

CONDITIONS FOR INDEXABILITY OF RESTLESS BANDITS AND AN $\mathcal{O}(K^3)$ ALGORITHM TO COMPUTE WHITTLE INDEX

NIMA AKBARZADEH * ** AND
ADITYA MAHAJAN,* *** *McGill University*

Abstract

Restless bandits are a class of sequential resource allocation problems concerned with allocating one or more resources among several alternative processes where the evolution of the processes depends on the resources allocated to them. Such models capture the fundamental trade-offs between exploration and exploitation. In 1988, Whittle developed an index heuristic for restless bandit problems which has emerged as a popular solution approach because of its simplicity and strong empirical performance. The Whittle index heuristic is applicable if the model satisfies a technical condition known as indexability. In this paper, we present two general sufficient conditions for indexability and identify simpler-to-verify refinements of these conditions. We then revisit a previously proposed algorithm called the adaptive greedy algorithm which is known to compute the Whittle index for a sub-class of restless bandits. We show that a generalization of the adaptive greedy algorithm computes the Whittle index for all indexable restless bandits. We present an efficient implementation of this algorithm which can compute the Whittle index of a restless bandit with K states in $\mathcal{O}(K^3)$ computations. Finally, we present a detailed numerical study which affirms the strong performance of the Whittle index heuristic.

Keywords: Multi-armed bandits; restless bandits; Whittle index; indexability; stochastic scheduling; resource allocation

2020 Mathematics Subject Classification: Primary 90C40
Secondary 90C39; 49M20; 91B32

1. Introduction

Restless bandits are a class of sequential resource allocation problems concerned with allocating one or more resources among several alternative processes where the evolution of the processes depends on the resources allocated to them. Such models arise in various applications, such as machine maintenance [18], congestion control [7], healthcare [11], finance [15], channel scheduling [21], smart grid [1], and others.

Restless bandits are a generalization of classical multi-armed bandits [13], where the processes remain frozen when resources are not allocated to them. Gittins [14] showed that when a single resource is to be allocated among multiple processes, the optimal policy has a simple

Received 24 March 2020; revision received 26 November 2021.

* Postal address: Department of Electrical and Computer Engineering, McGill University, 3480 Rue University, Montréal, QC H3A 0E9.

** Email address: nima.akbarzadeh@mail.mcgill.ca

*** Email address: aditya.mahajan@mcgill.ca

© The Author(s), 2022. Published by Cambridge University Press on behalf of Applied Probability Trust.

structure: compute an index for each process and allocate the resource to the process with the largest (or the lowest) index. In contrast, the general restless bandit problem is PSPACE-hard [27]. Whittle [32] showed that index-based policies are optimal for the Lagrangian relaxation of the restless bandit problem and argued that the corresponding index, now called the Whittle index, is a reasonable heuristic for restless bandit problems. Subsequently, it has been shown that the Whittle index heuristic is optimal under some conditions [22, 31] and performs well in practice [5, 16, 19].

The Whittle index heuristic is applicable if a technical condition known as *indexability* is satisfied. Sufficient conditions for indexability have been investigated under specific modeling assumptions: two-state restless bandits [7, 21]; monotone bandits [5, 7, 15]; models with right-skip-free transitions [18, 19]; models with monotone or convex cost/reward [5–8, 19, 33]; models satisfying partial conservation laws [23, 25]; and models arising in specific applications [6, 8, 16, 18, 19].

Niño-Mora [25, 26] proposed a generalization of the Whittle index called the marginal productivity index (MPI) for resource allocation problems where processes can be allocated fractional resources. In [25], he also proposed an algorithm called the adaptive greedy algorithm, to compute the MPI when the model satisfies a technical condition called partial conservation laws (PCL). For restless bandits which satisfy the PCL condition, the Whittle index can be computed using the adaptive greedy algorithm. However, for general restless bandits, there are no known efficient algorithms to exactly compute the Whittle index. It is possible to *approximately* compute the Whittle index by conducting a binary search over a penalty for the active action (or a subsidy for the passive action) [2, 29], but such a binary search is computationally expensive, because each step of the binary search requires solving a dynamic program.

In this paper, we revisit the restless bandit problem and present three contributions. Our first contribution is to provide general sufficient conditions for indexability which are based on an alternative characterization of the passive set. We also present easy-to-verify refinements of these sufficient conditions.

Our second contribution is to use a novel geometric interpretation of the Whittle index to show that a refinement of the adaptive greedy algorithm proposed by Niño-Mora [25] computes the Whittle index for all indexable restless bandits. We provide a computationally efficient implementation, which computes the Whittle indices of a restless bandit with K states in $\mathcal{O}(K^3)$ computations.

Our third contribution is to present three special cases: (i) *restless bandits with optimal threshold-based policy*, which were previously studied in [3, 5, 7, 15, 17, 30], (ii) *stochastic monotone bandits* which may be considered as a generalization of monotone bandits [5, 7, 15], and (iii) *restless bandits with controlled restarts*, similar to [3, 30], which are a generalization of the restart models [18, 19]. We show that these models are always indexable and the Whittle index can be computed in closed form.

Finally, we present a detailed numerical study comparing the performance of the Whittle index policy with that of the optimal and myopic policies. Our study shows that in general, the performance of Whittle index policy is comparable to that of the optimal policy and considerably better than that of the myopic policy.

Notation. Uppercase letters (X, Y , etc.) denote random variables, lowercase letters (x, y , etc.) denote their realization, and script letters (\mathcal{X}, \mathcal{Y} , etc.) denote their state spaces. Subscripts denote time: so, X_t denotes a system variable at time t and $X_{1:t}$ is shorthand for the system variables (X_1, \dots, X_t) . $\mathbb{P}(\cdot)$ denotes the probability of an event; $\mathbb{E}[\cdot]$ denotes the expectation of

a random variable. \mathbb{Z} and \mathbb{R} denote the sets of integers and real numbers. Given a matrix P , P_{ij} denotes its (i, j) th element.

2. Restless bandits: problem formulation and solution concept

2.1. Restless bandit process

A discrete-time restless bandit process (RB) is a controlled Markov process $(\mathcal{X}, \{0, 1\}, \{P^i(a)\}_{a \in \{0,1\}}, c, x_0)$ where \mathcal{X} denotes the state space, which is a finite or countable set; $\{0, 1\}$ denotes the action space, where the action 0 is called the *passive* action and the action 1 is the *active* action; $P^i(a)$, $a \in \{0, 1\}$, denotes the transition matrix when the action a is chosen; $c : \mathcal{X} \times \{0, 1\} \rightarrow \mathbb{R}$ denotes the cost function; and x_0 denotes the initial state. We use X_t and A_t to denote the action of the process at time t . The process evolves in a controlled Markov manner; i.e., for any realization $x_{0:t+1}$ of $X_{0:t+1}$ and $a_{0:t+1}$ of $A_{0:t+1}$, we have $\mathbb{P}(X_{t+1} = x_{t+1} | X_{0:t} = x_{0:t}, A_{0:t} = a_{0:t}) = \mathbb{P}(X_{t+1} = x_{t+1} | X_t = x_t, A_t = a_t)$, which we denote by $P_{x_t x_{t+1}}(a_t)$.

2.2. Restless multi-armed bandit problem

A restless multi-armed bandit is a collection of n independent RBs $(\mathcal{X}^i, \{0, 1\}, \{P^i(a)\}_{a \in \{0,1\}}, c^i, x_0^i)$, $i \in \mathcal{N} := \{1, \dots, n\}$. A decision-maker observes the state of all RBs, may choose to activate only $m < n$ of them, and incurs a cost equal to the sum of the costs incurred by the individual RBs.

Let $\mathcal{X} := \prod_{i \in \mathcal{N}} \mathcal{X}^i$ and $\mathcal{A}(m) := \{a = (a^1, \dots, a^n) \in \mathcal{A}^n : \sum_{i \in \mathcal{N}} a^i = m\}$ denote the joint state space and the feasible action space, respectively. Let $X_t := (X_t^1, \dots, X_t^n)$ and $A_t = (A_t^1, \dots, A_t^n)$ denote the joint state and actions at time t . As the RBs evolve independently, for any realization $x_{0:t}$ of $X_{0:t}$ and $a_{0:t}$ of $A_{0:t}$, we have

$$\mathbb{P}(X_{t+1} = x_{t+1} | X_{0:t} = x_{0:t}, A_{0:t} = a_{0:t}) = \prod_{i=1}^n \mathbb{P}(X_{t+1}^i = x_{t+1}^i | X_t^i = x_t^i, A_t^i = a_t^i).$$

When the system is in the state $x_t = (x_t^1, \dots, x_t^n)$ and the decision-maker chooses action $a_t = (a_t^1, \dots, a_t^n)$, the system incurs a cost $\bar{c}(x_t, a_t) := \sum_{i \in \mathcal{N}} c^i(x_t^i, a_t^i)$. The decision-maker chooses his actions using a time-homogeneous Markov policy $g : \mathcal{X} \rightarrow \mathcal{A}(m)$, i.e., chooses $A_t = g(X_t)$. The performance of any Markov policy g is given by

$$J^g(x_0) := (1 - \beta) \mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t \bar{c}(X_t, g(X_t)) \mid X_0 = x_0 \right],$$

where $\beta \in (0, 1)$ is the discount factor and x_0 is the initial state of the system.

We are interested in the following optimization problem.

Problem 1. Given the discount factor $\beta \in (0, 1)$, the total number n of arms, the number m of active arms, RBs $(\mathcal{X}^i, \{0, 1\}, \{P^i(a)\}_{a \in \{0,1\}}, c^i, x_0^i)$, $i \in \mathcal{N}$, and the initial state $x_0 \in \mathcal{X}$, choose a Markov policy $g : \mathcal{X} \rightarrow \mathcal{A}(m)$ that minimizes $J^g(x_0)$.

Problem 1 is a multi-stage stochastic control problem and one can obtain an optimal solution using dynamic programming. However, the dynamic programming solution is intractable for large n since the cardinality of the state space is $\prod_{i \in \mathcal{N}} |\mathcal{X}^i|$, which grows exponentially with n . In the next section, we describe a heuristic known as the *Whittle index* to efficiently obtain a suboptimal solution of the problem.

2.3. Indexability and the Whittle index

Consider an RB $(\mathcal{X}, \{0, 1\}, \{P(a)\}_{a \in \{0,1\}}, c, x_0)$. For any $\lambda \in \mathbb{R}$, we consider a Markov decision process $\{\mathcal{X}, \{0, 1\}, \{P(a)\}_{a \in \{0,1\}}, c_\lambda, x_0\}$, where

$$c_\lambda(x, a) := c(x, a) + \lambda a, \quad \forall x \in \mathcal{X}, \forall a \in \{0, 1\}. \tag{1}$$

The parameter λ may be viewed as a penalty for taking the active action. The performance of any time-homogeneous policy $g : \mathcal{X} \rightarrow \{0, 1\}$ is

$$J_\lambda^{(g)}(x_0) := (1 - \beta) \mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t c_\lambda(X_t, g(X_t)) \mid X_0 = x_0 \right]. \tag{2}$$

Consider the following optimization problem.

Problem 2. Given the RB $(\mathcal{X}, \{0, 1\}, \{P(a)\}_{a \in \{0,1\}}, c_\lambda, x_0)$ and the discount factor $\beta \in (0, 1)$, choose a Markov policy $g : \mathcal{X} \rightarrow \{0, 1\}$ to minimize $J_\lambda^{(g)}(x_0)$.

Problem 2 is also a Markov decision process and one can obtain an optimal solution using dynamic programming. Let $V_\lambda : \mathcal{X} \rightarrow \mathbb{R}$ be the unique fixed point of the following:

$$V_\lambda(x) = \min \{ H_\lambda(x, 0), H_\lambda(x, 1) \}, \quad \forall x \in \mathcal{X}, \tag{3}$$

where

$$H_\lambda(x, a) = (1 - \beta)c_\lambda(x, a) + \beta \sum_{y \in \mathcal{X}} P_{xy}(a)V_\lambda(y), \quad a \in \{0, 1\}. \tag{4}$$

Let $g_\lambda(x)$ denote the minimizer of the right-hand side of (3), where we set $g_\lambda(x) = 1$ if $H_\lambda(x, 0) = H_\lambda(x, 1)$. Then, from Markov decision theory [28], we know that the time-homogeneous policy g_λ is optimal for Problem 2.

Define the *passive set* Π_λ to be the set of states where the passive action is optimal, i.e.,

$$\Pi_\lambda := \{x \in \mathcal{X} : g_\lambda(x) = 0\}. \tag{5}$$

Definition 1 (Indexability). An RB is indexable if Π_λ is increasing in λ , i.e., for any $\lambda', \lambda'' \in \mathbb{R}$, $\lambda' \leq \lambda''$ implies that $\Pi_{\lambda'} \subseteq \Pi_{\lambda''}$.

