arXiv:0810.0558v1 [cs.DS] 3 Oct 2008

The Ratio Index for Budgeted Learning, with Applications

Ashish Goel [*]
Stanford University

Sanjeev Khanna [†]
University of Pennsylvania

Brad Null [‡]
Stanford University

February 22, 2013

Abstract

In the budgeted learning problem, we are allowed to experiment on a set of alternatives (given a fixed experimentation budget) with the goal of picking a single alternative with the largest possible expected payoff. Approximation algorithms for this problem were developed by Guha and Munagala by rounding a linear program that couples the various alternatives together. In this paper we present an index for this problem, which we call the ratio index, which also guarantees a constant factor approximation. Index-based policies have the advantage that a single number (i.e. the index) can be computed for each alternative irrespective of all other alternatives, and the alternative with the highest index is experimented upon. This is analogous to the famous Gittins index for the discounted multi-armed bandit problem.

The ratio index has several interesting structural properties. First, we show that it can be computed in strongly polynomial time. Second, we show that with the appropriate discount factor, the Gittins index and our ratio index are constant factor approximations of each other, and hence the Gittins index also gives a constant factor approximation to the budgeted learning problem. Finally, we show that the ratio index can be used to create an index-based policy that achieves an $O(1)$-approximation for the finite horizon version of the multi-armed bandit problem. Moreover, the policy does not require any knowledge of the horizon (whereas we compare its performance against an optimal strategy that is aware of the horizon). This yields the following surprising result: there is an index-based policy that achieves an $O(1)$-approximation for the multi-armed bandit problem, oblivious to the underlying discount factor.

[*]Departments of Management Science and Engineering and (by courtesy) Computer Science, Stanford University. Research supported by NSF ITR grant 0428868, NSF CAREER award 0339262, and gifts from Google, Microsoft, and Cisco. Email: ashishg@stanford.edu .

[†]Department of Computer and Information Science, University of Pennsylvania, Philadelphia PA. Email: sanjeev@cis.upenn.edu. Supported in part by a Guggenheim Fellowship and by NSF Award CCF-0635084.

[‡]Department of Management Science and Engineering, Stanford University. Research supported by NSF ITR grant 0428868. Email: null@stanford.edu .

1 Introduction

The classical multi-armed bandit problem provides an elegant model to study the tradeoff between collecting rewards in the present based on the current state of knowledge (exploitation) versus deferring rewards to the future in favor of gaining more knowledge (exploration[1]) [2]. Specifically, in this model a user has a choice of bandit-arms to play, and at each time step it must decide which arm to play. The expected reward from playing a bandit-arm depends on the state of the bandit-arm where the state represents a "prior" belief on the bandit-arm. Each time a bandit-arm is played, this prior gets updated according to some transition matrix defined on the state space. For instance, a typical assumption on the bandit-arms is that they have (α, β)-priors: the success probability of an (α, β)-bandit-arm is $\alpha/(\alpha + \beta)$; in case of a success a reward of 1 is obtained and α gets incremented, whereas in case of a failure no reward is obtained and β gets incremented. The user wishes to maximize the total expected discounted reward over time. This simple setting effectively models many applications. A canonical example is exploring the effectiveness of different treatments in clinical trials while maximizing the benefit received by patients.

The discount factor in a multi-armed bandit problem may be viewed as modulating the horizon over which the strategy explores to identify the bandit-arm with maximum expected reward, before switching to exploitation. This facet of the multi-armed bandit problem is explicitly captured by the *budgeted learning problem*, recently studied by Guha and Munagala [18]. The input to the budgeted learning problem is the same as for the multi-armed bandit problem, except the discount factor is replaced by a horizon h. The goal is to identify the bandit-arm with maximum expected reward using at most h steps of exploration. The work of [18] gives a constant factor approximation for the budgeted learning problem via a linear programming based approach that determines the allocation of exploration and exploitation budgets across the various arms. The budgeted learning problem is the main object of study in this paper.

The multi-armed bandit problem admits an elegant solution: compute a score for each bandit-arm using only the current state of the bandit-arm and the discount factor, *independent* of all other bandit-arms in the system, and then play the bandit-arm with the highest score. This score is known as the *Gittins index*, and many proofs are known to show this is an optimal strategy (e.g., see [9]). The optimality of this "index-based" strategy implies that this problem exhibits a "separability" property whereby the optimal decision at each step is obtained by computations performed *separately* for each bandit-arm. This structural insight translates into efficient decision making algorithms. In fact, for commonly used prior update rules and discount rates, extensive collections of pre-computed Gittins indices exist, enabling in principle, a simple lookup-based approach for optimal decision-making. There are multiple definitions of what it means for a problem to have an "index". We will use the term index in its strongest form, i.e., where the index of an arm depends *only on the state of that arm*. This is also sometimes called a decomposable index (eg. [3, 4]).

The inherent appeal and efficiency of index-based policies is the unifying theme underlying our work. We show that many interesting and non-trivial variations of the multi-armed bandit problem, including the budgeted learning problem and the finite horizon problem, can all be well-approximated by index-based policies. Moreover, our approach gives decision strategies that are *oblivious* to parameters such as the underlying horizon or the discount factor while being constant-factor competitive to optimal strategies that are fully aware of these parameters.

1.1 Our results

We will study this problem when the state space of each arm satisfies the "martingale property", i.e., if we play an arm multiple times, the sequence of expected rewards is a martingale. This is a natural assumption for multi-armed bandit and related problems, e.g. the commonly used (α, β) priors satisfy this property.

[1] We will use the terms experimentation and exploration interchangeably in this paper, depending on the context.

An Index for Budgeted Learning Problems: Our first result is that the budgeted learning problem admits an approximate index, which we call the *ratio index*. Informally speaking, given a single bandit-arm and an exploration budget of h steps, the ratio index for that arm is the maximum expected exploitation reward per unit of the exploration and exploitation budget utilized. The ratio index suggests the following natural algorithm: at each step, play the arm with the highest ratio index. We show that this simple greedy algorithm gives a constant factor approximation to the budgeted learning problem. An $O(1)$-approximation algorithm for this problem is already known [18]. However, the algorithm of [18] is based on solving a *coupled* LP over all the arms, whereas the ratio index can be computed for each arm in isolation, much like the Gittins index. The ratio index has many other interesting properties. For example:

(1) We show that the Gittins index with discount factor $(1 - 1/h)$ and the ratio index over horizon h are within a constant factor of each other. This gives the following surprising result:

Theorem 1.1 *Given an exploration budget h, playing at each step the arm with the highest Gittins index, with discount factor $1 - 1/h$, yields a constant factor approximation to the budgeted learning problem.*

The proof relies on comparing the "decision-trees" of the ratio index and Gittins index strategies. Even in retrospect, it is not clear to us how such a result could be derived using an LP-based formulation such as the one used by Guha and Munagala [18]. Interestingly, the policy described in theorem 1.1 is known to often work well in practice [23]. Nonetheless, before the work of Guha and Munagala [18], we do not know of any provable guarantees for polynomial time algorithms in this setting. And until now, we don't know of any formal guarantees that relate the exponential discounting approach (which yields the Gittins index) and the budgeted learning approach.

(2) The ratio index can be computed in time which is strongly polynomial in the size of the state space (independent of h) of each arm if the state space is acyclic, and strongly polynomial in the size of the state space and h if the state space is general. Our proof of this fact involves recursively analyzing the basic feasible solutions of an underlying LP for computing optimum single arm strategies and using the structure of the basic feasible solutions to prove that these strategies have a simple form.

Finite Horizon and Discount-Oblivious Multi-Armed Bandits: We next study an important and natural variation of the budgeted learning problem, called the *finite horizon* multi-armed bandit problem. We are given a finite horizon h, and the goal is to maximize the expected reward collected during the horizon. Thus, in contrast to the budgeted learning problem, the horizon h is being used for both exploration and exploitation, and no payoffs are obtained after time h. We show the following result using the ratio index:

Theorem 1.2 *There is an index-based policy that gives a constant factor approximation to the finite horizon multi-armed bandit problem.*

Finally, we study the role of the discount factor in the design of an optimal strategy for the exploration-exploitation tradeoff. Small variations in discount factors can alter the choice of bandit-arm played at any step, highlighting the sensitivity of the Gittins index to the discount rate. We study the "Discount-oblivious" multi-armed bandit problem where the underlying discount factor is not known, and in fact, may even vary from one time step to the next. A finite horizon problem can be viewed as a special case of this general setting where the discount factor is 1 for the first h steps and is 0 for all subsequent steps. There is a useful relationship between the finite horizon and discount oblivious versions of the multi-armed problem: a strategy is κ-approximate for the discount-oblivious multi-armed bandit problem iff it is κ-approximate (*simultaneously*) for all finite horizons. Using this connection, and building on Theorem 1.2, we show the following result:

Theorem 1.3 *There is an index-based policy that gives a constant factor approximation for the multi-armed bandit problem with respect to all possible discount factors simultaneously.*

Our proof of both of these results is based on the following easy consequence of the ratio index approach to the budgeted learning problem. For any constant β, the expected profit of the optimal h/β-horizon strategy is an $\Omega(1)$-fraction of the expected profit of an optimal h-horizon strategy. Using this result, we design an algorithm that alternates between budgeted exploration and exploitation, using geometrically increasing horizons; each increasing horizon competing against a lower discount rate on future rewards. It is worth noting that this result can also be shown using the LP-based proof of Guha and Munagala. However, the following corollary is a consequence of our index-based approach and the relation between ratio and Gittins indices.

Corollary 1.4 *The strategy that alternates between exploring the arm with the highest Gittins index, and exploiting the arm with the highest reward, in phases of geometrically increasing length (and discount factor $1 - 1/t$ during a phase of length t) provides a constant factor approximation to the multi-armed bandit problem* simultaneously *for all finite horizons and for all discount factors.*

1.2 Related Work and Organization

There are many sources for the canonical work on Gittins indices, particularly with reference to (α, β) bandits and Bernoulli bandit processes [10, 11, 12, 9]. Glazebrook and others have studied approximation algorithms for other extensions to multi-armed bandit problems [13, 14]. Their approach builds upon the concept of achievable regions and general conservation laws and a related linear programming approach built by Tsoucas, Bertsimas, Nino-Mora, and others [1, 3, 25]. Relaxed linear programming based approaches to extensions of the multi-armed bandit problem have also been developed, e.g. for restless bandits [27, 4]. Our work on the ratio index builds on the insights obtained from the LP relaxation based approach of Guha and Munagala [18] as well as related work in model-driven optimization [15, 17] and stochastic packing [7, 8, 16, 19]. Additionally, related LP formulations have been developed for multi-stage stochastic optimization [6, 24].

In the theoretical computer science community, multi-armed bandits have primarily been studied in an adversarial setting, with the goal being to minimize the regret (see [5] for a nice overview). A typical guarantee in these settings is that the *total regret* after T steps grows as $\tilde{O}(\sqrt{TN})$ where N is the number of alternatives, assuming the partial information model (i.e. only the reward for the alternative that is actually played is revealed), which corresponds well to our setting. These results assume no prior beliefs, unlike our decision theoretic framework. However, the regret based bounds in the adversarial setting are meaningless unless $T > N$. The decision theoretic framework which has a rich history (starting perhaps with Wald's work in 1947 [26]) is more suited to the situation where the number of exploration steps is drastically limited, as is often the case. A typical setting, for example, is one where an advertiser that can advertise on 100,000 possible phrases and is willing to pay for 100 clicks to decide which keyword attracts visitors that convert into paid customers. So a traditional regret based bound may not be very meaningful in this setting.

In section 2 we define the budgeted learning problem and the ratio index, and prove that the ratio index is a constant factor approximation to the budgeted learning problem. Section 3 establishes that the Gittins and ratio indices are constant factor approximations of each other. We also show here that playing the arm with the largest Gittins index (with a suitable discount factor), gives a constant factor approximation to the budgeted learning problem. Section 4 presents index-based policies for finite horizon and discount oblivious versions of the multi-armed bandit problem. In section 5, we present a strongly polynomial algorithm to compute the ratio index as well as several useful insights into its structural properties.

2 The Budgeted Learning Problem and the Ratio Index

2.1 The Budgeted Learning Problem.

We are given n arms. Arm i has state space T_i, with initial state ρ_i. Experimenting on an arm i in state $u \in T_i$ results in the arm entering state $v \in T_i$ with known probability P_{uv}. The payoff of state u is given as $\zeta(u)$. Given an experimentation budget h, we are interested in finding the optimal policy, π^*, so that $\mathbf{E}_{\pi^*}[\max_{i\in\{1,\ldots,n\}} \zeta(v_i)]$ is maximum among all policies, where v_i is the state of arm i after the policy has been executed (the number of experiments cannot exceed h).

We will use T to denote $\cup_i T_i$. For convenience, we will assume that the T_i are disjoint and that $P_{uv} = 0$ if u and v are in the state spaces of different bandit-arms; this can be easily enforced by duplicating any shared states. The initial states represent a prior belief on the payoff from the bandit-arms. We will assume that the expected payoff is a martingale, i.e., $\zeta(u) = \sum_{v\in T} P_{uv}\zeta(v)$; the martingale assumption is crucial to our results. We will also assume without loss of generality that the state space of any arm is acyclic and truncated at depth h.

The martingale property has some useful and easy consequences which we will use repeatedly:

1. For an arbitrary policy let $p(t)$ denote its expected payoff if it is terminated after t experiment steps. Then, $p(t)$ is non-decreasing in t. In other words, extra experiments can never hurt.

2. Given a *single arm*, no policy can have a higher expected payoff than the one which does no exploration and simply chooses the initial state as the winner; in other words, extra experiments can never help given just one arm.

