

Approximation Algorithms for Correlated Knapsacks and Non-Martingale Bandits

Anupam Gupta*

Ravishankar Krishnaswamy*

Marco Molinaro[†]

R. Ravi[†]

Abstract

In the stochastic knapsack problem, we are given a knapsack of size B , and a set of jobs whose sizes and rewards are drawn from a known probability distribution. However, the only way to know the actual size and reward is to schedule the job—when it completes, we get to know these values. How should we schedule jobs to maximize the expected total reward? We know constant-factor approximations for this problem when we assume that rewards and sizes are independent random variables, and that we cannot prematurely cancel jobs after we schedule them. What can we say when either or both of these assumptions are changed?

The stochastic knapsack problem is of interest in its own right, but techniques developed for it are applicable to other stochastic packing problems. Indeed, ideas for this problem have been useful for budgeted learning problems, where one is given several arms which evolve in a specified stochastic fashion with each pull, and the goal is to pull the arms a total of B times to maximize the reward obtained. Much recent work on this problem focus on the case when the evolution of the arms follows a martingale, i.e., when the expected reward from the future is the same as the reward at the current state. What can we say when the rewards do not form a martingale?

In this paper, we give constant-factor approximation algorithms for the stochastic knapsack problem with correlations and/or cancellations, and also for budgeted learning problems where the martingale condition is not satisfied, using similar ideas. Indeed, we can show that previously proposed linear programming relaxations for these problems have large integrality gaps. We propose new time-indexed LP relaxations; using a decomposition and “gap-filling” approach, we convert these fractional solutions to distributions over strategies, and then use the LP values and the time ordering information from these strategies to devise a randomized adaptive scheduling algorithm. We hope our LP formulation and decomposition methods may provide a new way to address other correlated bandit problems with more general contexts.

*Department of Computer Science, Carnegie Mellon University, Pittsburgh PA 15213.

[†]Tepper School of Business, Carnegie Mellon University, Pittsburgh PA 15213.

1 Introduction

Stochastic packing problems seem to be conceptually harder than their deterministic counterparts—imagine a situation where some rounding algorithm outputs a solution in which the budget constraint has been exceeded by a constant factor. For deterministic packing problems (with a single constraint), one can now simply pick the most profitable subset of the items which meets the packing constraint; this would give us a profit within a constant of the optimal value. The deterministic packing problems not well understood are those with multiple (potentially conflicting) packing constraints.

However, for the stochastic problems, even a single packing constraint is not simple to handle. Even though they arise in diverse situations, the first study from an approximations perspective was in an important paper of Dean et al. [DGV08] (see also [DGV05, Dea05]). They defined the stochastic knapsack problem, where each job has a random size and a random reward, and the goal is to give an adaptive strategy for irrevocably picking jobs in order to maximize the expected value of those fitting into a knapsack with size B —they gave an LP relaxation and rounding algorithm, which produced *non-adaptive* solutions whose performance was surprisingly within a constant-factor of the best *adaptive* ones (resulting in a constant adaptivity gap, a notion they also introduced). However, the results required that (a) the random rewards and sizes for items were independent of each other, and (b) once a job was placed, it could not be prematurely canceled—it is easy to see that these assumptions change the nature of the problem significantly.

The study of the stochastic knapsack problem was very influential—in particular, the ideas here were used to obtain approximation algorithms for *budgeted learning problems* studied by Guha and Munagala [GM07b, GM07a, GM09] and Goel et al. [GKN09], among others. They considered problems in the multi-armed bandit setting with k arms, each arm evolving according to an underlying state machine with probabilistic transitions when pulled. Given a budget B , the goal is to pull arms up to B times to maximize the reward—payoffs are associated with states, and the reward is some function of payoffs of the states seen during the evolution of the algorithm. (E.g., it could be the sum of the payoffs of all states seen, or the reward of the best final state, etc.) The above papers gave $O(1)$ -approximations, index-based policies and adaptivity gaps for several budgeted learning problems. However, these results all required the assumption that the rewards satisfied a *martingale property*, namely, if an arm is some state u , one pull of this arm would bring an expected payoff equal to the payoff of state u itself—the motivation for such an assumption comes from the fact that the different arms are assumed to be associated with a fixed (but unknown) reward, but we only begin with a prior distribution of possible rewards. Then, the expected reward from the next pull of the arm, *conditioned* on the previous pulls, forms a Doob martingale.

However, there are natural instances where the martingale property need not hold. For instance, the evolution of the prior could not just depend on the observations made but on external factors (such as time) as well. Or, in a marketing application, the evolution of a customer’s state may require repeated “pulls” (or marketing actions) before the customer transitions to a high reward state and makes a purchase, while the intermediate states may not yield any reward. These lead us to consider the following problem: there are a collection of n arms, each characterized by an arbitrary (known) Markov chain, and there are rewards associated with the different states. When we play an arm, it makes a state transition according to the associated Markov chain, and fetches the corresponding reward of the new state. What should our strategy be in order to maximize the expected total reward we can accrue by making at most B pulls in total?

1.1 Results

Our main results are the following: We give the first constant-factor approximations for the general version of the stochastic knapsack problem where rewards could be correlated with the sizes. Our techniques are general and also apply to the setting when jobs could be canceled arbitrarily. We then extend those ideas to give the first constant-factor approximation algorithms for a class of budgeted learning problems with Markovian transitions where the martingale property is not satisfied. We summarize these in Table 1.

Problem	Restrictions	Paper
Stochastic Knapsack	Fixed Rewards, No Cancellation	[DGV05]
	Correlated Rewards, No Cancellation	Section 2
	Correlated Rewards, Cancellation	Section 3
Multi-Armed Bandits	Martingale Assumption	[GM07b]
	No Martingale Assumption	Section 4

Table 1: Summary of Results

1.2 Why Previous Ideas Don’t Extend, and Our Techniques

One reason why stochastic packing problems are more difficult than their deterministic counterparts is that, unlike in the deterministic setting, here we cannot simply take a solution with expected reward R^* that packs into a knapsack of size $2B$ and convert it (by picking a subset of the items) into a solution which obtains a constant fraction of the reward R^* whilst packing into a knapsack of size B . In fact, there are examples where a budget of $2B$ can fetch much more reward than what a budget of size B can (see Appendix A.2). Another distinction from deterministic problems is that allowing cancellations can drastically increase the value of the solution (see Appendix A.1). The model used in previous works on stochastic knapsack and on budgeted learning circumvented both issues—in contrast, our model forces us to address them.

Stochastic Knapsack: Dean et al. [DGV08, Dea05] assume that the reward/profit of an item is independent of its stochastic size. Moreover, their model does not consider the possibility of canceling jobs in the middle. These assumptions simplify the structure of the decision tree and make it possible to formulate a (deterministic) knapsack-style LP, and round it. However, as shown in Appendix A, their LP relaxation performs poorly when either correlation or cancellation is allowed. This is the first issue we need to address.

Budgeted Learning: Obtaining approximations for budgeted learning problems is a more complicated task, since cancellations may be inherent in the problem formulation, i.e., any strategy would stop playing a particular arm and switch to another, and the rewards by playing any arm are naturally correlated with the (current) state and hence the number of previous pulls made on the item/arm. The first issue is often tackled by using more elaborate LPs with a flow-like structure that compute a probability distribution over the different times at which the LP stops playing an arm (e.g., [GM07a]), but the latter issue is less understood. Indeed, several papers on this topic present strategies that fetch an expected reward which is a constant-factor of an optimal solution’s reward, but which may violate the budget by a constant factor. In order to obtain an approximate solution without violating the budget, they critically make use of the *martingale property*—with this assumption at hand, they can truncate the last arm played to fit the budget without incurring any loss in expected reward. However, such an idea fails when the martingale property is not satisfied, and these LPs now have large integrality gaps (see Appendix A.2).

At a high level, a major drawback with previous LP relaxations for both problems is that the constraints are *local* for each arm/job, i.e., they track the probability distribution over how long each item/arm is processed (either till completion or cancellation), and there is an additional global constraint binding the total number of pulls/total size across items. This results in two different issues. For the (correlated) stochastic knapsack problem, these LPs do not capture the case when all the items have high contention, since they want to play early in order to collect profit. And for the general multi-armed bandit problem, we show that no local LP can be good since such LPs do not capture the notion of *preempting* an arm, namely switching from one arm to another, and possibly returning to the original arm later later. Indeed, we show cases when any near-optimal strategy must switch between different arms (see Appendix A.3)—this is a major difference from previous work with the martingale property where there exist near-optimal strategies that never return to any arm [GM09, Lemma 2.1]. At a high level, the lack of the martingale property means our algorithm needs to make adaptive decisions, where each move is a function of the previous outcomes; in particular this may involve revisiting a particular arm several times, with interruptions in the middle.

We resolve these issues in the following manner: incorporating cancellations into stochastic knapsack can be handled by just adapting the flow-like LPs from the multi-armed bandits case. To resolve the problems of contention and preemption, we formulate a *global time-indexed* relaxation that forces the LP solution to commit each job to begin at a time, and places constraints on the maximum expected reward that can be obtained if the algorithm begins an item a particular time. Furthermore, the time-indexing also enables our rounding scheme to extract information about when to preempt an arm and when to re-visit it based on the LP solution; in fact, these decisions will possibly be different for different (random) outcomes of any pull, but the LP encodes the information for each possibility. We believe that our rounding approach may be of interest in other applications in Stochastic optimization problems.

Another important version of budgeted learning is when we are allowed to make up to B plays as usual but now we can “exploit” at most K times: reward is only fetched when an arm is exploited and again depends on its current state. There is a further constraint that once an arm is exploited, it must then be discarded. The LP-based approach here can be easily extended to that case as well.

1.3 Roadmap

We begin in Section 2 by presenting a constant-factor approximation algorithm for the stochastic knapsack problem (Stock) when rewards could be correlated with the sizes, but decisions are irrevocable, i.e., job cancellations are not allowed. Then, we build on these ideas in Section 3, and present our results for the (correlated) stochastic knapsack problem, where job cancellation is allowed.

In Section 4, we move on to the more general class of multi-armed bandit (MAB) problems. For clarity in exposition, we present our algorithm for MAB, assuming that the transition graph for each arm is an *arborescence* (i.e., a directed tree), and then generalize it to arbitrary transition graphs in Section 5.

We remark that while our LP-based approach for the budgeted learning problem implies approximation algorithms for the stochastic knapsack problem as well, the knapsack problem provides a gentler introduction to the issues—it motivates and gives insight into our techniques for MAB. Similarly, it is easier to understand our techniques for the MAB problem when the transition graph of each arm’s Markov chain is a tree. Several illustrative examples are presented in Appendix A, e.g., illustrating why we need adaptive strategies for the non-martingale MAB problems, and why some natural ideas do not work. Finally, the extension of our algorithm for MAB for the case when rewards are available only when the arms are explicitly exploited with budgets on both the exploration and exploitation pulls appear in Appendix F. Note that this algorithm strictly generalizes the previous work on budgeted learning for MAB with the martingale property [GM07a].

1.4 Related Work

Stochastic scheduling problems have been long studied since the 1960s (e.g., [BL97, Pin95]); however, there are fewer papers on approximation algorithms for such problems. Kleinberg et al. [KRT00], and Goel and Indyk [GI99] consider stochastic knapsack problems with chance constraints: find the max-profit set which will overflow the knapsack with probability at most p . However, their results hold for deterministic profits and specific size distributions. Approximation algorithms for minimizing average completion times with arbitrary job-size distributions was studied by [MSU99, SU01]. The work most relevant to us is that of Dean, Goemans and Vondrák [DGV08, DGV05, Dea05] on stochastic knapsack and packing; apart from algorithms (for independent rewards and sizes), they show the problem to be PSPACE-hard when correlations are allowed. [CR06] study stochastic flow problems. Recent work of Bhalgat et al. [BGK11] presents a PTAS but violate the capacity by a factor $(1 + \epsilon)$; they also get better constant-factor approximations without violations.

The general area of learning with costs is a rich and diverse one (see, e.g., [Ber05, Git89]). Approximation algorithms start with the work of Guha and Munagala [GM07a], who gave LP-rounding algorithms for some problems. Further papers by these authors [GMS07, GM09] and by Goel et al. [GKN09] give improvements, relate LP-based techniques and index-based policies and also give new index policies. (See also [GGM06, GM07b].)

[GM09] considers switching costs, [GMP11] allows pulling many arms simultaneously, or when there is delayed feedback. All these papers assume the martingale condition.

2 The Correlated Stochastic Knapsack without Cancellation

We begin by considering the stochastic knapsack problem (StocK), when the job rewards may be correlated with its size. This generalizes the problem studied by Dean et al. [DGV05] who assume that the rewards are independent of the size of the job. We first explain why the LP of [DGV05] has a large integrality gap for our problem; this will naturally motivate our time-indexed formulation. We then present a simple randomized rounding algorithm which produces a non-adaptive strategy and show that it is an $O(1)$ -approximation.

2.1 Problem Definitions and Notation

We are given a knapsack of total budget B and a collection of n stochastic items. For any item $i \in [1, n]$, we are given a probability distribution over (size, reward) pairs specified as follows: for each integer value of $t \in [1, B]$, the tuple $(\pi_{i,t}, R_{i,t})$ denotes the probability $\pi_{i,t}$ that item i has a size t , and the corresponding reward is $R_{i,t}$. Note that the reward for a job is now correlated to its size; however, these quantities for two different jobs are still independent of each other.

An algorithm to *adaptively* process these items can do the following actions at the end of each timestep; (i) an item may complete at a certain size, giving us the corresponding reward, and the algorithm may choose a new item to start processing, or (ii) the knapsack becomes full, at which point the algorithm cannot process any more items, and any currently running job does not accrue any reward. The objective function is to maximize the total expected reward obtained from all completed items. Notice that we do not allow the algorithm to cancel an item before it completes. We relax this requirement in Section 3.

2.2 LP Relaxation

The LP relaxation in [DGV05] was (essentially) a knapsack LP where the sizes of items are replaced by the expected sizes, and the rewards are replaced by the expected rewards. While this was sufficient when an item's reward is fixed (or chosen randomly but independent of its size), we give an example in Appendix A.2 where such an LP (and in fact, the class of more general LPs used for approximating MAB problems) would have a large integrality gap. As mentioned in Section 1.2, the reason why local LPs don't work is that there could be high contention for being scheduled early (i.e., there could be a large number of items which all fetch reward if they instantiate to a large size, but these events occur with low probability). In order to capture this contention, we write a global time-indexed LP relaxation.

The variable $x_{i,t} \in [0, 1]$ indicates that item i is scheduled at (global) time t ; S_i denotes the random variable for the size of item i , and $\text{ER}_{i,t} = \sum_{s \leq B-t} \pi_{i,s} R'_{i,s}$ captures the expected reward that can be obtained from item i if it begins at time t ; (no reward is obtained for sizes that cannot fit the (remaining) budget.)

$$\begin{aligned} \max \sum_{i,t} \text{ER}_{i,t} \cdot x_{i,t} & & (\text{LP}_{\text{NoCancel}}) \\ \sum_t x_{i,t} & \leq 1 & \forall i \end{aligned} \tag{2.1}$$

$$\sum_{i,t' \leq t} x_{i,t'} \cdot \mathbb{E}[\min(S_i, t)] \leq 2t \quad \forall t \in [B] \tag{2.2}$$

$$x_{i,t} \in [0, 1] \quad \forall t \in [B], \forall i \tag{2.3}$$

While the size of the above LP (and the running time of the rounding algorithm below) polynomially depend on B , i.e., pseudo-polynomial, it is possible to write a compact (approximate) LP and then round it; details on the polynomial time implementation appear in Appendix B.2.

Notice the constraints involving the *truncated random variables* in equation (2.2): these are crucial for showing the correctness of the rounding algorithm **StocK-NoCancel**. Furthermore, the ideas used here will appear sub-

sequently in the MAB algorithm later; for MAB, even though we can't explicitly enforce such a constraint in the LP, we will end up inferring a similar family of inequalities from a near-optimal LP solution.

Lemma 2.1 *The above relaxation is valid for the Stock problem when cancellations are not permitted, and has objective value $\text{LP}_{\text{Opt}} \geq \text{Opt}$, where Opt is the expected profit of an optimal adaptive policy.*

Proof. Consider an optimal policy Opt and let $x_{i,t}^*$ denote the probability that item i is scheduled at time t . We first show that $\{x^*\}$ is a feasible solution for the LP relaxation $\text{LP}_{\text{NoCancel}}$. It is easy to see that constraints (2.1) and (2.3) are satisfied. To prove that (2.2) are also satisfied, consider some $t \in [B]$ and some run (over random choices of item sizes) of the optimal policy. Let $\mathbf{1}_{i,t'}^{\text{sched}}$ be indicator variable that item i is scheduled at time t' and let $\mathbf{1}_{i,s}^{\text{size}}$ be the indicator variable for whether the size of item i is s . Also, let L_t be the random variable indicating the last item scheduled at or before time t . Notice that L_t is the only item scheduled before or at time t whose execution may go over time t . Therefore, we get that

$$\sum_{i \neq L_t} \sum_{t' \leq t} \sum_{s \leq B} \mathbf{1}_{i,t'}^{\text{sched}} \cdot \mathbf{1}_{i,s}^{\text{size}} \cdot s \leq t.$$

Including L_t in the summation and truncating the sizes by t , we immediately obtain

$$\sum_i \sum_{t' \leq t} \sum_s \mathbf{1}_{i,t'}^{\text{sched}} \cdot \mathbf{1}_{i,s}^{\text{size}} \cdot \min(s, t) \leq 2t.$$

Now, taking expectation (over all of Opt 's sample paths) on both sides and using linearity of expectation we have

$$\sum_i \sum_{t' \leq t} \sum_s \mathbb{E} \left[\mathbf{1}_{i,t'}^{\text{sched}} \cdot \mathbf{1}_{i,s}^{\text{size}} \right] \cdot \min(s, t) \leq 2t.$$

However, because Opt decides whether to schedule an item before observing the size it instantiates to, we have that $\mathbf{1}_{i,t'}^{\text{sched}}$ and $\mathbf{1}_{i,s}^{\text{size}}$ are independent random variables; hence, the LHS above can be re-written as

$$\begin{aligned} & \sum_i \sum_{t' \leq t} \sum_s \Pr[\mathbf{1}_{i,t'}^{\text{sched}} = 1 \wedge \mathbf{1}_{i,s}^{\text{size}} = 1] \min(s, t) \\ &= \sum_i \sum_{t' \leq t} \Pr[\mathbf{1}_{i,t'}^{\text{sched}} = 1] \sum_s \Pr[\mathbf{1}_{i,s}^{\text{size}} = 1] \min(s, t) \\ &= \sum_i \sum_{t' \leq t} x_{i,t'}^* \cdot \mathbb{E}[\min(S_i, t)] \end{aligned}$$

Hence constraints (2.2) are satisfied. Now we argue that the expected reward of Opt is equal to the value of the solution x^* . Let O_i be the random variable denoting the reward obtained by Opt from item i . Again, due to the independence between Opt scheduling an item and the size it instantiates to, we get that the expected reward that Opt gets from executing item i at time t is

$$\mathbb{E}[O_i | \mathbf{1}_{i,t}^{\text{sched}} = 1] = \sum_{s \leq B-t} \pi_{i,s} R_{i,s} = \text{ER}_{i,t}.$$

Thus the expected reward from item i is obtained by considering all possible starting times for i :

$$\mathbb{E}[O_i] = \sum_t \Pr[\mathbf{1}_{i,t}^{\text{sched}} = 1] \cdot \mathbb{E}[O_i | \mathbf{1}_{i,t}^{\text{sched}} = 1] = \sum_t \text{ER}_{i,t} \cdot x_{i,t}^*.$$

This shows that $\text{LP}_{\text{NoCancel}}$ is a valid relaxation for our problem and completes the proof of the lemma. \blacksquare

Algorithm 2.1 Algorithm Stock-NoCancel

- 1: for each item i , **assign** a random start-time $D_i = t$ with probability $\frac{x_{i,t}^*}{4}$; with probability $1 - \sum_t \frac{x_{i,t}^*}{4}$, completely ignore item i ($D_i = \infty$ in this case).
 - 2: **for** j from 1 to n **do**
 - 3: Consider the item i which has the j th smallest deadline (and $D_i \neq \infty$)
 - 4: **if** the items added so far to the knapsack occupy at most D_i space **then**
 - 5: add i to the knapsack.
-

We are now ready to present our rounding algorithm **Stock-NoCancel** (Algorithm 2.1). It is a simple randomized rounding procedure which (i) picks the start time of each item according to the corresponding distribution in the optimal LP solution, and (ii) plays the items in order of the (random) start times. To ensure that the budget is not violated, we also drop each item independently with some constant probability.