Definition 2 (Whittle index). The Whittle index of state x of an indexable RB is the smallest value of λ for which x is part of the passive set Π_λ , i.e., $w(x) = \inf \{ \lambda \in \mathbb{R} : x \in \Pi_\lambda \}$.

Alternatively, the Whittle index $w(x)$ is a value of the penalty λ for which the optimal policy is indifferent between taking the active and the passive action when the RB is in state x .

2.4. Whittle index heuristic

A restless multi-armed bandit problem is said to be indexable if all RBs are indexable. For indexable problems, the Whittle index heuristic is as follows: *Compute the Whittle indices of all arms offline. Then, at each time, obtain the Whittle indices of the current state of all arms and play arms with the m largest Whittle indices.*

As mentioned earlier, the Whittle index policy is a popular approach for restless bandits because (i) its complexity is linear in the number of alternatives and (ii) its performance is often close to optimal in practice [5, 16, 19]. However, there are only a few general conditions to check indexability for general models.

2.5. Alternative characterizations of the passive set

We now present alternative characterizations of the passive set, which are important for the sufficient conditions for indexability that we provide later.

Let Σ denote the family of all stopping times with respect to the natural filtration of $\{X_t\}_{t \geq 0}$. For any state $x \in \mathcal{X}$, penalty $\lambda \in \mathbb{R}$, and stopping time $\tau \in \Sigma$, define

$$\begin{aligned}
 M(x, \tau) &:= \mathbb{E}[\beta^\tau | X_0 = x, \{A_t = 0\}_{t=0}^{\tau-1}], \\
 L(x, \tau) &:= \mathbb{E}\left[\sum_{t=0}^{\tau-1} \beta^t c(X_t, 0) + \beta^\tau c(X_\tau, 1) \mid X_0 = x, \{A_t = 0\}_{t=0}^{\tau-1}\right], \\
 W_\lambda(x) &:= (1 - \beta)\lambda + \beta \sum_{y \in \mathcal{X}} P_{xy}(1)V_\lambda(x).
 \end{aligned}
 \tag{6}$$

Let $h_{\tau,\lambda}$ denote the (history-dependent) policy that takes the passive action up to time $\tau - 1$, takes the active action at time τ , and then follows the optimal policy g_λ (for Problem 2). We now present different characterizations of the passive set.

Proposition 1. *The following characterizations of the passive set are equivalent:*

- $\Pi_\lambda^{(a)} = \{x \in \mathcal{X} : g_\lambda(x) = 0\}$,
- $\Pi_\lambda^{(b)} = \{x \in \mathcal{X} : H_\lambda(x, 0) < H_\lambda(x, 1)\}$,
- $\Pi_\lambda^{(c)} = \left\{x \in \mathcal{X} : \exists \sigma \in \Sigma, \sigma \neq 0, \text{ such that } J_\lambda^{(h_{\sigma,\lambda})}(x) < J_\lambda^{(h_0)}(x)\right\}$,
- $\Pi_\lambda^{(d)} = \{x \in \mathcal{X} : \exists \sigma \in \Sigma, \sigma \neq 0, \text{ such that } (1 - \beta)(L(x, \sigma) - c(x, 1)) < W_\lambda(x) - \mathbb{E}[\beta^\sigma W_\lambda(X_\sigma) | X_0 = x]\}$.

See Appendix A for the proof.

3. Sufficient conditions for indexability

In this section, we identify sufficient conditions for an RB to be indexable.

3.1. Preliminary results

Consider an RB $(\mathcal{X}, \{0, 1\}, \{P(a)\}_{a \in \{0,1\}}, c, x_0)$. For any Markov policy $g : \mathcal{X} \rightarrow \{0, 1\}$ and $\lambda \in \mathbb{R}$, we can write

$$J_\lambda^{(g)}(x) = D^{(g)}(x) + \lambda N^{(g)}(x),
 \tag{7}$$

where

$$D^{(g)}(x) := (1 - \beta)\mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t c(X_t, g(X_t)) \mid X_0 = x\right]$$

and

$$N^{(g)}(x) := (1 - \beta)\mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t g(X_t) \mid X_0 = x\right]$$

are the expected discounted total cost and the expected number of activations under the policy g starting at the initial state x . $D^{(g)}(\cdot)$ and $N^{(g)}(\cdot)$ can be computed using policy evaluation formulas. In particular, define $P^{(g)} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ and $c^{(g)} : \mathcal{X} \rightarrow \mathbb{R}$ as follows:




FIGURE 1. An illustration of the plot of $J_\lambda^{(\cdot)}(x)$ versus λ for $g \in \mathcal{G} := \{h_1, h_2, h_3\}$. Let λ_c^{ij} denote the λ -value of the intersection of $J_\lambda^{(h_i)}(x)$ and $J_\lambda^{(h_j)}(x)$. Note that in this plot, for all $\lambda \in (-\infty, \lambda_c^{12}]$ the policy h_1 is optimal; for all $\lambda \in [\lambda_c^{12}, \lambda_c^{23}]$ the policy h_2 is optimal; and for all $\lambda \in [\lambda_c^{23}, \infty)$ the policy h_3 is optimal. The lower concave envelope of $J_\lambda^{(h_i)}(x)$ (shown as a thick line) is the value function $V_\lambda(x)$, which is piecewise linear, concave, increasing, and continuous.

$P_{xy}^{(g)} = P_{xy}(g(x))$ and $c_\lambda^{(g)}(x) = c_\lambda(x, g(x)) = c^{(g)}(x, g(x)) + \lambda g(x)$ for any $x \in \mathcal{X}$. We also view g as an element in $\{0, 1\}^{|\mathcal{X}|}$. Then, using the policy evaluation formula for infinite-horizon Markov decision processes [28], we obtain

$$D^{(g)}(x) = (1 - \beta)[(I - \beta P^{(g)})^{-1} c^{(g)}](x) \text{ and } N^{(g)}(x) = (1 - \beta)[(I - \beta P^{(g)})^{-1} g](x). \quad (8)$$

We now provide a geometric interpretation of the value function $V_\lambda(x)$ as a function of λ . For any $g \in \mathcal{G}$, the plot of $J_\lambda^{(g)}(x) = D^{(g)}(x) + \lambda N^{(g)}(x)$ as a function of λ is a straight line with y-intercept $D^{(g)}(x)$ and slope $N^{(g)}(x)$. By definition, $V_\lambda(x) = \inf_{g \in \mathcal{G}} J_\lambda^{(g)}(x)$. Thus, $V_\lambda(x)$ is the lower concave envelope of the family of straight lines $\{J_\lambda^{(g)}(x)\}_{g \in \mathcal{G}}$. (See Figure 1 for an illustration.) Thus, we have the following.

Lemma 1. For any $x \in \mathcal{X}$, $V_\lambda(x)$ is continuous, increasing, piecewise linear, and concave in λ . Furthermore, when \mathcal{X} is finite, $V_\lambda(x)$ is piecewise linear.

Proof. For any Markov policy g , $N^{(g)}(x)$ is non-negative. Therefore, $J_\lambda^{(g)}(x) = D^{(g)}(x) + \lambda N^{(g)}(x)$ is increasing and continuous in λ . Since $V_\lambda(x)$ is an infimization of a family of linear functions, it is concave (see Figure 1). In addition, as monotonicity and continuity are preserved under infimization, the value function is also increasing and continuous in λ . Finally, when \mathcal{X} is finite, there are only finite number of pieces. Thus, $V_\lambda(x)$ is the minimum of finitely many linear functions, and hence piecewise linear. \square

Lemma 2. For any $\lambda', \lambda'' \in \mathbb{R}$,

$$(\lambda'' - \lambda')N^{(g_{\lambda''})}(x) \leq V_{\lambda''}(x) - V_{\lambda'}(x) \leq (\lambda'' - \lambda')N^{(g_{\lambda'})}(x), \quad \forall x \in \mathcal{X}.$$

Consequently, $N^{(g_\lambda)}(x)$ is non-increasing in λ .

Proof. Recall that $V_\lambda(x) = J_\lambda^{(g_\lambda)}(x) \leq J_\lambda^{(g_{\lambda'})}(x)$ for any $\lambda' \neq \lambda$. Thus,

$$\begin{aligned} V_{\lambda''}(x) - V_{\lambda'}(x) &= J_{\lambda''}^{(g_{\lambda''})}(x) - J_{\lambda'}^{(g_{\lambda'})}(x) \leq J_{\lambda''}^{(g_{\lambda'})}(x) - J_{\lambda'}^{(g_{\lambda'})}(x) \\ &\stackrel{(a)}{=} (\lambda'' - \lambda')N^{(g_{\lambda'})}(x), \end{aligned} \quad (9)$$

where (a) follows from (7). Similarly, we have

$$V_{\lambda''}(x) - V_{\lambda'}(x) = J_{\lambda''}^{(g_{\lambda''})}(x) - J_{\lambda'}^{(g_{\lambda'})}(x) \geq J_{\lambda''}^{(g_{\lambda''})}(x) - J_{\lambda'}^{(g_{\lambda''})}(x) \stackrel{(a)}{=} (\lambda'' - \lambda')N^{(g_{\lambda''})}(x), \tag{10}$$

where (a) follows from (7). The result follows from combining the above inequalities. □

3.2. Sufficient conditions for indexability

Theorem 1. Define $\mathcal{H} = \{(g, h) : g, h : \mathcal{X} \rightarrow \{0, 1\} \text{ such that for all } x \in \mathcal{X}, N^{(g)}(x) \geq N^{(h)}(x)\}$. Each of the following is a sufficient condition for Whittle indexability:

(a) For any $g, h \in \mathcal{H}$, we have that for every $x, z \in \mathcal{X}$,

$$\sum_{y \in \mathcal{X}} \left\{ [\beta P_{zy}(1) - P_{xy}(1)]^+ N^{(g)}(y) - [P_{xy}(1) - \beta P_{zy}(1)]^+ N^{(h)}(y) \right\} \leq \frac{(1 - \beta)^2}{\beta}. \tag{11}$$

(b) For any $g, h \in \mathcal{H}$, we have that for every $x \in \mathcal{X}$,

$$\sum_{y \in \mathcal{X}} \left\{ [P_{xy}(0) - P_{xy}(1)]^+ N^{(g)}(y) - [P_{xy}(1) - P_{xy}(0)]^+ N^{(h)}(y) \right\} \leq \frac{1 - \beta}{\beta}. \tag{12}$$

See Appendix B for the proof. The sufficient conditions of Theorem 1 can be difficult to verify. Simpler sufficient conditions are stated below.

Proposition 2. Each of the following is a sufficient condition for (11):

(a) $\max_{x,z \in \mathcal{X}} \sum_{y \in \mathcal{X}} [\beta P_{zy}(1) - P_{xy}(1)]^+ \leq (1 - \beta)^2 / \beta.$

(b) $P_{xy}(1) = P_{zy}(1)$, for any $x, y, z \in \mathcal{X}$.

In addition, each of the following is a sufficient condition for (12):

(c) $\max_{x \in \mathcal{X}} \sum_{y \in \mathcal{X}} [P_{xy}(0) - P_{xy}(1)]^+ \leq (1 - \beta) / \beta.$

(d) $\beta \leq 0.5.$

See Appendix C for the proof.

Some remarks

1. The sufficient conditions of Theorem 1 and of parts (a), (c), and (d) of Proposition 2 may be viewed as bounds on the discount factor β for which an RB is indexable. Numerical experiments to explore such a property are presented in [25]. A qualitatively similar result was established in [23, Corollary 5], which showed that a restless bandit process is generalized-conservation-laws (GCL-) indexable for sufficiently small discount factors. GCL-indexability is a sub-class of PCL-indexability, which is a sub-class of Whittle indexability. Thus, Proposition 2 provides a quantitative characterization of the qualitative observation made in [9] and generalizes it to a broader class of models.
2. We refer to models that satisfy the sufficient condition of Proposition 2(b) as *restless bandits with controlled restarts*. Such models arise in various scheduling problems

(e.g., machine maintenance, surveillance) where taking the active action resets the state according to a known probability distribution. Specific instances of such models are considered in [3, 30]. The special case when the active action resets to a specific (pristine) state are considered in [18, 19].

3. A different class of restart models have been considered in [7, 8, 20], where the passive action resets the state of the arm. Note that such models do not satisfy Proposition 2(b), and additional modeling assumptions are required to establish indexability. See [7, 8, 20] for details.