The proof of the following theorem is deferred to appendix A. It is conceivable (though not obvious to us) that this theorem can also be obtained via the "indexability characterization" of [3]. In any case, the proof is quite elementary and provides useful intuition.

Theorem 2.1 *There is no exact index for the budgeted learning problem.*

2.2 The ratio index

We will now define the ratio index, which is an approximate index for this problem. At any given time, the *current state* of the system is denoted by $\mathbf{S} = \{u_1, u_2, \ldots, u_n, \delta\}$, which captures the current states of all the arms, and the budget left (i.e. the number of experimentation steps that are still remaining), δ. The initial state of the system has all the arms in their initial states, and $\delta = h$. Since we use the term state for both the system and an arm, we will disambiguate where necessary by referring to these as "system-state" and "arm-state" respectively. A policy π is a function which takes as input a system-state $\mathbf{S}$ and either returns an arm i for experimentation (i.e. explores the arm-state u_i), or chooses an arm i as a winner and terminates (i.e. exploits the arm-state u_i), or simply terminates (abandons). If $\delta = 0$ then the only options are to abandon or exploit. The martingale property (see the comment at the end of section 2.1) implies that there always exists an optimal policy which explores some arm iff $\delta > 0$ and exploits some arm iff $\delta = 0$. We now introduce two vectors x^π and z^π. The probability that arm-state u is the final exploited state by policy π is given by x^π_u. The probability that arm-state u is explored by policy π is given by z^π_u. We define the *cost* of policy π as

$$\mathcal{C}(\pi) = \frac{\sum_{u\in T} z^\pi_u}{h} + \sum_{u\in T} x^\pi_u.$$

Observe that $\mathcal{C}(\pi) \leq 2$, for any policy π. The profit of policy π is defined as

$$\mathcal{P}(\pi) = \sum_{u\in T} x^\pi_u \zeta(u).$$

Observe that our definition of policy is an adaptive one; the decisions made in step $j > 1$ depend on the entire system-state at time j and hence on the outcome of previous experimentation steps. Further, it is easy to see that randomized strategies can not do any better than deterministic strategies.

If we drop the requirement that a policy must either exploit or abandon when the remaining budget δ is 0, we obtain what we call a *pseudo-policy*. A single arm policy is one which makes all its decisions based only on the state of a single pre-determined arm i, ignoring all other arms. We are now ready to define the ratio index and prove that it leads quite naturally to an approximation of the budgeted learning problem.

Ratio Index. *The ratio index $r(u, h)$ of a bandit-arm (say arm i) in initial state u and with experimentation budget h, is defined as*

$$\max_{\pi} \frac{\mathcal{P}(\pi)}{\mathcal{C}(\pi)},$$

where the max *is over all single arm pseudo-policies π which have initial arm-state u, budget h, state space T_i, and cost $\mathcal{C}(\pi) > 0$. We refer to a policy which yields the ratio index as a ratio index policy for state u, denoted $\pi_r(u, h)$.*

Even though we allow pseudo-policies in the definition of the ratio index, any ratio index policy respects the budget constraint:

Lemma 2.2 *Any ratio index policy for state u has cost at most* 1.

Proof: Because of the martingale property, no single arm policy starting from arm-state u can obtain profit more than $\zeta(u)$. Hence, any single arm policy π that has cost more than 1 must have a smaller ratio (of expected profit to expected cost) than the single arm policy which exploits in state u. ■

Greedy Algorithm. *Suppose the initial experimentation budget is h, and the current system-state is given by $\mathbf{S} = \{u_1, u_2, \ldots, u_n, \delta\}$. If $\delta > 0$, the greedy algorithm explores the arm i with the maximum ratio index, $r(u_i, h)$, with ties broken arbitrarily but consistently. If $\delta = 0$ the greedy algorithm exploits the arm i with maximum current expected reward $\zeta(u_i)$. We denote the greedy algorithm by G.*

Note: The greedy algorithm uses the same h at every step to compute the ratio index. Hence, given a table of the ratio index of every state in T (which can be pre-computed efficiently as specified in the section 5), we can implement this algorithm using a simple min-heap and the complexity of each step would be just $O(\log n)$, which is much better than solving a coupled LP with $3nh$ variables.

2.3 Analysis of the greedy algorithm

We now show that the greedy algorithm gives an $O(1)$-approximation to the budgeted learning problem.

Lemma 2.3 *A ratio index policy for arm-state u, $\pi_r(u, h)$, does not abandon any arm-state v with $r(v, h) > r(u, h)$ and does not explore or exploit any arm-state v with $r(v, h) < r(u, h)$.*

The proof of the above lemma is deferred to appendix D (as corollary D.6). Now consider the following algorithm, which we call the *persistent* algorithm, denoted G':

The persistent algorithm $\mathbf{G'}$: Given a system-state $\mathbf{S}$, let i be the arm with the highest ratio index $r(u_i, h)$ where u_i denotes the current state of arm i. Play arm i in accordance with the policy $\pi_r(u_i, h)$ until the policy chooses to exploit or abandon. If $\pi_r(u_i, h)$ abandons, let $\mathbf{S}'$ be the resulting system-state. Repeat the process starting with $\mathbf{S}'$. If at any time, the system-state is such that $\delta = 0$, immediately exploit the arm that has the highest current ratio index.

Observe that as for the greedy algorithm, the ratio index used by the persistent algorithm G' is computed using a fixed budget h; the number of remaining exploration steps δ is used only to terminate G'.

Lemma 2.4 *The expected profit of the greedy algorithm G is at least as much as the expected profit of the persistent algorithm G'.*

Proof: Couple the greedy and the persistent algorithms such that if they both explore an arm in a given state, that arm transitions to the same state for both strategies. Let $I = \langle i_1, i_2, \ldots, i_\kappa \rangle$ denote the sequence of arms explored by G' before exploitation; here $\kappa \leq h$ since G' can exploit early. Let $J = \langle j_1, j_2, \ldots, j_h \rangle$ denote the set of arms explored by G. By Lemma 2.3, we can conclude that I is a prefix of J. By the martingale property, early termination can never result in increased profit; the lemma follows. ■

Thus it suffices to analyze G'. Given two single arm pseudo-policies $\tilde{\pi}$ and π for an arm i, we say that $\tilde{\pi} \succeq \pi$, if for all arm-states $u \in T_i$ for which π explores (exploits) arm i, $\tilde{\pi}$ also explores (exploits) the arm i. Notice that $\tilde{\pi}$ might choose to continue exploration/exploitation when π abandons an arm-state. Informally, $\tilde{\pi} \succeq \pi$ means that policy $\tilde{\pi}$ can be played after policy π has been played to completion. We will now state a useful technical lemma; the proof is in appendix A:

Lemma 2.5 *Given two arbitrary single arm pseudo-policies π, π' for arm i in initial arm-state u, there exists another single arm pseudo-policy $\tilde{\pi}$ starting in the same initial arm-state u such that, (1) $\tilde{\pi} \succeq \pi$, (2) $\mathcal{C}(\tilde{\pi}) - \mathcal{C}(\pi) \leq \mathcal{C}(\pi')$, and (3) $\mathcal{P}(\tilde{\pi}) \geq \mathcal{P}(\pi')$.*

The above property is akin to submodularity. We now state our main theorem, which says that the greedy algorithm G gives a constant factor approximation to the optimal policy. Let $B_\pi(h, \mathbf{S})$ denote the expected profit obtained by strategy π for the budgeted learning problem run with budget h and initial system-state $\mathbf{S}$, and let $B^*(h, \mathbf{S})$ denote the expected profit obtained by an optimum strategy with the same parameters. We will omit the system-state when it is the same for all the strategies involved.

Theorem 2.6 $B_G(h) \geq 0.22 B^*(h)$

Proof: From Lemma 2.4, it suffices to analyze the persistent algorithm G' rather than the greedy algorithm G. We divide the persistent algorithm into stages, starting from stage 1. Let i_1 be the arm with the highest ratio index at the beginning of stage 1 (and hence the arm that will be played by G' at the first step). Since the arms evolve probabilistically, the first stage (as well as subsequent stages) will result in a distribution over system-states. Let $\mathbf{S}_j$ denote the system-state at the start of stage j, and let $\mathcal{D}_j$ denote the distribution of $\mathbf{S}_j$. Let u_j be the arm-state with the highest ratio index among the arm-states which have a non-zero probability, say γ_j, in $\mathcal{D}_j$, and let i_j be the corresponding arm. The j-th stage of G' is to simply move to the next stage if the arm i_j is not in state u_j (which happens with probability $1 - \gamma_j$); we call this stage "empty" in this case. If the arm i_j is in state u_j, then the j-th stage of G' is to mimic an optimum ratio index policy for state u_j. If the exploration budget gets exhausted during this mimicking process, then the j-th stage exploits arm i_j right away and the policy terminates; the cost of the extra exploitation is not charged to this stage of the policy. By the martingale property, this early termination can only increase the expected profit of the j-th stage. If the j-th stage exploits an arm, then the persistent algorithm terminates as well.

Let π_j denote the policy corresponding to the j-th stage. Let p_j and c_j be the cumulative expected profit and expected cost of the first j stages. Use $\Delta_p(j)$ and $\Delta_c(j)$ to denote the expected profit and the expected cost of the j-th stage, conditioned on this stage being played (i.e. the j-th stage being non-empty and the persistent algorithm not terminating before reaching the j-th stage). The following statement is a corollary of Lemma 2.5: the proof is a digression from the current theorem and is deferred to appendix A.

Corollary 2.7 *At the beginning of stage j, there exists a single arm pseudo-policy with profit to cost ratio at least $(\mathcal{P}(\pi^*) - p_{j-1})/2$.*

Since the persistent algorithm follows an optimum ratio index policy we are guaranteed that $\Delta_p(j)/\Delta_c(j) \geq (\mathcal{P}(\pi^*) - p_{j-1})/2$. By Markov's inequality, the probability that the budget has not been exhausted before stage j starts is at least $1 - c_{j-1}$. Also, recall that γ_j is the probability that the j-th stage is non-empty. The expected unconditioned profit of the j-th stage is at least $\gamma_j(1 - c_{j-1})\Delta_p(j)$. The expected unconditioned cost of the j-th stage is at most $\gamma_j\Delta_c(j)$. Hence, we get

$$\frac{p_j - p_{j-1}}{c_j - c_{j-1}} \geq \frac{\mathcal{P}(\pi^*) - p_{j-1}}{2} \cdot (1 - c_{j-1}).$$

Thus, the profit obtained by the persistent algorithm is more than the one attained by the following differential process, where p is the cumulative profit and c is the cumulative cost, and $p^* = \mathcal{P}(\pi^*)$ (view the process as increasing the expected cost from 0 to 1):

$$\frac{dp}{dc} = \frac{p^* - p}{2}(1 - c).$$

Integrating from $c = 0$ to $c = 1$, we get that the expected profit is at least $(1 - e^{-0.25})p^* \geq 0.22p^*$. ■

Thus, we have shown the existence of a simple index which yields almost as good an approximation ratio as the LP-based approach of Guha and Munagala. The results above assume that each exploration step has the same cost, but can easily be extended to the weighted exploration cost case. We can also modify the proof slightly to obtain the following corollary[2]:

Corollary 2.8 $B^*(h/2) \geq 0.17B^*(h)$.

Combining Corollary 2.8 and Theorem 2.6, we obtain the following corollary:

Corollary 2.9 $B_G(h/2) = \Omega(B^*(h))$.

3 Relating the Gittins and Ratio indices

We will use S_δ to denote a standard (stationary) bandit-arm with a fixed reward of δ. We will use $\mathcal{A}$ to denote a given bandit-arm in some initial state u. A *Gittins index strategy* $\mathcal{S}$ takes as input an arm $\mathcal{A}$ with an unknown reward distribution (but a known initial state) and a standard bandit-arm S_δ for some $\delta \geq 0$, and gives a strategy for maximizing the discounted reward for a multi-armed bandit with $\mathcal{A}$ and S_δ as its two bandit-arms. Thus each node in the decision tree of $\mathcal{S}$ is labeled as playing either the given arm $\mathcal{A}$ or the standard bandit S_δ. We can assume w.l.o.g. that once the strategy $\mathcal{S}$ plays the standard bandit at a node in the tree, it plays it forever from here onwards. The *Gittins index* of an arm $\mathcal{A}$ is defined to be the least δ such that the Gittins index strategy with input arms $\mathcal{A}$ and S_δ is indifferent between playing either one of them at time 0. We will assume u to be the initial state of $\mathcal{A}$ in the remainder of this section, and drop its explicit mention. Let $r(h)$ denote the ratio index for $\mathcal{A}$ when the horizon is limited to h. Let $\rho(\theta)$ denote the Gittins index for $\mathcal{A}$ when the discount factor is uniform for some $0 < \theta < 1$. The following lemmas show that the Gittins and the ratio indices are constant factor approximations of each other. The proofs involve transforming the Gittins index strategy to the ratio index strategy (and vice verse) and are in appendix B.

Lemma 3.1 *For any $h \geq 2$, $\rho(\theta) \geq r(h)\left(1 - \frac{1}{h}\right)^h$ where $\theta = (1 - \frac{1}{h})$. Thus as $h \to \infty$, $\rho(\theta) \geq r(h)/e$.*

Lemma 3.2 *For any $h \geq 2$, $\rho(\theta) \leq (2 + 4e)r(h)$ where $\theta = (1 - \frac{1}{h})$.*

[2]This corollary can also be obtained using the LP-based framework of Guha and Munagala [18].