Notice that the strategy obtained by the rounding procedure obtains reward from all items which are not dropped and which do not fail (i.e. they can start being scheduled before the sampled start-time D_i in Step 1); we now bound the failure probability.

Lemma 2.2 For every i , $\Pr(i \text{ fails} \mid D_i = t) \leq 1/2$.

Proof. Consider an item i and time $t \neq \infty$ and condition on the event that $D_i = t$. Let us consider the execution of the algorithm when it tries to add item i to the knapsack in steps 3-5. Now, let Z be a random variable denoting *how much of the interval* $[0, t]$ of the knapsack is occupied by previously scheduling items, at the time when i is considered for addition; since i does not fail when $Z < t$, it suffices to prove that $\Pr(Z \geq t) \leq 1/2$.

For some item $j \neq i$, let $\mathbf{1}_{D_j \leq t}$ be the indicator variable that $D_j \leq t$; notice that by the order in which algorithm **Stock-NoCancel** adds items into the knapsack, it is also the indicator that j was considered before i . In addition, let $\mathbf{1}_{j,s}^{\text{size}}$ be the indicator variable that $S_j = s$. Now, if Z_j denotes the total amount of the interval $[0, t]$ that j occupies, we have

$$Z_j \leq \mathbf{1}_{D_j \leq t} \sum_s \mathbf{1}_{j,s}^{\text{size}} \min(s, t).$$

Now, using the independence of $\mathbf{1}_{D_j \leq t}$ and $\mathbf{1}_{j,s}^{\text{size}}$, we have

$$\mathbb{E}[Z_j] \leq \mathbb{E}[\mathbf{1}_{D_j \leq t}] \cdot \mathbb{E}[\min(S_j, t)] = \frac{1}{4} \sum_{t' \leq t} x_{j,t'}^* \cdot \mathbb{E}[\min(S_j, t)] \quad (2.4)$$

Since $Z = \sum_j Z_j$, we can use linearity of expectation and the fact that $\{x^*\}$ satisfies LP constraint (2.2) to get

$$\mathbb{E}[Z] \leq \frac{1}{4} \sum_j \sum_{t' \leq t} x_{j,t'}^* \cdot \mathbb{E}[\min(S_j, t)] \leq \frac{t}{2}.$$

To conclude the proof of the lemma, we apply Markov's inequality to obtain $\Pr(Z \geq t) \leq 1/2$. ■

To complete the analysis, we use the fact that any item chooses a random start time $D_i = t$ with probability $x_{i,t}^*/4$, and conditioned on this event, it is added to the knapsack with probability at least $1/2$ from Lemma 2.2; in this case, we get an expected reward of at least $\text{ER}_{i,t}$. The theorem below (formally proved in Appendix B.1) then follows by linearity of expectations.

Theorem 2.3 The expected reward of our randomized algorithm is at least $\frac{1}{8}$ of LPOpt .

3 Stochastic Knapsack with Correlated Rewards and Cancellations

In this section, we present our algorithm for stochastic knapsack (**Stock**) where we allow correlations between rewards and sizes, and also allow cancellation of jobs. The example in Appendix A.1 shows that there can be an

arbitrarily large gap in the expected profit between strategies that can cancel jobs and those that can't. Hence we need to write new LPs to capture the benefit of cancellation, which we do in the following manner.

Consider any job j : we can create two jobs from it, the “early” version of the job, where we discard profits from any instantiation where the size of the job is more than $B/2$, and the “late” version of the job where we discard profits from instantiations of size at most $B/2$. Hence, we can get at least half the optimal value by flipping a fair coin and either collecting rewards from either the early or late versions of jobs, based on the outcome. In the next section, we show how to obtain a constant factor approximation for the first kind. For the second kind, we argue that cancellations don't help; we can then reduce it to Stock without cancellations (considered in Section 2).

3.1 Case I: Jobs with Early Rewards

We begin with the setting in which only small-size instantiations of items may fetch reward, i.e., the rewards $R_{i,t}$ of every item i are assumed to be 0 for $t > B/2$. In the following LP relaxation LP_S , $v_{i,t} \in [0, 1]$ tries to capture the probability with which Opt will process item i for *at least* t timesteps¹, $s_{i,t} \in [0, 1]$ is the probability that Opt stops processing item i *exactly* at t timesteps. The time-indexed formulation causes the algorithm to have running times of $\text{poly}(B)$ —however, it is easy to write compact (approximate) LPs and then round them; we describe the necessary changes to obtain an algorithm with running time $\text{poly}(n, \log B)$ in Appendix C.2.

$$\max \sum_{1 \leq t \leq B/2} \sum_{1 \leq i \leq n} v_{i,t} \cdot R_{i,t} \frac{\pi_{i,t}}{\sum_{t' \geq t} \pi_{i,t'}} \quad (\text{LP}_S)$$

$$v_{i,t} = s_{i,t} + v_{i,t+1} \quad \forall t \in [0, B], i \in [n] \quad (3.5)$$

$$s_{i,t} \geq \frac{\pi_{i,t}}{\sum_{t' \geq t} \pi_{i,t'}} \cdot v_{i,t} \quad \forall t \in [0, B], i \in [n] \quad (3.6)$$

$$\sum_{i \in [n]} \sum_{t \in [0, B]} t \cdot s_{i,t} \leq B \quad (3.7)$$

$$v_{i,0} = 1 \quad \forall i \quad (3.8)$$

$$v_{i,t}, s_{i,t} \in [0, 1] \quad \forall t \in [0, B], i \in [n] \quad (3.9)$$

Theorem 3.1 *The linear program (LP_S) is a valid relaxation for the Stock problem, and hence the optimal value LPOpt of the LP is at least the total expected reward Opt of an optimal solution.*

Proof. Consider an optimal solution Opt and let $v_{i,t}^*$ and $s_{i,t}^*$ denote the probability that Opt processes item i for at least t timesteps, and the probability that Opt stops processing item i at exactly t timesteps. We will now show that all the constraints of LP_S are satisfied one by one.

To this end, let R_i denote the random variable (over different executions of Opt) for the amount of processing done on job i . Notice that $\Pr[R_i \geq t] = \Pr[R_i \geq (t+1)] + \Pr[R_i = t]$. But now, by definition we have $\Pr[R_i \geq t] = v_{i,t}^*$ and $\Pr[R_i = t] = s_{i,t}^*$. This shows that $\{v^*, s^*\}$ satisfies these constraints.

For the next constraint, observe that conditioned on Opt running an item i for at least t time steps, the probability of item i stopping due to its size having instantiated to exactly equal to t is $\pi_{i,t} / \sum_{t' \geq t} \pi_{i,t'}$, i.e., $\Pr[R_i = t \mid R_i \geq t] \geq \pi_{i,t} / \sum_{t' \geq t} \pi_{i,t'}$. This shows that $\{v^*, s^*\}$ satisfies constraints (3.6).

Finally, to see why constraint (3.7) is satisfied, consider any particular run of the optimal algorithm and let $\mathbf{1}_{i,t}^{\text{stop}}$ denote the indicator random variable of the event $R_i = t$. Then we have

$$\sum_i \sum_t \mathbf{1}_{i,t}^{\text{stop}} \cdot t \leq B$$

Now, taking expectation over all runs of Opt and using linearity of expectation and the fact that $\mathbb{E}[\mathbf{1}_{i,t}^{\text{stop}}] = s_{i,t}^*$, we get constraint (3.7). As for the objective function, we again consider a particular run of the optimal algorithm and let $\mathbf{1}_{i,t}^{\text{proc}}$ now denote the indicator random variable for the event $(R_i \geq t)$, and $\mathbf{1}_{i,t}^{\text{size}}$ denote the indicator

¹In the following two sections, we use the word timestep to refer to processing one unit of some item.

variable for whether the size of item i is instantiated to exactly t in this run. Then we have the total reward collected by Opt in this run to be exactly

$$\sum_i \sum_t \mathbf{1}_{i,t}^{proc} \cdot \mathbf{1}_{i,t}^{size} \cdot R_{i,t}$$

Now, we simply take the expectation of the above random variable over all runs of Opt, and then use the following fact about $\mathbb{E}[\mathbf{1}_{i,t}^{proc} \mathbf{1}_{i,t}^{size}]$:

$$\begin{aligned} \mathbb{E}[\mathbf{1}_{i,t}^{proc} \mathbf{1}_{i,t}^{size}] &= \Pr[\mathbf{1}_{i,t}^{proc} = 1 \wedge \mathbf{1}_{i,t}^{size} = 1] \\ &= \Pr[\mathbf{1}_{i,t}^{proc} = 1] \Pr[\mathbf{1}_{i,t}^{size} = 1 \mid \mathbf{1}_{i,t}^{proc} = 1] \\ &= v_{i,t}^* \frac{\pi_{i,t}}{\sum_{t' \geq t} \pi_{i,t'}} \end{aligned}$$

We thus get that the expected reward collected by Opt is exactly equal to the objective function value of the LP formulation for the solution (v^*, s^*) . \blacksquare

Our rounding algorithm is very natural, and simply tries to mimic the probability distribution (over when to stop each item) as suggested by the optimal LP solution. To this end, let (v^*, s^*) denote an optimal fractional solution. The reason why we introduce some damping (in the selection probabilities) up-front is to make sure that we could appeal to Markov's inequality and ensure that the knapsack does not get violated with good probability.

Algorithm 3.1 Algorithm Stock-Small

- 1: **for** each item i **do**
 - 2: **ignore** i with probability $1 - 1/4$ (i.e., do not schedule it at all).
 - 3: **for** $0 \leq t \leq B/2$ **do**
 - 4: **cancel** item i at this step with probability $\frac{s_{i,t}^*}{v_{i,t}^*} - \frac{\pi_{i,t}}{\sum_{t' \geq t} \pi_{i,t'}}$ and **continue** to next item.
 - 5: process item i for its $(t + 1)^{st}$ timestep.
 - 6: **if** item i terminates after being processed for exactly $(t + 1)$ timesteps **then**
 - 7: **collect** a reward of $R_{i,t+1}$ from this item; **continue** onto next item;
-

Notice that while we let the algorithm proceed even if its budget is violated, we will collect reward only from items that complete before time B . This simplifies the analysis a fair bit, both here and for the MAB algorithm. In Lemma 3.2 below (proof in Appendix C), we show that for any item that is not dropped in step 2, its probability distribution over stopping times is identical to the optimal LP solution s^* . We then use this to argue that the expected reward of our algorithm is $\Omega(1)$ LPOpt.

Lemma 3.2 *Consider item i that was not dropped in step 2, Then, for any timestep $t \geq 0$, the following hold:*

- (i) *The probability (including cancellation & completion) of stopping at timestep t for item i is $s_{i,t}^*$.*
- (ii) *The probability that item i gets processed for its $(t + 1)^{st}$ timestep is exactly $v_{i,t+1}^*$*
- (iii) *If item i has been processed for $(t + 1)$ timesteps, the probability of completing successfully at timestep $(t + 1)$ is $\pi_{i,t+1} / \sum_{t' \geq t+1} \pi_{i,t'}$*

Theorem 3.3 *The expected reward of our randomized algorithm is at least $\frac{1}{8}$ of LPOpt.*

Proof. Consider any item i . In the worst case, we process i after all other items. Then the total expected size occupied thus far is at most $\sum_{i' \neq i} \mathbf{1}_{i'}^{keep} \sum_{t \geq 0} t \cdot s_{i',t}^*$, where $\mathbf{1}_{i'}^{keep}$ is the indicator random variable denoting whether item i' is not dropped in step 2. Here we have used Lemma 3.2 to argue that if an item i' is selected, its stopping-time distribution follows $s_{i',t}^*$. Taking expectation over the randomness in step 2, the expected space occupied by other jobs is at most $\sum_{i' \neq i} \frac{1}{3} \sum_{t \geq 0} t \cdot s_{i',t}^* \leq \frac{B}{4}$. Markov's inequality implies that this is at most $B/2$ with probability at least $1/2$. In this case, if item i is started (which happens w.p. $1/4$), it runs without violating the knapsack, with expected reward $\sum_{t \geq 1} v_{i,t}^* \cdot \pi_{i,t} / (\sum_{t' \geq t} \pi_{i,t'})$; the total expected reward is then at least $\sum_i \frac{1}{8} \sum_t v_{i,t}^* \pi_{i,t} / (\sum_{t' \geq t} \pi_{i,t'}) \geq \frac{\text{LPOpt}}{8}$. \blacksquare

3.2 Case II: Jobs with Late Rewards

Now we handle instances in which only large-size instantiations of items may fetch reward, i.e., the rewards $R_{i,t}$ of every item i are assumed to be 0 for $t \leq B/2$. For such instances, we now argue that *cancellation is not helpful*. As a consequence, we can use the results of Section 2 and obtain a constant-factor approximation algorithm!

To see why, intuitively, as an algorithm processes a job for its t^{th} timestep for $t < B/2$, it gets no more information about the reward than when starting (since all rewards are at large sizes). Furthermore, there is no benefit of canceling a job once it has run for at least $B/2$ timesteps – we can't get any reward by starting some other item.

More formally, consider a (deterministic) strategy S which in some state makes the decision of scheduling item i and halting its execution if it takes more than t timesteps. First suppose that $t \leq B/2$; since this job does will not be able to reach size larger than $B/2$, no reward will be accrued from it and hence we can change this strategy by skipping the scheduling of i without altering its total reward. Now consider the case where $t > B/2$. Consider the strategy S' which behaves as S except that it does not preempt i in this state but lets i run to completion. We claim that S' obtains at least as much expected reward as S . First, whenever item i has size at most t then S and S' obtain the same reward. Now suppose that we are in a scenario where i reached size $t > B/2$. Then item i is halted and S cannot obtain any other reward in the future, since no item that can fetch any reward would complete before the budget runs out; in the same situation, strategy S' obtains non-negative rewards. Using this argument we can eliminate all the cancellations of a strategy without decreasing its expected reward.

Lemma 3.4 *There is an optimal solution in this case which does not cancel.*

As mentioned earlier, we can now appeal to the results of Section 2 and obtain a constant-factor approximation for the large-size instances. Now we can combine the algorithms that handle the two different scenarios (or choose one at random and run it), and get a constant fraction of the expected reward that an optimal policy fetches.

4 Multi-Armed Bandits

We now turn our attention to the more general Multi-Armed Bandits problem (MAB). In this framework, there are n arms: arm i has a collection of states denoted by \mathcal{S}_i , a starting state $\rho_i \in \mathcal{S}_i$; Without loss of generality, we assume that $\mathcal{S}_i \cap \mathcal{S}_j = \emptyset$ for $i \neq j$. Each arm also has a *transition graph* T_i , which is given as a polynomial-size (weighted) directed tree rooted at ρ_i ; we will relax the tree assumption later. If there is an edge $u \rightarrow v$ in T_i , then the edge weight $p_{u,v}$ denotes the probability of making a transition from u to v if we play arm i when its current state is node u ; hence $\sum_{v:(u,v) \in T_i} p_{u,v} = 1$. Each time we play an arm, we get a reward whose value depends on the state from which the arm is played. Let us denote the reward at a state u by r_u . Recall that the martingale property on rewards requires that $\sum_{v:(u,v) \in T_i} p_{u,v} r_v = r_u$ for all states u .

Problem Definition. For a concrete example, we consider the following budgeted learning problem on *tree transition graphs*. Each of the arms starts at the start state $\rho_i \in \mathcal{S}_i$. We get a reward from each of the states we play, and the goal is to maximize the total expected reward, while not exceeding a pre-specified allowed number of plays B across all arms. The framework described below can handle other problems (like the explore/exploit kind) as well, and we discuss this in Appendix F.

Note that the Stochastic Knapsack problem considered in the previous section is a special case of this problem where each item corresponds to an arm, where the evolution of the states corresponds to the explored size for the item. Rewards are associated with each stopping size, which can be modeled by end states that can be reached from the states of the corresponding size with the probability of this transition being the probability of the item taking this size. Thus the resulting trees are paths of length up to the maximum size B with transitions to end states with reward for each item size. For example, the transition graph in Figure 4.1 corresponds to an item which instantiates to a size of 1 with probability $1/2$ (and fetches a reward R_1), takes size 3 with probability $1/4$ (with reward R_3), and size 4 with the remaining probability $1/4$ (reward is R_4). Notice that the reward on stopping at all intermediate nodes is 0 and such an instance therefore does not satisfy the martingale property.

Even though the rewards are obtained in this example on reaching a state rather than playing it, it is not hard to modify our methods for this version as well.




Figure 4.1: Reducing Stochastic Knapsack to MAB

Notation. The transition graph T_i for arm i is an out-arborescence defined on the states \mathcal{S}_i rooted at ρ_i . Let $\text{depth}(u)$ of a node $u \in \mathcal{S}_i$ be the depth of node u in tree T_i , where the root ρ_i has depth 0. The unique parent of node u in T_i is denoted by $\text{parent}(u)$. Let $\mathcal{S} = \cup_i \mathcal{S}_i$ denote the set of all states in the instance, and $\text{arm}(u)$ denote the arm to which state u belongs, i.e., the index i such that $u \in \mathcal{S}_i$. Finally, for $u \in \mathcal{S}_i$, we refer to the act of playing arm i when it is in state u as “playing state $u \in \mathcal{S}_i$ ”, or “playing state u ” if the arm is clear in context.