4. An algorithm to compute the Whittle index

Given an indexable RB, a naive method to compute the Whittle index at state x is to do a binary search over the penalty λ and find the critical penalty $w(x)$ such that for $\lambda \in (-\infty, w(x))$, $g_\lambda(x) = 0$, and for $\lambda \in [w(x), \infty)$, $g_\lambda(x) = 1$. Although such an approach has been used in the literature [2, 29], it is not efficient as it requires a separate binary search for each state. For a sub-class of restless bandits which satisfy an additional technical condition called partial conservation laws (PCL), Niño-Mora [24, 25] presented an algorithm called the adaptive greedy algorithm to compute the Whittle index. In this section, we present an algorithm that may be viewed as a refinement of the adaptive greedy algorithm and show that it computes the Whittle index for all indexable RBs. The result of this section are restricted to the case of finite \mathcal{X} .

Let K denote the number of states (i.e., $K = |\mathcal{X}|$), and let $K_D (\leq K)$ denote the number of distinct Whittle indices. Let $\Lambda^* = \{\lambda_1, \dots, \lambda_{K_D}\}$, where $\lambda_1 < \lambda_2 < \dots < \lambda_{K_D}$ is the sorted list of distinct Whittle indices. Also, let $\lambda_0 = -\infty$. For any $d \in \{0, \dots, K_D\}$, let $\mathcal{P}_d := \{x \in \mathcal{X} : w(x) \leq \lambda_d\}$ denote the set of states with Whittle index less than or equal to λ_d . Note that $\mathcal{P}_0 = \emptyset$ and $\mathcal{P}_{K_D} = \mathcal{X}$. Let $\Gamma_{d+1} = \mathcal{P}_{d+1} \setminus \mathcal{P}_d$ denote the set of states with Whittle index λ_{d+1} .

For any subset $\mathcal{S} \subseteq \mathcal{X}$, define the policy $\bar{g}^{(\mathcal{S})} : \mathcal{X} \rightarrow \{0, 1\}$ as

$$\bar{g}^{(\mathcal{S})}(x) = \begin{cases} 0, & \text{if } x \in \mathcal{S}, \\ 1, & \text{if } x \in \mathcal{X} \setminus \mathcal{S}. \end{cases} \tag{13}$$

Thus, the policy $\bar{g}^{(\mathcal{S})}$ takes the passive action in the set \mathcal{S} and the active action in the set $\mathcal{X} \setminus \mathcal{S}$.

Now, for any $d \in \{0, \dots, K_D - 1\}$ and all states $y \in \mathcal{X} \setminus \mathcal{P}_d$, define $h_d = \bar{g}^{(\mathcal{P}_d)}$, $h_{d,y} = \bar{g}^{(\mathcal{P}_d \cup \{y\})}$, and for all $x \in \Lambda_{d,y}$,

$$\Lambda_{d,y} = \{x \in \mathcal{X} : N^{(h_d)}(x) \neq N^{(h_{d,y})}(x)\}, \mu_{d,y}(x) = \frac{D^{(h_{d,y})}(x) - D^{(h_d)}(x)}{N^{(h_d)}(x) - N^{(h_{d,y})}(x)}. \tag{14}$$

Lemma 3. For an indexable RB with $d \in \{0, \dots, K_D - 1\}$, we have the following:

1. For all $y \in \Gamma_{d+1}$, we have $w(y) = \lambda_{d+1}$.
2. For all $y \in \mathcal{X} \setminus \mathcal{P}_d$ and $\lambda \in (\lambda_d, \lambda_{d+1}]$, we have $J_\lambda^{(h_{d,y})}(x) \geq J_\lambda^{(h_d)}(x)$ for all $x \in \mathcal{X}$, with equality if and only if $y \in \Gamma_{d+1}$ and $\lambda = \lambda_{d+1}$.

Proof. See Appendix D. □

Theorem 2. *For an indexable RB, the following properties hold:*

1. For any $y \in \Gamma_{d+1}$, the set $\Lambda_{d,y}$ is non-empty.
2. For any $x \in \Lambda_{d,y}$, $\mu_{d,y}(x) \geq \lambda_{d+1}$, with equality if and only if $y \in \Gamma_{d+1}$.

Proof. The proof of each part is as follows:

1. We prove the result by contradiction. Suppose that there exists a $y \in \Gamma_{d+1}$ such that $\Lambda_{d,y} = \emptyset$, which means $N^{(h_d)}(x) = N^{(h_{d,y})}(x)$ for all $x \in \mathcal{X}$. By Lemma 3, we have that

$$J_{\lambda_{d+1}}^{(h_d)}(x) = J_{\lambda_{d+1}}^{(h_{d,y})}(x).$$

Therefore, from (7) we infer $D^{(h_d)}(x) = D^{(h_{d,y})}(x)$ for all $x \in \mathcal{X}$. Since both $D^{(g)}(x)$ and $N^{(g)}(x)$ do not depend on λ , (7) implies that for any λ and $x \in \mathcal{X}$, we have $J_{\lambda}^{(h_d)}(x) = J_{\lambda}^{(h_{d,y})}(x)$. This implies that the policies h_d and $h_{d,y}$ will be optimal for the same set of λ . Now, since the policy h_d is optimal for all $\lambda \in (\lambda_d, \lambda_{d+1}]$ (by definition), so is $h_{d,y}$. Hence $y \in \mathcal{P}_d$. But we started by assuming that $y \notin \mathcal{P}_d$, so we have a contradiction.

2. By Lemma 3, Part 2, for all $y \in \mathcal{X} \setminus \mathcal{P}_d$, $\lambda \in (\lambda_d, \lambda_{d+1}]$, and for all $x \in \Lambda_{d,y}$, we have $J_{\lambda}^{(h_{d,y})}(x) \geq J_{\lambda}^{(h_d)}(x)$. Then, by (7), we infer

$$D^{(h_{d,y})}(x) + \lambda N^{(h_{d,y})}(x) \geq D^{(h_d)}(x) + \lambda N^{(h_d)}(x).$$

Finally, we have $\mu_{d,y}(x) \geq \lambda$ and thus $\mu_{d,y}(x) \geq \lambda_{d+1}$ for all $x \in \Lambda_{d,y}$. This proves the first part of the statement. To prove the second part, note that the policy h_d is an optimal policy for $\lambda \in (\lambda_d, \lambda_{d+1}]$, and for any $y \in \mathcal{P}_{d+1}$, the policy $h_{d,y}$ is an optimal policy for $\lambda \in (\lambda_{d+1}, \lambda_{d+2}]$. From Lemma 1, we know that $V_{\lambda}(x)$ is continuous in λ for all $x \in \mathcal{X}$. So, for all $x \in \mathcal{X}$,

$$\lim_{\lambda \uparrow \lambda_{d+1}} J_{\lambda}^{(h_d)}(x) = \lim_{\lambda \uparrow \lambda_{d+1}} V_{\lambda}(x) = \lim_{\lambda \downarrow \lambda_{d+1}} V_{\lambda}(x) = \lim_{\lambda \downarrow \lambda_{d+1}} J_{\lambda}^{(h_{d,y})}(x).$$

Thus, for all $x \in \mathcal{X}$, $J_{\lambda_{d+1}}^{(h_d)}(x) = J_{\lambda_{d+1}}^{(h_{d,y})}(x)$, and therefore,

$$D^{(h_d)}(x) + \lambda_{d+1} N^{(h_d)}(x) = D^{(h_{d,y})}(x) + \lambda_{d+1} N^{(h_{d,y})}(x).$$

As a result, $\lambda_{d+1} = \mu_{d,y}(x)$ for all $x \in \Lambda_{d,y}$. □

Theorem 2 suggests a method for identifying the Whittle index of any indexable RB by iteratively identifying the set \mathcal{P}_d and the Whittle index λ_d . By definition, $\mathcal{P}_0 = \emptyset$ and $\lambda_0 = -\infty$. Now suppose $\mathcal{P}_0 \subset \mathcal{P}_1 \subset \dots \subset \mathcal{P}_d$ and $\lambda_0 < \lambda_1 < \dots < \lambda_d$ have been identified. We will describe how to determine \mathcal{P}_{d+1} and λ_{d+1} :

1. For $h_d = \bar{g}(\mathcal{P}_d)$, compute $N^{(h_d)}$ by solving (8).
2. For all $y \in \mathcal{X} \setminus \mathcal{P}_d$, compute $N^{(h_{d,y})}$ where $h_{d,y} = \bar{g}(\mathcal{P} \cup \{y\})$ by solving (8), and compute $\Lambda_{d,y}$. Let $\mu_{d,y}^* = \min_{x \in \Lambda_{d,y}} \mu_{d,y}(x)$, where $\mu_{d,y}(x)$ is given by Theorem 2. Then $\lambda_{d+1} = \min_{y \in \mathcal{X} \setminus \mathcal{P}_d} \mu_{d,y}^*$, $\Gamma_{d+1} = \arg \min_{y \in \mathcal{X} \setminus \mathcal{P}_d} \mu_{d,y}^*$, and we get $\mathcal{P}_{d+1} = \mathcal{P}_d \cup \Gamma_{d+1}$ (recall that argmin denotes the set of all minimizers) and $w(x) = \lambda_{d+1}$ for all $x \in \Gamma_{d+1}$.

Algorithm 1: Computing Whittle index of all states of an indexable RB

input : RB $(\mathcal{X}, \{0, 1\}, P(a)_{a \in \{0,1\}}, c, x_0)$, discount factor β .

Initialize $d = 0$ and $\mathcal{P}_0 = \emptyset$.

while $\mathcal{P}_d \neq \mathcal{X}$ **do**

Compute $\Lambda_{d,y}$ and $\mu_{d,y}(x)$ using (14), $\forall y \in \mathcal{X} \setminus \mathcal{P}_d$.

Compute $\mu_{d,y}^* = \min_{x \in \Lambda_{d,y}} \mu_{d,y}(x)$, $\forall y \in \mathcal{X} \setminus \mathcal{P}_d$.

Compute $\lambda_{d+1} = \min_{y \in \mathcal{X} \setminus \mathcal{P}_d} \mu_{d,y}^*$.

Compute $\Gamma_{d+1} = \arg \min_{y \in \mathcal{X} \setminus \mathcal{P}_d} \mu_{d,y}^*$.

Set $w(z) = \lambda_{d+1}$, $\forall z \in \Gamma_{d+1}$.

Set $\mathcal{P}_{d+1} = \mathcal{P}_d \cup \Gamma_{d+1}$.

Set $d = d + 1$.

Iteratively proceeding in this way, we can compute the Whittle index for all states. The detailed algorithm is presented in Algorithm 1.

4.1. An efficient implementation using the Sherman–Morrison formula

We now present an efficient implementation of Algorithm 1 using the Sherman–Morrison inverse formula. Suppose $A \in \mathbb{R}^{n \times n}$ is an invertible square matrix and $u, v \in \mathbb{R}^n$ are column vectors such that $A + uv^T$ is invertible. Then the Sherman–Morrison inverse formula is

$$(A + uv^T)^{-1} = A^{-1} - \frac{A^{-1}uv^T A^{-1}}{1 + v^T A^{-1}u}. \tag{15}$$

Furthermore, given $b \in \mathbb{R}^n$, if x is the solution of $Ax = b$ and y is the solution $Ay = u$, then the solution of $(A + uv^T)\tilde{x} = b$ is given by the following (see [12, Corollary 2]):

$$\tilde{x} = x - \frac{v^T x}{1 + v^T y} y. \tag{16}$$

Let $\Phi^{(g)} = (I - \beta P^{(g)})^{-1}$. Note that for any Markov policy g , $(I - \beta P^{(g)})$ is invertible, because $\beta P^{(g)}$ is a sub-stochastic matrix and has a spectral radius less than 1. Therefore, the conditions for using the Sherman–Morrison formula are satisfied. Hence, using (16), Equation (8) may be written (in matrix form) as

$$D^{(g)} = (1 - \beta)\Phi^{(g)}c^{(g)} \text{ and } N^{(g)} = (1 - \beta)\Phi^{(g)}g. \tag{17}$$

Now, for any $d \in \{0, \dots, K_D - 1\}$ and a state $y \in \mathcal{X} \setminus \mathcal{P}_d$, consider the policies $h_d = \bar{g}^{(\mathcal{P}_d)}$ and $h_{d,y} = \bar{g}^{(\mathcal{P}_{d,y})}$. Let e_y denote the unit vector with 1 in the y th location, and let ρ_y be a vector given by $[\rho_y]_x = P_{yx}(1) - P_{yx}(0)$, for all $x \in \mathcal{X}$. Then $P^{(h_{d,y})} = P^{(h_d)} - e_y \rho_y^T$. Therefore,

$$I - \beta P^{(h_{d,y})} = (I - \beta P^{(h_d)}) + \beta e_y \rho_y^T.$$

Let $\Phi_{\cdot y}^{(h_d)}$ denote the y th column of $\Phi^{(h_d)}$, i.e., $\Phi_{\cdot y}^{(h_d)} = \Phi^{(h_d)} e_y$. Then, by the Sherman–Morrison inverse formula ((15) and (16)), we have

$$\Phi^{(h_{d,y})} = \Phi^{(h_d)} - \frac{\beta \Phi^{(h_d)} e_y \rho_y^T \Phi^{(h_d)}}{1 + \beta \rho_y^T \Phi_{\cdot y}^{(h_d)}}, \tag{18}$$

$$D^{(h_{d,y})} = D^{(h_d)} - \frac{\beta \rho_y^T D^{(h_d)}}{1 + \beta \rho_y^T \Phi_{\cdot y}^{(h_d)}} \Phi_{\cdot y}^{(h_d)}, \quad N^{(h_{d,y})} = N^{(h_d)} - \frac{\beta \rho_y^T N^{(h_d)}}{1 + \beta \rho_y^T \Phi_{\cdot y}^{(h_d)}} \Phi_{\cdot y}^{(h_d)}. \tag{19}$$

Thus, if $\Phi^{(h_d)}$ has been computed, then $\Phi^{(h_{d,y})}$ can be computed in $\mathcal{O}(K^2)$ computations. In addition, if $\Phi^{(h_d)}$, $D^{(h_d)}$, and $N^{(h_d)}$ have been computed, then $D^{(h_{d,y})}$ and $N^{(h_{d,y})}$ can be computed in $\mathcal{O}(K)$ computations.