Lemma 3.3 *Let $\rho_i(t)$ denote the Gittins index of arm i at time t, where the discount factor θ is $1 - 1/h$. Playing the arm with the highest value of $\rho_i(t)$ for $t = 1, 2, \ldots, h$ and then picking the arm with the highest expected payoff at time t results in a constant factor approximation to the budgeted learning problem.*

While the constants in our proofs are large, the algorithms are simple and intuitive. For instance, Schnieder and Moore [23] and Madani, Lizotte and Greiner [21] have studied the policy defined in lemma 3.3 and other similar policies, and found that they often work well in practice.

4 Finite Horizon and Discount Oblivious Multi-armed Bandits

In the traditional multi-armed bandit problem, we are given a fixed discount factor $\theta \in (0, 1)$ and allowed to play one arm at each time. If the reward at time t is $r(t)$ then the total discounted reward is $\sum_{t \geq 0} \theta^t r(t)$. Always playing the arm with the currently highest Gittins index maximizes the expected total discounted reward; however the Gittins index of an arm depends crucially on the parameter θ. In this section, we discuss both *finite horizon* and *discount oblivious* versions of the multi-armed bandit problem.

In the finite horizon multi-armed bandit problem, we are given a fixed number of steps, h, as in the budgeted learning problem. However, unlike the budgeted learning problem, the objective of the finite horizon problem is to maximize the total (undiscounted) expected reward obtained *during* the first h steps. This models many important problems such as optimally placing bets with a fixed number of chips, and optimally assigning impressions to advertisers [22].

In the discount oblivious multi-armed bandit problem, we want to find a strategy that provides a constant factor approximation to the optimum reward for all $\theta \in (0, 1)$ *simultaneously*. It is not clear up front that such a strategy exists. In fact, we will allow the discounts to be even more general. Let $\Lambda = \langle \Lambda_0, \Lambda_1, \Lambda_2, ...$ be an infinite sequence of discount factors that satisfies the property $1 = \Lambda_0 \geq \Lambda_1 \geq \Lambda_2 \geq ...$ and where $\Lambda_t \to 0$ as $t \to \infty$. We will call such a sequence a discount factor sequence. Let the system-state $\mathbf{S}$ denote the vector of all the arm-states. We will use $D_\pi(\Lambda, \mathbf{S})$ to denote the total expected discounted reward of any strategy π for discount factor sequence Λ starting from system-state $\mathbf{S}$. If strategy π obtains reward $r(t)$ at time t when started in initial system-state $\mathbf{S}$, then $D_\pi(\Lambda, \mathbf{S}) = \sum_{t=0}^{\infty} \Lambda_t r(t)$. Setting $\Lambda_t = \theta^t$ leads to the standard multi-armed bandit problem. Setting $\Lambda_t = 1$ for $t < h$ and $\Lambda_t = 0$ otherwise leads to a fixed horizon problem where we only get the reward from the first h time steps.

We will use $F_\pi(h, \mathbf{S})$ to denote the total (undiscounted) expected reward over a window of h steps of any strategy π, starting from $\mathbf{S}$. We will use $D^*(\Lambda, \mathbf{S})$, $F^*(h, \mathbf{S})$ to denote the optimum values for the two problems. We will omit the parameter $\mathbf{S}$ when it is the same for all strategies under discussion. All proofs are deferred to appendix C.

4.1 An approximate index for the finite horizon problem

Recall (from section 2) that $B_G(h)$ and $B^*(h)$ denote the expected profit of the greedy algorithm (which always explores the arm with the largest ratio index) and the optimum strategy respectively, for the budgeted multi-armed bandit problem. We first relate the budgeted learning and finite horizon problems:

Lemma 4.1 *For any positive integer h, we have $\lceil \frac{h}{2} \rceil \cdot B^* \left(\lfloor \frac{h}{2} \rfloor \right) \leq F^*(h) \leq h \cdot B^*(h)$.*

We will now define two index-based strategies for the finite horizon problem, assuming horizon h and initial system-state $\mathbf{S}$:.

1. For the first $\lfloor h/2 \rfloor$ steps, play the arm with the highest ratio index, where the ratio index is computed assuming a budget of $\lfloor h/2 \rfloor$. For the remaining $\lceil h/2 \rceil$ steps, play the arm with the highest expected reward. We will denote this strategy as RATIOSWITCH$(h, \mathbf{S})$ since it switches from using the ratio index (in the first half) to using the expected profit as an index in the second half.

2. Similarly, define GITTINSSWITCH$(h, \mathbf{S})$ as the strategy which plays the arm with the highest Gittins index (assuming a discount factor $1 - 1/\lfloor h/2 \rfloor$) during the first $\lfloor h/2 \rfloor$ steps, and then switches to using the arm with the highest expected reward.

Observe that both strategies use an index at each step, and the choice of index does not depend on the state of the system; it only depends on the time step. As before, we will omit the system-state parameter $\mathbf{S}$ when is it is the same for all strategies under discussion.

Theorem 4.2 $F_{\text{RATIOSWITCH}(h)}(h) = \Omega(F^*(h))$.

Combining lemma 3.3 with the proof of theorem 4.2, we obtain:

Theorem 4.3 $F_{\text{GITTINSSWITCH}(h)}(h) = \Omega(F^*(h))$.

To the best of our knowledge, this is the first index for the finite horizon problem with provable approximation guarantees. It would be interesting to obtain a smooth version of GITTINSSWITCH(h) which does not need to make the discrete jump from a discount factor of $1 - 2/h$ in the first half to playing the arm with the highest expected reward (i.e. to a discount factor of 0) in the second half.

4.2 An approximate index for the discount oblivious problem

We will first establish a connection between the discount oblivious and finite horizon problems and then use this connection to obtain a simple index-based approximation algorithm for the discount oblivious problem.

Lemma 4.4 *For any κ, a strategy gives a κ-approximation simultaneously for all discount factor sequences Λ iff it gives a κ-approximation simultaneously to the fixed horizon problems with all horizons $h \geq 0$.*

Let RATIOSCALE be the following discount oblivious strategy: play in sequence the strategies RATIOSWITCH$(1, \mathbf{S_0})$, RATIOSWITCH$(2, \mathbf{S_1})$, RATIOSWITCH$(4, \mathbf{S_3})$, RATIOSWITCH$(8, \mathbf{S_7}), \ldots$, where each RATIOSWITCH$(2^k)$ is started from the state of the system after time $2^k - 1$, denoted $\mathbf{S_{2^k-1}}$; this is the state in which the arms are left by the previous RATIOSWITCH strategy. $\mathbf{S_0}$ is the initial state of the system.

Like RATIOSWITCH, RATIOSCALE is also an index-based strategy; the index used at any time step t depends only on t. Analogously, GITTINSSCALE plays the sequence GITTINSSWITCH$(1, \mathbf{S_0})$, GITTINSSWITCH$(2, \mathbf{S_1})$, GITTINSSWITCH$(4, \mathbf{S_3})$, GITTINSSWITCH$(8, \mathbf{S_7}), \ldots$.

Since the state of the system at the start of RATIOSWITCH(2^i) depends on the outcomes of the previous steps, the following technical lemma, which is an easy consequence of the Martingale property, will be useful. This lemma states that performing an arbitrary sequence of extra explorations at the beginning cannot hurt the optimum solution for the budgeted learning problem. Observe that the state $\mathbf{T}$ is itself a random variable in this lemma; the expectation is over all values of $\mathbf{T}$.

Lemma 4.5 *Let π_1 be any arbitrary finite sequence of explorations starting from system-state $\mathbf{S}$. Let $\mathbf{T}$ be the system-state at the end of π_1. Let π_2 be an optimum h step strategy for the budgeted learning problem starting from the system-state $\mathbf{T}$. Then $\mathbf{E}[B_{\pi_2}(h, \mathbf{T})] \geq B^*(h, \mathbf{S})$.*

Using the above lemma, we can show the following:

Lemma 4.6 *For any positive integer $h \geq 1$, the expected reward of the discount oblivious strategy* RATIOSCALE *in the first h steps is $\Omega(F^*(h, \mathbf{S_0}))$.*

Lemma 4.7 *For any positive integer $h \geq 1$, the expected reward of the discount oblivious strategy* GITTINSSCALE *in the first h steps is $\Omega(F^*(h, \mathbf{S_0}))$.*

Invoking lemma 4.4 now gives us:

Theorem 4.8 *Strategies* RATIOSCALE *and* GITTINSSCALE *both give a constant factor approximation to the multi-armed bandit problem simultaneously for all discount factor sequences* $\mathbb{\Lambda}$.

5 Computing the Ratio Index

We will now sketch how the ratio index can be computed. In the process, we will also get several useful insights into its structural properties. Given a single bandit-arm i, an initial state ρ for i, an exploration budget of h, and a state space T_i truncated to depth h, we view T_i as a layered DAG of depth h, which is to say that for any arm-state, u, in layer j, if $P_{uv} > 0$, then v must be in layer $j+1$. As explained in section 5.1, this is without loss of generality. We let Σ be the number of nodes in the layered DAG. Additionally, for any state u in T_i, we use T_i^u to denote the sub-DAG of T_i with root u; thus $T_i^\rho = T_i$.

For the purposes of this section, we require the use of *randomized single arm policies*. Whereas a *deterministic* single arm policy (corresponding to arm-state v) will always either explore v, exploit v, or abandon with probability 1, a randomized policy, π, selects $e_v, p_v : e_v, p_v \geq 0, e_v + p_v \leq 1$ where e_v represents the probability π explores in this state, p_v represents the probability π exploits in this state, and $1 - e_v - p_v$ represents the probability π abandons in this state. The vectors x^π and z^π are defined for randomized policies as for deterministic policies, as are the profit $\mathcal{P}(\pi)$ and cost $\mathcal{C}(\pi)$ of the policy. Our approach below will calculate the ratio index $r(u, h)$ for all $u \in T_i$ as well as the entire *profit curve* $\mathcal{P}_u(\cdot)$ for all u where $\mathcal{P}_u(\mathcal{C}_u) = \max_\pi \mathcal{P}(\pi)$ where the $\max$ is over all randomized single arm policies π with initial state u and $\mathcal{C}(\pi) \leq \mathcal{C}_u$. We show that there exists a deterministic policy that induces the maximum $\mathcal{P}_u(\mathcal{C}_u)/\mathcal{C}_u$ over all $\mathcal{C}_u > 0$ (and in fact our algorithm will find such a policy). Thus, the value $\max \mathcal{P}_u(\mathcal{C}_u)/\mathcal{C}_u$ is the ratio index for u given h, i.e., $r(u, h)$. Our algorithm relies heavily upon the following theorem on the structure of the profit curve.

Theorem 5.1 *The profit curve, $\mathcal{P}_u(\cdot)$, for any given state u is concave and piecewise linear with at most $2\Sigma_u$ segments where Σ_u represents the number of states in T_i^u.*

The proof of this theorem involves several steps and is deferred to appendix D. Towards proving the theorem, we show that as the budget increases along the profit curve for u, a *monotonicity property* holds that for every state $v \in T_i^u$, both p_v and $e_v + p_v$ are non-decreasing.

Lemma 5.2 *For any $C^{(1)}, C^{(2)}$, with $C^{(2)} > C^{(1)}$, there exist optimal solutions $\langle e^{(1)}, p^{(1)} \rangle$ and $\langle e^{(2)}, p^{(2)} \rangle$ to $LP_u(C^{(1)})$ and $LP_u(C^{(2)})$ respectively such that $p_v^{(1)} \leq p_v^{(2)}$ and $e_v^{(1)} + p_v^{(1)} \leq e_v^{(2)} + p_v^{(2)}$ for all v in T_i^u.*

We further characterize the intersection of line segments of the profit curve as "corner" solutions and show that at these points $p_v \in \{0, 1\}$ and $e_v \in \{0, 1\}$ for *all* states in T_i^u. Thus, these points of the curve are induced by deterministic policies. Thus, the policy which induces the "corner" solution at the end of the first segment of the profit curve is a deterministic ratio index policy.

5.1 Algorithm for Computing the Profit Curve

The algorithm for computing the profit curve (and hence the ratio index) involves recursively calculating the profit curve for a state u given the profit curves for all of its successor states. We begin by constructing an *exploration profit curve* for u, $\mathcal{X}_u(\cdot)$, which denotes the optimal profit for any given cost conditioned on the fact that we are exploring at u (i.e. $e_u = 1$). We then take the concave envelope over this curve combined with the abandonment policy and the exploitation policy. Figure 1 in appendix D shows a typical example of the relationship between these two curves.

Superficially, it might seem that the number of segments of the profit curves could increase exponentially as we perform this process up the DAG. However, Theorem 5.1 guarantees that the number of segments remains bounded and the entire curve for u can be computed in time $O(d\Sigma_u \log \Sigma_u)$ given the successor curves, where d represents the maximum number of immediate descendants for any node. Thus, this algorithm is strongly polynomial (in Σ) for computing the entire profit curve of a state in the layered DAG, and hence, the ratio index. If the underlying state space of the bandit-arm is an unlayered DAG, we can make it layered by multiplying the number of states by at most Σ, so the algorithm is still strongly polynomial in Σ. If the underlying state space is not a DAG, we can convert it into a layered DAG by multiplying the number of states by at most h. Details of the algorithm and the analysis are in appendix D.

Acknowledgements

The authors would like to thank Rajat Bhattacharjee and Sudipto Guha for helpful discussions, as well as anonymous referees for pointing us to the citations [23, 21, 26].

References

[1] P. Bhattacharya, L. Georgiadis, and P. Tsoucas. Extended Polymatroids, Properties, and Optimization. *Integer Programming and Combinatorial Optimization, IPCO2*, ed. E. Bala, G Cornnejols, and R. Kannan, Carnegie-Mellon University 298-315, 1992.

[2] R. Bellman. A Problem in the Sequential Design of Experiments. *Sankhia*, 16:221-229, 1956.