4.1 Global Time-indexed LP

In the following, the variable $z_{u,t} \in [0, 1]$ indicates that the algorithm plays state $u \in \mathcal{S}_i$ at time t . For state $u \in \mathcal{S}_i$ and time t , $w_{u,t} \in [0, 1]$ indicates that arm i *first enters* state u at time t : this happens if and only if the algorithm *played* $\text{parent}(u)$ at time $t - 1$ and the arm made a transition into state u .

$$\max \sum_{u,t} r_u \cdot z_{u,t} \tag{LP_{\text{mab}}}$$

$$w_{u,t} = z_{\text{parent}(u),t-1} \cdot p_{\text{parent}(u),u} \quad \forall t \in [2, B], u \in \mathcal{S} \setminus \cup_i \{\rho_i\} \tag{4.10}$$

$$\sum_{t' \leq t} w_{u,t'} \geq \sum_{t' \leq t} z_{u,t'} \quad \forall t \in [1, B], u \in \mathcal{S} \tag{4.11}$$

$$\sum_{u \in \mathcal{S}} z_{u,t} \leq 1 \quad \forall t \in [1, B] \tag{4.12}$$

$$w_{\rho_i,1} = 1 \quad \forall i \in [1, n] \tag{4.13}$$

Lemma 4.1 *The value of an optimal LP solution LP_{Opt} is at least Opt , the expected reward of an optimal adaptive strategy.*

Proof. We convention that Opt starts playing at time 1. Let $z_{u,t}^*$ denote the probability that Opt plays state u at time t , namely, the probability that arm $\text{arm}(u)$ is in state u at time t and is played at time t . Also let $w_{u,t}^*$ denote the probability that Opt “enters” state u at time t , and further let $w_{\rho_i,1}^* = 1$ for all i .

We first show that $\{z^*, w^*\}$ is a feasible solution for LP_{mab} and later argue that its LP objective is at least Opt . Consider constraint (4.10) for some $t \in [2, B]$ and $u \in \mathcal{S}$. The probability of entering state u at time t conditioned on Opt playing state $\text{parent}(u)$ at time $t - 1$ is $p_{\text{parent}(u),u}$. In addition, the probability of entering state u at time t conditioning on Opt not playing state $\text{parent}(u)$ at time $t - 1$ is zero. Since $z_{\text{parent}(u),t-1}^*$ is the probability that Opt plays state $\text{parent}(u)$ at time $t - 1$, we remove the conditioning to obtain $w_{u,t}^* = z_{\text{parent}(u),t-1}^* \cdot p_{\text{parent}(u),u}$.

Now consider constraint (4.11) for some $t \in [1, B]$ and $u \in \mathcal{S}$. For any outcome of the algorithm (denoted by a sample path σ), let $\mathbf{1}_{u',t'}^{\text{enter}}$ be the indicator variable that Opt enters state u' at time t' and let $\mathbf{1}_{u',t'}^{\text{play}}$ be the indicator variable that Opt plays state u' at time t' . Since T_i is acyclic, state u is played at most once in σ and is also entered at most once in σ . Moreover, whenever u is played before or at time t , it must be that u was also entered before or at time t , and hence $\sum_{t' \leq t} \mathbf{1}_{u,t'}^{\text{play}} \leq \sum_{t' \leq t} \mathbf{1}_{u,t'}^{\text{enter}}$. Taking expectation on both sides and using the fact that $\mathbb{E}[\mathbf{1}_{u,t'}^{\text{play}}] = z_{u,t'}^*$ and $\mathbb{E}[\mathbf{1}_{u,t'}^{\text{enter}}] = w_{u,t'}^*$, linearity of expectation gives $\sum_{t' \leq t} z_{u,t'}^* \leq \sum_{t' \leq t} w_{u,t'}^*$.

To see that constraints (4.12) are satisfied, notice that we can play at most one arm (or alternatively one state) in each time step, hence $\sum_{u \in \mathcal{S}} \mathbf{1}_{u,t}^{\text{play}} \leq 1$ holds for all $t \in [1, B]$; the claim then follows by taking expectation on both sides as in the previous paragraph. Finally, constraints (4.13) is satisfied by definition of the start states.

To conclude the proof of the lemma, it suffices to show that $\text{Opt} = \sum_{u,t} r_u \cdot z_{u,t}^*$. Since Opt obtains reward r_u whenever it plays state u , it follows that Opt 's reward is given by $\sum_{u,t} r_u \cdot \mathbf{1}_{u,t}^{\text{play}}$; by taking expectation we get $\sum_{u,t} r_u z_{u,t}^* = \text{Opt}$, and hence $\text{LPOpt} \geq \text{Opt}$. ■

4.2 The Rounding Algorithm

In order to best understand the motivation behind our rounding algorithm, it would be useful to go over the example which illustrates the necessity of preemption (repeatedly switching back and forth between the different arms) in Appendix A.3.

At a high level, the rounding algorithm proceeds as follows. In Phase I, given an optimal LP solution, we decompose the fractional solution for each arm into a convex² combination of integral “strategy forests” (which are depicted in Figure 4.2): each of these tells us at what times to play the arm, and in which states to abandon the arm. Now, if we sample a random strategy forest for each arm from this distribution, we may end up scheduling multiple arms to play at some of the timesteps, and hence we need to resolve these conflicts. A natural first approach might be to (i) sample a strategy forest for each arm, (ii) play these arms in a random order, and (iii) for any arm follow the decisions (about whether to abort or continue playing) as suggested by the sampled strategy forest. In essence, we are ignoring the times at which the sampled strategy forest has scheduled the plays of this arm and instead playing this arm continually until the sampled forest abandons it. While such a non-preemptive strategy works when the martingale property holds, the example in Appendix A.3 shows that preemption is unavoidable.

Another approach would be to try to play the sampled forests at their prescribed times; if multiple forests want to play at the same time slot, we round-robin over them. The expected number of plays in each timestep is 1, and the hope is that round-robin will not hurt us much. However, if some arm needs B contiguous steps to get to a state with high reward, and a single play of some other arm gets scheduled by bad luck in some timestep, we would end up getting nothing!

Guided by these bad examples, we try to use the continuity information in the sampled strategy forests—once we start playing some contiguous component (where the strategy forest plays the arm in every consecutive time step), we play it to the end of the component. The naïve implementation does not work, so we first alter the LP solution to get convex combinations of “nice” forests—loosely, these are forests where the strategy forest plays contiguously in almost all timesteps, or in at least half the timesteps. This alteration is done in Phase II, and then the actual rounding in Phase III, and the analysis appears in Section 4.2.3.

4.2.1 Phase I: Convex Decomposition

In this step, we decompose the fractional solution into a convex combination of “forest-like strategies” $\{\mathbb{T}(i, j)\}_{i,j}$, corresponding to the j^{th} forest for arm i . We first formally define what these forests look like: The j^{th} strategy forest $\mathbb{T}(i, j)$ for arm i is an assignment of values $\text{time}(i, j, u)$ and $\text{prob}(i, j, u)$ to each state $u \in \mathcal{S}_i$ such that:

- (i) For $u \in \mathcal{S}_i$ and $v = \text{parent}(u)$, it holds that $\text{time}(i, j, u) \geq 1 + \text{time}(i, j, v)$, and
- (ii) For $u \in \mathcal{S}_i$ and $v = \text{parent}(u)$, if $\text{time}(i, j, u) \neq \infty$ then $\text{prob}(i, j, u) = p_{v,u} \text{prob}(i, j, v)$; else if $\text{time}(i, j, u) = \infty$ then $\text{prob}(i, j, u) = 0$.

We call a triple (i, j, u) a *tree-node* of $\mathbb{T}(i, j)$. When i and j are understood from the context, we identify the tree-node (i, j, u) with the state u .

For any state u , the values $\text{time}(i, j, u)$ and $\text{prob}(i, j, u)$ denote the time at which the arm i is played at state u , and the probability with which the arm is played, according to the strategy forest $\mathbb{T}(i, j)$.³ The probability values are particularly simple: if $\text{time}(i, j, u) = \infty$ then this strategy does not play the arm at u , and hence the probability

²Strictly speaking, we do not get convex combinations that sum to one; our combinations sum to $\sum_t z_{\rho_i, t}$, the value the LP assigned to pick to play the root of the arm over all possible start times, which is at most one.

³When i and j are clear from the context, we will just refer to state u instead of the triple (i, j, u) .

is zero, else $\text{prob}(i, j, u)$ is equal to the probability of reaching u over the random transitions according to T_i if we play the root with probability $\text{prob}(i, j, \rho_i)$. Hence, we can compute $\text{prob}(i, j, u)$ just given $\text{prob}(i, j, \rho_i)$ and whether or not $\text{time}(i, j, u) = \infty$. Note that the time values are not necessarily consecutive, plotting these on the timeline and connecting a state to its parents only when they are in consecutive timesteps (as in Figure 4.2) gives us forests, hence the name.

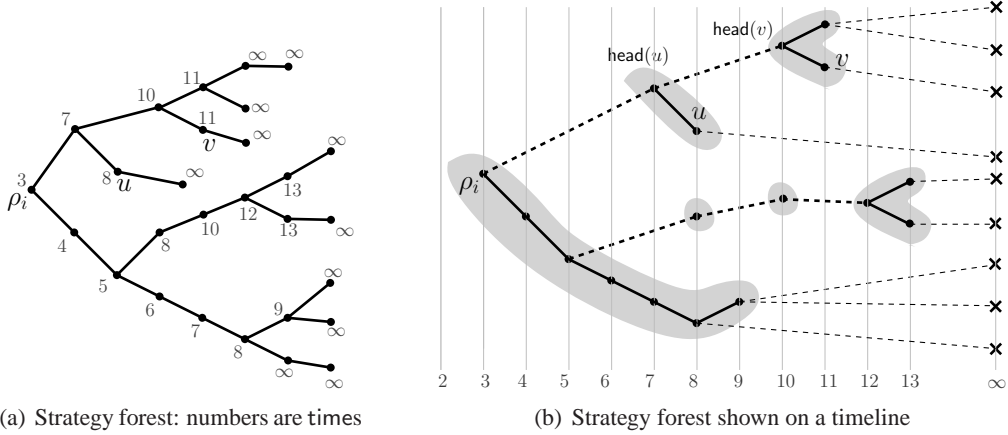


Figure 4.2: Strategy forests and how to visualize them: grey blobs are connected components.

The algorithm to construct such a decomposition proceeds in rounds for each arm i ; in a particular round, it “peels” off such a strategy as described above, and ensures that the residual fractional solution continues to satisfy the LP constraints, guaranteeing that we can repeat this process, which is similar to (but slightly more involved than) performing flow-decompositions. The decomposition lemma is proved in Appendix D.1:

Lemma 4.2 *Given a solution to (LP_{mab}) , there exists a collection of at most $nB|\mathcal{S}|$ strategy forests $\{\mathbb{T}(i, j)\}$ such that $z_{u,t} = \sum_{j:\text{time}(i,j,u)=t} \text{prob}(i, j, u)$.⁴ Hence, $\sum_{(i,j,u):\text{time}(i,j,u)=t} \text{prob}(i, j, u) \leq 1$ for all t .*

For any $\mathbb{T}(i, j)$, these prob values satisfy a “preflow” condition: the in-flow at any node v is always at least the out-flow, namely $\text{prob}(i, j, v) \geq \sum_{u:\text{parent}(u)=v} \text{prob}(i, j, u)$. This leads to the following simple but crucial observation.

Observation 4.3 *For any arm i , for any set of states $X \subseteq \mathcal{S}_i$ such that no state in X is an ancestor of another state in X in the transition tree T_i , and for any $z \in \mathcal{S}_i$ that is an ancestor of all states in X , $\text{prob}(i, j, z) \geq \sum_{x \in X} \text{prob}(i, j, x)$.*

More generally, given similar conditions on X , if Z is a set of states such that for any $x \in X$, there exists $z \in Z$ such that z is an ancestor of x , we have $\sum_{z \in Z} \text{prob}(i, j, z) \geq \sum_{x \in X} \text{prob}(i, j, x)$

4.2.2 Phase II: Eliminating Small Gaps

While Appendix A.3 shows that preemption is necessary to remain competitive with respect to Opt, we also should not get “tricked” into switching arms during very short breaks taken by the LP. For example, say, an arm of length $(B - 1)$ was played in two continuous segments with a gap in the middle. In this case, we should not lose out on profit from this arm by starting some other arms’ plays during the break. To handle this issue, whenever some path on the strategy tree is almost contiguous—i.e., gaps on it are relatively small—we make these portions completely contiguous. Note that we will not make the entire tree contiguous, but just combine some sections together.

⁴To reiterate, even though we call this a convex decomposition, the sum of the probability values of the root state of any arm is at most one by constraint 4.12, and hence the sum of the probabilities of the root over the decomposition could be less than one in general.

Before we make this formal, here is some useful notation: Given $u \in \mathcal{S}_i$, let $\text{Head}(i, j, u)$ be its ancestor node $v \in \mathcal{S}_i$ of least depth such that the plays from v through u occur in consecutive time values. More formally, the path $v = v_1, v_2, \dots, v_l = u$ in T_i is such that $\text{time}(i, j, v_l) = \text{time}(i, j, v_{l-1}) + 1$ for all $l' \in [2, l]$. We also define the *connected component* of a node u , denoted by $\text{comp}(i, j, u)$, as the set of all nodes u' such that $\text{Head}(i, j, u) = \text{Head}(i, j, u')$. Figure 4.2 shows the connected components and heads.

The main idea of our *gap-filling* procedure is the following: if a head state $v = \text{Head}(i, j, u)$ is played at time $t = \text{time}(i, j, v)$ s.t. $t < 2 \cdot \text{depth}(v)$, then we “advance” the $\text{comp}(i, j, v)$ and get rid of the gap between v and its parent (and recursively apply this rule)⁵. The procedure can be described in more detail as follows.

Algorithm 4.1 Gap Filling Algorithm GapFill

```

1: for  $\tau = B$  to 1 do
2:   while there exists a tree-node  $u \in \mathbb{T}(i, j)$  such that  $\tau = \text{time}(\text{Head}(u)) < 2 \cdot \text{depth}(\text{Head}(u))$  do
3:     let  $v = \text{Head}(u)$ .
4:     if  $v$  is not the root of  $\mathbb{T}(i, j)$  then
5:       let  $v' = \text{parent}(v)$ .
6:       advance the component  $\text{comp}(v)$  rooted at  $v$  such that  $\text{time}(v) \leftarrow \text{time}(v') + 1$ , to make  $\text{comp}(v)$ 
contiguous with the ancestor forming one larger component. Also alter the times of  $w \in \text{comp}(v)$ 
appropriately to maintain contiguity with  $v$  (and now with  $v'$ ).

```

One crucial property is that these “advances” do not increase by much the number of plays that occur at any given time t . Essentially this is because if for some time slot t we “advance” a set of components that were originally scheduled after t to now cross time slot t , these components moved because their ancestor paths (fractionally) used up at least $t/2$ of the time slots before t ; since there are t time slots to be used up, each to unit extent, there can be at most 2 units of components being moved up. Hence, in the following, we assume that our \mathbb{T} ’s satisfy the properties in the following lemma:

Lemma 4.4 *Algorithm GapFill produces a modified collection of \mathbb{T} ’s such that*

- (i) For each i, j, u such that $r_u > 0$, $\text{time}(\text{Head}(i, j, u)) \geq 2 \cdot \text{depth}(\text{Head}(i, j, u))$.
- (ii) The total extent of plays at any time t , i.e., $\sum_{(i,j,u):\text{time}(i,j,u)=t} \text{prob}(i, j, u)$ is at most 3.

The proof appears in Appendix D.2.

4.2.3 Phase III: Scheduling the Arms

Having done the preprocessing, the rounding algorithm is simple: it first randomly selects at most one strategy forest from the collection $\{\mathbb{T}(i, j)\}_j$ for each arm i . It then picks an arm with the earliest connected component (i.e., that with smallest $\text{time}(\text{Head}(i, j, u))$) that contains the current state (the root states, to begin with), plays it to the end—which either results in terminating the arm, or making a transition to a state played much later in time, and repeats. The formal description appears in Algorithm 4.2. (If there are ties in Step 5, we choose the smallest index.) Note that the algorithm runs as long as there is some active node, regardless of whether or not we have run out of plays (i.e., the budget is exceeded)—however, we only count the profit from the first B plays in the analysis.

Observe that Steps 7-9 play a connected component of a strategy forest contiguously. In particular, this means that all $\text{currstate}(i)$ ’s considered in Step 5 are head vertices of the corresponding strategy forests. These facts will be crucial in the analysis.

Lemma 4.5 *For arm i and strategy $\mathbb{T}(i, j)$, conditioned on $\sigma(i) = j$ after Step 1 of AlgMAB, the probability of playing state $u \in \mathcal{S}_i$ is $\text{prob}(i, j, u) / \text{prob}(i, j, \rho_i)$, where the probability is over the random transitions of arm i .*

⁵The intuition is that such vertices have only a small gap in their play and should rather be played contiguously.

Algorithm 4.2 Scheduling the Connected Components: Algorithm AlgMAB

- 1: for arm i , **sample** strategy $\mathbb{T}(i, j)$ with probability $\frac{\text{prob}(i, j, \rho_i)}{24}$; ignore arm i w.p. $1 - \sum_j \frac{\text{prob}(i, j, \rho_i)}{24}$.
 - 2: let $A \leftarrow$ set of “active” arms which chose a strategy in the random process.
 - 3: for each $i \in A$, **let** $\sigma(i) \leftarrow$ index j of the chosen $\mathbb{T}(i, j)$ and **let** $\text{currstate}(i) \leftarrow$ root ρ_i .
 - 4: **while** active arms $A \neq \emptyset$ **do**
 - 5: **let** $i^* \leftarrow$ arm with state played earliest in the LP (i.e., $i^* \leftarrow \text{argmin}_{i \in A} \{\text{time}(i, \sigma(i), \text{currstate}(i))\}$).
 - 6: **let** $\tau \leftarrow \text{time}(i^*, \sigma(i^*), \text{currstate}(i^*))$.
 - 7: **while** $\text{time}(i^*, \sigma(i^*), \text{currstate}(i^*)) \neq \infty$ **and** $\text{time}(i^*, \sigma(i^*), \text{currstate}(i^*)) = \tau$ **do**
 - 8: **play** arm i^* at state $\text{currstate}(i^*)$
 - 9: **update** $\text{currstate}(i^*)$ be the new state of arm i^* ; **let** $\tau \leftarrow \tau + 1$.
 - 10: **if** $\text{time}(i^*, \sigma(i^*), \text{currstate}(i^*)) = \infty$ **then**
 - 11: **let** $A \leftarrow A \setminus \{i^*\}$
-

The above lemma is relatively simple, and proved in Appendix D.3. The rest of the section proves that in expectation, we collect a constant factor of the LP reward of each strategy $\mathbb{T}(i, j)$ before running out of budget; the analysis is inspired by our Stock rounding procedure. We mainly focus on the following lemma.

Lemma 4.6 *Consider any arm i and strategy $\mathbb{T}(i, j)$. Then, conditioned on $\sigma(i) = j$ and on the algorithm playing state $u \in \mathcal{S}_i$, the probability that this play happens before time $\text{time}(i, j, u)$ is at least $1/2$.*

Proof. Fix an arm i and an index j for the rest of the proof. Given a state $u \in \mathcal{S}_i$, let \mathcal{E}_{iju} denote the event $(\sigma(i) = j) \wedge$ (state u is played). Also, let $\mathbf{v} = \text{Head}(i, j, u)$ be the head of the connected component containing u in $\mathbb{T}(i, j)$. Let r.v. τ_u (respectively $\tau_{\mathbf{v}}$) be the actual time at which state u (respectively state \mathbf{v}) is played—these random variables take value ∞ if the arm is not played in these states. Then

$$\Pr[\tau_u \leq \text{time}(i, j, u) \mid \mathcal{E}_{iju}] \geq \frac{1}{2} \iff \Pr[\tau_{\mathbf{v}} \leq \text{time}(i, j, \mathbf{v}) \mid \mathcal{E}_{iju}] \geq \frac{1}{2}, \quad (4.14)$$

because the time between playing u and \mathbf{v} is exactly $\text{time}(i, j, u) - \text{time}(i, j, \mathbf{v})$ since the algorithm plays connected components continuously (and we have conditioned on \mathcal{E}_{iju}). Hence, we can just focus on proving the right inequality in (4.14) for vertex \mathbf{v} .