So we can use (18) and (19) to implement Algorithm 1 in a more efficient manner. However, there is one additional step that needs to be handled, which we explain next.

As $h_{d+1} = \bar{g}^{(\mathcal{P}_d \cup \Gamma_{d+1})}$, we get

$$I - \beta P^{(h_{d+1})} = (I - \beta P^{(h_d)}) + \beta \sum_{y \in \Gamma_{d+1}} e_y \rho_y^T. \tag{20}$$

Thus, $I - \beta P^{(h_{d+1})}$ is a rank- $|\Gamma_{d+1}|$ update of $I - \beta P^{(h_d)}$. When $|\Gamma_{d+1}| > 1$, we can either sequentially apply Equations (18) and (19) for all $y \in \Gamma_{d+1}$ or use the Woodbury formula to compute $\Phi^{(h_{d+1})}$ and compute $D^{(h_{d+1})}$ and $N^{(h_{d+1})}$ using (17). The Woodbury formula is a generalization of the Sherman–Morrison formula. Suppose $A \in \mathbb{R}^{n \times n}$ is an invertible matrix and $U, V \in \mathbb{R}^{n \times m}$ are such that $A + UV^T$ is invertible. Then the Woodbury formula is

$$(A + UV^T)^{-1} = A^{-1} - A^{-1}U(I + V^T U)^{-1}V^T A^{-1}.$$

The complexity of sequentially applying the Sherman–Morrison formula is $\mathcal{O}(|\Gamma_{d+1}|K^2)$ to compute $\Phi^{(h_{d+1})}$ and $\mathcal{O}(|\Gamma_{d+1}|K)$ to compute $D^{(h_{d+1})}$ and $N^{(h_{d+1})}$. The complexity of using the Woodbury formula is $\mathcal{O}(|\Gamma_{d+1}|^{2.807} + K^2)$ to compute $\Phi^{(h_{d+1})}$ and $\mathcal{O}(K^2)$ to compute $D^{(h_{d+1})}$ and $N^{(h_{d+1})}$.

We show the complete algorithm to efficiently compute the Whittle index in Algorithm 2, where we use sequential application of Sherman–Morrison formula to compute $\Phi^{(h_{d+1})}$, $D^{(h_{d+1})}$, and $N^{(h_{d+1})}$.

Some remarks

1. The idea of computing the index by iteratively sorting the states according to their index is commonly used in the algorithms to compute the Gittins index; for example, the largest-remaining-index algorithm, the state-elimination algorithm, the triangularization algorithm, and the fast-pivoting algorithm use variations of this idea. See [10] for details.
2. The computational complexity of Algorithm 2 is $\mathcal{O}(K^3)$, which can be characterized as follows. The algorithm starts with computing $\Phi^{(h_0)}$, which requires $\mathcal{O}(K^{2.807})$ computations (using Strassen’s algorithm), and $D^{(h_0)}$ and $N^{(h_0)}$, each of which requires $\mathcal{O}(K^2)$ computations. Then, in the inner for loop, computing each of $D^{(h_{d,y})}$, $N^{(h_{d,y})}$, and $\mu_{d,y}^*$ requires $\mathcal{O}(K)$ computations, and the inner loop is executed $|\mathcal{X} \setminus \mathcal{P}_d|$ times.

Algorithm 2: Computing Whittle index of all states of an indexable RB

input : RB $(\mathcal{X}, \{0, 1\}, P(a)_{a \in \{0,1\}}, c, x_0)$, discount factor β .

Initialize $d = 0, \mathcal{P}_0 = \emptyset, h_0 = \mathbf{1}_K$.

Compute $\Phi^{(h_0)} = (I - \beta P^{(h_0)})^{-1}$ and $[D^{(h_0)} N^{h_0}] = (1 - \beta)\Phi^{(h_0)}[c^{(h_0)} h_0]$

while $\mathcal{P}_d \neq \mathcal{X}$ **do**

forall $y \in \mathcal{X} \setminus \mathcal{P}_d$ **do**

 Compute $D^{(h_{d,y})}$ and $N^{(h_{d,y})}$ using (19).

 Compute $\Lambda_{d,y}$ and $\mu_{d,y}(x)$ for all $x \in \Lambda_{d,y}$ using (14).

 Compute $\mu_{d,y}^* = \min_{x \in \Lambda_{d,y}} \mu_{d,y}(x)$.

 Compute $\lambda_{d+1} = \min_{y \in \mathcal{X} \setminus \mathcal{P}_d} \mu_{d,y}^*$.

 Compute $\Gamma_{d+1} = \arg \min_{y \in \mathcal{X} \setminus \mathcal{P}_d} \mu_{d,y}^*$.

 Set $w(z) = \lambda_{d+1}, \forall z \in \Gamma_{d+1}$.

 Set $\mathcal{P}_{d+1} = \mathcal{P}_d \cup \Gamma_{d+1}$.

 Initialize $\Phi^{(h_{d+1})} = \Phi^{(h_d)}, D^{(h_{d+1})} = D^{(h_d)}$ and $N^{(h_{d+1})} = N^{(h_d)}$.

forall $z \in \Gamma_{d+1}$ **do**

 Compute $\Phi^{(h_{d+1,z})}, D^{(h_{d+1,z})}$ and $N^{(h_{d+1,z})}$ by using (18) and (19).

 Update $\Phi^{(h_{d+1})} = \Phi^{(h_{d+1,z})}, D^{(h_{d+1})} = D^{(h_{d+1,z})}$ and $N^{(h_{d+1})} = N^{(h_{d+1,z})}$.

 Set $d = d + 1$.

Afterwards, updating $\Phi^{(h_{d+1})}, D^{(h_{d+1})}$, and $N^{(h_{d+1})}$ requires $\mathcal{O}(|\Gamma_{d+1}|K^2)$, $\mathcal{O}(|\Gamma_{d+1}|K)$, and $\mathcal{O}(|\Gamma_{d+1}|K)$ computations, respectively. Therefore, the computational complexity of the algorithm is

$$\begin{aligned} & \mathcal{O}(K^{2.807}) + \mathcal{O}(K^2) + \sum_{d=1}^{K_D} (\mathcal{O}(|\mathcal{X} \setminus \mathcal{P}_d|K) + \mathcal{O}(|\Gamma_{d+1}|K^2) + \mathcal{O}(2|\Gamma_{d+1}|K)) \\ & \leq \mathcal{O}(K^{2.807}) + \sum_{d=1}^{K_D} \mathcal{O}(K^2) + \mathcal{O}\left(\left[\sum_{d=1}^{K_D} |\Gamma_{d+1}| \right] K^2\right) \\ & \leq \mathcal{O}(K^{2.807}) + \mathcal{O}(K^3) + \mathcal{O}(K^3) \leq \mathcal{O}(K^3), \end{aligned}$$

where the first inequality uses the fact that $|\mathcal{X} \setminus \mathcal{P}_d| \leq K$ and the second inequality uses the fact that $\sum_{d=1}^{K_D} |\Gamma_{d+1}| = K$.

- Note that Algorithm 2 computes the Whittle index exactly. In contrast, using binary search [2] computes the Whittle index approximately. Let C_{\max} and C_{\min} denote the upper and lower bound on the per-step cost. Then we know that for any state $x, w(x) \in [C_{\min}, C_{\max}]$. Now, suppose we want to compute the Whittle index to an accuracy of δ . Then the interval $[C_{\min}, C_{\max}]$ needs to be divided into $\log_2((C_{\max} - C_{\min})/\delta)$ steps.

For each step of the binary search, we need to solve the dynamic program (3). There are two main algorithms to solve dynamic programs [28]: policy iteration, which gives an exact solution, and value iteration, which gives an approximate solution.

Each iteration of policy iteration has two steps: (i) policy evaluation, which is done by solving a linear system and has complexity $\mathcal{O}(K^3)$, and (ii) policy improvement, which requires $\mathcal{O}(K^2)$ computations. Thus, the overall complexity of policy iteration is $\mathcal{O}(N_{PI}K^3)$, where N_{PI} is the number of iterations after which policy iteration converges.

Each iteration of value iteration requires $\mathcal{O}(K^2)$ computations. Let N_{VI} denote the number of iterations for value iteration (typically, the iteration stops when a stopping criterion is met). Then the complexity of approximately solving the dynamic program is $\mathcal{O}(N_{VI}K^2)$.

Note that the binary search needs to be repeated for each state. Thus, using binary search to compute the Whittle index to an accuracy of δ has a complexity of $\mathcal{O}(\log_2((C_{\max} - C_{\min})/\delta)N_{PI}K^4)$ if the dynamic program at each step is solved exactly and has a complexity of $\mathcal{O}(\log_2((C_{\max} - C_{\min})/\delta)N_{VI}K^3)$ if the dynamic program at each step is solved approximately.

4.2. Discussion on PCL-indexability

As mentioned earlier, an algorithm very similar to Algorithm 1 was proposed in [25] for computing the Whittle index for RBs that satisfy a technical condition known as PCL-indexability. The analysis in [25] is done under the assumption that the system starts from a designated start state distribution π_0 . For any policy g , define $N^{(g)} = \sum_{x \in \mathcal{X}} N^{(g)}(x)\pi_0(x)$, and define $D^{(g)} = \sum_{x \in \mathcal{X}} D^{(g)}(x)\pi_0(x)$. Let x_1, \dots, x_K be a permutation of the state space such that the corresponding Whittle indices are non-decreasing: $\lambda_1 \leq \dots \leq \lambda_K$. For any $k \in \{1, \dots, K\}$, let $\bar{\mathcal{P}}_k$ denote the set $\{x_1, \dots, x_k\}$.

Now, for any $k \in \{1, \dots, K\}$ and all states $y \in \mathcal{X} \setminus \bar{\mathcal{P}}_k$, define $\bar{h}_k = \bar{g}(\bar{\mathcal{P}}_k)$, $\bar{h}_{k,y} = \bar{g}(\bar{\mathcal{P}}_k \cup \{y\})$, and define

$$\bar{\mu}_{k,y} = \frac{D(\bar{h}_{k,y}) - D(\bar{h}_k)}{N(\bar{h}_k) - N(\bar{h}_{k,y})}. \tag{21}$$

In [25] an algorithm called the adaptive greedy algorithm is presented to iteratively identify the sets $\bar{\mathcal{P}}_k$ and compute the corresponding Whittle indices. This algorithm is shown in Algorithm 3.

An RB is said to be PCL-indexable [25] if it satisfies the following conditions:

1. For any $\mathcal{S} \subseteq \mathcal{X}$ and $y \in \mathcal{X} \setminus \mathcal{S}$, we have $N(\bar{g}(\mathcal{S})) - N(\bar{g}(\mathcal{S} \cup \{y\})) > 0$.
2. The sequence of index values produced by the adaptive greedy algorithm is monotonically non-decreasing.

Finally, the following result is established.

Theorem 3 (Theorem 1 of [25]). *A PCL-indexable RB is indexable, and the adaptive greedy algorithm gives its Whittle indices.*

The main differences between our result and that of [25] are as follows:

1. An implication of the first condition in the definition of PCL indexability is that the denominator in (21) is never zero. In contrast, we do not impose such a restriction and work with the non-empty subset of states for which the denominator in (14) is non-zero.

Algorithm 3: The adaptive greedy algorithm of [25]

input : RB $(\mathcal{X}, \{0, 1\}, P(a)_{a \in \{0,1\}}, c, x_0)$, discount factor β .