[3] D. Bertsimas and J. Nino-Mora. Conservation Laws, Extended Polymatroids and Multi-armed Bandit Problems. *Mathematics of Operations Research*, 21:257-306, 1996.

[4] D. Bertsimas and J. Nino-Mora. Restless Bandits, Linear Programming Relaxations, and a Primal-Dual Index Heuristic. *Operations Research*, 48:80-90, 2000.

[5] A. Blum and Y. Mansour. Learning, Regret Minimization, and Equilibria. *Algorithmic Game Theory*, ed. N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani. 79-102, 2007.

[6] M. Charikar, C. Chekuri, and M. Pal. Sampling Bounds for Stochastic Optimization. *APPROX-RANDOM* 257-269, 2005.

[7] B. C. Dean, M. X. Goemans, and J. Vondrak. Approximation the Stochastic Knapsack Problem: The Benefit of Adaptivity. *FOCS '04: Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, 208-217, 2004.

[8] B. C. Dean, M. X. Goemans, and J. Vondrak. Adaptivity and Approximation for Stochastic Packing Problems. *Porc. 16th ACM-SIAM Symp. on Discrete Algorithms*, 395-404, 2005.

[9] E. Frostig and G. Weiss. Four Proofs of Gittins' Multiarmed Bandit Theorem. *Applied Probability Trust*, November 10, 1999.

[10] J.C. Gittins and D. M. Jones. A Dynamic Allocation Index for the Sequential Design of Experiments. *Progress in Statistics: European Meeting of Statisticians, Budapest, 1972*, ed. J. Ganu, K. Sarkadi, and I. Vince. 241-266, 1974.

[11] J. C. Gittins. Bandit Processes and Dynamic Allocation Indices. *J Royal Statistical Societe Series B*, 14:148-167, 1979.

[12] J. C. Gittins. *Multiarmed Bandits Allocation Indices*, Wiley, New York, 1989.

[13] K. D. Glazebrook and R. Garbe. Almost Optimal Policies for Stochastic Systems which Almost Satisfy Conservation Laws. *Annals of Operations Research*, 92:19-43, 1999.

[14] K. D. Glazebrook and D. J. Wilkinson. Index-Based Policies for Discounted Multi-armed Bandits on Parallel Machines. *The Annals of Applied Probability*, 10:3:877-896, 2000.

[15] A. Goel, S. Guha, and K. Munagala. Asking the Right Questions: Model-driven Optimization Using Prbes. *Proc. ACM Symp. on Principles of Database Systems*, 2006.

[16] A. Goel and P. Indyk. Stochastic Load Balancing and Related Problems. *Proc. Symp. on Foundations of Computer Science*, 1999.

[17] S. Guha and K. Munagala. Model Driven Optimization Using Adaptive Probes. *Proc. ACM-SIAM Symp. on Discrete Algorithms*, 2007.

[18] S. Guha and K. Munagala. Approximation Algorithms for Budgeted Learning Problems. *STOC*, 2007.

[19] J. Kleinberg, Y. Rabini, and E. Tardos. Allocation Bandwidth for Bursty Connections. *SIAM J. Comput* 30(1), 2000.

[20] D. Luenberger. *Linear and Nonlinear Programming*, Reading, MA, 1984.

[21] O. Madani, D. Lizotte, and R. Greiner. Active Model Selection. *Proceedings of the 20th conference on Uncertainty in artificial intelligence*, 357-365, 2004.

[22] P. Rusmevichientong and D. Williamson. An Adaptive Algorithm for Selecting Profitable Keywords for Search-based Advertising Services. *Proceedings of the 7th ACM conference on Electronic commerce,* 260-269, 2006.

[23] J. Schneider and A. Moore. Active Learning in Discrete Input Spaces. *Proceedings of the 34th Interface Symposium*, 2002.

[24] D. Shmoys and C. Swamy. Stochastic Optimization is (Almost) as Easy as Discrete Optimization. *Proc. 45th Symp. on Foundations of Computer Science* 228-237, 2004.

[25] P. Tsoucas. The Region of Achievable Performance in a Model of Klimov. Research Report RC16543, IBM T.J. Watson Research Center, Yorktown Heights, New York, 1991.

[26] A. Wald. *Sequential Analysis*, J. Wiley & Sons, New York, 212p, 1947.

[27] P. Whittle. Restless Bandits: Activity Allocation in a Changing World. *A Celebration of Applied Probability. J. Applied Probability*, 25A:287-298, 1988.

A The Budgeted Learning Problem

We will now provide the missing proofs from section 2:

Proof of Theorem 2.1: Consider three arms A, B, C with (α, β) priors: $\alpha_A = 5, \beta_A = 4$, $\alpha_B = 28, \beta_B = 19$, $\alpha_C = 28, \beta_C = 19$. Assume the horizon h is just 1.

First consider the scenario where arms A and B are the only ones present. Observe that $28/48 > 5/9$. So if we play arm B once and the trial is a failure, arm B will still be more profitable than arm A. Hence, playing arm B gives an expected profit of $28/47 = 0.5957$ since arm B will be chosen as the winner regardless of the outcome. Also, $28/47 < 6/10$ so if we play arm A once and the trial is a success, arm A becomes more profitable than arm B. Hence, exploring A first gives an expected profit of $(5/9) \times (6/10) + (4/9) \times (28/47) = 0.5981$. Therefore, if there exists an index for the budgeted learning problem, the index of arm A must be higher than that of arm B.

Now consider the scenario where arms A, B, C are all present. Playing arm A first gives the same expected profit as before: 0.5981. If we play arm B, and the trial is a failure, arm C will be chosen as the winner, giving an expected profit of $(28/47) \times (29/48) + (19/47) \times (28/47) = 0.6008$. Hence, arm B must have a higher index than arm A, which is a contradiction. ■

Proof of Lemma 2.5: Without loss of generality, assume that the state space T_i is a tree. Define $\tilde{\pi}$ as follows: $\tilde{\pi}$ makes the same choices as π for all arm-states which are either explored or exploited by π. This ensures that condition 1 in the lemma is satisfied. Further, the total cost of $\tilde{\pi}$ on these states is merely the total cost of π.

For an arm-state u that is abandoned by π, $\tilde{\pi}$ does the following:

(a) If any ancestor of u in T_i gets abandoned by π', then $\tilde{\pi}$ abandons u.

(b) If any ancestor of u gets exploited by π', then $\tilde{\pi}$ exploits u.

(c) Else, $\tilde{\pi}$ makes the same choice as π' on u and all the descendants of u.

In order to prove condition 2 in the lemma, we have to bound (by charging to π') the cost incurred (by $\tilde{\pi}$) in exploring/exploiting those arm-states u (and their descendants) that are abandoned by π. In case (a) above, u is abandoned and no extra cost is incurred. In case (b), let v be the ancestor arm-state of u that was exploited by π'. The cost $x_u^{\tilde{\pi}}$ of exploiting u is the cost $x_v^{\pi'}$ of exploiting v times the probability of π reaching u conditioned on π reaching v. Since the total probability that π abandons a descendant of v conditioned on π reaching v is at most 1, the incremental cost for $\tilde{\pi}$ for all descendants of v can be no more than $x_v^{\pi'}$ and thus there can be no overcharging. In case (c), the charging is quite straight-forward since $\tilde{\pi}$ just mimics π'.

In order to prove condition 3 in the lemma, consider any arm-state u that is exploited by π'. If that state is exploited by π, then it is also exploited by $\tilde{\pi}$. If that state is explored by π, then eventually, π must either abandon or exploit a descendant along every path in the state space starting from u. Descendants that get abandoned by π will get exploited by $\tilde{\pi}$ by property (b). By the martingale property, the total profit obtained by $\tilde{\pi}$ from all the descendant states of u is the same as the profit obtained by π' from u. If the arm-state u is abandoned or not reached by π, then $\tilde{\pi}$ mimics π' according to property (c). Hence, $\tilde{\pi}$ gets at least as much profit as π'. ■

Proof of Corollary 2.7: Let π^* denote the optimum policy. Define a new policy that we call the restriction of π^* to arm i, denoted L_i, as follows: L_i follows π^* when it explores/exploits arm i, and simulates π^* (without really playing it) on other states. A simple coupling argument shows that the total expected cost of all these single arm policies is equal to the expected cost of π^*, and the total expected profit

of all these single arm policies is equal to the expected profit of π^*. Similarly, let M denote the greedy algorithm over the first $j-1$ stages, and let N_i denote the restriction of M to arm i.

The following is now immediate: $\sum_i \mathcal{P}_{L_i} - \sum_i \mathcal{P}_{N_i} = \mathcal{P}(\pi^*) - p_{j-1}$ and $\sum_i \mathcal{C}_{L_i} = \mathcal{C}(\pi^*) \leq 2$. Hence, there exists some i such that $(\mathcal{P}_{L_i} - \mathcal{P}_{N_i})/\mathcal{C}_{L_i} \geq (\mathcal{P}(\pi^*) - p_{j-1})/2$. Applying Lemma 2.5 with N_i serving the role of π and L_i serving the role of π' proves the corollary. ■

Proof of Corollary 2.8: If we take the integral in the proof of the above theorem with c going from 0 to $1/2$ instead of from 0 to 1, we get an approximation factor of $(1 - e^{-3/16}) \approx 0.17$. Hence, the optimum reward from the budgeted learning problem with budget $h/2$ is at least 0.17 times the optimum reward with budget h. ■

B Relating the Gittins and ratio indices

We will now prove the lemmas from section 3.

Lemma B.1 (same as Lemma 3.1) *For any $h \geq 2$,*

$$\rho(\theta) \geq r(h) \left(1 - \frac{1}{h}\right)^h$$

where $\theta = (1 - \frac{1}{h})$. Thus as $h \to \infty$, $\rho(\theta) \geq r(h)/e$.

Proof: Consider a strategy $\mathcal{S}_1$ that computes the ratio index $r(h)$ for arm $\mathcal{A}$ with horizon limited to h. Let $r = r(h)$. We will modify the strategy $\mathcal{S}_1$ in to another strategy $\mathcal{S}_2$ that will certify that the Gittins index for arm $\mathcal{A}$ is at least r/c for some suitable constant $c > 1$. The decision tree for strategy $\mathcal{S}_2$ is obtained by modifying the decision tree for $\mathcal{S}_1$ as follows:

1. When the decision tree for strategy $\mathcal{S}_1$ visits a node labeled "Abandon", the tree for $\mathcal{S}_2$ changes it to playing the standard bandit-arm $S_{r/c}$ and terminates this branch of the tree.
2. When the decision tree for strategy $\mathcal{S}_1$ visits a node labeled "Exploit", the tree for $\mathcal{S}_2$ replaces it to playing the arm $\mathcal{A}$ forever.

Let γ denote the expected reward of the bandit-arm $\mathcal{A}$ conditioned on exploitation by the strategy $\mathcal{S}_1$. Also, let b_R and b_T denote the exploration and the exploitation budget, respectively, of the strategy $\mathcal{S}_1$. Thus, by definition, the ratio index is given by

$$r = \frac{\gamma b_T}{b_R + b_T}.$$

We can now lower bound the difference Δ in the expected reward generated by the strategy $\mathcal{S}_2$ over the standard arm $S_{r/c}$ to be at least

$$\Delta \geq b_T \left(\gamma - \frac{r}{c}\right) \theta^h \sum_{i \geq 0} \theta^i - b_R h \left(\frac{r}{c}\right).$$

Substituting $\gamma = \frac{b_R + b_T}{b_T}$, we get

$$\Delta \geq \left((b_R + b_T)r - b_T \left(\frac{r}{c}\right)\right) \frac{\theta^h}{1-\theta} - b_R h \left(\frac{r}{c}\right) \geq (b_R h r) \left(\theta^h - \frac{1}{c}\right).$$

Thus $\Delta \geq 0$ if $c \geq (1 - \frac{1}{h})^{-h}$, establishing that $\rho \geq r \left(1 - \frac{1}{h}\right)^h$.

■

Lemma B.2 (same as Lemma 3.2) *For any $h \geq 2$,*

$$\rho(\theta) \leq 18r(h)$$

where $\theta = (1 - \frac{1}{h})$.

Proof: Let $r = r(h)$ and $\rho = \rho(\theta)$. The optimum discounted reward strategy given a bandit-arm $\mathcal{A}$ and the standard bandit-arm S_ρ is indifferent at time 0 between playing the two bandit-arms. Let $\mathcal{S}_1$ be a Gittins index strategy for arm $\mathcal{A}$ and the standard bandit-arm S_ρ such that $\mathcal{S}_1$ plays $\mathcal{A}$ at time 0. Each node in the decision tree of $\mathcal{S}_1$ is labeled as either playing the standard bandit S_ρ or playing the arm $\mathcal{A}$.

We will modify the strategy $\mathcal{S}_1$ into another strategy $\mathcal{S}_2$ that will certify that the ratio index for arm $\mathcal{A}$ is at least ρ/c when horizon is restricted to h; here $c > 1$ is a suitable constant to be specified later. Note that unlike strategy $\mathcal{S}_1$, the strategy $\mathcal{S}_2$ does not have access to the standard bandit S_ρ. However, since the standard bandit S_ρ is in a stationary state, when we say a state s in $\mathcal{S}_1$, we will assume that it refers only to the state of the arm $\mathcal{A}$. The decision tree for strategy $\mathcal{S}_2$ is obtained by modifying the decision tree for $\mathcal{S}_1$ as follows:

1. When the decision tree for strategy $\mathcal{S}_1$ visits a state that is labeled as playing S_ρ and the depth is at most h, the strategy $\mathcal{S}_2$ abandons.