For brevity of notation, let $t_{\mathbf{v}} = \text{time}(i, j, \mathbf{v})$. In addition, we define the order \preceq to indicate which states can be played before \mathbf{v} . That is, again making use of the fact that the algorithm plays connected components contiguously, we say that $(i', j', v') \preceq (i, j, \mathbf{v})$ iff $\text{time}(\text{Head}(i', j', v')) \leq \text{time}(\text{Head}(i, j, \mathbf{v}))$. Notice that this order is independent of the run of the algorithm.

For each arm $i' \neq i$ and index j' , we define random variables $Z_{i', j'}$ used to count the number of plays that can possibly occur before the algorithm plays state \mathbf{v} . If $\mathbf{1}_{(i', j', v')}$ is the indicator variable of event $\mathcal{E}_{i', j', v'}$, define

$$Z_{i', j'} = \min(t_{\mathbf{v}}, \sum_{v': (i', j', v') \preceq (i, j, \mathbf{v})} \mathbf{1}_{(i', j', v')}). \quad (4.15)$$

We truncate $Z_{i', j'}$ at $t_{\mathbf{v}}$ because we just want to capture how much time up to $t_{\mathbf{v}}$ is being used. Now consider the sum $Z = \sum_{i' \neq i} \sum_{j'} Z_{i', j'}$. Note that for arm i' , at most one of the $Z_{i', j'}$ values will be non-zero in any scenario, namely the index $\sigma(i')$ sampled in Step 1. The first claim below shows that it suffices to consider the upper tail of Z , and show that $\Pr[Z \geq t_{\mathbf{v}}/2] \leq 1/2$, and the second gives a bound on the conditional expectation of $Z_{i', j'}$.

Claim 4.7 $\Pr[\tau_{\mathbf{v}} \leq t_{\mathbf{v}} \mid \mathcal{E}_{iju}] \geq \Pr[Z \leq t_{\mathbf{v}}/2]$.

Proof. We first claim that $\Pr[\tau_{\mathbf{v}} \leq t_{\mathbf{v}} \mid \mathcal{E}_{iju}] \geq \Pr[Z \leq t_{\mathbf{v}}/2 \mid \mathcal{E}_{iju}]$. So, let us condition on \mathcal{E}_{iju} . Then if $Z \leq t_{\mathbf{v}}/2$, none of the $Z_{i', j'}$ variables were truncated at $t_{\mathbf{v}}$, and hence Z exactly counts the total number of plays (by all other arms $i' \neq i$, from any state) that could possibly be played before the algorithm plays v in strategy $\mathbb{T}(i, j)$. Therefore, if Z is smaller than $t_{\mathbf{v}}/2$, then combining this with the fact that $\text{depth}(v) \leq t_{\mathbf{v}}/2$ (from

Lemma 4.4(i)), we can infer that all the plays (including those of v 's ancestors) that can be made before playing v can indeed be completed within t_v . In this case the algorithm will definitely play v before t_v ; hence we get that conditioning on \mathcal{E}_{iju} , the event $\tau_v \leq t_v$ holds when $Z \leq t_v/2$.

Finally, to remove the conditioning: note that $Z_{i',j'}$ is just a function of (i) the random variables $\mathbf{1}_{(i',j',v')}$, i.e., the random choices made by playing $\mathbb{T}(i',j')$, and (ii) the constant $t_v = \text{time}(i,j,v)$. However, the r.v.s $\mathbf{1}_{(i',j',v')}$ are clearly independent of the event \mathcal{E}_{iju} for $i' \neq i$ since the plays of **AlgMAB** in one arm are independent of the others, and $\text{time}(i,j,v)$ is a constant determined once the strategy forests are created in Phase II. Hence the event $Z \leq t_v/2$ is independent of \mathcal{E}_{iju} ; hence $\Pr[Z \leq t_v/2 \mid \mathcal{E}_{iju}] = \Pr[Z \leq t_v/2]$, which completes the proof. ■

Claim 4.8

$$\mathbb{E}[Z_{i',j'} \mid \sigma(i') = j'] \leq \sum_{v' \text{ s.t. } \text{time}(i',j',v') \leq t_v} \frac{\text{prob}(i',j',v')}{\text{prob}(i',j',\rho_{i'})} + t_v \left(\sum_{v' \text{ s.t. } \text{time}(i',j',v') = t_v} \frac{\text{prob}(i',j',v')}{\text{prob}(i',j',\rho_{i'})} \right)$$

Proof. Recall the definition of $Z_{i',j'}$ in Eq (4.15): any state v' with $\text{time}(i',j',v') > t_v$ may contribute to the sum only if it is part of a connected component with head $\text{Head}(i',j',v')$ such that $\text{time}(\text{Head}(i',j',v')) \leq t_v$, by the definition of the ordering \preceq . Even among such states, if $\text{time}(i',j',v') > 2t_v$, then the truncation implies that $Z_{i',j'}$ is unchanged whether or not we include $\mathbf{1}_{(i',j',v')}$ in the sum. Indeed, if $\mathbf{1}_{(i',j',v')} = 1$ then all of v' 's ancestors will have their indicator variables at value 1; moreover $\text{depth}(v') > t_v$ since there is a contiguous collection of nodes that are played from this tree $\mathbb{T}(i',j')$ from time t_v onwards till $\text{time}(i',j',v') > 2t_v$; so the sum would be truncated at value t_v whenever $\mathbf{1}_{(i',j',v')} = 1$. Therefore, we can write

$$Z_{i',j'} \leq \sum_{v': \text{time}(i',j',v') \leq t_v} \mathbf{1}_{(i',j',v')} + \sum_{\substack{v': t_v < \text{time}(i',j',v') \leq 2t_v \\ (i',j',v') \preceq (i,j,v)}} \mathbf{1}_{(i',j',v')} \quad (4.16)$$

Recall we are interested in the conditional expectation given $\sigma(i') = j'$. Note that $\Pr[\mathbf{1}_{(i',j',v')} \mid \sigma(i') = j'] = \text{prob}(i',j',v')/\text{prob}(i',j',\rho_{i'})$ by Lemma 4.5, hence the first sum in (4.16) gives the first part of the claimed bound. Now the second part: observe that for any arm i' , any fixed value of $\sigma(i') = j'$, and any value of $t' \geq t_v$,

$$\sum_{\substack{v' \text{ s.t. } \text{time}(i',j',v') = t' \\ (i',j',v') \preceq (i,j,v)}} \text{prob}(i',j',v') \leq \sum_{v' \text{ s.t. } \text{time}(i',j',v') = t_v} \text{prob}(i',j',v')$$

This is because of the following argument: Any state that appears on the LHS of the sum above is part of a connected component which crosses t_v , they must have an ancestor which is played at t_v . Also, since all states which appear in the LHS are played at t' , no state can be an ancestor of another. Hence, we can apply the second part of Observation 4.3 and get the above inequality. Combining this with the fact that $\Pr[\mathbf{1}_{(i',j',v')} \mid \sigma(i') = j'] = \text{prob}(i',j',v')/\text{prob}(i',j',\rho_{i'})$, and applying it for each value of $t' \in (t_v, 2t_v]$, gives us the second term. ■

Equipped with the above claims, we are ready to complete the proof of Lemma 4.6. Employing Claim 4.8 we get

$$\begin{aligned} \mathbb{E}[Z] &= \sum_{i' \neq i} \sum_{j'} \mathbb{E}[Z_{i',j'}] = \sum_{i' \neq i} \sum_{j'} \mathbb{E}[Z_{i',j'} \mid \sigma(i') = j'] \cdot \Pr[\sigma(i') = j'] \\ &= \frac{1}{24} \sum_{i' \neq i} \sum_{j'} \left\{ \sum_{v': \text{time}(i',j',v') \leq t_v} \text{prob}(i',j',v') + t_v \left(\sum_{v': \text{time}(i',j',v') = t_v} \text{prob}(i',j',v') \right) \right\} \quad (4.17) \end{aligned}$$

$$= \frac{1}{24} (3 \cdot t_v + 3 \cdot t_v) \leq \frac{1}{4} t_v. \quad (4.18)$$

Equation (4.17) follows from the fact that each tree $\mathbb{T}(i,j)$ is sampled with probability $\frac{\text{prob}(i,j,\rho_i)}{24}$ and (4.18) follows from Lemma 4.4. Applying Markov's inequality, we have that $\Pr[Z \geq t_v/2] \leq 1/2$. Finally, Claim 4.7 says that $\Pr[\tau_v \leq t_v \mid \mathcal{E}_{iju}] \geq \Pr[Z \leq t_v/2] \geq 1/2$, which completes the proof. ■

Theorem 4.9 *The reward obtained by the algorithm AlgMAB is at least $\Omega(\text{LPOpt})$.*

Proof. The theorem follows by a simple linearity of expectation. Indeed, the expected reward obtained from any state $u \in \mathcal{S}_i$ is at least $\sum_j \Pr[\sigma(i) = j] \Pr[\text{state } u \text{ is played} \mid \sigma(i) = j] \Pr[\tau_u \leq t_u \mid \mathcal{E}_{iju}] \cdot R_u \geq \sum_j \frac{\text{prob}(i,j,u)}{24} \frac{1}{2} \cdot R_u$. Here, we have used Lemmas 4.5 and 4.6 for the second and third probabilities. But now we can use Lemma 4.2 to infer that $\sum_j \text{prob}(i,j,u) = \sum_t z_{u,t}$; Making this substitution and summing over all states $u \in \mathcal{S}_i$ and arms i completes the proof. ■

5 MABs with Arbitrary Transition Graphs

We now show how we can use techniques akin to those we described for the case when the transition graph is a tree, to handle the case when it can be an arbitrary directed graph. A naïve way to do this is to expand out the transition graph as a tree, but this incurs an exponential blowup of the state space which we want to avoid. We can assume we have a layered DAGs, though, since the conversion from a digraph to a layered DAG only increases the state space by a factor of the horizon B ; this standard reduction appears in Appendix E.1.

While we can again write an LP relaxation of the problem for layered DAGs, the challenge arises in the rounding algorithm: specifically, in (i) obtaining the convex decomposition of the LP solution as in Phase I, and (ii) eliminating small gaps as in Phase II by advancing forests in the strategy.

- We handle the first difficulty by considering convex decompositions not just over strategy forests, but over slightly more sophisticated strategy DAGs. Recall (from Figure 4.2) that in the tree case, each state in a strategy forest was labeled by a unique time and a unique probability associated with that time step. As the name suggests, we now have labeled DAGs—but the change is more than just that. Now each state has a copy associated with *each* time step in $\{1, \dots, B\}$. This change tries to capture the fact that our strategy may play from a particular state u at different times depending on the path taken by the random transitions used to reach this state. (This path was unique in the tree case.)
- Now having sampled a strategy DAG for each arm, one can expand them out into strategy forests (albeit with an exponential blow-up in the size), and use Phases II and III from our previous algorithm—it is not difficult to prove that this algorithm is a constant-factor approximation. However, the above is not a poly-time algorithm, since the size of the strategy forests may be exponentially large. If we don't expand the DAG, then we do not see how to define gap elimination for Phase II. But we observe that instead of explicitly performing the advance steps in Phase II, it suffices to perform them as a *thought experiment*—i.e., to not alter the strategy forest at all, but merely to infer when these advances would have happened, and play accordingly in the Phase III ⁶. Using this, we can give an algorithm that plays just on the DAG, and argue that the sequence of plays made by our DAG algorithm faithfully mimics the execution if we had constructed the exponential-size tree from the DAG, and executed Phases II and III on that tree.

The details of the LP rounding algorithm for layered DAGs follows in Sections 5.1-5.3.

⁶This is similar to the idea of lazy evaluation of strategies. The DAG contains an implicit randomized strategy which we make explicit as we toss coins of the various outcomes using an algorithm.

5.1 LP Relaxation

There is only one change in the LP—constraint (5.19) now says that if a state u is visited at time t , then one of its ancestors must have been pulled at time $t - 1$; this ancestor was unique in the case of trees.

$$\begin{aligned} \max \sum_{u,t} r_u \cdot z_{u,t} & & (\text{LP}_{\text{mabdag}}) \\ w_{u,t} = \sum_v z_{v,t-1} \cdot p_{v,u} & \quad \forall t \in [2, B], u \in \mathcal{S} \setminus \cup_i \{\rho_i\}, v \in \mathcal{S} & (5.19) \end{aligned}$$

$$\sum_{t' \leq t} w_{u,t'} \geq \sum_{t' \leq t} z_{u,t'} \quad \forall t \in [1, B], u \in \mathcal{S} \quad (5.20)$$

$$\sum_{u \in \mathcal{S}} z_{u,t} \leq 1 \quad \forall t \in [1, B] \quad (5.21)$$

$$w_{\rho_i,1} = 1 \quad \forall i \in [1, n] \quad (5.22)$$

Again, a similar analysis to the tree case shows that this is a valid relaxation, and hence the LP value is at least the optimal expected reward.

5.2 Convex Decomposition: The Altered Phase I

This is the step which changes the most—we need to incorporate the notion of peeling out a “strategy DAG” instead of just a tree. The main complication arises from the fact that a play of a state u may occur at different times in the LP solution, depending on the path to reach state u in the transition DAG. However, we don’t need to keep track of the entire history used to reach u , just how much time has elapsed so far. With this in mind, we create B copies of each state u (which will be our nodes in the strategy DAG), indexed by (u, t) for $1 \leq t \leq B$.

The j^{th} strategy dag $\mathbb{D}(i, j)$ for arm i is an assignment of values $\text{prob}(i, j, u, t)$ and a relation ‘ \rightarrow ’ from 4-tuples to 4-tuples of the form $(i, j, u, t) \rightarrow (i, j, v, t')$ such that the following properties hold:

- (i) For $u, v \in \mathcal{S}_i$ such that $p_{u,v} > 0$ and any time t , there is exactly one time $t' \geq t + 1$ such that $(i, j, u, t) \rightarrow (i, j, v, t')$. Intuitively, this says if the arm is played from state u at time t it transitions to state v , then it is played from v at a unique time t' , if it played at all. If $t' = \infty$, the play from v never happens.
- (ii) For any $u \in \mathcal{S}_i$ and time $t \neq \infty$, $\text{prob}(i, j, u, t) = \sum_{(v,t') \text{ s.t. } (i,j,v,t') \rightarrow (i,j,u,t)} \text{prob}(i, j, v, t') \cdot p_{v,u}$.

For clarity, we use the following notation throughout the remainder of the section: *states* refer to the states in the original transition DAG, and *nodes* correspond to the tuples (i, j, u, t) in the strategy DAGs. When i and j are clear in context, we may simply refer to a node of the strategy DAG by (u, t) .

Equipped with the above definition, our convex decomposition procedure appears in Algorithm 5.2. The main subroutine involved is presented first (Algorithm 5.1). This subroutine, given a fractional solution, identifies the structure of the DAG that will be peeled out, depending on when the different states are first played fractionally in the LP solution. Since we have a layered DAG, the notion of the *depth* of a state is well-defined as the number of hops from the root to this state in the DAG, with the depth of the root being 0.

Algorithm 5.1 Sub-Routine PeelStrat (i, j)

- 1: **mark** (ρ_i, t) where t is the earliest time s.t. $z_{\rho_i, t} > 0$ and set $\text{peelProb}(\rho_i, t) = 1$. All other nodes are un-marked and have $\text{peelProb}(v, t') = 0$.
 - 2: **while** \exists a marked unvisited node **do**
 - 3: **let** (u, t) denote the marked node of smallest depth and earliest time; **update** its status to visited.
 - 4: **for** every v s.t. $p_{u,v} > 0$ **do**
 - 5: **if** there is t' such that $z_{v,t'} > 0$, consider the earliest such t' and **then**
 - 6: **mark** (v, t') and **set** $(i, j, u, t) \rightarrow (i, j, v, t')$; **update** $\text{peelProb}(v, t') := \text{peelProb}(v, t') + \text{peelProb}(u, t) \cdot p_{u,v}$.
 - 7: **else**
 - 8: **set** $(i, j, u, t) \rightarrow (i, j, v, \infty)$ and leave $\text{peelProb}(v, \infty) = 0$.
-

The convex decomposition algorithm is now very easy to describe with the sub-routine in Algorithm 5.1 in hand.

Algorithm 5.2 Convex Decomposition of Arm i

```

1: set  $\mathcal{C}_i \leftarrow \emptyset$  and set loop index  $j \leftarrow 1$ .
2: while  $\exists$  a state  $u \in \mathcal{S}_i$  s.t.  $\sum_t z_{u,t}^{j-1} > 0$  do
3:   run sub-routine PeelStrat to extract a DAG  $\mathbb{D}(i, j)$  with the appropriate  $\text{peelProb}(u, t)$  values.
4:   let  $A \leftarrow \{(u, t) \text{ s.t. } \text{peelProb}(u, t) \neq 0\}$ .
5:   let  $\epsilon = \min_{(u,t) \in A} z_{u,t}^{j-1} / \text{peelProb}(u, t)$ .
6:   for every  $(u, t)$  do
7:     set  $\text{prob}(i, j, u, t) = \epsilon \cdot \text{peelProb}(u, t)$ .
8:     update  $z_{u,t}^j = z_{u,t}^{j-1} - \text{prob}(i, j, u, t)$ .
9:     update  $w_{v,t+1}^j = w_{v,t+1}^{j-1} - \text{prob}(i, j, u, t) \cdot p_{u,v}$  for all  $v$ .
10:  set  $\mathcal{C}_i \leftarrow \mathcal{C}_i \cup \mathbb{D}(i, j)$ .
11:  increment  $j \leftarrow j + 1$ .

```

An illustration of a particular DAG and a strategy dag $\mathbb{D}(i, j)$ peeled off is given in Figure 5.3 (notice that the states w, y and z appear more than once depending on the path taken to reach them).




Figure 5.3: Strategy dags and how to visualize them: notice the same state played at different times.

Now we analyze the solutions $\{z^j, w^j\}$ created by Algorithm 5.2.

Lemma 5.1 Consider an integer j and suppose that $\{z^{j-1}, w^{j-1}\}$ satisfies constraints (4.10)-(4.12) of $\text{LP}_{\text{mabdag}}$. Then after iteration j of Step 2, the following properties hold:

- (a) $\mathbb{D}(i, j)$ (along with the associated $\text{prob}(i, j, \cdot, \cdot)$ values) is a valid strategy dag, i.e., satisfies the conditions (i) and (ii) presented above.
- (b) The residual solution $\{z^j, w^j\}$ satisfies constraints (5.19)-(5.21).
- (c) For any time t and state $u \in \mathcal{S}_i$, $z_{u,t}^{j-1} - z_{u,t}^j = \text{prob}(i, j, u, t)$.

Proof. We show the properties stated above one by one.