Initialize $k = 0$ and $\mathcal{P}_0 = \emptyset$

while $k \neq K$ **do**

forall $y \in \mathcal{X} \setminus \mathcal{P}_k$ **do**

Compute $\bar{\mu}_{d,y}$ using (21)

Pick $x_{k+1} \in \arg \min_{y \in \mathcal{X} \setminus \bar{\mathcal{P}}_k} \bar{\mu}_{k,y}$

Set $w(s_{k+1}) = \min_{y \in \mathcal{X} \setminus \bar{\mathcal{P}}_k} \bar{\mu}_{k,y}$ and $\bar{\mathcal{P}}_{k+1} = \bar{\mathcal{P}} \cup \{x_k\}$

$k = k + 1$

2. In Algorithm 3, the sets $\{\bar{\mathcal{P}}_k\}_{k=1}^K$ are constructed by adding states one by one, even when $\bar{\mu}_{k,y}$ has multiple argmins. In contrast, in Algorithm 1, the sets $\{\mathcal{P}_d\}_{d=1}^{K_D}$ are constructed by adding all states which have the same Whittle index at once.
3. In Algorithm 3, one has to check that the indices are generated in a non-decreasing order (which is the second condition of PCL-indexability). In contrast, in Algorithm 1, the indices are always generated in an increasing order, and therefore Condition 2 for PCL-indexability is always satisfied.
4. Finally, Theorem 3 only guarantees that Algorithm 3 computes the Whittle index for RBs which satisfy PCL-indexability. Moreover, the second condition for PCL-indexability can only be checked after running Algorithm 3. In contrast, Theorem 2 guarantees that Algorithm 1 computes the Whittle index for all indexable RBs.

We conclude this discussion by revisiting an example from [25] which is an indexable RB that is not PCL-indexable. For this example, $\mathcal{X} = \{1, 2, 3\}$; the transition matrices are

$$P(0) = \begin{bmatrix} 0.3629 & 0.5028 & 0.1343 \\ 0.0823 & 0.7534 & 0.1643 \\ 0.2460 & 0.0294 & 0.7246 \end{bmatrix}, \quad P(1) = \begin{bmatrix} 0.1719 & 0.1749 & 0.6532 \\ 0.0547 & 0.9317 & 0.0136 \\ 0.1547 & 0.6271 & 0.2182 \end{bmatrix};$$

the per-step cost is $c(x, 0) = 0$ for all $x \in \mathcal{X}$ and $c(1, 1) = -0.44138$, $c(2, 1) = -0.8033$, $c(3, 1) = -0.14257$; and $\beta = 0.9$. The corresponding Whittle indices are $[0.18, 0.8, 0.57]$.

This model is not PCL-indexable, since if $g = [1, 1, 0]$ and $h = [0, 1, 0]$, then $N^{(g)} = [5.66, 8.24, 4.23]$ and $N^{(h)} = [6.65, 8.59, 4.88]$. Therefore, for any initial state distribution π_0 , $N^{(g)} < N^{(h)}$.

However, as the problem is indexable, we can still apply Algorithm 1 to compute Whittle indices without any limitations. The steps are as follows:

1. Initialize $d = 0$ and have $\mathcal{P}_0 = \emptyset$. Thus $\bar{g}^{(\mathcal{P}_0)} = [1, 1, 1]$ and we compute $N^{(\bar{g}^{(\mathcal{P}_0)})} = [10, 10, 10]$ and $D^{(\bar{g}^{(\mathcal{P}_0)})} = [-6.43, -7.43, -6.51]$.

2. There are three possibilities for $y \in \mathcal{X} \setminus \mathcal{P}_0 = \{1, 2, 3\}$:

- For $y = 1$, $h_{0,1} = [0, 1, 1]$. We compute $N^{(h_{0,1})} = [7.88, 9.29, 9.13]$ and $D^{(h_{0,1})} = [-6.05, -7.30, -6.35]$. Therefore, $\Lambda_{0,1} = \{x \in \mathcal{X} : N^{(\bar{g}^{(\mathcal{P}_0)}(x))} \neq N^{(h_{0,1})}(x)\} = \{1, 2, 3\}$. Now for each $x \in \Lambda_{0,1}$, we compute $\mu_{0,1}(1) = \mu_{0,1}(2) = \mu_{0,1}(3) = 0.18$. Therefore, $\mu_{0,1}^* = 0.18$.
- For $y = 2$, $h_{0,2} = [1, 0, 1]$. We compute $N^{(h_{0,2})} = [4.58, 2.93, 4.10]$ and $D^{(h_{0,2})} = [-1.27, -0.7, -0.89]$. Therefore, $\Lambda_{0,2} = \{x \in \mathcal{X} : N^{(\bar{g}^{(\mathcal{P}_0)}(x))} \neq N^{(h_{0,2})}(x)\} = \{1, 2, 3\}$. Now for each $x \in \Lambda_{0,2}$, we compute $\mu_{0,2}(1) = \mu_{0,2}(2) = \mu_{0,2}(3) = 0.95$. Therefore, $\mu_{0,2}^* = 0.95$.
- For $y = 3$, $h_{0,3} = [1, 1, 0]$. We compute $N^{(h_{0,3})} = [5.66, 8.24, 4.23]$ and $D^{(h_{0,3})} = [-3.64, -6.30, -2.79]$. Therefore, $\Lambda_{0,3} = \{x \in \mathcal{X} : N^{(\bar{g}^{(\mathcal{P}_0)}(x))} \neq N^{(h_{0,3})}(x)\} = \{1, 2, 3\}$. Now for each $x \in \Lambda_{0,3}$, we compute $\mu_{0,3}(1) = \mu_{0,3}(2) = \mu_{0,3}(3) = 0.64$. Therefore, $\mu_{0,3}^* = 0.64$.

Now $\lambda_1 = \min \{\mu_{0,1}^*, \mu_{0,2}^*, \mu_{0,3}^*\} = 0.18$. Therefore, $\mathcal{P}_1 = \{1\}$, $w(1) = 0.18$, $\bar{g}^{(\mathcal{P}_1)} = [0, 1, 1]$. We have already computed $N^{(\bar{g}^{(\mathcal{P}_1)})} = [7.88, 9.29, 9.13]$ and $D^{(\bar{g}^{(\mathcal{P}_1)})} = [-6.05, -7.30, -6.35]$.

3. There are two possibilities for $y \in \mathcal{X} \setminus \mathcal{P}_1 = \{2, 3\}$:

- For $y = 2$, $h_{1,2} = [0, 0, 1]$. We compute $N^{(h_{1,2})} = [1.48, 1.52, 2.57]$ and $D^{(h_{1,2})} = [-0.21, -0.22, -0.37]$. Therefore, $\Lambda_{1,2} = \{x \in \mathcal{X} : N^{(\bar{g}^{(\mathcal{P}_1)}(x))} \neq N^{(h_{1,2})}(x)\} = \{1, 2, 3\}$. Now for each $x \in \Lambda_{1,2}$, we compute $\mu_{1,2}(1) = \mu_{1,2}(2) = \mu_{1,2}(3) = 0.91$. Therefore, $\mu_{1,2}^* = 0.91$.
- For $y = 3$, $h_{1,3} = [0, 1, 0]$. We compute $N^{(h_{1,3})} = [6.65, 8.59, 4.88]$ and $D^{(h_{1,3})} = [-1.22, -0.66, -0.83]$. Therefore, $\Lambda_{1,3} = \{x \in \mathcal{X} : N^{(\bar{g}^{(\mathcal{P}_1)}(x))} \neq N^{(h_{1,3})}(x)\} = \{1, 2, 3\}$. Now for each $x \in \Lambda_{1,3}$, we compute $\mu_{1,3}(1) = \mu_{1,3}(2) = \mu_{1,3}(3) = 0.57$. Therefore, $\mu_{1,3}^* = 0.57$.

Now $\lambda_2 = \min \{\mu_{1,2}^*, \mu_{1,3}^*\} = 0.57$. Therefore, $\mathcal{P}_2 = \{1, 3\}$, $w(3) = 0.57$, $\bar{g}^{(\mathcal{P}_2)} = [0, 1, 0]$. We have already computed $N^{(\bar{g}^{(\mathcal{P}_2)})} = [6.65, 8.59, 4.88]$ and $D^{(\bar{g}^{(\mathcal{P}_2)})} = [-1.22, -0.66, -0.83]$.

4. There is only one possibility for $y \in \mathcal{X} \setminus \mathcal{P}_2 = \{2\}$:

- For $y = 3$, $h_{3,2} = [0, 0, 0]$, $N^{(h_{3,2})} = [0, 0, 0]$, and $D^{(h_{3,2})} = [-0.21, -0.22, -0.37]$. Therefore, $\Lambda_{3,2} = \{x \in \mathcal{X} : N^{(\bar{g}^{(\mathcal{P}_2)}(x))} \neq N^{(h_{3,2})}(x)\} = \{1, 2, 3\}$. Now for each $x \in \Lambda_{3,2}$, we compute $\mu_{3,2}(1) = \mu_{3,2}(2) = \mu_{3,2}(3) = 0.8$. Therefore, $\mu_{3,2}^* = 0.8$.

Now $\mu_{3,2}^* = 0.57$. Therefore, $\mathcal{P}_3 = \{1, 2, 3\}$ and $w(2) = 0.8$.

Finally, the Whittle indices are $[0.18, 0.8, 0.57]$.

5. Some special cases

In this section, we refine the results developed in this paper to some special cases.

5.1. Restless bandits with optimal threshold-based policy

Consider an RB $(\mathcal{X}, \{0, 1\}, \{P(a)\}_{a \in \{0,1\}}, c, x_0)$ where the state space \mathcal{X} is a totally ordered set, i.e., $\mathcal{X} = \{1, \dots, K\}$. Let $\mathcal{X}_0 = \{0, \dots, K\}$; let $\mathcal{X}_{\geq \ell}$ denote the set of states greater than or equal to the state ℓ , and let $\mathcal{X}_{\leq \ell}$ denote the set of states less than or equal to the state ℓ . We suppose that the model satisfies the following assumption:

- (P) There exists a non-decreasing family of thresholds $\{\ell_\lambda\}_{\lambda \in \mathbb{R}}$, $\ell_\lambda \in \mathcal{X}_0$, such that the threshold policy $g^{(\ell_\lambda)}$ is optimal for Problem 2 with activation cost λ .

Several models where (P) holds have been considered in the literature [3, 5, 7, 15, 17, 30]. A key implication of Property (P) is the following.

Lemma 4. *Suppose an RB defined on a totally ordered state space satisfies Property (P). Then the RB is indexable and the Whittle index $w(\ell)$ is non-decreasing in $\ell \in \mathcal{X}$.*

Proof. Note that Property (P) implies that the passive set $\Pi_\lambda = \{x \in \mathcal{X} : g_\lambda(x) = 0\} = \mathcal{X}_{\leq \ell_\lambda}$, which is increasing in λ . Hence the RB is indexable. Moreover, for any state ℓ , the Whittle index $w(\ell)$ is the smallest value of λ such that $\ell_\lambda = \ell$. Therefore, by Property (P), $w(\ell)$ is non-decreasing in ℓ . □

As in Section 4, we assume that there are $K_D (\leq K)$ distinct Whittle indices given by $\Lambda^* = \{\lambda_1, \dots, \lambda_{K_D}\}$, where $\lambda_1 < \lambda_2 < \dots < \lambda_{K_D}$. We also let $\lambda_0 = -\infty$, and for any $d \in \{0, \dots, K_D\}$, we let $\mathcal{P}_d = \{x \in \mathcal{X} : w(x) \leq \lambda_d\}$. As stated in the proof of Lemma 4, Property (P) implies that $\mathcal{P}_d = \mathcal{X}_{\leq \ell_{\lambda_d}}$. Therefore, $\Gamma_{d+1} = \{\ell_{\lambda_d} + 1, \dots, \ell_{\lambda_{d+1}}\}$. Thus, Theorem 2 simplifies as follows.

Corollary 1. *Suppose an RB defined on a totally ordered state space satisfies Property (P). Then the following properties hold:*

1. For any $y \in \Gamma_{d+1}$, the set $\Lambda_{d,y}$ is non-empty.
2. For any $x \in \Lambda_{d,y}$, $\mu_{d,y}(x) \geq \lambda_{d+1}$, with equality if and only if $y \in \Gamma_{d+1}$.

Thus, based on Corollary 1, for models that satisfy Property (P), we can simplify Algorithm 2 as shown in Algorithm 4. Instead of computing $\mu_{d,y}^*$ for all $y \in \mathcal{X} \setminus \mathcal{P}_d$, we can compute it sequentially and break the loop when $\mu_{d,y}^* \neq \lambda_{d+1}$. Note that this simplification does not change the asymptotic complexity of the algorithm, which is still $\mathcal{O}(K^3)$.