2. When the decision tree for strategy $\mathcal{S}_1$ visits a state s labeled as playing the arm $\mathcal{A}$, then

   (a) If the depth is less than h, then the strategy $\mathcal{S}_2$ explores if the expected reward of *playing the arm $\mathcal{A}$ in state s* is less than $2\rho/c$, and it exploits otherwise.

   (b) if the depth is exactly h, strategy $\mathcal{S}_2$ exploits.

Clearly, the strategy $\mathcal{S}_2$ as defined above has its horizon restricted to h.

We will use $\mathsf{label}(s)$ and $\dot{(s)}$ to denote respectively the label and depth of a state s in the decision tree of strategy $\mathcal{S}_2$. A state s is assigned a label of "R" if $\mathcal{S}_2$ explores in this state, and a label of "T" if it exploits in this state. Finally, we will denote the probability of reaching a state s in $\mathcal{S}_1$ by $p_1(s)$, and the reward at a state s in the strategies $\mathcal{S}_1$ and $\mathcal{S}_2$ by $\pi_1(s)$ and $\pi_2(s)$ respectively. Note that $\pi_1(s) = \pi_2(s)$ for any state s where strategy $\mathcal{S}_1$ plays the arm $\mathcal{A}$.

The expected reward γ collected by $\mathcal{S}_2$, as well as the exploration and the exploitation budget of $\mathcal{S}_2$, denoted by b_R and b_T respectively, are given by

$$\gamma = \sum_{\{s \in \mathcal{S}_2 \mid \mathsf{label}(s) = T\}} p_1(s)\pi_2(s) \qquad b_R = \sum_{\{s \in \mathcal{S}_2 \mid \mathsf{label}(s) = R\}} \frac{p_1(s)}{h} \qquad b_T = \sum_{\{s \in \mathcal{S}_2 \mid \mathsf{label}(s) = T\}} p_1(s).$$

We now lower bound the total reward *deficit* that is accumulated by the strategy $\mathcal{S}_1$ in the states that are labeled as "exploring" in the strategy $\mathcal{S}_2$, as compared to the reward of playing the standard bandit S_ρ.

$$\sum_{\{s \in \mathcal{S}_2 \mid \mathsf{label}(s) = R\}} p_1(s)(\rho - \pi_1(s))\theta^{\dot{(s)}} = \sum_{\{s \in \mathcal{S}_2 \mid \mathsf{label}(s) = R\}} p_1(s)(\rho - \pi_2(s))\theta^{\dot{(s)}} \geq \frac{hb_R}{4}\left(\rho - \frac{2\rho}{c}\right).$$

since $\theta^{\dot{(s)}} \geq 1/4$ for $\dot{(s)} \leq h$.

On the other hand, the total reward *surplus* accumulated by the strategy $\mathcal{S}_1$ above the reward of playing the standard bandit S_ρ in the subtrees of the states labeled as "exploiting" in $\mathcal{S}_2$, can be upper bounded by

$$\sum_{\{s\in\mathcal{S}_2 \mid \mathsf{label}(s)=T\}} p_1(s)(\pi_1(s)-\rho)\left(\sum_{i\geq 0}\theta^i\right) \leq h\left(\sum_{\{s\in\mathcal{S}_2 \mid \mathsf{label}(s)=T\}} p_1(s)\pi_2(s)\right).$$

The bound above follows from the martingale property which implies that for any state s, we must satisfy the relationship $\pi_1(s) \leq \pi_2(s)+\rho$. To see this, observe that for any state s, by martingale property we have $\pi_2(s) = \sum_{v\in N(s)} P_{sv}\pi_2(v)$, where $N(s)$ denotes the set of descendants of a states s, and P_{sv} denotes the probability of transitioning from state s to state v when arm $\mathcal{A}$ is played in states s. On the other hand,

$$\pi_1(s) = \sum_{v\in N(s)} P_{sv}\max\{\pi_2(v),\rho\} = \sum_{\{v\in N(s) \mid \pi_2(v)>\rho\}} P_{sv}\cdot\pi_2(s) + \sum_{\{v\in N(s) \mid \pi_2(v)\leq\rho\}} P_{sv}\cdot\rho \leq \pi_2(s)+\rho.$$

Since the total reward surplus in $\mathcal{S}_1$ w.r.t. the standard bandit S_ρ must be at least as large as the total reward deficit w.r.t. to S_ρ, it follows that

$$h\left(\sum_{\{s\in\mathcal{S}_2 \mid \mathsf{label}(s)=T\}} p_1(s)\pi_2(s)\right) \geq \frac{hb_R}{4}\left(\rho-\frac{2\rho}{c}\right).$$

Thus

$$\gamma \geq \frac{b_R}{4}\left(\rho-\frac{2\rho}{c}\right) = \frac{b_R}{4c}(c-2)\rho. \tag{1}$$

Let $b_{T_1} = \sum_{\{s\in\mathcal{S}_2 \mid \mathsf{label}(s)=T \text{ and } \pi_2(s)\geq 2\rho/c\}} p_1(s)$, and $b_{T_2} = b_T - b_{T_1}$. Clearly,

$$\gamma \geq \left(\frac{2\rho}{c}\right) b_{T_1}. \tag{2}$$

Finally, we observe that the Gittins index policy $\mathcal{S}_1$ never plays the bandit-arm $\mathcal{A}$ in a state s if it ever plays the standard bandit in any ancestor of state s. Therefore, for any $i \leq h$,

$$b_{T_2} \leq \sum_{\{s\in\mathcal{S}_2 \mid \mathsf{label}(s)=R \text{ and } \dot{(}s)=i\}} p_1(s).$$

Averaging over all $i \leq h$, we get

$$b_{T_2} \leq \frac{1}{h}\left(\sum_{\{s\in\mathcal{S}_2 \mid \mathsf{label}(s)=R\}} p_1(s)\right) = b_R. \tag{3}$$

We can now lower bound the ratio index r by the expected reward to exploration and exploitation budget ratio achieved by the policy $\mathcal{S}_2$:

$$
\begin{aligned}
r &\geq \frac{\gamma}{b_R + b_{T_1} + b_{T_2}} \\
&\geq \frac{\gamma}{2b_R + b_{T_1}} && \text{(using (3))} \\
&\geq \frac{\gamma}{\frac{8\gamma c}{(c-2)\rho} + \frac{\gamma c}{2\rho}} && \text{(using (1) and (2))} \\
&\geq \frac{\rho}{\frac{8c}{c-2} + \frac{c}{2}} \\
&\geq \frac{\rho}{c} && \text{(when } c \geq 18\text{)}
\end{aligned}
$$

■

Lemma B.3 (same as Lemma 3.3) *Let $\rho_i(t)$ denote the Gittins index of arm i at time t, where the discount factor θ is $1 - 1/h$. Playing the arm with the highest value of $\rho_i(t)$ for $t = 1, 2, \ldots, h$ and then picking the arm with the highest expected payoff at time t results in a constant factor approximation to the budgeted learning problem.*

Proof: Consider the strategy S_2 described in the proof of Lemma 3.2 above. By the proof of Lemma 3.2, the profit to cost ratio of this strategy is $\Omega(\rho(\theta))$, which by Lemma 3.1 is $\Omega(r(h))$. We will call S_2 the truncated Gittins policy. Using the truncated Gittins policy instead of the ratio index policy in the proof of theorem 2.6, we can conclude that repeatedly choosing the arm with the highest Gittins index and playing the truncated Gittins policy for this arm (the entire policy, not just for a single step) till it abandons or exploits gives a constant factor approximation to the budgeted learning problem. We will call this the TG algorithm.

Now consider the strategy which repeatedly plays the arm with the highest Gittins index for one step and then at time h, picks the arm with the highest payoff; we will call this the Gittins strategy. The Gittins strategy is identical to the TG algorithm as long as the TG algorithm continues exploring some arm. The TG algorithm may choose to exploit before the entire budget is exhausted, in which case making an arbitrary set of additional exploration steps can not hurt (by the martingale property). The TG algorithm may also choose to exploit an arm with a smaller current expected payoff than some other arm; again, choosing the arm with the highest expected current payoff can not hurt. In either case, the expected profit of the Gittins strategy is no worse than that of the TG algorithm, and hence, the Gittins strategy is also a constant factor approximation to the budgeted learning problem. ■

C Finite Horizon and Discount Oblivious Multi-armed Bandits

We will now provide the proofs from section 4.

Proof of Lemma 4.1: Consider the fixed horizon strategy (for horizon h) that first solves the budgeted learning problem with budget $\lfloor h/2 \rfloor$ and then exploits the winner from this budgeted learning problem for the remaining $\lceil h/2 \rceil$ time steps. This strategy has expected pay off at least $\left\lceil \frac{h}{2} \right\rceil \cdot B^* \left(\left\lfloor \frac{h}{2} \right\rfloor \right)$ and hence, $F^*(h) \geq \left\lceil \frac{h}{2} \right\rceil \cdot B^* \left(\left\lfloor \frac{h}{2} \right\rfloor \right)$.

Now consider the budgeted learning strategy (with budget h) that emulates the optimum fixed horizon strategy (for horizon h) but only for t steps, where t is an integer chosen uniformly at random from the set $\{0, 1, \ldots, h-1\}$, and then declares the arm that the optimum fixed horizon strategy was about to play in step $t+1$ as the winner. This budgeted learning strategy has expected payoff exactly $F^*(h)/h$ and hence, $F^*(h) \leq h \cdot B^*(h)$. ■

Proof of Theorem 4.2: Assume, for simplicity, that h is even; the odd case is very similar. Now,

$$\begin{aligned} F_{\text{RATIOSWITCH}(h)}(h) &\geq (h/2)B_G(h/2) \\ &= \Omega(hB^*(h)) && \text{[from lemma 2.9]} \\ &= \Omega(F^*(h)) && \text{[from lemma 4.1]} \end{aligned}$$

■

Proof of Lemma 4.4: The "only if" part is trivial since a fixed horizon problem can be modeled using a discount factor sequence, as explained earlier. For the "if" part, consider a strategy π that offers expected reward $r(t)$ at time t. Then the expected discounted reward, $D_\pi(\Lambda)$ of π given a discount factor Λ can be written as $\sum_{h=0}^{\infty}\left((\Lambda_h - \Lambda_{h+1})\sum_{t=0}^{h} r(t)\right)$. Observe that $\Lambda_h - \Lambda_{h+1} \geq 0$ by the definition of a discount factor sequence. If we simultaneously κ-approximate each of the finite horizon rewards $\sum_{t=0}^{h} r(t)$ we also simultaneously κ-approximate any non-negative linear combination, and hence the optimum expected discounted reward for any discount factor sequence. ■

Proof of Lemma 4.6: If $h = 1$, $\text{RATIOSWITCH}(1, \mathbf{S_0})$ simply plays the arm with the highest expected profit, and hence obtains profit $F^*(1)$. If $h = 2$, it is easy to see that $F^*(2) \leq 3F^*(1)$ and hence the lemma holds. We will now assume that $h \geq 3$.

Let t be the largest power of 2 which is no larger than $(h+1)/2$; hence $t > h/4$. Since $h \geq 3$, we have $t \geq 2$. The strategy RATIOSCALE is guaranteed to execute the strategy $\text{RATIOSWITCH}(t, \mathbf{S_{t-1}})$ sometime during the first h steps (in fact during steps $t, \ldots, 2t-1$). Using lemma 4.5, we can repeat the steps in the proof of theorem 4.2 to claim that $\mathbf{E}[F_{\text{RATIOSWITCH}(t,\mathbf{S_{t-1}})}(t, \mathbf{S_{t-1}})] = \Omega(F^*(t, \mathbf{S_0})) = \Omega(F^*(h, \mathbf{S_0}))$. ■

The proof of Lemma 4.7 is similar to the above proof.

D Calculating the profit curve

We will now present the full details of the proof of theorem 5.1 as well as the full algorithm to compute the profit curves and ratio indices for all states up to depth h for a given bandit-arm.

D.1 Proof of Theorem 5.1

Below, we prove a series of claims that together imply theorem 5.1. We begin by considering two methods of calculating $\mathcal{P}_u(\mathcal{C}_u)$ that will be used in our discussions. The first is a recursive equation that can be used to calculate $\mathcal{P}_u(\mathcal{C}_u)$ for a given state u and budget $\mathcal{C}_u$ provided that we have the entire profit curves of all successor states. This equation is

$$\mathcal{P}_u(\mathcal{C}_u) = \max_{p_u, e_u, E^u} p_u\zeta(u) + e_u \sum_{v \in D(u)} P_{uv}\mathcal{P}_v(E^u_v)$$

the constraints are:

$$\begin{gathered} p_u + e_u \leq 1 \\ p_u + e_u(1/h + \sum_{v \in D(u)} P_{uv}E^u_v) \leq \mathcal{C}_u \\ p_u, e_u, E^u \geq 0 \end{gathered}$$

$D(u)$ is the set of immediate descendants of u. The decision variables p_u and e_u represent the probability of exploiting and exploring in u respectively (as in the definition of a randomized policy). The vector E^u

represents the budgets we would allocate to each of the immediate descendants of u should we visit them. Recall that P_{uv} is the probability of transitioning to state v given we are experimenting in state u. We assume $E_v^u = 0$ if $v \notin D(u)$.