Property (a): This follows from the construction of Algorithm 5.1. More precisely, condition (i) is satisfied because in Algorithm 5.1 each (u, t) is visited at most once and that is the only time when a pair $(u, t) \rightarrow (v, t')$ (with $t' \geq t + 1$) is added to the relation. For condition (ii), notice that every time a pair $(u, t) \rightarrow (v, t')$ is added to the relation we keep the invariant $\text{peelProb}(v, t') = \sum_{(w,\tau) \text{ s.t. } (i,j,w,\tau) \rightarrow (i,j,v,t')} \text{peelProb}(w, \tau) \cdot p_{w,v}$; condition (ii) then follows since $\text{prob}(\cdot)$ is a scaling of $\text{peelProb}(\cdot)$.

Property (b): Constraint (5.19) of $\text{LP}_{\text{mabdag}}$ is clearly satisfied by the new LP solution $\{z^j, w^j\}$ because of the two updates performed in Steps 8 and 9: if we decrease the z value of any state at any time, the w of all children are appropriately reduced for the subsequent timestep.

Before showing that the solution $\{z^j, w^j\}$ satisfies constraint (5.20), we first argue that after every round of the procedure they remain non-negative. By the choice of ϵ in step 5, we have $\text{prob}(i, j, u, t) = \epsilon \cdot \text{peelProb}(u, t) \leq \frac{z_{u,t}^{j-1}}{\text{peelProb}(u,t)} \text{peelProb}(u, t) = z_{u,t}^{j-1}$ (notice that this inequality holds even if $\text{peelProb}(u, t) = 0$); consequently even after the update in step 8, $z_{u,t}^j \geq 0$ for all u, t . This and the fact that the constraints (5.19) are satisfied implies that $\{z^j, w^j\}$ satisfies the non-negativity requirement.

We now show that constraint (5.20) is satisfied. Suppose for the sake of contradiction there exist some $u \in \mathcal{S}$ and $t \in [1, B]$ such that $\{z^j, w^j\}$ violates this constraint. Then, let us consider any such u and the earliest time t_u such that the constraint is violated. For such a u , let $t'_u \leq t_u$ be the latest time before t_u where $z_{u,t'}^{j-1} > 0$. We now consider two cases.

Case (i): $t'_u < t_u$. This is the simpler case of the two. Because t_u was the earliest time where constraint (5.20) was violated, we know that $\sum_{t' \leq t_u} w_{u,t'}^j \geq \sum_{t' \leq t_u} z_{u,t'}^j$. Furthermore, since $z_{u,t}$ is never increased during the course of the algorithm we know that $\sum_{t'=t'_u+1}^{t_u} z_{u,t'}^j = 0$. This fact coupled with the non-negativity of $w_{u,t}^j$ implies that the constraint in fact is not violated, which contradicts our assumption about the tuple u, t_u .

Case (ii): $t'_u = t_u$. In this case, observe that there cannot be any pair of tuples $(v, t_1) \rightarrow (u, t_2)$ s.t. $t_1 < t_u$ and $t_2 > t_u$, because any copy of v (some ancestor of u) that is played before t_u , will mark a copy of u that occurs before t_u or the one being played at t_u in Step 6 of PeelStrat. We will now show that summed over all $t' \leq t_u$, the decrease in the LHS is counter-balanced by a corresponding drop in the RHS, between the solutions $\{z^{j-1}, w^{j-1}\}$ and $\{z^j, w^j\}$ for this constraint (5.20) corresponding to u and t_u . To this end, notice that the only times when $w_{u,t'}$ is updated (in Step 9) for $t' \leq t_u$, are when considering some (v, t_1) in Step 6 such that $(v, t_1) \rightarrow (u, t_2)$ and $t_1 < t_2 \leq t_u$. The value of w_{u,t_1+1} is dropped by exactly $\text{prob}(i, j, v, t_1) \cdot p_{v,u}$. But notice that the corresponding term z_{u,t_2} drops by $\text{prob}(i, j, u, t_2) = \sum_{(v'', t'') \text{ s.t. } (v'', t'') \rightarrow (u, t_2)} \text{prob}(i, j, v'', t'') \cdot p_{v'',u}$. Therefore, the total drop in w is balanced by a commensurate drop in z on the RHS.

Finally, constraint (5.21) is also satisfied as the z variables only decrease in value.

Property (c): This is an immediate consequence of the Step 8 of the convex decomposition algorithm. ■

As a consequence of the above lemma, we get the following.

Lemma 5.2 *Given a solution to $(\text{LP}_{\text{mabdag}})$, there exists a collection of at most $nB^2|\mathcal{S}|$ strategy dags $\{\mathbb{D}(i, j)\}$ such that $z_{u,t} = \sum_j \text{prob}(i, j, u, t)$. Hence, $\sum_{(i,j,u)} \text{prob}(i, j, u, t) \leq 1$ for all t .*

5.3 Phases II and III

We now show how to execute the strategy dags $\mathbb{D}(i, j)$. At a high level, the development of the plays mirrors that of Sections 4.2.2 and 4.2.3. First we transform $\mathbb{D}(i, j)$ into a (possibly exponentially large) blown-up tree and show how this playing these exactly captures playing the strategy dags. Hence (if running time is not a concern), we can simply perform the gap-filling algorithm and make plays on these blown-up trees following Phases II and III in Sections 4.2.2 and 4.2.3. To achieve polynomial running time, we then show that we can *implicitly execute* the gap-filling phase while playing this tree, thus getting rid of actually performing Phase 4.2.2. Finally, to complete our argument, we show how we do not need to explicitly construct the blown-up tree, and can generate the required portions depending on the transitions made thus far *on demand*.

5.3.1 Transforming the DAG into a Tree

Consider any strategy dag $\mathbb{D}(i, j)$. We first transform this dag into a (possibly exponential) tree by making as many copies of a node (i, j, u, t) as there are paths from the root to (i, j, u, t) in $\mathbb{D}(i, j)$. More formally, define $\mathbb{DT}(i, j)$ as the tree whose vertices are the simple paths in $\mathbb{D}(i, j)$ which start at the root. To avoid confusion, we will explicitly refer to vertices of the tree \mathbb{DT} as tree-nodes, as distinguished from the *nodes* in \mathbb{D} ; to simplify the notation we identify each tree-node in \mathbb{DT} with its corresponding path in \mathbb{D} . Given two tree-nodes P, P' in $\mathbb{DT}(i, j)$, add an arc from P to P' if P' is an immediate extension of P , i.e., if P corresponds to some

path $(i, j, u_1, t_1) \rightarrow \dots \rightarrow (i, j, u_k, t_k)$ in $\mathbb{D}(i, j)$, then P' is a path $(i, j, u_1, t_1) \rightarrow \dots \rightarrow (i, j, u_k, t, k) \rightarrow (i, j, u_{k+1}, t_{k+1})$ for some node (i, j, u_{k+1}, t_{k+1}) .

For a tree-node $P \in \mathbb{DT}(i, j)$ which corresponds to the path $(i, j, u_1, t_1) \rightarrow \dots \rightarrow (i, j, u_k, t_k)$ in $\mathbb{D}(i, j)$, we define $\text{state}(P) = u_k$, i.e., $\text{state}(\cdot)$ denotes the final state (in \mathcal{S}_i) in the path P . Now, for tree-node $P \in \mathbb{DT}(i, j)$, if u_1, \dots, u_k are the children of $\text{state}(P)$ in \mathcal{S}_i with positive transition probability from $\text{state}(P)$, then P has exactly k children P_1, \dots, P_k with $\text{state}(P_l)$ equal to u_l for all $l \in [k]$. The *depth* of a tree-node P is defined as the depth of $\text{state}(P)$.

We now define the quantities *time* and *prob* for tree-nodes in $\mathbb{DT}(i, j)$. Let P be a path in $\mathbb{D}(i, j)$ from ρ_i to node (i, j, u, t) . We define $\text{time}(P) := t$ and $\text{prob}(P) := \text{prob}(P')p_{(\text{state}(P'), u)}$, where P' is obtained by dropping the last node from P . The blown-up tree $\mathbb{DT}(i, j)$ of our running example $\mathbb{D}(i, j)$ (Figure 5.3) is given in Figure 5.4.

Lemma 5.3 *For any state u and time t , $\sum_{P: \text{time}(P)=t \text{ and } \text{state}(P)=u} \text{prob}(P) = \text{prob}(i, j, u, t)$.*




Figure 5.4: Blown-up Strategy Forest $\mathbb{DT}(i, j)$

Now that we have a tree labeled with *prob* and *time* values, the notions of connected components and heads from Section 4.2.2 carry over. Specifically, we define $\text{Head}(P)$ to be the ancestor P' of P in $\mathbb{DT}(i, j)$ with least depth such that there is a path $(P' = P_1 \rightarrow \dots \rightarrow P_l = P)$ satisfying $\text{time}(P_i) = \text{time}(P_{i-1}) + 1$ for all $i \in [2, l]$, i.e., the plays are made contiguously from $\text{Head}(P)$ to P in the blown-up tree. We also define $\text{comp}(P)$ as the set of all tree-nodes P' such that $\text{Head}(P) = \text{Head}(P')$.

In order to play the strategies $\mathbb{DT}(i, j)$ we first eliminate small gaps. The algorithm *GapFill* presented in Section 4.2.2 can be employed for this purpose and returns trees $\mathbb{DT}'(i, j)$ which satisfy the analog of Lemma 4.4.

Lemma 5.4 *The trees returned by *GapFill* satisfy the followings properties.*

- (i) *For each tree-node P such that $r_{\text{state}(P)} > 0$, $\text{time}(\text{Head}(P)) \geq 2 \cdot \text{depth}(\text{Head}(P))$.*
- (ii) *The total extent of plays at any time t , i.e., $\sum_{P: \text{time}(P)=t} \text{prob}(P)$ is at most 3.*

Now we use Algorithm 4.2 to play the trees $\mathbb{DT}(i, j)$. We restate the algorithm to conform with the notation used in the trees $\mathbb{DT}(i, j)$.

Now an argument identical to that for Theorem 4.9 gives us the following:

Theorem 5.5 *The reward obtained by the algorithm *AlgDAG* is at least a constant fraction of the optimum for $(\text{LP}_{\text{mabdag}})$.*

5.3.2 Implicit gap filling

Our next goal is to execute *GapFill* implicitly, that is, to incorporate the gap-filling within Algorithm *AlgDAG* without having to explicitly perform the advances.

To do this, let us review some properties of the trees returned by *GapFill*. For a tree-node P in $\mathbb{DT}(i, j)$, let $\text{time}(P)$ denote the associated time in the original tree (i.e., before the application of *GapFill*) and let $\text{time}'(P)$ denote the time in the modified tree (i.e., after $\mathbb{DT}(i, j)$ is modified by *GapFill*).

Algorithm 5.3 Scheduling the Connected Components: Algorithm AlgDAG

```
1: for arm  $i$ , sample strategy  $\mathbb{DT}(i, j)$  with probability  $\frac{\text{prob}(\text{root}(\mathbb{DT}(i, j)))}{24}$ ; ignore arm  $i$  w.p.  $1 - \sum_j \frac{\text{prob}(\text{root}(\mathbb{DT}(i, j)))}{24}$ .
2: let  $A \leftarrow$  set of “active” arms which chose a strategy in the random process.
3: for each  $i \in A$ , let  $\sigma(i) \leftarrow$  index  $j$  of the chosen  $\mathbb{DT}(i, j)$  and let  $\text{currnode}(i) \leftarrow$  root of  $\mathbb{DT}(i, \sigma(i))$ .
4: while active arms  $A \neq \emptyset$  do
5:   let  $i^* \leftarrow$  arm with tree-node played earliest (i.e.,  $i^* \leftarrow \text{argmin}_{i \in A} \{\text{time}(\text{currnode}(i))\}$ ).
6:   let  $\tau \leftarrow \text{time}(\text{currnode}(i^*))$ .
7:   while  $\text{time}(\text{currnode}(i^*)) \neq \infty$  and  $\text{time}(\text{currnode}(i^*)) = \tau$  do
8:     play arm  $i^*$  at state  $\text{state}(\text{currnode}(i^*))$ 
9:     let  $u$  be the new state of arm  $i^*$  and let  $P$  be the child of  $\text{currnode}(i^*)$  satisfying  $\text{state}(P) = u$ .
10:    update  $\text{currnode}(i^*)$  to be  $P$ ; let  $\tau \leftarrow \tau + 1$ .
11:   if  $\text{time}(\text{currnode}(i^*)) = \infty$  then
12:     let  $A \leftarrow A \setminus \{i^*\}$ 
```

Claim 5.6 For a non-root tree-node P and its parent P' , $\text{time}'(P) = \text{time}'(P') + 1$ if and only if, either $\text{time}(P) = \text{time}(P') + 1$ or $2 \cdot \text{depth}(P) > \text{time}(P)$.

Proof. Let us consider the forward direction. Suppose $\text{time}'(P) = \text{time}'(P') + 1$ but $\text{time}(P) > \text{time}(P') + 1$. Then P must have been the head of its component in the original tree and an **advance** was performed on it, so we must have $2 \cdot \text{depth}(P) > \text{time}(P)$.

For the reverse direction, if $\text{time}(P) = \text{time}(P') + 1$ then P could not have been a head since it belongs to the same component as P' and hence it will always remain in the same component as P' (as **GapFill** only merges components and never breaks them apart). Therefore, $\text{time}'(P) = \text{time}'(P') + 1$. On the other hand, if $\text{time}(P) > \text{time}(P') + 1$ and $2 \cdot \text{depth}(P) > \text{time}(P)$, then P was a head in the original tree, and because of the above criterion, **GapFill** must have made an advance on P' thereby including it in the same component as P ; so again it is easy to see that $\text{time}'(P) = \text{time}'(P') + 1$. ■

The crucial point here is that whether or not P is in the same component as its predecessor after the gap-filling (and, consequently, whether it was played contiguously along with its predecessor should that transition happen in **AlgDAG**) can be inferred from the time values of P, P' before gap-filling and from the depth of P —it does not depend on any other **advances** that happen during the gap-filling.

Algorithm 5.4 is a procedure which plays the original trees $\mathbb{DT}(i, j)$ while implicitly performing the **advance** steps of **GapFill** (by checking if the properties of Claim 5.6 hold). This change is reflected in Step 7 where we may play a node even if it is not contiguous, so long it satisfies the above stated properties. Therefore, as a consequence of Claim 5.6, we get the following Lemma that the plays made by **ImplicitFill** are identical to those made by **AlgDAG** after running **GapFill**.

Lemma 5.7 Algorithm *ImplicitFill* obtains the same reward as algorithm $\text{AlgDAG} \circ \text{GapFill}$.

5.3.3 Running **ImplicitFill** in Polynomial Time

With the description of **ImplicitFill**, we are almost complete with our proof with the exception of handling the exponential blow-up incurred in moving from \mathbb{D} to \mathbb{DT} . To resolve this, we now argue that while the blown-up \mathbb{DT} made it easy to visualize the transitions and plays made, all of it can be done implicitly from the strategy DAG \mathbb{D} . Recall that the tree-nodes in $\mathbb{DT}(i, j)$ correspond to simple paths in $\mathbb{D}(i, j)$. In the following, the final algorithm we employ (called **ImplicitPlay**) is simply the algorithm **ImplicitFill**, but with the exponentially blown-up trees $\mathbb{DT}(i, \sigma(i))$ being generated *on-demand*, as the different transitions are made. We now describe how this can be done.

Algorithm 5.4 Filling gaps implicitly: Algorithm ImplicitFill

```
1: for arm  $i$ , sample strategy  $\mathbb{DT}(i, j)$  with probability  $\frac{\text{prob}(\text{root}(\mathbb{DT}(i, j)))}{24}$ ; ignore arm  $i$  w.p.  $1 - \sum_j \frac{\text{prob}(\text{root}(\mathbb{DT}(i, j)))}{24}$ .
2: let  $A \leftarrow$  set of “active” arms which chose a strategy in the random process.
3: for each  $i \in A$ , let  $\sigma(i) \leftarrow$  index  $j$  of the chosen  $\mathbb{DT}(i, j)$  and let  $\text{currnode}(i) \leftarrow$  root of  $\mathbb{DT}(i, \sigma(i))$ .
4: while active arms  $A \neq \emptyset$  do
5:   let  $i^* \leftarrow$  arm with state played earliest (i.e.,  $i^* \leftarrow \text{argmin}_{i \in A} \{\text{time}(\text{currnode}(i))\}$ ).
6:   let  $\tau \leftarrow \text{time}(\text{currnode}(i^*))$ .
7:   while  $\text{time}(\text{currnode}(i^*)) \neq \infty$  and  $(\text{time}(\text{currnode}(i^*)) = \tau$  or  $2 \cdot \text{depth}(\text{currnode}(i^*)) > \text{time}(\text{currnode}(i^*)))$  do
8:     play arm  $i^*$  at state  $\text{state}(\text{currnode}(i^*))$ 
9:     let  $u$  be the new state of arm  $i^*$  and let  $P$  be the child of  $\text{currnode}(i^*)$  satisfying  $\text{state}(P) = u$ .
10:    update  $\text{currnode}(i^*)$  to be  $P$ ; let  $\tau \leftarrow \tau + 1$ .
11:    if  $\text{time}(\text{currnode}(i^*)) = \infty$  then
12:      let  $A \leftarrow A \setminus \{i^*\}$ 
```

In Step 3 of ImplicitFill, we start off at the roots of the trees $\mathbb{DT}(i, \sigma(i))$, which corresponds to the single-node path corresponding to the root of $\mathbb{D}(i, \sigma(i))$. Now, at some point in time in the execution of ImplicitFill, suppose we are at the tree-node $\text{currnode}(i^*)$, which corresponds to a path Q in $\mathbb{D}(i, \sigma(i))$ that ends at $(i, \sigma(i), v, t)$ for some v and t . The invariant we maintain is that, in our algorithm ImplicitPlay, we are at node $(i, \sigma(i), v, t)$ in $\mathbb{D}(i, \sigma(i))$. Establishing this invariant would show that the two runs ImplicitPlay and ImplicitFill would be identical, which when coupled with Theorem 5.5 would complete the proof—the information that ImplicitFill uses of Q , namely $\text{time}(Q)$ and $\text{depth}(Q)$, can be obtained from $(i, \sigma(i), v, t)$.

The invariant is clearly satisfied at the beginning, for the different root nodes. Suppose it is true for some tree-node $\text{currnode}(i)$, which corresponds to a path Q in $\mathbb{D}(i, \sigma(i))$ that ends at $(i, \sigma(i), v, t)$ for some v and t . Now, suppose upon playing the arm i at state v (in Step 8), we make a transition to state u (say), then ImplicitFill would find the unique child tree-node P of Q in $\mathbb{DT}(i, \sigma(i))$ with $\text{state}(P) = u$. Then let $(i, \sigma(i), u, t')$ be the last node of the path P , so that P equals Q followed by $(i, \sigma(i), u, t')$.

But, since the tree $\mathbb{DT}(i, \sigma(i))$ is just an expansion of $\mathbb{D}(i, \sigma(i))$, the unique child P in $\mathbb{DT}(i, \sigma(i))$ of tree-node Q which has $\text{state}(P) = u$, is (by definition of \mathbb{DT}) the unique node $(i, \sigma(i), u, t')$ of $\mathbb{D}(i, \sigma(i))$ such that $(i, \sigma(i), v, t) \rightarrow (i, \sigma(i), u, t')$. Hence, just as ImplicitFill transitions to P in $\mathbb{DT}(i, \sigma(i))$ (in Step 9), we can transition to the state $(i, \sigma(i), u, t')$ with just \mathbb{D} at our disposal, thus establishing the invariant.