Remark 1. Note that if the model satisfies additional assumptions such that it is known up front that no two states have the same Whittle index, then we do not need the inner for loop (over y) in Algorithm 4, and can simply compute the Whittle index of the state ℓ as

$$w(\ell) = \min \frac{D(\bar{g}^{(\mathcal{X}_{\leq \ell+1})})(x) - D(g^{(\mathcal{X}_{\leq \ell})})(x)}{N(\bar{g}^{(\mathcal{X}_{\leq \ell})})(x) - N(g^{(\mathcal{X}_{\leq \ell+1})})(x)},$$

where the minimum is over all x such that the denominator is non-zero.

Algorithm 4: Whittle index for RB with optimal threshold-based policy

input : RB $(\mathcal{X}, \{0, 1\}, P(a)_{a \in \{0,1\}}, c, x_0)$, discount factor β .

Initialize $d = 0, \ell = 0, h_0 = \mathbf{1}_K$.

Compute $\Phi^{(h_0)} = (I - \beta P^{(h_0)})^{-1}, [D^{(h_0)} N^{h_0}] = (1 - \beta)\Phi^{(h_0)}[c^{(h_0)} h_0]$

while $\ell \leq K$ **do**

forall $y \in \{\ell + 1, \dots, K\}$ **do**

Compute $D^{(h_{d,y})}$ and $N^{(h_{d,y})}$ using (19).

Compute $\Lambda_{d,y}$ and $\mu_{d,y}(x)$ for all $x \in \Lambda_{d,y}$ using (14).

Compute $\mu_{d,y}^* = \min_{x \in \Lambda_{d,y}} \mu_{d,y}(x)$.

if $y = \ell + 1$ **then**

Set $\lambda_{d+1} = \mu_{d,y}^*$ and $\Gamma_{d+1} = \{y\}$

Set $w(y) = \lambda_{d+1}$

else

if $\lambda_{d+1} = \mu_{d,y}^*$ **then**

Update $\Gamma_{d+1} = \Gamma_{d+1} \cup \{y\}$

Set $w(y) = \lambda_{d+1}$

else

Set $\ell = y$

break

Initialize $\Phi^{(h_{d+1})} = \Phi^{(h_d)}, D^{(h_{d+1})} = D^{(h_d)}$ and $N^{(h_{d+1})} = N^{(h_d)}$.

forall $z \in \Gamma_{d+1}$ **do**

Compute $\Phi^{(h_{d+1,z})}, D^{(h_{d+1,z})}$ and $N^{(h_{d+1,z})}$ by using (18) and (19).

Update $\Phi^{(h_{d+1})} = \Phi^{(h_{d+1,z})}, D^{(h_{d+1})} = D^{(h_{d+1,z})}$ and $N^{(h_{d+1})} = N^{(h_{d+1,z})}$.

Set $d = d + 1$.

In the next section, we present a new model called *stochastic monotone bandits*, which may be considered as a generalization of monotone bandits [5, 7, 15], and show that these models satisfy Property (P).

5.2. Stochastic monotone bandits

We first start with the definition of a submodular function. Given ordered sets \mathcal{X} and \mathcal{Y} , a function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ is called submodular if for any $x_1, x_2 \in \mathcal{X}$ and $y_1, y_2 \in \mathcal{Y}$ such that $x_2 \geq x_1$ and $y_2 \geq y_1$, we have $f(x_1, y_2) - f(x_1, y_1) \geq f(x_2, y_2) - f(x_2, y_1)$.

We say that the RB is *stochastic monotone* if it satisfies the following conditions:

(D1) For any $a \in \{0, 1\}$, $P(a)$ is stochastically monotone; i.e., for any $x, y \in \mathcal{X}$ such that $x < y$, we have

$$\sum_{w \in \mathcal{X}_{\geq z}} P_{xw}(a) \leq \sum_{w \in \mathcal{X}_{\geq z}} P_{yw}(a)$$

for any $z \in \mathcal{X}$.

(D2) For any $z \in \mathcal{X}$, $S_{zx}(a) := \sum_{w \in \mathcal{X}_{\geq z}} P_{xw}(a)$ is submodular in (x, a) .

(D3) For any $a \in \{0, 1\}$, $c(x, a)$ is non-decreasing in x .

(D4) The cost $c(x, a)$ is submodular in (x, a) .

For ease of notation, for any $\ell \in \mathcal{X}_0$, we let $g^{(\ell)} = \bar{g}^{(\mathcal{X}_{\leq \ell})}$ denote a policy with threshold ℓ (where $\bar{g}^{(S)}$ is as defined in (13)).

Lemma 5. *A stochastic monotone RB satisfies the following properties:*

1. For any $\lambda \in \mathbb{R}$, there exists a threshold $\ell_\lambda \in \mathcal{X}^*$ such that the threshold policy $g^{(\ell_\lambda)}$ is optimal for Problem 2. If there are multiple such thresholds, we use ℓ_λ to denote the largest threshold.
2. If, for any $x \in \mathcal{X}$, $N^{(g^{(\ell)})}(x)$ is non-increasing in ℓ , then ℓ_λ is non-decreasing with λ . Hence, the model satisfies Property (P) and is therefore indexable.

Proof. For the first part, we note that Conditions (D1)–(D4) are the same as the properties of [28, Theorem 4.7.4], which implies that there exists a threshold as stated.

For the second part, we first show that for any $\ell \in \mathcal{X}^*$, $J_\lambda^{(g^{(\ell)})}(x)$ is submodular in (ℓ, λ) for all $x \in \mathcal{X}$. In particular, for any $k < \ell$, we have

$$J_\lambda^{(g^{(\ell)})}(x) - J_\lambda^{(g^{(k)})}(x) = D_\lambda^{(g^{(\ell)})}(x) - D_\lambda^{(g^{(k)})}(x) + \lambda(N_\lambda^{(g^{(\ell)})}(x) - N_\lambda^{(g^{(k)})}(x)).$$

Now (D5) implies that the difference $J_\lambda^{(g^{(\ell)})}(x) - J_\lambda^{(g^{(k)})}(x)$ is non-increasing in λ . Therefore, $J_\lambda^{(g^{(\ell)})}(x)$ is submodular in (ℓ, λ) . Consequently, from [28, Theorem 2.8.2],

$$\ell_\lambda = \max \left\{ \ell' \in \arg \min_{\ell \in \mathcal{X}^*} J_\lambda^{(g^{(\ell)})}(x) \right\}$$

is non-decreasing in λ . □

5.3. Restless bandits with controlled restarts

Consider restless bandits with controlled restarts (i.e., models where $P_{xy}(1)$ does not depend on x). By Proposition 2(c), such models are indexable. In this section, we explain how to simplify the computation of the Whittle index for such models. For ease of notation, we use P_{xy} to denote $P_{xy}(0)$ and Q_y to denote $P_{xy}(1)$.

Define $D^{(g)} = \sum_{x \in \mathcal{X}} Q_x D^{(g)}(x)$ and $N^{(g)} = \sum_{x \in \mathcal{X}} Q_x N^{(g)}(x)$. Now, following the discussion of Section 4, we can show that the result of Theorem 2 continues to hold when $\mu_{d,y}$ is replaced by

$$\hat{\mu}_{d,y} = \frac{D^{(h_y)} - D^{(\bar{g}^{(\mathcal{P}_d)})}}{N^{(\bar{g}^{(\mathcal{P}_d)})} - N^{(h_y)}}.$$

Therefore, we can replace $\mu_{d,y}(x)$ in Algorithm 1 by $\hat{\mu}_{d,y}$. Our key result for this section is that $D^{(g)}$ and $N^{(g)}$ can be computed efficiently for models with controlled restarts.

For that matter, given any policy g , let τ_g denote the hitting time of the set $\Pi^{(g)} = \{x \in \mathcal{X} : g(x) = 1\}$. Let

$$L^{(g)} := \mathbb{E} \left[\sum_{t=0}^{\tau_g} \beta^t c(X_t, g(X_t)) \mid X_0 \sim Q \right] \quad \text{and} \quad M^{(g)} := \mathbb{E} \left[\sum_{t=0}^{\tau_g} \beta^t \mid X_0 \sim Q \right]$$

denote the expected discounted cost and expected discounted time for hitting $\Pi^{(g)}$ starting with an initial state distribution of Q . Then, using ideas from renewal theory, we can show the following.

Theorem 4. *For any policy g ,*

$$D^{(g)} = \frac{L^{(g)}}{M^{(g)}} \quad \text{and} \quad N^{(g)} = \frac{1}{\beta M^{(g)}} - \frac{1 - \beta}{\beta}.$$

Proof. The proof follows from standard ideas in renewal theory. By the strong Markov property, we have

$$\begin{aligned} D^{(g)} &= \mathbb{E} \left[(1 - \beta) \sum_{t=0}^{\tau_g} \beta^t c(X_t, g(X_t)) + \beta^{\tau_g+1} D^{(g)} \mid X_0 \sim Q \right] \\ &= (1 - \beta)L^{(g)} + \mathbb{E}[\beta^{\tau_g+1} | X_0 \sim Q] D^{(g)}. \end{aligned} \tag{22}$$

Using the definition of $M^{(g)}$, we have $\mathbb{E}[\beta^{\tau_g+1} | X_0 \sim Q] = 1 - (1 - \beta)M^{(g)}$. Substituting this in (22) and rearranging the terms we get $D^{(g)} = L^{(g)}/M^{(g)}$.

For $N^{(g)}$, by the strong Markov property we have

$$\begin{aligned} N^{(g)} &= \mathbb{E} \left[(1 - \beta)\beta^{\tau_g} + \beta^{\tau_g+1} N^{(g)} \mid X_0 \sim Q \right] \\ &= \mathbb{E}[\beta^{\tau_g} | X_0 \sim Q] (1 - \beta + \beta N^{(g)}) = \frac{1 - (1 - \beta)M^{(g)}}{\beta} (1 - \beta + \beta N^{(g)}). \end{aligned}$$

Therefore, we get $N^{(g)} = (1 - (1 - \beta)M^{(g)})/\beta M^{(g)}$. □

Given any policy g , we can efficiently compute $L^{(g)}$ and $M^{(g)}$ using standard formulas for truncated Markov chains. For any vector v , let $v^{(g)}$ denote the vector with components indexed by the set $\{x \in \mathcal{X} : g(x) = 0\}$, and let $\tilde{v}^{(g)}$ denote the remaining components. For example, if $\mathcal{X} = \{1, 2, 3, 4\}$, $g = (1, 0, 1, 0)$, and $v = [1, 2, 3, 4]$, then $v^{(g)} = (2, 4)$ and $\tilde{v}^{(g)} = (1, 3)$. Similarly, for any square matrix Z , let $Z^{[g]}$ denote the square sub-matrix corresponding to

Algorithm 5: Myopic heuristic

input : Set \mathcal{N} of arms; arms m to be activated

foreach time t **do**

let $\ell = 0$, $\mathcal{M} = \emptyset$, and $\mathcal{Z} = \mathcal{N}$.

foreach $\ell \in \{0, \dots, m\}$ **do**

$i_\ell^* \in \arg \min_{i \in \mathcal{Z}} \sum_{j \in \mathcal{Z} \setminus \{i\}} \{c^j(X_t^j, 0) + c^i(X_t^i, 1)\}$ // Pick any arg min

let $\mathcal{M} = \mathcal{M} \cup \{i_\ell^*\}$, $\mathcal{Z} = \mathcal{Z} \setminus \{i_\ell^*\}$

Activate arms in \mathcal{Z}

the elements $\{x \in \mathcal{X} : g(x) = 0\}$, and let $\tilde{Z}^{[g]}$ denote the sub-matrix with rows $\{x \in \mathcal{X} : g(x) = 0\}$ and columns $\{x \in \mathcal{X} : g(x) = 1\}$. For example, if $g = [1, 0, 1, 0]$ and if

$$Z = \begin{bmatrix} 1 & 2 & 3 & 4 \\ 5 & 6 & 8 & 8 \\ 9 & 10 & 11 & 12 \\ 13 & 14 & 15 & 16 \end{bmatrix},$$

then $Z^{[g]} = \begin{bmatrix} 6 & 8 \\ 14 & 16 \end{bmatrix}$ and $Z^{[g]} = \begin{bmatrix} 5 & 8 \\ 13 & 15 \end{bmatrix}$. Then, from standard formulas for truncated Markov chains, we have the following.

Proposition 3. For any policy g , let c_0 and c_1 denote column vectors corresponding to $c(\cdot, 0)$ and $c(\cdot, 1)$. Then

$$\begin{aligned} L^{(g)} &= Q^{(g)}(I - \beta P^{[g]})^{-1} \left(c_0^{(g)} + \beta \tilde{P}^{[g]} \tilde{c}_1^{(g)} \right) + \tilde{Q}^{(g)} \tilde{c}_1^{(g)}, \\ M^{(g)} &= Q^{(g)}(I - \beta P^{[g]})^{-1} \left(\mathbf{1}^{(g)} + \beta \tilde{P}^{[g]} \tilde{\mathbf{1}}^{(g)} \right) + \tilde{Q}^{(g)} \tilde{\mathbf{1}}^{(g)}. \end{aligned}$$

This gives us an efficient method to compute $L^{(g)}$ and $M^{(g)}$, which can in turn be used to compute $D^{(g)}$ and $N^{(g)}$ and used in a modified version of Algorithm 1 as explained.