Alternatively, the following LP ($LP_u(\mathcal{C}_u)$) (similar to the one in [18]) for a given state u of bandit-arm i reveals a policy, $\langle w, x, z\rangle$ which induces $\mathcal{P}_u(\mathcal{C}_u)$ for a given $\mathcal{C}_u$.

$$\max_{w,x,z} \sum_{v\in T_i^u} x_v \zeta_i(v)$$

s.t.

$$\begin{aligned}
\frac{\sum_{v\in T_i^u} z_v}{h} + \sum_{v\in T_i^u} x_v &\le \mathcal{C}_u \\
\sum_{y:v\in D(y)} z_y P_{yv} &= w_v \qquad \forall v \in T_i^u \setminus \{u\} \\
w_u &= 1 \\
x_v + z_v &\le w_v \qquad \forall v \in T_i^u \\
x_v, z_v &\ge 0 \qquad \forall v \in T_i^u
\end{aligned}$$

For any state $v \in T_i^u$, w_v represents the probability the bandit-arm enters v, x_v represents the unconditional probability of exploiting in v, and z_v represents the unconditional probability of experimenting in v. Given the stochastic nature of the P matrix and the fact that all of the z values are less than or equal to their corresponding w values, we can see that each element of w will be bounded between 0 and 1. Thus, x and z will also be automatically bounded above by 1. Thus, we do not need further constraints bounding these variables.

Using both the recursive equation and the LP, we can show the following.

Claim D.1 *$\mathcal{P}_u(\cdot)$ is a concave, nondecreasing function and if $\mathcal{C}_u \ge 1$, $\mathcal{P}_u(\mathcal{C}_u) = \mathcal{P}_u(1) = \zeta(u)$.*

Proof: Given any $\mathcal{C}_u$ and associated policy $\{p_u, e_u, E^u\}$ for the recursive equation that realizes $\mathcal{P}_u(\mathcal{C}_u)$, observe first that this policy is also feasible for any cost greater than or equal to $\mathcal{C}_u$. Thus, $\mathcal{P}_u(\mathcal{C}_u)$ is a nondecreasing function.

We prove the remainder of the claim by induction. Looking at $\mathcal{P}_u(\mathcal{C}_u)$ for a state u that is at depth h it is easy to see that, $\mathcal{P}_u(\mathcal{C}_u) = \mathcal{C}_u * \zeta(u)$ if $\mathcal{C}_u \le 1$ and $\zeta(u)$ if $\mathcal{C}_u > 1$. Thus, at depth h if $\mathcal{C}_u \ge 1$, $\mathcal{P}_u(\mathcal{C}_u) = \mathcal{P}_u(1) = \zeta(u)$ and $\mathcal{P}_u()$ is concave.

Now assume these properties hold for all states at depth $i+1$ and look at a state u at depth i. Our profit is obviously non-decreasing in each of the decision variables $\{p_u, e_u, E^u\}$. Set $E_v^u = 1\ \forall v \in D(u)$. From the induction hypothesis, $\sum_{v\in D(u)} P_{uv}\mathcal{P}_v(1) = \sum_{v\in D(u)} P_{uv}\zeta(v)$ and by the martingale property, $\sum_{v\in D(u)} P_{uv}\zeta(v) = \zeta(u)$, so our objective becomes $\max p_u\zeta(u) + e_u\zeta(u)$ or equivalently $\max(p_u + e_u)\zeta(u)$. Since $(p_u + e_u) \le 1$, this can be no larger that $\zeta(u)$. But clearly $\zeta(u)$ can be achieved with a cost of 1 by setting $p_u = 1$ and $e_u = 0$. Henceforth, $\mathcal{P}_u(\mathcal{C}_u) = \zeta(u)\ \forall \mathcal{C}_u \ge 1$.

With respect to concavity, for any two points on the profit curve of u, $\mathcal{P}_u(C^{(1)})$ and $\mathcal{P}_u(C^{(2)})$, corresponding to $C^{(1)} < C^{(2)}$ let $\langle w^{(1)}, x^{(1)}, z^{(1)}\rangle$ and $\langle w^{(2)}, x^{(2)}, z^{(2)}\rangle$ be their associated policies respectively as defined on $LP_u(\mathcal{C}_u)$. Consider the policy $\langle \frac{w^{(1)}+w^{(2)}}{2}, \frac{x^{(1)}+x^{(2)}}{2}, \frac{z^{(1)}+z^{(2)}}{2}\rangle$. This policy is feasible for the problem of finding the profit associated with budget $\frac{C^{(1)}+C^{(2)}}{2}$, so $\mathcal{P}_u(\frac{C^{(1)}+C^{(2)}}{2}) \ge \sum_{v\in T_i^u} \frac{x_v^{(1)}+x_v^{(2)}}{2}\zeta_i(v) = \frac{\mathcal{P}_u(C^{(1)})+\mathcal{P}_u(C^{(2)})}{2}$. Thus the profit curve for u is concave.

■

With respect to $LP_u(\mathcal{C}_u)$, we define the vectors e and p where $e_v = z_v/w_v$ and $p_v = x_v/w_v$ if $w_v > 0$. Otherwise, we require only that $e_v, p_v \ge 0$ and $e_v + p_v \le 1$. Thus e_v and p_v are the conditional probability of exploring and exploiting, respectively, given that we are in state v. Thus, the two vectors $\langle e, p\rangle$ define

a randomized policy that induces a point on $\mathcal{P}_u()$. Alternatively, we could define the same policy with the three vectors $\langle w, x, z \rangle$. In what follows, we will freely interchange between the two notations. Note that the one thing we must be careful to observe is that for any policy and state where $w_v = 0$, there are infinitely many equivalent assignments of e_v and p_v.

Appealing to linear programming theory with respect to $LP_u(\mathcal{C}_u)$, we can derive several interesting properties of the profit curve. We begin with the following:

Claim D.2 *For any $\mathcal{C}_u \in [0, 1]$ there exists an optimal policy with respect to $LP_u(\mathcal{C}_u)$ such that $p_v \in \{0, 1\}$ and $e_v \in \{0, 1\}$ for all but at most one state in T_i^u.*

Proof: Let's consider a basic feasible optimal solution to $LP_u(\mathcal{C}_u)$, $\langle w^*, x^*, z^* \rangle$. We know such a solution exists since the LP is bounded and a basic feasible solution to the LP exists (for instance the solution that corresponds to setting $x_v = z_v = 0 \; \forall v$). Let us create $\hat{T_i^u}$ by removing all states v from T_i^u for which $x_u^* = z_u^* = 0$. (Note: Since we may be removing some children of states remaining in $\hat{T_i^u}$, the martingale property may no longer hold with respect to $\hat{T_i^u}$.) Let us create $\hat{LP}_u(\mathcal{C}_u)$ by replacing T_i^u in $LP_u(\mathcal{C}_u)$ with $\hat{T_i^u}$. $\hat{LP}_u(\mathcal{C}_u)$ will have the same objective value as $LP_u(\mathcal{C}_u)$ and an optimal solution $\langle \hat{w^*}, \hat{x^*}, \hat{z^*} \rangle$ such that $\hat{w_v^*} = w_v^*$ and $\hat{x_v^*} = x_v^*$ for every state $v \in \hat{T_i^u}$.

Let us define the number of states in $\hat{T_i^u}$ as $\hat{\Sigma_u}$. $\hat{LP}_u(\mathcal{C}_u)$ has $3\hat{\Sigma_u}$ variables, $2\hat{\Sigma_u}$ non-negativity constraints, and $2\hat{\Sigma_u} + 1$ other constraints. Thus $\hat{LP}_u(\mathcal{C}_u)$ has a basic feasible optimal solution, which will have at least $\hat{\Sigma_u} - 1$ variables equal to zero. (For discussion of LP theory and the role of basic feasible solutions, see [20] or a similar resource.)

If exactly $\hat{\Sigma_u} - 1$ variables are equal to zero, all constraints of the type $x_v + z_v \leq w_v$ must be tight. By virtue of how we created $\hat{T_i^u}$, we know that each of these zero variables must correspond to a distinct state. Thus, for these $\hat{\Sigma_u} - 1$ states either $\hat{x_v^*} = \hat{w_v^*} > 0$ or $\hat{z_v^*} = \hat{w_v^*} > 0$, and for the only remaining state, (call it y), $\hat{x_y^*} > 0$ and $\hat{z_y^*} > 0$.

Alternatively there could be $\hat{\Sigma_u}$ variables equal to zero (but no more since for all states $\hat{x_v^*} + \hat{z_v^*} > 0$), in which case, for at least $\hat{\Sigma_u} - 1$ states either $\hat{x_v^*} = \hat{w_v^*} > 0$ or $\hat{z_v^*} = \hat{w_v^*} > 0$.

Looking at this in terms of $\hat{p}$ and $\hat{e}$, in at least $\hat{\Sigma_u} - 1$ states either $\hat{p_v^*} = 1$ or $\hat{e_v^*} = 1$. For all states v in $T_i^u / \hat{T_i^u}$ we could arbitrarily assign p_v and e_v, so this property holds with respect to $LP_u(\mathcal{C}_u)$ as well.

■

From the analysis above, we can see that for a given $\mathcal{C}_u$, $LP_u(\mathcal{C}_u)$ has an optimal basic feasible solution $\langle w^*, x^*, z^* \rangle$ that takes one of the following three forms:

1. For every state v where $x_v^* + z_v^* > 0$, $x_v^* + z_v^* = w_v^*$, and in exactly one state y, $x_y^* > 0$ and $z_y^* > 0$. In this case exactly $3\hat{\Sigma_u}$ constraints of $\hat{LP}_u(\mathcal{C}_u)$ are binding (including the cost constraint).

2. For every state v where $x_v^* + z_v^* > 0$, either $x_v^* = 0$ or $z_v^* = 0$, and in exactly one state y, $x_y^* + z_y^* < w_y^*$. In this case exactly $3\hat{\Sigma_u}$ constraints of $\hat{LP}_u(\mathcal{C}_u)$ are binding (including the cost constraint).

3. For every state v where $x_v^* + z_v^* > 0$, either $x_v^* = 0$ or $z_v^* = 0$, and $x_v^* + z_v^* = w_v^*$. In this case either $3\hat{\Sigma_u} + 1$ constraints of $\hat{LP}_u(\mathcal{C}_u)$ are binding, or else the cost constraint is not binding which implies the slope of the profit curve at this point is zero.

In the first two types of policies, we will call the state y for which it does not hold that $p_y^*, e_y^* \in \{0, 1\}$ the *transitional state*. Note that policy type 3 does not have a transitional state, and thus corresponds to a deterministic policy.

Further leveraging our understanding of the basic feasible solutions of $LP_u()$ and linear programming theory, we make the following pair of claims.

Claim D.3 *Let $\langle e^*, p^* \rangle$ be a basic feasible optimal policy for $LP_u(\mathcal{C}_u)$. If $\langle e^*, p^* \rangle$ has a transitional state, then there exists an $\epsilon > 0$ such that there exists an optimal policy $\langle e^{*(1)}, p^{*(1)} \rangle$ for $LP_u(\mathcal{C}_u^{(1)})$ where $\mathcal{C}_u^{(1)} \in [\mathcal{C}_u - \epsilon, \mathcal{C}_u + \epsilon]$ that has the same transitional state y and for all states $v \neq y$ $e_v^{*(1)} = e_v^*$ and $p_v^{*(1)} = p_v^*$.*

Proof: This result follows directly from linear programming theory. Given that we have a transitional state, we have a non-degenerate solution to $\hat{LP}_u()$ (with exactly $3\hat{\Sigma}_u$ constraints at equality). Thus, small changes in the budget will result in a solution that has the same set of binding constraints. These results carry over to $LP_u()$. (We refer the reader to [20] or another suitable optimization text for more discussion of linear programming theory and sensitivity analysis.)

With respect to our problem, this implies that for every non-transitional state v, if $x_v^* = 0, x_v^{*(1)} = 0$, if $z_v^* = 0, z_v^{*(1)} = 0$, if $x_v^* = w_v^*, x_v^{*(1)} = w_v^{*(1)}$, and if $z_v^* = w_v^*, z_v^{*(1)} = w_v^{*(1)}$. Or equivalently, $e_v^{*(1)} = e_v^*, p_v^{*(1)} = p_v^*$.

■

This leads naturally to the following result:

Claim D.4 *The profit curve is piecewise linear. Further, each "corner" solution or point connecting two segments of the curve can be achieved by a deterministic policy.*

Proof: Again by linear programming theory, beginning at a non-degenerate solution to $\hat{LP}_u()$ and making incremental changes to the budget will change the optimal objective value at a constant rate (the shadow price of the budget constraint) until a new constraint becomes binding. With respect to $LP_u(\mathcal{C}_u)$ this implies that if the optimal solution to this LP has a transitional state, then there exist some $C^{(1)} < \mathcal{C}_u$ and $C^{(2)} > \mathcal{C}_u$ such that the segment of the profit curve from $\mathcal{P}_u(C^{(1)})$ to $\mathcal{P}_u(C^{(2)})$ is linear.

Furthermore, since we know that each end of this line segment must have an additional binding constraint, the optimal policy associated with these points on the curve must be of type 3 above (i.e. a deterministic policy).

■

This result, combined with the fact that profit curves are concave, implies that for any state, u, the ratio index can always be calculated with infinitesimal cost, i.e. $r(u,h) = \mathcal{P}'_{+u}(0)$, where $\mathcal{P}'_{+u}(\mathcal{C}_u)$ denotes the right-sided derivative of the profit curve evaluated at $\mathcal{C}_u$. Further, the policy at the end of the first line segment of $\mathcal{P}_u(\cdot)$ is a deterministic ratio index policy.

Before we prove Lemma 5.2, we will need to establish the following four properties. In the following, recall that $\mathcal{P}'_{+u}(\mathcal{C}_u)$ represents the right-sided derivative of $\mathcal{P}_u()$ evaluated at $\mathcal{C}_u$. Correspondingly, $\mathcal{P}'_{-u}(\mathcal{C}_u)$ represents the left-sided derivative of $\mathcal{P}_u()$ evaluated at $\mathcal{C}_u$.