For completeness, we present the implicit algorithm below.

6 Concluding Remarks

We presented the first constant-factor approximations for the stochastic knapsack problem with cancellations and correlated size/reward pairs, and for the budgeted learning problem without the martingale property. We showed that existing LPs for the restricted versions of the problems have large integrality gaps, which required us to give new LP relaxations, and well as new rounding algorithms for these problems.

Acknowledgments. We thank Kamesh Munagala and Sudipto Guha for useful conversations.

References

[Ber05] Dimitri P. Bertsekas. *Dynamic programming and optimal control*. Athena Scientific, Belmont, MA, third edition, 2005.

Algorithm 5.5 Algorithm ImplicitPlay

```
1: for arm  $i$ , sample strategy  $\mathbb{D}(i, j)$  with probability  $\frac{\text{prob}(\text{root}(\mathbb{D}(i, j)))}{24}$ ; ignore arm  $i$  w.p.  $1 - \sum_j \frac{\text{prob}(\text{root}(\mathbb{D}(i, j)))}{24}$ .
2: let  $A \leftarrow$  set of “active” arms which chose a strategy in the random process.
3: for each  $i \in A$ , let  $\sigma(i) \leftarrow$  index  $j$  of the chosen  $\mathbb{D}(i, j)$  and let  $\text{currnode}(i) \leftarrow$  root of  $\mathbb{D}(i, \sigma(i))$ .
4: while active arms  $A \neq \emptyset$  do
5:   let  $i^* \leftarrow$  arm with state played earliest (i.e.,  $i^* \leftarrow \text{argmin}_{i \in A} \{\text{time}(\text{currnode}(i))\}$ ).
6:   let  $\tau \leftarrow \text{time}(\text{currnode}(i^*))$ .
7:   while  $\text{time}(\text{currnode}(i^*)) \neq \infty$  and  $(\text{time}(\text{currnode}(i^*)) = \tau$  or  $2 \cdot \text{depth}(\text{currnode}(i^*)) >$ 
    $\text{time}(\text{currnode}(i^*)))$  do
8:     play arm  $i^*$  at state  $\text{state}(\text{currnode}(i^*))$ 
9:     let  $u$  be the new state of arm  $i^*$ .
10:    update  $\text{currnode}(i^*)$  to be  $u$ ; let  $\tau \leftarrow \tau + 1$ .
11:   if  $\text{time}(\text{currnode}(i^*)) = \infty$  then
12:     let  $A \leftarrow A \setminus \{i^*\}$ 
```

- [BGK11] Anand Bhalgat, Ashish Goel, and Sanjeev Khanna. Improved approximation results for stochastic knapsack problems. In *SODA '11*. Society for Industrial and Applied Mathematics, 2011.
- [BL97] John R. Birge and François Louveaux. *Introduction to stochastic programming*. Springer Series in Operations Research. Springer-Verlag, New York, 1997.
- [CR06] Shuchi Chawla and Tim Roughgarden. Single-source stochastic routing. In *Proceedings of APPROX*, pages 82–94. 2006.
- [Dea05] Brian C. Dean. *Approximation Algorithms for Stochastic Scheduling Problems*. PhD thesis, MIT, 2005.
- [DGV05] Brian C. Dean, Michel X. Goemans, and Jan Vondrák. Adaptivity and approximation for stochastic packing problems. In *SODA*, pages 395–404, 2005.
- [DGV08] Brian C. Dean, Michel X. Goemans, and Jan Vondrák. Approximating the stochastic knapsack problem: The benefit of adaptivity. *Math. Oper. Res.*, 33(4):945–964, 2008.
- [GGM06] Ashish Goel, Sudipto Guha, and Kamesh Munagala. Asking the right questions: model-driven optimization using probes. In *PODS*, pages 203–212, 2006.
- [GI99] Ashish Goel and Piotr Indyk. Stochastic load balancing and related problems. In *40th Annual Symposium on Foundations of Computer Science (New York, 1999)*, pages 579–586. IEEE Computer Soc., Los Alamitos, CA, 1999.
- [Git89] J. C. Gittins. *Multi-armed bandit allocation indices*. Wiley-Interscience Series in Systems and Optimization. John Wiley & Sons Ltd., Chichester, 1989. With a foreword by Peter Whittle.
- [GKN09] Ashish Goel, Sanjeev Khanna, and Brad Null. The ratio index for budgeted learning, with applications. In *SODA '09: Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 18–27, Philadelphia, PA, USA, 2009. Society for Industrial and Applied Mathematics.
- [GM07a] Sudipto Guha and Kamesh Munagala. Approximation algorithms for budgeted learning problems. In *STOC'07—Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, pages 104–113. ACM, New York, 2007. Full version as *Sequential Design of Experiments via Linear Programming*, <http://arxiv.org/abs/0805.2630v1>.
- [GM07b] Sudipto Guha and Kamesh Munagala. Model-driven optimization using adaptive probes. In *SODA '07: Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 308–317, Philadelphia, PA, USA, 2007. Society for Industrial and Applied Mathematics. Full version as *Adaptive Uncertainty Resolution in Bayesian Combinatorial Optimization Problems*, <http://arxiv.org/abs/0812.1012v1>.
- [GM09] Sudipto Guha and Kamesh Munagala. Multi-armed bandits with metric switching costs. In *ICALP*, pages 496–507, 2009.

- [GMP11] Sudipto Guha, Kamesh Munagala, and Martin Pal. Iterated allocations with delayed feedback. *ArXiv*, arxiv:abs/1011.1161, 2011.
- [GMS07] Sudipto Guha, Kamesh Munagala, and Peng Shi. On index policies for restless bandit problems. *CoRR*, abs/0711.3861, 2007. <http://arxiv.org/abs/0711.3861>. Full version of *Approximation algorithms for partial-information based stochastic control with Markovian rewards* (FOCS’07), and *Approximation algorithms for restless bandit problems*, (SODA’09).
- [KRT00] Jon Kleinberg, Yuval Rabani, and Éva Tardos. Allocating bandwidth for bursty connections. *SIAM J. Comput.*, 30(1):191–217 (electronic), 2000.
- [MSU99] Rolf H. Möhring, Andreas S. Schulz, and Marc Uetz. Approximation in stochastic scheduling: the power of lp-based priority policies. *Journal of the ACM (JACM)*, 46(6):924–942, 1999.
- [Pin95] Michael Pinedo. *Scheduling: Theory, Algorithms, and Systems*. Prentice Hall, 1995.
- [SU01] Martin Skutella and Marc Uetz. Scheduling precedence-constrained jobs with stochastic processing times on parallel machines. In *Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*, pages 589–590. Society for Industrial and Applied Mathematics, 2001.

A Some Bad Examples

A.1 Badness Due to Cancellations

We first observe that the LP relaxation for the Stock problem used in [DGV08] has a large integrality gap in the model where cancellations are allowed, *even when the rewards are fixed for any item*. This was also noted in [Dea05]. Consider the following example: there are n items, every item instantiates to a size of 1 with probability 0.5 or a size of $n/2$ with probability 0.5, and its reward is always 1. Let the total size of the knapsack be $B = n$. For such an instance, a good solution would cancel any item that does not terminate at size 1; this way, it can collect a reward of at least $n/2$ in expectation, because an average of $n/2$ items will instantiate with a size 1 and these will all contribute to the reward. On the other hand, the LP from [DGV08] has value $O(1)$, since the mean size of any item is at least $n/4$. In fact, any strategy that does not cancel jobs will also accrue only $O(1)$ reward.

A.2 Badness Due to Correlated Rewards

While the LP relaxations used for MAB (e.g., the formulation in [GM07a]) can handle the issue explained above w.r.t cancellations, we now present an example of stochastic knapsack (where the reward is correlated with the actual size) for which the existing MAB LP formulations all have a large integrality gap.

Consider the following example: there are n items, every item instantiates to a size of 1 with probability $1 - 1/n$ or a size of n with probability $1/n$, and its reward is 1 only if its size is n , and 0 otherwise. Let the total size of the knapsack be $B = n$. Clearly, any integral solution can fetch an expected reward of $1/n$ — if the first item it schedules instantiates to a large size, then it gives us a reward. Otherwise, no subsequent item can be fit within our budget even if it instantiates to its large size. The issue with the existing LPs is that the *arm-pull* constraints are ensured locally, and there is one global budget. That is, even if we play each arm to completion individually, the expected size (i.e., number of pulls) they occupy is $1 \cdot (1 - 1/n) + n \cdot (1/n) \leq 2$. Therefore, such LPs can accommodate $n/2$ jobs, fetching a total reward of $\Omega(1)$. This example brings to attention the fact that all these items are competing to be pulled in the first time slot (if we begin an item in any later time slot it fetches zero reward), thus naturally motivating our time-indexed LP formulation in Section 3.2.

In fact, the above example also shows that if we allow ourselves a budget of $2B$, i.e., $2n$ in this case, we can in fact achieve an expected reward of $O(1)$ (much higher than what is possible with a budget of B) — keep playing all items one by one, until one of them does not step after size 1 and then play that to completion; this event happens with probability $\Omega(1)$.

A.3 Badness Due to the Non-Martingale Property in MAB: The Benefit of Preemption

Not only do cancelations help in our problems (as can be seen from the example in Appendix A.1), we now show that even *preemption* is necessary in the case of MAB where the rewards do not satisfy the martingale property. In fact, this brings forward another key difference between our rounding scheme and earlier algorithms for MAB—the necessity of preempting arms is not an artifact of our algorithm/analysis but, rather, is unavoidable.

Consider the following instance. There are n identical arms, each of them with the following (recursively defined) transition tree starting at $\rho(0)$:

When the root $\rho(j)$ is pulled for $j < m$, the following two transitions can happen:

- (i) with probability $1/(n \cdot n^{m-j})$, the arm transitions to the “right-side”, where if it makes $B - n(\sum_{k=0}^j L^k)$ plays, it will deterministically reach a state with reward n^{m-j} . All intermediate states have 0 reward.
- (ii) with probability $1 - 1/(n \cdot n^{m-j})$, the arm transitions to the “left-side”, where if it makes $L^{j+1} - 1$ plays, it will deterministically reach the state $\rho(j+1)$. No state along this path fetches any reward.

Finally, node $\rho(m)$ makes the following transitions when played: (i) with probability $1/n$, to a leaf state that has a reward of 1 and the arm ends there; (ii) with probability $1 - 1/n$, to a leaf state with reward of 0.

For the following calculations, assume that $B \gg L > n$ and $m \gg 0$.

Preempting Solutions. We first exhibit a preempting solution with expected reward $\Omega(m)$. The strategy plays $\rho(0)$ of all the arms until one of them transitions to the “right-side”, in which case it continues to play this until it fetches a reward of n^m . Notice that any root which transitioned to the right-side can be played to completion, because the number of pulls we have used thus far is at most n (only those at the $\rho(0)$ nodes for each arm), and the size of the right-side is exactly $B - n$. Now, if all the arms transitioned to the left-side, then it plays the $\rho(1)$ of each arm until one of them transitioned to the right-side, in which case it continues playing this arm and gets a reward of n^{m-1} . Again, any root $\rho(1)$ which transitioned to the right-side *can be played* to completion, because the number of pulls we have used thus far is at most $n(1 + L)$ (for each arm, we have pulled the root $\rho(0)$, transitioned the walk of length $L - 1$ to $\rho(1)$ and then pulled $\rho(1)$), and the size of the right-side is exactly $B - n(1 + L)$. This strategy is similarly defined, recursively.

We now calculate the expected reward: if any of the roots $\rho(0)$ made a transition to the right-side, we get a reward of n^m . This happens with probability roughly $1/n^m$, giving us an expected reward of 1 in this case. If all the roots made the transition to the left-side, then at least one of the $\rho(1)$ states will make a transition to their right-side with probability $\approx 1/n^{m-1}$ in which case will will get reward of n^{m-1} , and so on. Thus, summing over the first $m/2$ such rounds, our expected reward is at least

$$\frac{1}{n^m} n^m + \left(1 - \frac{1}{n^m}\right) \frac{1}{n^{m-1}} n^{m-1} + \left(1 - \frac{1}{n^m}\right) \left(1 - \frac{1}{n^{m-1}}\right) \frac{1}{n^{m-2}} n^{m-2} + \dots$$

Each term above is $\Omega(1)$ giving us a total of $\Omega(m)$ expected reward.

Non-Preempting Solutions. Consider any non-preempting solution. Once it has played the first node of an arm and it has transitioned to the left-side, it has to irrevocably decide if it abandons this arm or continues playing. But if it has continued to play (and made the transition of $L - 1$ steps), then it cannot get any reward from the right-side of $\rho(0)$ of any of the other arms, because $L > n$ and the right-side requires $B - n$ pulls before reaching a reward-state. Likewise, if it has decided to move from $\rho(i)$ to $\rho(i+1)$ on any arm, it cannot get *any* reward from the right-sides of $\rho(0), \rho(1), \dots, \rho(i)$ on *any* arm due to budget constraints. Indeed, for any $i \geq 1$, to have reached $\rho(i+1)$ on any particular arm, it must have utilized $(1 + L - 1) + (1 + L^2 - 1) + \dots + (1 + L^{i+1} - 1)$ pulls in total, which exceeds $n(1 + L + L^2 + \dots + L^i)$ since $L > n$. Finally, notice that if the strategy has decided to move from $\rho(i)$ to $\rho(i+1)$ on any arm, the maximum reward that it can obtain is n^{m-i-1} , namely, the reward from the right-side transition of $\rho(i+1)$.

Using these properties, we observe that an optimal non-preempting strategy proceeds in rounds as described next.

Strategy at round i . Choose a set N_i of n_i available arms and play them as follows: pick one of these arms, play until reaching state $\rho(i)$ and then play once more. If there is a right-side transition before reaching state $\rho(i)$, discard this arm since there is not enough budget to play until reaching a state with positive reward. If there is a right-side transition at state $\rho(i)$, play this arm until it gives reward of n^{m-i} . If there is no right-side transition and there is another arm in N_i which is still to be played, discard the current arm and pick the next arm in N_i .

In round i , at least $\max(0, n_i - 1)$ arms are discarded, hence $\sum_i n_i \leq 2n$. Therefore, the expected reward can be at most

$$\frac{n_1}{n \cdot n^m} n^m + \frac{n_2}{n \cdot n^{m-1}} n^{m-1} + \dots + \frac{n_m}{n} \leq 2$$

B Proofs from Section 2

B.1 Proof of Theorem 2.3

Let add_i denote the event that item i was added to the knapsack in Step 5. Also, let V_i denote the random variable corresponding to the reward that our algorithm gets from item i .

Clearly if item i has $D_i = t$ and was added, then it is added to the knapsack before time t . In this case it is easy to see that $\mathbb{E}[V_i \mid \text{add}_i \wedge (D_i = t)] \geq R_{i,t}$ (because its random size is independent of when the algorithm started it). Moreover, from the previous lemma we have that $\Pr(\text{add}_i \mid (D_i = t)) \geq 1/2$ and from Step 1 we have $\Pr(D_i = t) = \frac{x_{i,t}^*}{4}$; hence $\Pr(\text{add}_i \wedge (D_i = t)) \geq x_{i,t}^*/8$. Finally adding over all possibilities of t , we lower bound the expected value of V_i by

$$\mathbb{E}[V_i] \geq \sum_t \mathbb{E}[V_i \mid \text{add}_i \wedge (D_i = t)] \cdot \Pr(\text{add}_i \wedge (D_i = t)) \geq \frac{1}{8} \sum_t x_{i,t}^* R_{i,t}.$$

Finally, linearity of expectation over all items shows that the total expected reward of our algorithm is at least $\frac{1}{8} \cdot \sum_{i,t} x_{i,t}^* R_{i,t} = \text{LPOpt}/8$, thus completing the proof.

B.2 Making Stock-NoCancel Fully Polynomial

Recall that our LP relaxation $\text{LP}_{\text{NoCancel}}$ in Section 2 uses a global time-indexed LP. In order to make it compact, our approach will be to group the B timeslots in $\text{LP}_{\text{NoCancel}}$ and show that the grouped LP has optimal value within constant factor of $\text{LP}_{\text{NoCancel}}$; furthermore, we show also that it can be rounded and analyzed almost identically to the original LP. To this end, consider the following LP relaxation:

$$\max \sum_i \sum_{j=0}^{\log B} \mathbb{E} R_{i,2^{j+1}} \cdot x_{i,2^j} \quad (\text{PolyLP}_L)$$

$$\sum_{j=0}^{\log B} x_{i,2^j} \leq 1 \quad \forall i \quad (\text{B.23})$$

$$\sum_{i,j' \leq j} x_{i,2^{j'}} \cdot \mathbb{E}[\min(S_i, 2^{j+1})] \leq 2 \cdot 2^j \quad \forall j \in [0, \log B] \quad (\text{B.24})$$

$$x_{i,2^j} \in [0, 1] \quad \forall j \in [0, \log B], \forall i \quad (\text{B.25})$$

The next two lemmas relate the value of (PolyLP_L) to that of the original LP ($\text{LP}_{\text{NoCancel}}$).

Lemma B.1 *The optimum of (PolyLP_L) is at least half of the optimum of $(\text{LP}_{\text{NoCancel}})$.*

Proof. Consider a solution x for $(\text{LP}_{\text{NoCancel}})$ and define $\bar{x}_{i,1} = x_{i,1}/2 + \sum_{t \in [2,4]} x_{i,t}/2$ and $\bar{x}_{i,2^j} = \sum_{t \in [2^{j+1}, 2^{j+2}]} x_{i,t}/2$ for $1 < j \leq \log B$. It suffices to show that \bar{x} is a feasible solution to (PolyLP_L) with value greater than or equal to half of the value of x .

For constraints (B.23) we have $\sum_{j=0}^{\log B} \bar{x}_{i,2^j} = \sum_{t \geq 1} x_{i,t}/2 \leq 1/2$; these constraints are therefore easily satisfied. We now show that $\{\bar{x}\}$ also satisfies constraints (B.24):

$$\begin{aligned} \sum_{i,j' \leq j} x_{i,2^{j'}} \cdot \mathbb{E}[\min(S_i, 2^{j+1})] &= \sum_i \sum_{t=1}^{2^{j+2}-1} \frac{x_{i,t} \mathbb{E}[\min(S_i, 2^{j+1})]}{2} \\ &\leq \sum_i \sum_{t=1}^{2^{j+2}-1} \frac{x_{i,t} \mathbb{E}[\min(S_i, 2^{j+2} - 1)]}{2} \leq 2^{j+2} - 1, \end{aligned}$$

where the last inequality follows from feasibility of $\{x\}$.