6. Numerical experiments

In this section, we evaluate how well the Whittle index policy (WIP) performs compared to the optimal policy (OPT) as well as to a baseline policy known as the myopic policy (MYP) (shown in Algorithm 5). The code is available at [4].

6.1. Experimental setup

In our experiments, we consider restart bandits with $P(1) = [\mathbf{1}, \mathbf{0}, \dots, \mathbf{0}]$. There are two other components of the model: the transition matrix $P(0)$ and the cost function c . We choose these components as follows.




FIGURE 2. Relative performance α_{OPT} of WIP versus OPT for Experiment 2.

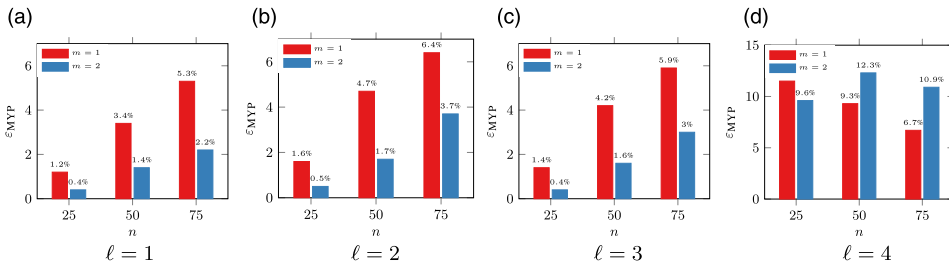


FIGURE 3. Relative improvement ϵ_{MYP} of WIP versus MYP for Experiment 3.

6.1.1. *The choice of transition matrices.* We have three setups for choosing $P(0)$. The first setup is a family of four types of structured stochastic monotone matrices, which we denote by $\mathcal{P}_\ell(p)$, $\ell \in \{1, \dots, 4\}$, where $p \in [0, 1]$ is a parameter of the model. The second setup is a randomly generated stochastic monotone matrix which we denote by $\mathcal{R}(d)$, where $d \in [0, 1]$ is a parameter of the model. In the third setup, we generate random stochastic matrices using the Lévy distribution. The details of these models are presented in the supplementary material.

6.1.2. *The choice of the cost function.* For all our experiments we choose $c(x, 0) = (x - 1)^2$ and $c(x, 1) = 0.5(|\mathcal{X}| - 1)^2$.

6.2. Experimental details and results

We conduct different experiments to compare the performance of the Whittle index with the optimal policy and the myopic policy for different setups (described in Section 6.1) and for different choices of the size $|\mathcal{X}|$ of the state space, the number n of arms, and the number m of active arms. For all experiments we choose the discount factor $\beta = 0.95$.

We evaluate the performance of a policy via Monte Carlo simulations over S trajectories, where each trajectory is of length T . In all our experiments, we choose $S = 2500$ and $T = 250$.

Experiment 1. *Comparison of Whittle index with the optimal policy for structured models.* The optimal policy is computed by solving the Markov decision process for Problem 1, which is feasible only for small values of $|\mathcal{X}|$ and n . We choose $|\mathcal{X}| = 5$ and $n = 5$ and compare the two policies for the model $\mathcal{P}_\ell(\cdot)$, $\ell \in \{1, \dots, 4\}$ and $m \in \{1, 2\}$.

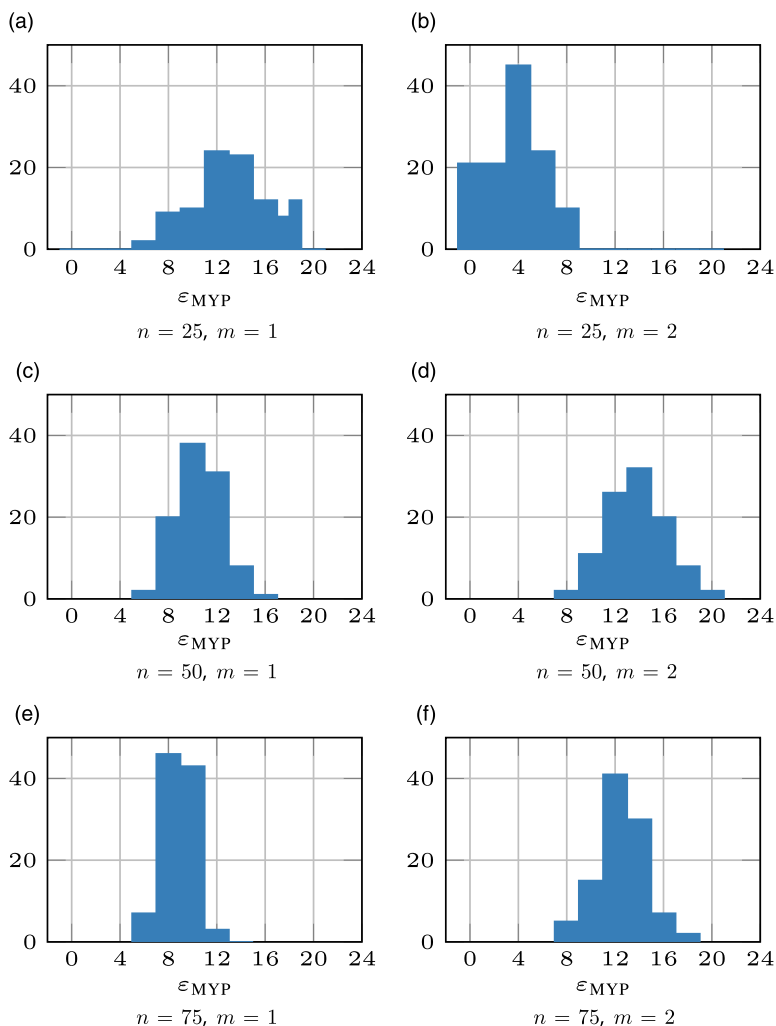


FIGURE 4. Relative improvement ε_{MYP} of WIP versus MYP for Experiment 4.

For a given value of n and ℓ , we pick n equispaced points (p_1, \dots, p_n) in the interval $[0.35, 1]$ and choose $\mathcal{P}_\ell(p_i)$ as the transition matrix of arm i . We observed that $\alpha_{\text{OPT}} = J(\text{OPT})/J(\text{WIP})$, the relative (percentage) performance improvement of WIP compared to OPT, was in the range of 99.95%–100% for all parameters.

Experiment 2. Comparison of Whittle index with the optimal policy for randomly sampled models. As before, we pick $|\mathcal{X}| = 5$ and $n = 5$ so that it is feasible to calculate the optimal policy. For each arm, we sample the transition matrix from $\mathcal{R}(5/|\mathcal{X}|)$ and repeat the experiment 250 times. Histograms of α_{OPT} over experiments for $m \in \{1, 2\}$ appear in Figure 2, which shows that WIP performs close to OPT in all cases.

Experiment 3. Comparison of Whittle index with the myopic policy for structured models. We generate the structured models as in Experiment 1 but for $|\mathcal{X}| = 25, n \in \{25, 50, 75\}$, and

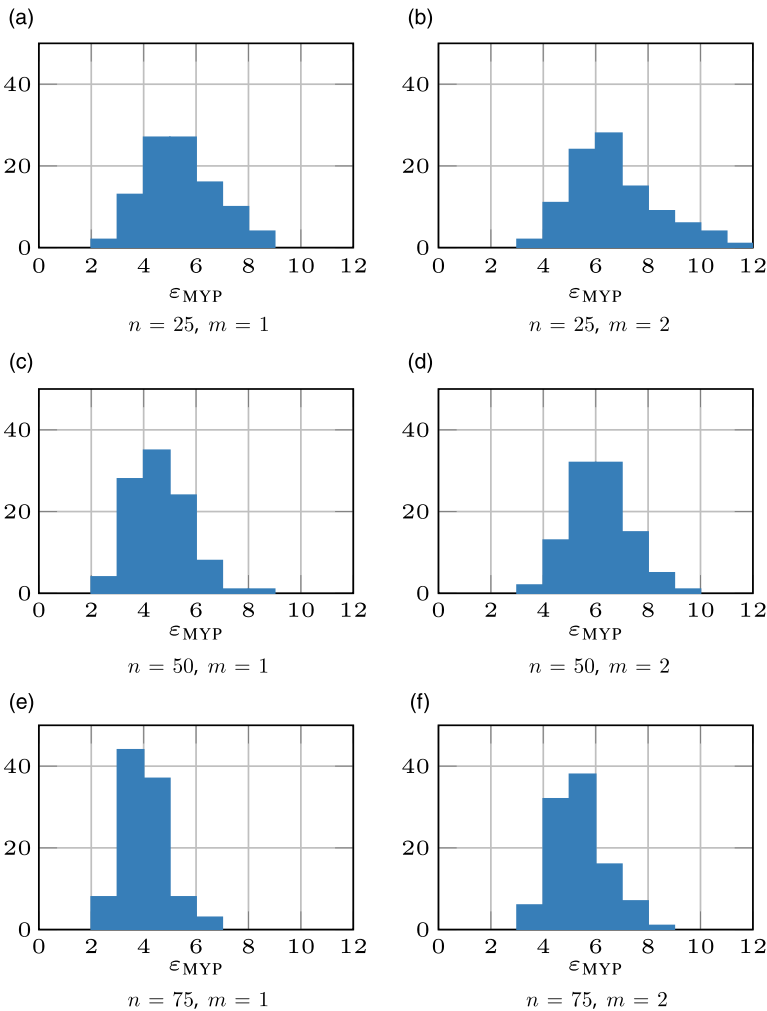


FIGURE 5. Relative improvement ε_{MYP} of WIP versus MYP for Experiment 5.

$m \in \{1, 2, 5\}$. In this case, let $\varepsilon_{MYP} = (J(MYP) - J(WIP))/J(MYP)$ denote the relative improvement of WIP compared to MYP. The results of ε_{MYP} for different choices of the parameters are shown in Figure 3.

In Figure 3, we observe that WIP performs considerably better than MYP. In addition to that, the performance of WIP is better with respect to MYP when $\ell = 4$, which is more complicated than models where $\ell \in \{1, 2, 3\}$. However, increasing m does not necessarily contribute to better ε_{MYP} , as overlap between the choices of the two policies may increase. Note that as $\mathcal{P}_4(\cdot)$ is very different from the rest of the models, the trend of bars in Figure 3d with respect to n varies differently from the rest of the models.

Experiment 4. Comparison of Whittle index with the myopic policy for randomly sampled models. We generate 250 random models as described in Experiment 2 but for $|\mathcal{X}| = 25$ and

larger values of n . For each case, ε_{MYP} is computed. Histograms of ε_{MYP} for different choices of the parameters are shown in Figure 4.

The results show that on average, WIP performs considerably better than MYP, and this improvement is guaranteed as the concentration of data for the sampled models is mostly on positive values of ε_{MYP} .

Experiment 5. Comparison of Whittle index with the myopic policy for restart models. We generate 250 random stochastic matrices for $P(0)$. Each row of the matrix is generate according to Section 1.3 of the supplementary material. We set $|\mathcal{X}| = 25, n \in \{25, 50, 75\}$, and $m \in \{1, 2\}$. For each case, ε_{MYP} is computed; histograms of ε_{MYP} for different choices of the parameters are shown in Figure 5.

7. Conclusion

We present two general sufficient conditions for restless bandit processes to be indexable. The first condition depends only on the transition matrix $P(1)$, while the second condition depends on both $P(0)$ and $P(1)$. These sufficient conditions are based on alternative characterizations of the passive set, which might be useful in general as well. We also present refinements of these sufficient conditions that are simpler to verify. Two of these simpler conditions are worth highlighting: models where the active action resets the state according to a known distribution, and models where the discount factor is less than 0.5.

We then present a generalization of a previously proposed adaptive greedy algorithm, which was developed to compute the Whittle index for a sub-class of restless bandits known as PCL-indexable bandits. We show that the generalized adaptive greedy algorithm computes the Whittle index for all indexable bandits. We provide a computationally efficient implementation of our algorithm, which computes the Whittle indices of a restless bandit with K states in $\mathcal{O}(K^3)$ computations.

Finally, we show how to refine the results for two classes of restless bandits: stochastic monotone bandits and restless bandits with controlled restarts. We also present a detailed numerical study which shows that the Whittle index policy performs close to the optimal policy and is considerably better than a myopic policy.