Claim D.5 *There exists an optimal solution $\langle e^*, p^* \rangle$ to $\mathrm{LP}_u(C)$, such that for any state v where $w_v^* > 0$:*

1. *If $p_v^* > 0, \mathcal{P}'_{-v}(1) \geq \mathcal{P}'_{-u}(C)$*
2. *If $e_v^* > 0, \mathcal{P}'_{+u}(C) \geq \mathcal{P}'_{-v}(1)$*
3. *If $e_v^* > 0, r(v,h) \geq \mathcal{P}'_{-u}(C)$*
4. *If $1 - e_v^* - p_v^* > 0, \mathcal{P}'_{+u}(C) \geq r(v,h)$*

Proof: First, consider the set of all optimal solutions to $LP_u(C)$. Arbitrarily observe one optimal solution $\langle \bar{e}, \bar{p} \rangle$. Use $\mathcal{S}$ to denote the set of all optimal solutions $\langle e, p \rangle$ to $LP_u(C)$ such that $w_v = \bar{w}_v$. Select $\langle e^*, p^* \rangle \in \mathcal{S}$ such that $p_v^* = \max_{\mathcal{S}} p_v, p_v^* + e_v^* = \max_{\mathcal{S}} p_v + e_v$. We now prove the four properties.

1. If $p_v^* > 0, \mathcal{P}'_{-v}(1) \geq \mathcal{P}'_{-u}(C)$: Given v such that $w_v^* > 0$ and $p_v^* > 0$ assume $\mathcal{P}'_{-v}(1) < \mathcal{P}'_{-u}(C)$. Denote by E_v^* the amount of budget devoted to v and its descendants. Thus, we know that total profit garnered from v and its descendants is $\mathcal{P}_v(E_v^*)$. Let $\langle \tilde{e}, \tilde{p} \rangle_v$ be the policy that induces $\mathcal{P}_v(E_v^* - \epsilon)$ on T_i^v. Let $\langle \tilde{e}, \tilde{p} \rangle$ be a policy that follows $\langle \tilde{e}, \tilde{p} \rangle_v$ in the subtree of v and follows $\langle e^*, p^* \rangle$ otherwise. The policy $\langle \tilde{e}, \tilde{p} \rangle$ is a feasible solution to the problem of finding the point on the profit curve of u with cost $C - w_v^* \epsilon$. This policy has profit $\mathcal{P}_u(C) - w_v^* \epsilon \mathcal{P}'_{-v}(E_v^*)$. Further, as $p_v^* > 0$, we know that either $p_v^* = 1$ (in which case $E_v^* = 1$) or else v is the transitional state for $LP_v(E_v^*)$. In the latter case, one end of the line segment of $\mathcal{P}_v()$ which contains $\mathcal{P}_v(E_v^*)$ must have $p_v^* = 1$. This obviously corresponds to $\mathcal{P}_v(1)$. Thus, $\mathcal{P}_v(E_v^*)$ must be on the last segment of the profit curve of v, so $\mathcal{P}'_{-v}(E_v^*) = \mathcal{P}'_{-v}(1)$. Thus,

$$\begin{aligned} \mathcal{P}_u(C - w_v^* \epsilon) &\geq \mathcal{P}_u(C) - w_v^* \epsilon \mathcal{P}'_{-v}(1) \\ &> \mathcal{P}_u(C) - w_v^* \epsilon \mathcal{P}'_{-u}(C) \end{aligned}$$

   Thus, $\mathcal{P}_u(C) - \mathcal{P}_u(C - w_v^* \epsilon) < w_v^* \epsilon \mathcal{P}'_u(C)$. By concavity, this cannot be true, so it must be true that $\mathcal{P}'_{-v}(1) \geq \mathcal{P}'_{-u}(C)$.

2. If $e_v^* > 0, \mathcal{P}'_{+u}(C) \geq \mathcal{P}'_{-v}(1)$: Given v such that $w_v^* > 0$ and $e_v^* > 0$ assume $\mathcal{P}'_{-v}(1) > \mathcal{P}'_{+u}(C)$. Denote by E_v^* the amount of budget devoted to v and its descendants. Observe that it must be the case that $E_v^* < 1$. Otherwise there exists some other optimal solution $\langle e^{*(1)}, p^{*(1)} \rangle$ to $LP_u(C)$ with $p_v^{*(1)} = 1, e_v^{*(1)} = 0$. Thus, we know that total profit garnered from v and its descendants is $\mathcal{P}_v(E_v^*)$. Let $\langle \tilde{e}, \tilde{p} \rangle_v$ be the policy that induces $\mathcal{P}_v(E_v^* + \epsilon)$ on T_i^v. Let $\langle \tilde{e}, \tilde{p} \rangle$ be a policy that follows $\langle \tilde{e}, \tilde{p} \rangle_v$ in the subtree of v and follows $\langle e^*, p^* \rangle$ otherwise. The policy $\langle \tilde{e}, \tilde{p} \rangle$ is a feasible solution to the problem of finding the point on the profit curve of u with cost $C + w_v^* \epsilon$. This policy has profit $\mathcal{P}_u(C) + w_v^* \epsilon \mathcal{P}'_{+v}(E_v^*)$. Thus,

$$\mathcal{P}_u(C + w_v^* \epsilon) \geq \mathcal{P}_u(C) + w_v^* \epsilon \mathcal{P}'_{+v}(E_v^*)$$

   Further, as $E_v^* < 1$, $\mathcal{P}'_{+v}(E_v^*) \geq \mathcal{P}'_{-v}(1)$. Thus,

$$\mathcal{P}_u(C + w_v^* \epsilon) \geq \mathcal{P}_u(C) + w_v^* \epsilon \mathcal{P}'_{-v}(1)$$

$$\Rightarrow \mathcal{P}_u(C + w_v^* \epsilon) - \mathcal{P}_u(C) > w_v^* \epsilon \mathcal{P}'_u(C)$$

   By concavity, this cannot be true, so it must be true that $\mathcal{P}'_{+u}(C) \geq \mathcal{P}'_{-v}(1)$.

3. If $e_v^* > 0, r(v, h) \geq \mathcal{P}'_{-u}(C)$: Given v such that $w_v^* > 0$ and $e_v^* > 0$ assume $r(v, h) < \mathcal{P}'_{-u}(C)$. Denote by E_v^* the amount of budget devoted to v and its descendants. Thus, we know that total profit garnered from v and its descendants is $\mathcal{P}_v(E_v^*)$. Let $\langle \tilde{e}, \tilde{p} \rangle_v$ be the policy that induces $\mathcal{P}_v(E_v^* - \epsilon)$ on T_i^v. Let $\langle \tilde{e}, \tilde{p} \rangle$ be a policy that follows $\langle \tilde{e}, \tilde{p} \rangle_v$ in the subtree of v and follows $\langle e^*, p^* \rangle$ otherwise. The policy $\langle \tilde{e}, \tilde{p} \rangle$ is a feasible solution to the problem of finding the point on the profit curve of u with cost

$C - w_v^* \epsilon$. This policy has profit $\mathcal{P}_u(C) - w_v^* \epsilon \mathcal{P}'_{-v}(E_v^*)$. Further, as $E_v^* > 0$ (since $e_v^* > 0$), we know that $\mathcal{P}'_{-v}(E_v^*) \leq r(v,h)$. Thus,

$$\begin{aligned}\mathcal{P}_u(C - w_v^* \epsilon) &\geq \mathcal{P}_u(C) - w_v^* \epsilon r(v,h) \\ &> \mathcal{P}_u(C) - w_v^* \epsilon \mathcal{P}'_{-u}(C)\end{aligned}$$

Thus, $\mathcal{P}_u(C) - \mathcal{P}_u(C - w_v^* \epsilon) < w_v^* \epsilon \mathcal{P}'_u(C)$. By concavity, this cannot be true, so it must be true that $\mathcal{P}'_{-v}(1) \geq \mathcal{P}'_{-u}(C)$.

4. If $1 - e_v^* - p_v^* > 0, \mathcal{P}'_{+u}(C) \geq r(v,h)$: Given v such that $w_v^* > 0$ and $1 - e_v^* - p_v^* > 0$ assume $r(v,h) > \mathcal{P}'_{+u}(C)$. Denote by E_v^* the amount of budget devoted to v and its descendants. Thus, we know that total profit garnered from v and its descendants is $\mathcal{P}_v(E_v^*)$. Let $\langle \tilde{e}, \tilde{p} \rangle_v$ be the policy that induces $\mathcal{P}_v(E_v^* + \epsilon)$ on T_i^v. Let $\langle \tilde{e}, \tilde{p} \rangle$ be a policy that follows $\langle \tilde{e}, \tilde{p} \rangle_v$ in the subtree of v and follows $\langle e^*, p^* \rangle$ otherwise. The policy $\langle \tilde{e}, \tilde{p} \rangle$ is a feasible solution to the problem of finding the point on the profit curve of u with cost $C + w_v^* \epsilon$. This policy has profit $\mathcal{P}_u(C) + w_v^* \epsilon \mathcal{P}'_{+v}(E_v^*)$. Thus,

$$\mathcal{P}_u(C + w_v^* \epsilon) \geq \mathcal{P}_u(C) + w_v^* \epsilon \mathcal{P}'_{+v}(E_v^*)$$

Further, as $1 - e_v^* - p_v^* > 0$, $\mathcal{P}'_{+v}(E_v^*) = r(v,h)$. (Since neither $e_v^* = 1$ nor $p_v^* = 1$, at least one must be zero and the other is either zero or the transitional state. Thus, at one end of this segment of $\mathcal{P}_v()$ must be the abandonment policy, so we are on the first segment of the profit curve of v.) Thus,

$$\mathcal{P}_u(C + w_v^* \epsilon) \geq \mathcal{P}_u(C) + w_v^* \epsilon r(v,h)$$

$$\Rightarrow \mathcal{P}_u(C + w_v^* \epsilon) - \mathcal{P}_u(C) > w_v^* \epsilon \mathcal{P}'_u(C)$$

By concavity, this cannot be true, so it must be true that $\mathcal{P}'_{+u}(C) \geq r(v,h)$.

■

This result directly implies Lemma 2.3, which we now prove.

Corollary D.6 (same as Lemma 2.3) *A ratio index policy for arm-state u, $\pi_r(u,h)$, does not abandon any arm-state v with $r(v,h) > r(u,h)$ and does not explore or exploit any arm-state v with $r(v,h) < r(u,h)$.*

Proof: With respect to Lemma D.5, select C such that $0 < C < \mathcal{C}(\pi_r(u,h)$. Thus, $\mathcal{P}'_{-u}(C) = \mathcal{P}'_{+u}(C) = r(u,h)$. since $\mathcal{P}'_{-v}(1) \leq r(v,h)$, the first, third and fourth properties proved in Lemma D.5 respectively imply that for the optimal randomized policy corresponding to this point on the profit curve: (1) if $r(u,h) > r(v,h)$, then $p_v^* = 0$; (3) if $r(u,h) > r(v,h)$, then $e_v^* = 0$; and (4) if $r(v,h) > r(u,h)$, then $1 - e_v^* - p_v^* = 0$. As v cannot be a transitional state under any of these conditions, then by the argument in Claim D.4 these properties must hold for the ratio index policy as well. ■

We are now ready to prove the following monotonicity result, which follows from the concavity of the profit curve.

Claim D.7 (same as Lemma 5.2) *For any state u and $C^{(1)}$ and $C^{(2)}$, where $0 \leq C^{(1)} < C^{(2)} \leq 1$, there exist optimal solutions $\langle e^{(1)}, p^{(1)} \rangle$ and $\langle e^{(2)}, p^{(2)} \rangle$ to $LP_u(C^{(1)})$ and $LP_u(C^{(2)})$ respectively such that $p_v^{(1)} \leq p_v^{(2)}$ and $e_v^{(1)} + p_v^{(1)} \leq e_v^{(2)} + p_v^{(2)}$ for all v in T_i^u.*

Proof: Given the four properties from Lemma D.5, we can see that for any state v, as the slope of $\mathcal{P}_u(C)$ decreases with increasing C, there will be up to three regions where we would first abandon at v, then we would explore at v, then we would exploit. The only remaining technical issue that remains is if the slope of the profit curve of u exactly equals $\mathcal{P}'_{-v}(1)$ or $r(v)$. In these cases, we are of necessity on a line segment of $\mathcal{P}_u()$ where v is at some point the transitional state.[3] In these cases, at the two endpoints of the segment we must have one of the three following pairs of values for v:

1. $p_v = e_v = 0$; $p_v = 1$
2. $p_v = e_v = 0$; $e_v = 1$
3. $e_v = 1$; $p_v = 1$

For the first two of these cases, it is obvious that the state on the right corresponds to higher cost. Thus, our monotonicity property holds. For the third case, we know by the martingale property that setting $p_v = 1$ yields maximum possible profit at the node. Thus, the monotonicity property must hold in this case as well. ■

We index the B_u "corner" solutions on $\mathcal{P}_u(\mathcal{C}_u)$ as $s_i, ..., s_{B_u}$ corresponding to the budget associated with their rightmost end point. The above claim further implies that if for every state $v \in T_i^u$ and every "corner" solution, s_i on $\mathcal{P}_u(\mathcal{C}_u)$ we apply a label $L_{s_i}(v) \in \{A, E, P\}$ where A corresponds to "abandoning" at v ($p_v = e_v = 0$), E corresponds to "exploring" ($e_v = 1$), and P corresponds to "exploiting" ($p_v = 1$), then for any v and any i, j such that $i \leq j$, if $L_{s_i}(v) = P$, then $L_{s_j}(v) = P$, and if $L_{s_i}(v) = E$, then $L_{s_j}(v) \neq A$. Thus, we can find a set of solutions such that as we increase the budget, once a state is labeled "P" it will always remain a "P" and once it becomes an "E" it will never become an"A". Thus, each state can only change labels at most twice. As every successive "corner" solution must involve changing the label of at least one state, the profit curve for a state can have at most $2\Sigma_u$ segments (where Σ_u represents the number of states in T_i^u). This establishes the following claim which completes the proof of theorem 5.1.