Finally, noticing that $\text{ER}_{i,t}$ is non-increasing with respect to t , it is easy to see that $\sum_i \sum_{j=0}^{\log B} \text{ER}_{i,2^{j+1}} \cdot \bar{x}_{i,2^j} \geq \sum_{i,t} \text{ER}_{i,t} \cdot x_{i,t}/2$ and hence \bar{x} has value greater than or equal to half of the value of x as desired. ■

Lemma B.2 *Let $\{\bar{x}\}$ be a feasible solution for (PolyLP_L) . Define $\{\hat{x}\}$ satisfying $\hat{x}_{i,t} = \bar{x}_{i,2^j}/2^j$ for all $t \in [2^j, 2^{j+1})$ and $i \in [n]$. Then $\{\hat{x}\}$ is feasible for $(\text{LP}_{\text{NoCancel}})$ and has value at least as large as $\{\bar{x}\}$.*

Proof. The feasibility of $\{\bar{x}\}$ directly implies that $\{\hat{x}\}$ satisfies constraints (2.1). For constraints (2.2), consider $t \in [2^j, 2^{j+1})$; then we have the following:

$$\begin{aligned} \sum_{i,t' \leq t} \hat{x}_{i,t'} \cdot \mathbb{E}[\min(S_i, t)] &\leq \sum_i \sum_{j' \leq j} \sum_{t \in [2^{j'}, 2^{j'+1})} \frac{\bar{x}_{i,2^{j'}}}{2^{j'}} \mathbb{E}[\min(S_i, 2^{j+1})] \\ &= \sum_i \sum_{j' \leq j} \bar{x}_{i,2^{j'}} \mathbb{E}[\min(S_i, 2^{j+1})] \leq 2 \cdot 2^j \leq 2t. \end{aligned}$$

Finally, again using the fact that $\text{ER}_{i,t}$ is non-increasing in t we get that the value of $\{\hat{x}\}$ is

$$\sum_{i,t} \text{ER}_{i,t} \cdot \hat{x}_{i,t} = \sum_i \sum_{j=0}^{\log B} \sum_{t \in [2^j, 2^{j+1})} \text{ER}_{i,t} \frac{\bar{x}_{i,2^j}}{2^j} \geq \sum_i \sum_{j=0}^{\log B} \sum_{t \in [2^j, 2^{j+1})} \text{ER}_{i,2^{j+1}} \frac{\bar{x}_{i,2^j}}{2^j} = \sum_i \sum_{j=0}^{\log B} \text{ER}_{i,2^{j+1}} \bar{x}_{i,2^j},$$

which is then at least as large as the value of $\{\bar{x}\}$. This concludes the proof of the lemma. ■

The above two lemmas show that the PolyLP_L has value close to that of $\text{LP}_{\text{NoCancel}}$: let's now show that we can simulate the execution of Algorithm **Stock-Large** just given an optimal solution $\{\bar{x}\}$ for (PolyLP_L) . Let $\{\hat{x}\}$ be defined as in the above lemma, and consider the Algorithm **Stock-Large** applied to $\{\hat{x}\}$. By the definition of $\{\hat{x}\}$, here's how to execute Step 1 (and hence the whole algorithm) in polynomial time: we obtain $D_i = t$ by picking $j \in [0, \log B]$ with probability $\bar{x}_{i,2^j}$ and then selecting $t \in [2^j, 2^{j+1})$ uniformly; notice that indeed $D_i = t$ (with $t \in [2^j, 2^{j+1})$) with probability $\bar{x}_{i,2^j}/2^j = \hat{x}_{i,t}$.

Using this observation we can obtain a $1/16$ approximation for our instance \mathcal{I} in polynomial time by finding the optimal solution $\{\bar{x}\}$ for (PolyLP_L) and then running Algorithm **Stock-Large** over $\{\hat{x}\}$ as described in the previous paragraph. Using a direct modification of Theorem 2.3 we have that the strategy obtained has expected reward at least as large as $1/8$ of the value of $\{\hat{x}\}$, which by Lemmas B.1 and B.2 (and Lemma 2.1) is within a factor of $1/16$ of the optimal solution for \mathcal{I} .

C Proofs from Section 3

C.1 Proof of Lemma 3.2

The proof works by induction. For the base case, consider $t = 0$. Clearly, this item is forcefully canceled in step 4 of Algorithm 3.1 **Stock-Small** (in the iteration with $t = 0$) with probability $s_{i,0}^*/v_{i,0}^* - \pi_{i,0}/\sum_{t \geq 0} \pi_{i,t}$.

But since $\pi_{i,0}$ was assumed to be 0 and $v_{i,0}^*$ is 1, this quantity is exactly $s_{i,0}^*$, and this proves property (i). For property (ii), item i is processed for its 1^{st} timestep if it did not get forcefully canceled in step 4. This therefore happens with probability $1 - s_{i,0}^* = v_{i,0}^* - s_{i,0}^* = v_{i,1}^*$. For property (iii), conditioned on the fact that it has been processed for its 1^{st} timestep, clearly the probability that its (unknown) size has instantiated to 1 is exactly $\pi_{i,1} / \sum_{t' \geq 1} \pi_{i,t'}$. When this happens, the job stops in step 7, thereby establishing the base case.

Assuming this property holds for every timestep until some fixed value $t - 1$, we show that it holds for t ; the proofs are very similar to the base case. Assume item i was processed for the t^{th} timestep (this happens w.p $v_{i,t}^*$ from property (ii) of the induction hypothesis). Then from property (iii), the probability that this item completes at this timestep is exactly $\pi_{i,t} / \sum_{t' \geq t} \pi_{i,t'}$. Furthermore, it gets forcefully canceled in step 4 with probability $s_{i,t}^* / v_{i,t}^* - \pi_{i,t} / \sum_{t' \geq t} \pi_{i,t'}$. Thus the total probability of stopping at time t , assuming it has been processed for its t^{th} timestep is exactly $s_{i,t}^* / v_{i,t}^*$; unconditionally, the probability of stopping at time t is hence $s_{i,t}^*$.

Property (ii) follows as a consequence of Property (i), because the item is processed for its $(t + 1)^{st}$ timestep only if it did not stop at timestep t . Therefore, conditioned on being processed for the t^{th} timestep, it continues to be processed with probability $1 - s_{i,t}^* / v_{i,t}^*$. Therefore, removing the conditioning, we get the probability of processing the item for its $(t + 1)^{st}$ timestep is $v_{i,t}^* - s_{i,t}^* = v_{i,t+1}^*$. Finally, for property (iii), conditioned on the fact that it has been processed for its $(t + 1)^{st}$ timestep, clearly the probability that its (unknown) size has instantiated to exactly $(t + 1)$ is $\pi_{i,t+1} / \sum_{t' \geq t+1} \pi_{i,t'}$. When this happens, the job stops in step 7 of the algorithm.

C.2 Stock with Small Sizes: A Fully Polytime Algorithm

The idea is to quantize the possible sizes of the items in order to ensure that LP LP_S has polynomial size, then obtain a good strategy (via Algorithm Stock-Small) for the transformed instance, and finally to show that this strategy is actually almost as good for the original instance.

Consider an instance $\mathcal{I} = (\pi, R)$ where $R_{i,t} = 0$ for all $t > B/2$. Suppose we start scheduling an item at some time; instead of making decisions of whether to continue or cancel an item at each subsequent time step, we are going to do it in time steps which are powers of 2. To make this formal, define instance $\bar{\mathcal{I}} = (\bar{\pi}, \bar{R})$ as follows: set $\bar{\pi}_{i,2^j} = \sum_{t \in [2^j, 2^{j+1})} \pi_{i,t}$ and $\bar{R}_{i,2^j} = (\sum_{t \in [2^j, 2^{j+1})} \pi_{i,t} R_{i,t}) / \bar{\pi}_{i,2^j}$ for all $i \in [n]$ and $j \in \{0, 1, \dots, \lfloor \log B \rfloor\}$. The instances are coupled in the natural way: the size of item i in the instance $\bar{\mathcal{I}}$ is 2^j iff the size of item i in the instance \mathcal{I} lies in the interval $[2^j, 2^{j+1})$.

In Section 3.1, a *timestep* of an item has duration of 1 time unit. However, due to the construction of $\bar{\mathcal{I}}$, it is useful to consider that the t^{th} time step of an item has duration 2^t ; thus, an item can only complete at its 0^{th} , 1^{st} , 2^{nd} , etc. timesteps. With this in mind, we can write an LP analogous to (LP_S) :

$$\max \sum_{1 \leq j \leq \log(B/2)} \sum_{1 \leq i \leq n} v_{i,2^j} \cdot \bar{R}_{i,2^j} \frac{\bar{\pi}_{i,2^j}}{\sum_{j' \geq j} \bar{\pi}_{i,2^{j'}}} \quad (\text{PolyLP}_S)$$

$$v_{i,2^j} = s_{i,2^j} + v_{i,2^{j+1}} \quad \forall j \in [0, \log B], i \in [n] \quad (\text{C.26})$$

$$s_{i,2^j} \geq \frac{\bar{\pi}_{i,2^j}}{\sum_{j' \geq j} \bar{\pi}_{i,2^{j'}}} \cdot v_{i,2^j} \quad \forall t \in [0, \log B], i \in [n] \quad (\text{C.27})$$

$$\sum_{i \in [n]} \sum_{j \in [0, \log B]} 2^j \cdot s_{i,2^j} \leq B \quad (\text{C.28})$$

$$v_{i,0} = 1 \quad \forall i \quad (\text{C.29})$$

$$v_{i,2^j}, s_{i,2^j} \in [0, 1] \quad \forall j \in [0, \log B], i \in [n] \quad (\text{C.30})$$

Notice that this LP has size polynomial in the size of the instance \mathcal{I} .

Consider the LP (LP_S) with respect to the instance \mathcal{I} and let (v, s) be a feasible solution for it with objective value z . Then define (\bar{v}, \bar{s}) as follows: $\bar{v}_{i,2^j} = v_{i,2^j}$ and $\bar{s}_{i,2^j} = \sum_{t \in [2^j, 2^{j+1})} s_{i,t}$. It is easy to check that (\bar{v}, \bar{s}) is a feasible solution for (PolyLP_S) with value at least z , where the latter uses the fact that $v_{i,t}$ is non-increasing in t . Using Theorem 3.1 it then follows that the optimum of (PolyLP_S) with respect to $(\bar{\pi}, \bar{R})$ is at least as large as the reward obtained by the optimal solution for the stochastic knapsack instance (π, R) .

Let (\bar{v}, \bar{s}) denote an optimal solution of (PolyLP_S) . Notice that with the redefined notion of timesteps we can naturally apply Algorithm **Stock-Small** to the LP solution (\bar{v}, \bar{s}) . Moreover, Lemma 3.2 still holds in this setting. Finally, modify Algorithm **Stock-Small** by ignoring items with probability $1 - 1/8 = 7/8$ (instead of $3/4$) in Step 2 (we abuse notation slightly and shall refer to the modified algorithm also as **Stock-Small**) and notice that Lemma 3.2 still holds.

Consider the strategy $\bar{\mathbb{S}}$ for $\bar{\mathcal{I}}$ obtained from Algorithm **Stock-Small**. We can obtain a strategy \mathbb{S} for \mathcal{I} as follows: whenever \mathbb{S} decides to process item i of $\bar{\mathcal{I}}$ for its j th timestep, we decide to continue item i of \mathcal{I} while it has size from 2^j to $2^{j+1} - 1$.

Lemma C.1 *Strategy \mathbb{S} is a $1/16$ approximation for \mathcal{I} .*

Proof. Consider an item i . Let \bar{O} be the random variable denoting the total size occupied before strategy $\bar{\mathbb{S}}$ starts processing item i and similarly let O denote the total size occupied before strategy \mathbb{S} starts processing item i . Since Lemma 3.2 still holds for the modified algorithm **Stock-Small**, we can proceed as in Theorem 3.3 and obtain that $\mathbb{E}[\bar{O}] \leq B/8$. Due to the definition of \mathbb{S} we can see that $O \leq 2\bar{O}$ and hence $\mathbb{E}[O] \leq B/4$. From Markov's inequality we obtain that $\Pr(O \geq B/2) \leq 1/2$. Noticing that i is started by \mathbb{S} with probability $1/8$ we get that the probability that i is started and there is at least $B/2$ space left on the knapsack at this point is at least $1/16$. Finally, notice that in this case $\bar{\mathbb{S}}$ and \mathbb{S} obtain the same expected value from item i , namely $\sum_j \bar{v}_{i,2^j} \cdot \bar{R}_{i,2^j} \frac{\bar{\pi}_{i,2^j}}{\sum_{j' \geq j} \bar{\pi}_{i,2^{j'}}$. Thus \mathbb{S} get expected value at least that of the optimum of (PolyLP_S) , which is at least the value of the optimal solution for \mathcal{I} as argued previously. ■

D Details from Section 4

D.1 Details of Phase I (from Section 4.2.1)

We first begin with some notation that will be useful in the algorithm below. For any state $u \in \mathcal{S}_i$ such that the path from ρ_i to u follows the states $u_1 = \rho_i, u_2, \dots, u_k = u$, let $\pi_u = \prod_{l=1}^{k-1} p_{u_l, u_{l+1}}$.

Fix an arm i , for which we will perform the decomposition. Let $\{z, w\}$ be a feasible solution to LP_{mab} and set $z_{u,t}^0 = z_{u,t}$ and $w_{u,t}^0 = w_{u,t}$ for all $u \in \mathcal{S}_i, t \in [B]$. We will gradually alter the fractional solution as we build the different forests. We note that in a particular iteration with index j , all z^{j-1}, w^{j-1} values that are not updated in Steps 12 and 13 are retained in z^j, w^j respectively. For brevity of notation, we shall use “iteration j of step 2” to

Algorithm D.1 Convex Decomposition of Arm i

- 1: **set** $\mathcal{C}_i \leftarrow \emptyset$ and **set loop index** $j \leftarrow 1$.
 - 2: **while** \exists a node $u \in \mathcal{S}_i$ s.t. $\sum_t z_{u,t}^{j-1} > 0$ **do**
 - 3: **initialize** a new tree $\mathbb{T}(i, j) = \emptyset$.
 - 4: **set** $A \leftarrow \{u \in \mathcal{S}_i \text{ s.t. } \sum_t z_{u,t}^{j-1} > 0\}$.
 - 5: for all $u \in \mathcal{S}_i$, **set** $\text{time}(i, j, u) \leftarrow \infty$, $\text{prob}(i, j, u) \leftarrow 0$, and **set** $\epsilon_u \leftarrow \infty$.
 - 6: **for** every $u \in A$ **do**
 - 7: **update** $\text{time}(i, j, u)$ to the smallest time t s.t. $z_{u,t}^{j-1} > 0$.
 - 8: **update** $\epsilon_u = z_{u, \text{time}(i, j, u)}^{j-1} / \pi_u$
 - 9: **let** $\epsilon = \min_u \epsilon_u$.
 - 10: **for** every $u \in A$ **do**
 - 11: **set** $\text{prob}(i, j, u) = \epsilon \cdot \pi_u$.
 - 12: **update** $z_{u, \text{time}(i, j, u)}^j = z_{u, \text{time}(i, j, u)}^{j-1} - \text{prob}(i, j, u)$.
 - 13: **update** $w_{v, \text{time}(i, j, u)+1}^j = w_{v, \text{time}(i, j, u)+1}^{j-1} - \text{prob}(i, j, u) \cdot p_{u,v}$ for all v s.t. $\text{parent}(v) = u$.
 - 14: **set** $\mathcal{C}_i \leftarrow \mathcal{C}_i \cup \mathbb{T}(i, j)$.
 - 15: **increment** $j \leftarrow j + 1$.
-

denote the execution of the entire block (steps 3 – 14) which constructs strategy forest $\mathbb{T}(i, j)$.

Lemma D.1 Consider an integer j and suppose that $\{z^{j-1}, w^{j-1}\}$ satisfies constraints (4.10)-(4.12) of LP_{mab} . Then after iteration j of Step 2, the following properties hold:

- (a) $\mathbb{T}(i, j)$ (along with the associated $\text{prob}(i, j, \cdot)$ and $\text{time}(i, j, \cdot)$ values) is a valid strategy forest, i.e., satisfies the conditions (i) and (ii) presented in Section 4.2.1.
- (b) The residual solution $\{z^j, w^j\}$ satisfies constraints (4.10)-(4.12).
- (c) For any time t and state $u \in \mathcal{S}_i$, $z_{u,t}^{j-1} - z_{u,t}^j = \text{prob}(i, j, u) \mathbf{1}_{\text{time}(i,j,u)=t}$.

Proof. We show the properties stated above one by one.

Property (a): We first show that the time values satisfy $\text{time}(i, j, u) \geq \text{time}(i, j, \text{parent}(u)) + 1$, i.e. condition (i) of strategy forests. For sake of contradiction, assume that there exists $u \in \mathcal{S}_i$ with $v = \text{parent}(u)$ where $\text{time}(i, j, u) \leq \text{time}(i, j, v)$. Define $t_u = \text{time}(i, j, u)$ and $t_v = \text{time}(i, j, \text{parent}(u))$; the way we updated $\text{time}(i, j, u)$ in step 7 gives that $z_{u,t_u}^{j-1} > 0$.

Then, constraint (4.11) of the LP implies that $\sum_{t' \leq t_u} w_{u,t'}^{j-1} > 0$. In particular, there exists a time $t' \leq t_u \leq t_v$ such that $w_{u,t'}^{j-1} > 0$. But now, constraint (4.10) enforces that $z_{v,t'-1}^{j-1} = w_{u,t'}^{j-1} / p_{v,u} > 0$ as well. But this contradicts the fact that t_v was the first time s.t $z_{v,t}^{j-1} > 0$. Hence we have $\text{time}(i, j, u) \geq \text{time}(i, j, \text{parent}(u)) + 1$.

As for condition (ii) about $\text{prob}(i, j, \cdot)$, notice that if $\text{time}(i, j, u) \neq \infty$, then $\text{prob}(i, j, u)$ is set to $\epsilon \cdot \pi_u$ in step 11. It is now easy to see from the definition of π_u (and from the fact that $\text{time}(i, j, u) \neq \infty \Rightarrow \text{time}(i, j, \text{parent}(u)) \neq \infty$) that $\text{prob}(i, j, u) = \text{prob}(i, j, \text{parent}(u)) \cdot p_{\text{parent}(u),u}$.

Property (b): Constraint (4.10) of LP_{mab} is clearly satisfied by the new LP solution $\{z^j, w^j\}$ because of the two updates performed in Steps 12 and 13: if we decrease the z value of any node at any time, the w of all children are appropriately reduced (for the subsequent timestep).

Before showing that the solution $\{z^j, w^j\}$ satisfies constraint (4.11), we first argue that they remain non-negative. By the choice of ϵ in step 9, we have $\text{prob}(i, j, u) = \epsilon \pi_u \leq \epsilon_u \pi_u = z_{u, \text{time}(i,j,u)}^{j-1}$ (where ϵ_u was computed in Step 8); consequently even after the update in step 12, $z_{u, \text{time}(i,j,u)}^j \geq 0$ for all u . This and the fact that the constraints (4.10) are satisfied implies that $\{z^j, w^j\}$ satisfies the non-negativity requirement.

We now show that constraint (4.11) is satisfied. For any time t and state $u \notin A$ (where A is the set computed in step 4 for iteration j), clearly it must be that $\sum_{t' \leq t} z_{u,t'}^{j-1} = 0$ by definition of the set A ; hence just the non-negativity of w^j implies that these constraints are trivially satisfied.