Claim D.8 *The profit curve $\mathcal{P}_u(\mathcal{C}_u)$ for any state u can have at most $2\Sigma_u$ segments, where Σ_u is the number of states in T_i^u.*

D.2 Algorithm for Computing the Profit Curve

The algorithm for computing the profit curve (and hence the ratio index) involves recursively calculating the profit curve for a state given the profit curves for all of its successor states. In this process, given the profit curve for all subsequent states of a given state u, we begin by constructing an intermediate curve, the *exploration profit curve* of u ($\mathcal{X}_u()$). This curve denotes the optimal profit for any given cost contingent upon exploring at u (i.e. $e_u = 1$). Once we have found the exploration profit curve, we then take the concave envelope over this curve combined with the abandonment policy and the exploitation policy to find the profit curve for a state (see figure 1 for an illustration).

Under these conditions, we can modify the recursive equation we introduced earlier to calculate the profit curve for u ($\mathcal{P}_u()$) into a simpler equation to calculate the exploration profit curve

$$\mathcal{X}_u(\mathcal{C}_u) = \max_{E^u} \sum_{v \in D(u)} P_{uv} \mathcal{P}_v(E_v^u)$$

[3]It is possible to have a line segment of the profit curve with multiple transitional states (e.g. in the case of two states with identical ratio indices). In this case, we can always look at a sequence of basic feasible solutions such that the transitions are ordered and there is at most one transitional state at any point on the segment.

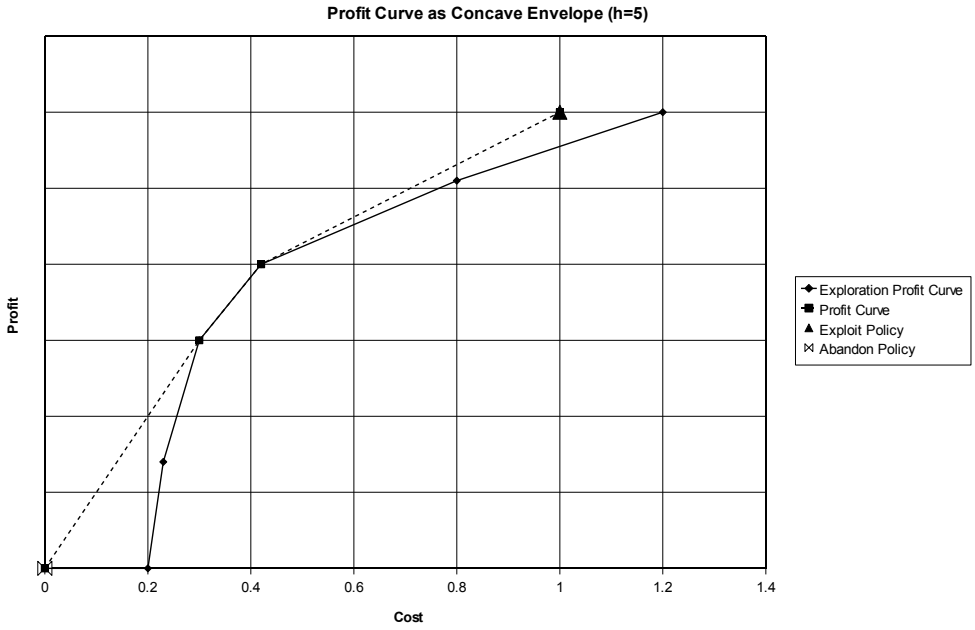

Figure 1: The relation between the profit and the exploration-profit curves

such that:

$$1/h + \sum_{v \in D(u)} P_{uv} E_v^u \leq \mathcal{C}_u$$
$$E^u \geq 0$$

Recall that E_v^u is the budget allocated to v should we visit it immediately after u, and E^u is the vector of these budgets for each of the immediate descendants of u.

For each $v \in D(u)$, let B_v represent the number of "corner" solutions on $\mathcal{P}_v()$ and denote the cost of the ith such "corner" solution as s_i^v (where $s_0^v = 0$ and $s_{B_v}^v = 1$). We can then create the following modified recursive equation to calculate the exploration profit curve.

$$\max_{\epsilon^u} \textstyle\sum_{v \in D(u)} P_{uv} \sum_{i=1}^{B_v} [\mathcal{P}_v(s_i^v) - \mathcal{P}_v(s_{i-1}^v)] \epsilon_{v,i}^u$$

s.t.

$$1/h + \textstyle\sum_{v \in D(u)} P_{uv} \sum_{i=1}^{B_v} (s_i^v - s_{i-1}^v) \epsilon_{v,i}^u \leq \mathcal{C}_u$$
$$0 \leq \epsilon_{v,i}^u \leq 1 \qquad \forall v \in D(u), i \in \{1, \ldots, B_v\}$$

This is a linear program where each decision variable, $\epsilon_{v,i}^u$, selects some fraction of the ith segment of the profit curve of v. ϵ^u represents the collection of all such variables. As the profit curve is concave (by claim D.1), $\frac{\mathcal{P}_v(s_i^v) - \mathcal{P}_v(s_{i-1}^v)}{s_i^v - s_{i-1}^v} \leq \frac{\mathcal{P}_v(s_j^v) - \mathcal{P}_v(s_{j-1}^v)}{s_j^v - s_{j-1}^v}$ $\forall i \geq j$. Thus, there exists an optimal solution which only assigns $\epsilon_{v,i}^u > 0$ if $\epsilon_{v,i-1}^u = 1$ for any $i > 1$.

Further, through inspection we can see that the optimal solution for any $\mathcal{C}_u$ is to select the segments in order of decreasing slope (where the slope of the ith segment of $\mathcal{P}_v()$ is $\frac{\mathcal{P}_v(s_i^v) - \mathcal{P}_v(s_{i-1}^v)}{s_i^v - s_{i-1}^v}$) until all budget is exhausted. By ordering segments thus, we can easily construct $\mathcal{X}_u()$. The algorithm below orders

the segments of the elements of $D(u)$ and builds $\mathcal{X}_u()$, storing the costs (c_i) and budgets allocated to all descendants ($E_v^u = \sum_i \epsilon_{v,i}^u$) for each "corner" solution of $\mathcal{X}_u()$.

Algorithm: COMPUTEEXPLORATIONPROFITCURVE

1) For all $v \in D(u), i \in \{1, ..., B_v\}$
   /*Compute profit, cost, and slope for each line segment of $\mathcal{P}_v()$*/
   Set $\pi_i^v = \mathcal{P}_v(s_i^v) - \mathcal{P}_v(s_{i-1}^v)$
   Set $c_i^v = s_i^v - s_{i-1}^v$
   Set $M_i^v = \pi_i^v / c_i^v$
2) Sort the $\sum_{v \in D(u)} B_v$ elements of the form M_i^v from largest to smallest
   Let $d(j)$ index the node for the jth largest element in the list
   Let $t(j)$ indicate which segment of V_j this is
   /*$M_{t(j)}^{d(j)}$ = slope of the jth line segment in the ordered list.*/
3) Set $\hat{S}_0 = \frac{1}{h}$ /*fixed cost*/, $\hat{X}_0 = 0$ /*initial profit*/, $\hat{\mathcal{E}}_{v,0} = 0 \ \forall v \in D(u)$ /*initial budgets for descendants*/
4) For $i = 1$ to $\sum_{v \in D(u)} B_v$
   /*Add next segment to the curve*/
   /*Compute current total profit ($\hat{X}_i$) and cost ($\hat{S}_i$)*/
   Let $\hat{X}_i = \hat{X}_{i-1} + P_{u,d(i)} \pi_{t(i)}^{d(i)}$
   Let $\hat{S}_i = \hat{S}_{i-1} + P_{u,d(i)} c_{t(i)}^{d(i)}$
   /*Compute budgets allocated to descendants ($\hat{\mathcal{E}}_{v,i}$) at the end*/
   /*of each segment (needed to represent the policy)*/
   Let $\hat{\mathcal{E}}_{d(i),i} = \hat{\mathcal{E}}_{d(i),i-1} + c_{t(i)}^{d(i)}$
   Let $\hat{\mathcal{E}}_{v,i} = \hat{\mathcal{E}}_{v,i-1} \ \forall v \neq d(i)$
5) Set $n = 1$, $k_1 = 1$
6) For $i = 2$ to $\sum_{v \in D(u)} B_v$
   /*Find changes in slope on the exploration profit curve*/
   If $\frac{\hat{X}_i - \hat{X}_{i-1}}{\hat{S}_i - \hat{S}_{i-1}} \neq \frac{\hat{X}_{i-1} - \hat{X}_{i-2}}{\hat{S}_{i-1} - \hat{S}_{i-2}}$
      $n++$
      $k_n = i$
7) For $m = 1$ to n
   /*Merge together segments of the exploration profit*/
   /*curve with the same slope*/
   $S_m = \hat{S}_{k_m}, \mathcal{X}_u(c_m) = \hat{X}_{k_m}$, $\mathcal{E}_{v,m}^u = \hat{\mathcal{E}}_{v,k_m}$
8) $B_u^{\mathcal{X}} = n$

Algorithm COMPUTEEXPLORATIONPROFITCURVE represents the exploration profit curve for state u by returning the number of line segments of the curve (not including the zero slope segment from cost of 0 to $S_0 = 1/h$), $B_u^{\mathcal{X}}$, as well as the cost (S_i), profit ($\mathcal{X}_u(S_i)$), and vector of budgets to allocate to all immediate descendants ($\mathcal{E}_{\cdot,i}^u$) corresponding to the endpoint of each segment.

Given that the profit curves of all descendants are concave, the sorting of line segments in step 2) equates to simply interleaving the segments of the different states and can be performed in $O(d\Sigma_u \log \Sigma_u)$ time using a simple min-heap, where d is the maximum number of immediate descendants for a node. After sorting these segments, steps 3) and 4) then determine the cost and profit associated with adding each segment to

the exploration profit curve. Finally, as there may and likely will be duplicates in the sorted list of slopes, steps 5) through 7) merge all segments of the curve with the same slope.

Given the exploration profit curve, it is much easier to calculate the profit curve for a state. From lemma D.1, we know that $\mathcal{P}_u(1) = \zeta(u)$. Further, this must be the only "corner" solution corresponding to exploiting at u ($p_u = 1$). All other "corner" solutions must thus correspond to exploring at u ($e_u = 1$) and thus must correspond to points on $\mathcal{X}_u()$. As we know $\mathcal{P}_u()$ is concave, we can simply take the concave envelope of the points $(0, \mathcal{P}_u(0) = 0)$, $(S_i, \mathcal{X}_u(S_i))\forall i \in \{1, \ldots, B_u^{\mathcal{X}}\}$, and $(1, \mathcal{P}_u(1) = \zeta(u))$. The algorithm below does this.

Algorithm: COMPUTEPROFITCURVE

1) Find $j^* = \arg\max_{j \in \{1,\ldots,B_u^{\mathcal{X}}\}} R_j$ where $R_j = \frac{\mathcal{X}_u(S_j)}{S_j}$

2a) If $R_{j^*} \leq \zeta(u)$ /*The ratio index policy exploits immediately*/

Set $r(u) = \zeta(u), s_1^u = 1, \mathcal{P}_u(s_1^u) = \zeta(u), E_v^u(1) = 0 \ \ \forall v \in D(u), B_u = 1$

2b) Else /*The ratio index policy explores*/

Set $r(u) = R_{j^*}, s_1^u = S_{j^*}, \mathcal{P}_u(s_1^u) = \mathcal{X}_u(S_{j^*}), E^u(1) = \mathcal{E}^u_{\cdot,j^*}$

Set $i = 1, j = j^* + 1$.

While $\frac{\mathcal{X}_u(S_j) - \mathcal{X}_u(S_{j-1})}{S_j - S_{j-1}} > \frac{\zeta(u) - \mathcal{X}_u(S_{j-1})}{1 - S_{j-1}}$

/*Greater marginal return to explore than exploit*/

$s_{i+1}^u = S_j$

$\mathcal{P}_u(s_{i+1}^u) = \mathcal{X}_u(S_j)$

$E^u(i+1) = \mathcal{E}^u_{\cdot,j}$

$i = i + 1, j = j + 1$

Set $s_{i+1}^u = 1, \mathcal{P}(s_{i+1}^u) = \zeta(u), B_u = i + 1$.

Algorithm COMPUTEPROFITCURVE runs in $O(d\Sigma_u)$ time. Step 1) computes the value R_j for each segment of $\mathcal{X}_u()$, which represents to slope of the line segment from the origin to the point $(c_i, \mathcal{X}_u(c_i))$. In the event that $R_{j^*} \leq \zeta(u)$, the ratio index policy is to exploit immediately and we are done determining the profit curve for u (step 2a). Otherwise, once we have found the ratio index policy, we continue to look at higher budget exploration policies to determine subsequent segments of the profit curve (step 2b). The quantity $E^u(i)$ represents the budgets allocated to each of the immediate descendants at the end of the ith segment of the profit curve. These values are only required to represent the actual policy, not calculate the ratio index or profit curve of any state. The quantity $\frac{\zeta(u) - \mathcal{X}_u(c_{j-1})}{1 - c_{j-1}}$ is the marginal ratio of transitioning from s_i^u to exploitation at u. Once the slopes of the segments of $\mathcal{X}_u()$ are no larger than this, it is optimal to transition to exploitation at u.

As each step above takes at most $O(d\Sigma_u \log \Sigma_u)$ time, the total time to compute the profit curves for all states in the layered DAG is $O(d\Sigma^2 \log \Sigma)$.