Therefore consider some $t \in [B]$ and a state $u \in A$. We know from step 7 that $\text{time}(i, j, u) \neq \infty$. If $t < \text{time}(i, j, u)$, then the way $\text{time}(i, j, u)$ is updated in step 7 implies that $\sum_{t' \leq t} z_{u,t'}^j = \sum_{t' \leq t} z_{u,t'}^{j-1} = 0$, so the constraint is trivially satisfied because w^j is non-negative. If $t \geq \text{time}(i, j, u)$, we claim that the change in the left hand side and right hand side (between the solutions $\{z^{j-1}, w^{j-1}\}$ and $\{z^j, w^j\}$) of the constraint under consideration is the same, implying that it will be still satisfied by $\{z^j, w^j\}$.

To prove this claim, observe that the right hand side has decreased by exactly $z_{u, \text{time}(i,j,u)}^{j-1} - z_{u, \text{time}(i,j,u)}^j = \text{prob}(i, j, u)$. But the only value which has been modified in the left hand side is $w_{u, \text{time}(i,j, \text{parent}(u))+1}^{j-1}$, which has gone down by $\text{prob}(i, j, \text{parent}(u)) \cdot p_{\text{parent}(u),u}$. Because $\mathbb{T}(i, j)$ forms a valid strategy forest, we have $\text{prob}(i, j, u) = \text{prob}(i, j, \text{parent}(u)) \cdot p_{\text{parent}(u),u}$, and thus the claim follows.

Finally, constraint (4.12) are also satisfied as the z variables only decrease in value over iterations.

Property (c): This is an immediate consequence of the Step 12. ■

To prove Lemma 4.2, firstly notice that since $\{z^0, w^0\}$ satisfies constraints (4.10)-(4.12), we can proceed by induction and infer that the properties in the previous lemma hold for every strategy forest in the decomposition; in particular, each of them is a valid strategy forest.

In order to show that the marginals are preserved, observe that in the last iteration j^* of procedure we have $z_{u,t}^{j^*} = 0$ for all u, t . Therefore, adding the last property in the previous lemma over all j gives

$$z_{u,t} = \sum_{j \geq 1} (z_{u,t}^{j-1} - z_{u,t}^j) = \sum_{j \geq 1} \text{prob}(i, j, u) \mathbf{1}_{\text{time}(i,j,u)=t} = \sum_{j: \text{time}(i,j,u)=t} \text{prob}(i, j, u).$$

Finally, since some $z_{u,t}^j$ gets altered to 0 since in each iteration of the above algorithm, the number of strategies for each arm in the decomposition is upper bounded by $B|\mathcal{S}|$. This completes the proof of Lemma 4.2.

D.2 Details of Phase II (from Section 4.2.2)

Proof of Lemma 4.4: Let $\text{time}^t(u)$ denote the time assigned to node u by the end of round $\tau = t$ of the algorithm; $\text{time}^{B+1}(u)$ is the initial time of u . Since the algorithm works backwards in time, our round index will start at B and end up at 1. To prove property (i) of the statement of the lemma, notice that the algorithm only converts head nodes to non-head nodes and not the other way around. Moreover, heads which survive the algorithm have the same time as originally. So it suffices to show that heads which originally did not satisfy property (i)—namely, those with $\text{time}^{B+1}(v) < 2 \cdot \text{depth}(v)$ —do not survive the algorithm; but this is clear from the definition of Step 2.

To prove property (ii), fix a time t , and consider the execution of **GapFill** at the end of round $\tau = t$. We claim that the total extent of fractional play at time t does not increase as we continue the execution of the algorithm from round $\tau = t$ to round 1. To see why, let C be a connected component at the end of round $\tau = t$ and let h denote its head. If $\text{time}^t(h) > t$ then no further **advance** affects C and hence it does not contribute to an increase in the number of plays at time t . On the other hand, if $\text{time}^t(h) \leq t$, then even if C is advanced in a subsequent round, each node w of C which ends up being played at t , i.e., has $\text{time}^1(w) = t$ must have an ancestor w' satisfying $\text{time}^t(w') = t$, by the contiguity of C . Thus, Observation 4.3 gives that $\sum_{u \in C: \text{time}^1(u)=t} \text{prob}(u) \leq \sum_{u \in C: \text{time}^t(u)=t} \text{prob}(u)$. Applying this for each connected component C , proves the claim. Intuitively, any component which advances forward in time is only reducing its load/total fractional play at any fixed time t .

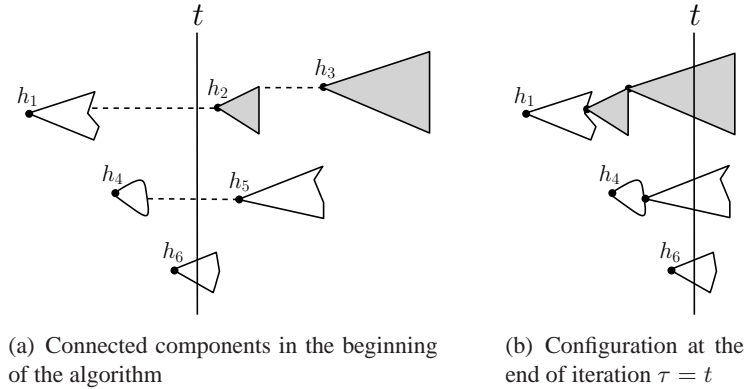


Figure D.5: Depiction of a strategy forest $\mathbb{T}(i, j)$ on a timeline, where each triangle is a connected component. In this example, $H = \{h_2, h_5\}$ and C_{h_2} consists of the grey nodes. From Observation 4.3 the number of plays at t do not increase as components are moved to the left.

Then consider the end of iteration $\tau = t$ and we now prove that the fractional extent of play at time t is at most 3. Due to Lemma 4.2, it suffices to prove that $\sum_{u \in U} \text{prob}(u) \leq 2$, where U is the set of nodes which caused an increase in the number of plays at time t , namely, $U = \{u : \text{time}^{B+1}(u) > t \text{ and } \text{time}^t(u) = t\}$.

Notice that a connected component of the original forest can only contribute to this increase if its head h crossed time t , that is $\text{time}^{B+1}(h) > t$ and $\text{time}^t(h) \leq t$. However, it may be that this crossing was not directly caused

by an **advance** on h (i.e. h advanced till $\text{time}^{B+1}(\text{parent}(h)) \geq t$), but an **advance** to a head h' in a subsequent round was responsible for h crossing over t . But in this case h must be part of the connected component of h' when the latter **advance** happens, and we can use h' 's advance to bound the congestion.

To make this more formal, let H be the set of heads of the original forest whose **advances** made them cross time t , namely, $h \in H$ iff $\text{time}^{B+1}(h) > t$, $\text{time}^t(h) \leq t$ and $\text{time}^{B+1}(\text{parent}(h)) < t$. Moreover, for $h \in H$ let C_h denote the connected component of h in the beginning of the iteration where an **advance** was executed on h , that is, when v was set to h in Step 3. The above argument shows that these components C_h 's contain all the nodes in U , hence it suffices to see how they increase the congestion at time t .

In fact, it is sufficient to focus just on the heads in H . To see this, consider $h \in H$ and notice that no node in $U \cap C_h$ is an ancestor of another. Then Observation 4.3 gives $\sum_{u \in U \cap C_h} \text{prob}(u) \leq \text{prob}(h)$, and adding over all h in H gives $\sum_{u \in U} \text{prob}(u) \leq \sum_{h \in H} \text{prob}(h)$.

To conclude the proof, we upper bound the right hand side of the previous inequality. The idea now is that the play probabilities on the nodes in H cannot be too large since their parents have $\text{time}^{B+1} < t$ (and each head has a large number of ancestors in $[1, t]$ because it was considered for an advance). More formally, fix i, j and consider a head h in $H \cap \mathbb{T}(i, j)$. From Step 2 of the algorithm, we obtain that $\text{depth}(h) > (1/2)\text{time}^{B+1}(h) \geq t/2$. Since $\text{time}^{B+1}(\text{parent}(h)) < t$, it follows that for every $d \leq \lfloor t/2 \rfloor$, h has an ancestor $u \in \mathbb{T}(i, j)$ with $\text{depth}(u) = d$ and $\text{time}^{B+1}(u) \leq t$. Moreover, the definition of H implies that no head in $H \cap \mathbb{T}(i, j)$ can be an ancestor of another. Then again employing Observation 4.3 we obtain

$$\sum_{h \in H \cap \mathbb{T}(i, j)} \text{prob}(h) \leq \sum_{u \in \mathbb{T}(i, j): \text{depth}(u)=d, \text{time}^{B+1}(u) \leq t} \text{prob}(u) \quad (\forall d \leq \lfloor t/2 \rfloor).$$

Adding over all i, j and $d \leq \lfloor t/2 \rfloor$ leads to the bound $(t/2) \cdot \sum_{h \in H} \text{prob}(h) \leq \sum_{u: \text{time}^{B+1}(u) \leq t} \text{prob}(u)$. Finally, using Lemma 4.2 we can upper bound the right hand side by t , which gives $\sum_{u \in U} \text{prob}(u) \leq \sum_{h \in H} \text{prob}(h) \leq 2$ as desired. \blacksquare

D.3 Details of Phase III (from Section 4.2.3)

Proof of Lemma 4.5: The proof is quite straightforward. Intuitively, it is because AlgMAB (Algorithm 4.2) simply follows the probabilities according to the transition tree T_i (unless $\text{time}(i, j, u) = \infty$ in which case it abandons the arm). Consider an arm i such that $\sigma(i) = j$, and any state $u \in \mathcal{S}_i$. Let $\langle v_1 = \rho_i, v_2, \dots, v_t = u \rangle$ denote the unique path in the transition tree for arm i from ρ_i to u . Then, if $\text{time}(i, j, u) \neq \infty$ the probability that state u is played is exactly the probability of the transitions reaching u (because in steps 8 and 9, the algorithm just keeps playing the states⁷ and making the transitions, unless $\text{time}(i, j, u) = \infty$). But this is precisely $\prod_{k=1}^{t-1} p_{v_k, v_{k+1}} = \text{prob}(i, j, u) / \text{prob}(i, j, \rho_i)$ (from the properties of each strategy in the convex decomposition). If $\text{time}(i, j, u) = \infty$ however, then the algorithm terminates the arm in Step 10 without playing u , and so the probability of playing u is $0 = \text{prob}(i, j, u) / \text{prob}(i, j, \rho_i)$. This completes the proof.

E Proofs from Section 5

E.1 Layered DAGs capture all Graphs

We first show that *layered DAGs* can capture all transition graphs, with a blow-up of a factor of B in the state space. For each arm i , for each state u in the transition graph \mathcal{S}_i , create B copies of it indexed by (v, t) for all $1 \leq t \leq B$. Then for each u and v such that $p_{u,v} > 0$ and for each $1 \leq t < B$, place an arc $(u, t) \rightarrow (v, t+1)$. Finally, delete all vertices that are not reachable from the state $(\rho_i, 1)$ where ρ_i is the starting state of arm i . There is a clear correspondence between the transitions in \mathcal{S}_i and the ones in this layered graph: whenever state u is played at time t and \mathcal{S}_i transitions to state v , we have the transition from (u, t) to $(v, t+1)$ in the layered DAG. Henceforth, we shall assume that the layered graph created in this manner is the transition graph for each arm.

⁷We remark that while the plays just follow the transition probabilities, they may not be made contiguously.

F MABs with Budgeted Exploitation

As we remarked before, we now explain how to generalize the argument from Section 4 to the presence of “exploits”. A strategy in this model needs to choose an arm in each time step and perform one of two actions: either it *pulls* the arm, which makes it transition to another state (this corresponds to *playing* in the previous model), or *exploits* it. If an arm is in state u and is exploited, it fetches reward r_u , and cannot be pulled any more. As in the previous case, there is a budget B on the total number of pulls that a strategy can make and an additional budget of K on the total number of exploits allowed. (We remark that the same analysis handles the case when pulling an arm also fetches reward, but for a clearer presentation we do not consider such rewards here.)

Our algorithm in Section 4 can be, for the large part, directly applied in this situation as well; we now explain the small changes that need to be done in the various steps, beginning with the new LP relaxation. The additional variable in the LP, denoted by $x_{u,t}$ (for $u \in \mathcal{S}_i, t \in [B]$) corresponds to the probability of exploiting state u at time t .

$$\max \sum_{u,t} r_u \cdot x_{u,t} \tag{LP4}$$

$$w_{u,t} = z_{\text{parent}(u),t-1} \cdot p_{\text{parent}(u),u} \quad \forall t \in [2, B], u \in \mathcal{S} \tag{F.31}$$

$$\sum_{t' \leq t} w_{u,t'} \geq \sum_{t' \leq t} (z_{u,t'} + x_{u,t'}) \quad \forall t \in [1, B], u \in \mathcal{S} \tag{F.32}$$

$$\sum_{u \in \mathcal{S}} z_{u,t} \leq 1 \quad \forall t \in [1, B] \tag{F.33}$$

$$\sum_{u \in \mathcal{S}, t \in [B]} x_{u,t} \leq K \quad \forall t \in [1, B] \tag{F.34}$$

$$w_{\rho_i,1} = 1 \quad \forall i \in [1, n] \tag{F.35}$$

F.1 Changes to the Algorithm

Phase I: Convex Decomposition

This is the step where most of the changes happen, to incorporate the notion of exploitation. For an arm i , its strategy forest $\times\mathbb{T}(i, j)$ (the “ \times ” to emphasize the “exploit”) is an assignment of values $\text{time}(i, j, u)$, $\text{pull}(i, j, u)$ and $\text{exploit}(i, j, u)$ to each state $u \in \mathcal{S}_i$ such that:

- (i) For $u \in \mathcal{S}_i$ and $v = \text{parent}(u)$, it holds that $\text{time}(i, j, u) \geq 1 + \text{time}(i, j, v)$, and
- (ii) For $u \in \mathcal{S}_i$ and $v = \text{parent}(u)$ s.t $\text{time}(i, j, u) \neq \infty$, then one of $\text{pull}(i, j, u)$ or $\text{exploit}(i, j, u)$ is equal to $p_{v,u} \text{pull}(i, j, v)$ and the other is 0; if $\text{time}(i, j, u) = \infty$ then $\text{pull}(i, j, u) = \text{exploit}(i, j, u) = 0$.

For any state u , the value $\text{time}(i, j, u)$ denotes the time at which arm i is *played* (i.e., pulled or exploited) at state u , and $\text{pull}(i, j, u)$ (resp. $\text{exploit}(i, j, u)$) denotes the probability that the state u is pulled (resp. exploited). With the new definition, if $\text{time}(i, j, u) = \infty$ then this strategy does not play the arm at u . If state u satisfies $\text{exploit}(i, j, u) \neq 0$, then strategy $\times\mathbb{T}(i, j)$ *always exploits* u upon reaching it and hence none of its descendants can be reached. For states u which have $\text{time}(i, j, u) \neq \infty$ and have $\text{exploit}(i, j, u) = 0$, this strategy *always pulls* u upon reaching it. In essence, if $\text{time}(i, j, u) \neq \infty$, either $\text{pull}(i, j, u) = \text{pull}(i, j, \rho_i) \cdot \pi_u$, or $\text{exploit}(i, j, u) = \text{pull}(i, j, \rho_i) \cdot \pi_u$.

Furthermore, these strategy forests are such that the following are also true.

- (i) $\sum_j \text{s.t } \text{time}(i, j, u) = t \text{ pull}(i, j, u) = z_{u,t}$,
- (ii) $\sum_j \text{s.t } \text{time}(i, j, u) = t \text{ exploit}(i, j, u) = x_{u,t}$.

For convenience, let us define $\text{prob}(i, j, u) = \text{pull}(i, j, u) + \text{exploit}(i, j, u)$, which denotes the probability of some play happening at u .

The algorithm to construct such a decomposition is very similar to the one presented in Section D.1. The only change is that in Step 7 of Algorithm D.1, instead of looking at the first time when $z_{u,t} > 0$, we look at the first time when either $z_{u,t} > 0$ or $x_{u,t} > 0$. If $x_{u,t} > 0$, we ignore all of u 's descendants in the current forest we

plan to peel off. Once we have such a collection, we again appropriately select the largest ϵ which preserves non-negativity of the x 's and z 's. Finally, we update the fractional solution to preserve feasibility. The same analysis can be used to prove the analogous of Lemma D.1 for this case, which in turn gives the desired properties for the strategy forests.

Phase II: Eliminating Small Gaps

This is identical to the Section 4.2.2.

Phase III: Scheduling the Arms

The algorithm is also identical to that in Section 4.2.3. We sample a strategy forest $\times\mathbb{T}(i, j)$ for each arm i and simply play connected components contiguously. Each time we finish playing a connected component, we play the next component that begins earliest in the LP. The only difference is that a play may now be either a *pull* or an *exploit* (which is deterministically determined once we fix a strategy forest); if this play is an exploit, the arm does not proceed to other states and is dropped. Again we let the algorithm run ignoring the pull and exploit budgets, but in the analysis we only collect reward from exploits which happen before either budget is exceeded.

The lower bound on the expected reward collected is again very similar to the previous model; the only change is to the statement of Lemma 4.6, which now becomes the following.

Lemma F.1 *For arm i and strategy $\times\mathbb{T}(i, j)$, suppose arm i samples strategy j in step 1 of AlgMAB (i.e., $\sigma(i) = j$). Given that the algorithm plays the arm i in state u during this run, the probability that this play happens before time $\text{time}(i, j, u)$ **and** the number of exploits before this play is smaller than K , is at least $11/24$.*

In Section 4, we showed Lemma 4.6 by showing that

$$\Pr[\tau_u > \text{time}(i, j, u) \mid \mathcal{E}_{iju}] \leq \frac{1}{2}$$

Additionally, suppose we can also show that

$$\Pr[\text{number of exploits before } u > (K - 1) \mid \mathcal{E}_{iju}] \leq \frac{1}{24} \tag{F.36}$$

Then we would have

$$\Pr[(\text{number of exploits before } u > (K - 1)) \vee (\tau_u > \text{time}(i, j, u)) \mid \mathcal{E}_{iju}] \leq 13/24,$$

which would imply the Lemma.

To show Equation F.36 we start with an analog of Lemma 4.5 for bounding arm exploitations: conditioned on $\mathcal{E}_{i,j,u}$ and $\sigma(i') = j'$, the probability that arm i' is exploited at state u' before u is exploited is at most $\text{exploit}(i', j', u') / \text{prob}(i', j', \rho_{i'})$. This holds even when $i' = i$: in this case the probability of arm i being exploited before reaching u is zero, since an arm is abandoned after its first exploit. Since $\sigma(i') = j'$ with probability $\text{prob}(i', j', \rho_{i'})/24$, it follows that the probability of exploiting arm i' in state u' conditioned on $\mathcal{E}_{i,j,u}$ is at most $\sum_{j'} \text{exploit}(i', j', u')/24$. By linearity of expectation, the expected number of exploits before u conditioned on $\mathcal{E}_{i,j,u}$ is at most $\sum_{(i',j',u')} \text{exploit}(i', j', u')/24 = \sum_{u',t} x_{u',t}/24$, which is upper bounded by $K/24$ due to LP feasibility. Then Equation F.36 follows from Markov inequality.

The rest of the argument is identical to that in Section 4 giving us the following.

Theorem F.2 *There is a randomized $O(1)$ -approximation algorithm for the MAB problem with an exploration budget of B and an exploitation budget of K .*