Appendix A. Proof of Proposition 1

We first present a preliminary result.

Lemma 6. For $\tau = 0$, the policy h_0 satisfies $J_\lambda^{(h_0)}(x) = H_\lambda(x, 1) = (1 - \beta)c(x, 1) + W_\lambda$.

Proof. Consider the stopping time $\tau = 0$. The policy h_0 takes the active action at time 0 and follows the optimal policy afterwards. Thus, for any $x \in \mathcal{X}$,

$$J^{(h_0)}(x) = (1 - \beta)(c(x, 1) + \lambda) + \beta \sum_{y \in \mathcal{X}} P_{xy}(1)V_\lambda(y) = H_\lambda(x, 1).$$

By (4) and (6) we have $H_\lambda(x, 1) = (1 - \beta)c(x, 1) + W_\lambda(x)$. □

We now proceed with the proof of Proposition 1. By definition, $\Pi_\lambda^{(a)} = \Pi_\lambda$. We establish the equality of other characterizations:

- (i) $\Pi_\lambda^{(a)} = \Pi_\lambda^{(b)}$. We have $x \in \Pi_\lambda \stackrel{(a)}{\iff} g_\lambda(x) = 0 \stackrel{(b)}{\iff} H_\lambda(x, 0) < H_\lambda(x, 1)$, where (a) follows from (5) and (b) follows from the dynamic program (3).

(ii) $\Pi_\lambda^{(b)} \subseteq \Pi_\lambda^{(c)}$. Let σ denote the hitting time of $\mathcal{X} \setminus \Pi_\lambda$. If we start in the state $x \in \Pi_\lambda^{(b)} = \Pi_\lambda$, then the policy $h_{\sigma,\lambda}$ is same as the optimal policy. Hence, $J_\lambda^{(h_{\sigma,\lambda})}(x) = H_\lambda(x, 0)$. Thus, for any $x \in \Pi_\lambda^{(b)} = \Pi_\lambda$,

$$J_\lambda^{(h_{\sigma,\lambda})}(x) = H_\lambda(x, 0) \stackrel{(a)}{<} H_\lambda(x, 1) \stackrel{(b)}{=} J_\lambda^{(h_0)}(x),$$

where (a) follows from fact that $x \in \Pi_\lambda^{(b)}$ and (b) from Lemma 6.

(iii) $\Pi_\lambda^{(c)} \subseteq \Pi_\lambda^{(b)}$. Let $x \in \Pi_\lambda^{(c)}$, and let $\sigma \in \Sigma$ denote a stopping time such that $J_\lambda^{(h_{\sigma,\lambda})}(x) < J_\lambda^{(h_0)}(x)$. Now, the optimal policy performs at least as well as the policy $h_{\sigma,\lambda}$. Therefore, $V_\lambda(x) \leq J_\lambda^{(h_{\sigma,\lambda})}(x)$. Combining this result with Lemma 6 we have $V_\lambda(x) < H_\lambda(x, 1)$. Thus, we must have $V_\lambda(x) = H_\lambda(x, 0)$, which results in $H_\lambda(x, 0) < H_\lambda(x, 1)$, which implies $x \in \Pi_\lambda^{(b)}$.

(iv) $\Pi_\lambda^{(c)} = \Pi_\lambda^{(d)}$. According to the definitions of $L(x, \tau)$ and $W_\lambda(x)$ we have

$$J_\lambda^{(h_{\tau,\lambda})}(x) = (1 - \beta)L(x, \tau) + \mathbb{E}[\beta^\tau W_\lambda(X_\tau) | X_0 = x]. \tag{23}$$

Thus, $J_\lambda^{(h_{\sigma,\lambda})}(x) < J_\lambda^{(h_0)}(x)$ if and only if

$$(1 - \beta)L(x, \sigma) + \mathbb{E}[\beta^\sigma W_\lambda(X_\sigma) | X_0 = x] < (1 - \beta)c(x, 1) + W_\lambda(x), \tag{24}$$

where we have used (23) for $J_\lambda^{(h_{\sigma,\lambda})}(x)$ and Lemma 6 for $J_\lambda^{(h_0)}(x)$. Rearranging the terms of (24) we get the expression in $\Pi_\lambda^{(d)}$. Hence, $\Pi_\lambda^{(c)} = \Pi_\lambda^{(d)}$.

Appendix B. Proof of Theorem 1

B.1. Proof of Theorem 1(a)

We first present a preliminary result. Let $\Delta_\lambda := \lambda'' - \lambda'$.

Lemma 7. Under (11), for any $\lambda'' > \lambda'$ and $\sigma \in \Sigma, \sigma \neq 0$, we have that for any $x \in \mathcal{X}$,

$$W_{\lambda'}(x) - \mathbb{E}[\beta^\sigma W_{\lambda'}(X_\sigma) | X_0 = x] \leq W_{\lambda''}(x) - \mathbb{E}[\beta^\sigma W_{\lambda''}(X_\sigma) | X_0 = x].$$

Proof. By (6), for any $x \in \mathcal{X}$ we have

$$\begin{aligned} & (W_{\lambda''}(x) - \mathbb{E}[\beta^\sigma W_{\lambda''}(X_\sigma) | X_0 = x]) - (W_{\lambda'}(x) - \mathbb{E}[\beta^\sigma W_{\lambda'}(X_\sigma) | X_0 = x]) \\ &= (1 - \beta)\Delta_\lambda(1 - M(x, \sigma)) + \beta \mathbb{E} \left[\sum_{y \in \mathcal{X}} (P_{xy}(1) - \beta^\sigma P_{X_\sigma y}(1))(V_{\lambda''}(y) - V_{\lambda'}(y)) \mid X_0 = x \right]. \end{aligned} \tag{25}$$

Now, since $\sigma \geq 1$, $M(x, \sigma) \leq \beta$, and

$$(1 - \beta)\Delta_\lambda(1 - M(x, \sigma)) \geq \Delta_\lambda(1 - \beta)^2. \tag{26}$$

Now consider

$$\begin{aligned} & \beta \mathbb{E} \left[\sum_{y \in \mathcal{X}} (P_{xy}(1) - \beta^\sigma P_{X_\sigma y}(1))(V_{\lambda''}(y) - V_{\lambda'}(y)) \mid X_0 = x \right] \\ & \stackrel{(a)}{\geq} \beta \mathbb{E} \left[\sum_{y \in \mathcal{X}} (P_{xy}(1) - \beta P_{X_\sigma y}(1))(V_{\lambda''}(y) - V_{\lambda'}(y)) \mid X_0 = x \right] \\ & \stackrel{(b)}{\geq} \beta \Delta_\lambda \mathbb{E} \left[\sum_{y \in \mathcal{X}} \left\{ [P_{xy}(1) - \beta P_{X_\sigma y}(1)]^+ N^{(g_{\lambda''})}(y) \right. \right. \\ & \qquad \qquad \qquad \left. \left. + [P_{xy}(1) - \beta P_{X_\sigma y}(1)]^- N^{(g_{\lambda'})}(y) \right\} \mid X_0 = x \right] \\ & \stackrel{(c)}{\geq} -\Delta_\lambda(1 - \beta)^2, \end{aligned} \tag{27}$$

where (a) holds because $\sigma \geq 1$, (b) holds by Lemma 2, and (c) follows from (11). Substituting (26) and (27) in (25), we get the result of the lemma. \square

We now proceed with the proof of Theorem 1(a). Consider $\lambda' < \lambda''$. Suppose $x \in \Pi_{\lambda'}$. By Proposition 1(d), there exists a $\sigma \neq 0$ such that $(1 - \beta)(L(x, \sigma) - c(x, 1)) < W_{\lambda'}(x) - \mathbb{E}[\beta^\sigma W_{\lambda'}(X_\sigma) \mid X_0 = x]$. Combining this result with the result of Lemma 7, we infer $(1 - \beta)(L(x, \sigma) - c(x, 1)) < W_{\lambda''}(x) - \mathbb{E}[\beta^\sigma W_{\lambda''}(X_\sigma) \mid X_0 = x]$. Thus, $x \in \Pi_{\lambda''}$. Hence, $\Pi_{\lambda'} \subseteq \Pi_{\lambda''}$ and the RB is indexable.

B.2. Proof of Theorem 1.b

Consider $\lambda' < \lambda''$. An RB is indexable if $\Pi_{\lambda'} \subseteq \Pi_{\lambda''}$, or equivalently, for any x such that $H_{\lambda'}(x, 0) < H_{\lambda'}(x, 1)$, we have $H_{\lambda''}(x, 0) < H_{\lambda''}(x, 1)$. A sufficient condition for this is to show that $H_{\lambda'}(x, 1) - H_{\lambda'}(x, 0) \leq H_{\lambda''}(x, 1) - H_{\lambda''}(x, 0)$, or equivalently, that $H_{\lambda''}(x, 0) - H_{\lambda'}(x, 0) \leq H_{\lambda''}(x, 1) - H_{\lambda'}(x, 1)$. We prove this inequality as follows.

Let $\Delta_\lambda = \lambda'' - \lambda'$. By (4), for any $x \in \mathcal{X}$ we have

$$\begin{aligned} & (H_{\lambda''}(x, 1) - H_{\lambda'}(x, 1)) - (H_{\lambda''}(x, 0) - H_{\lambda'}(x, 0)) \\ & = \Delta_\lambda(1 - \beta) + \beta \sum_{y \in \mathcal{X}} (P_{xy}(1) - P_{xy}(0))(V_{\lambda''}(y) - V_{\lambda'}(y)) \\ & \stackrel{(a)}{\geq} \Delta_\lambda \left(1 - \beta + \beta \sum_{y \in \mathcal{X}} [P_{xy}(1) - P_{xy}(0)]^+ N^{(g_{\lambda''})}(y) + [P_{xy}(1) - P_{xy}(0)]^- N^{(g_{\lambda'})}(y) \right) \stackrel{(b)}{\geq} 0, \end{aligned}$$

where (a) follows from Lemma 2 and (b) holds by (12). Therefore the RB is indexable.

Appendix C. Proof of Proposition 2

We prove the result of each part separately.

(a) This follows from observing that

$$\begin{aligned} & \sum_{y \in \mathcal{X}} \left\{ [\beta P_{zy}(1) - P_{xy}(1)]^+ N^{(g)}(y) - [P_{xy}(1) - \beta P_{zy}(1)]^+ N^{(h)}(y) \right\} \\ & \stackrel{(a)}{\leq} \sum_{y \in \mathcal{X}} [\beta P_{zy}(1) - P_{xy}(1)]^+ N^{(g)}(y) \\ & \stackrel{(b)}{\leq} \sum_{y \in \mathcal{X}} [\beta P_{zy}(1) - P_{xy}(1)]^+ \leq \max_{x, z \in \mathcal{X}} \sum_{y \in \mathcal{X}} [\beta P_{zy}(1) - P_{xy}(1)]^+, \end{aligned}$$

where we are ignoring negative terms in (a) and using $N^{(g)}(x) \leq 1$ in (b).

(b) For any $x, y, z \in \mathcal{X}$, $P_{xy}(1) - \beta P_{zy}(1) = (1 - \beta)P_{xy}(1)$. Thus,

$$\begin{aligned} & \sum_{y \in \mathcal{X}} \left\{ [\beta P_{zy}(1) - P_{xy}(1)]^+ N^{(g)}(y) - [P_{xy}(1) - \beta P_{zy}(1)]^+ N^{(h)}(y) \right\} \\ & = - \sum_{y \in \mathcal{X}} (1 - \beta)P_{xy}(1)N^{(h)}(y) \leq 0 < \frac{(1 - \beta)^2}{\beta}. \end{aligned}$$

(c) This follows from observing that

$$\begin{aligned} & \sum_{y \in \mathcal{X}} \left\{ [P_{xy}(0) - P_{xy}(1)]^+ N^{(g)}(y) - [P_{xy}(1) - P_{xy}(0)]^+ N^{(h)}(y) \right\} \\ & \stackrel{(a)}{\leq} \sum_{y \in \mathcal{X}} [P_{xy}(0) - P_{xy}(1)]^+ N^{(g)}(y) \\ & \stackrel{(b)}{\leq} \sum_{y \in \mathcal{X}} [P_{xy}(0) - P_{xy}(1)]^+ \leq \max_{x \in \mathcal{X}} \sum_{y \in \mathcal{X}} [P_{xy}(0) - P_{xy}(1)]^+, \end{aligned}$$

where we are ignoring negative terms in (a) and using $N^{(g)}(x) \leq 1$ in (b).

(d) The fact that $\beta \leq 0.5$ implies that

$$\frac{1 - \beta}{\beta} \geq 1 \geq \max_{x \in \mathcal{X}} [P_{xy}(0) - P_{xy}(1)]^+,$$

which is the same as the sufficient condition (c) established above.

Appendix D. Proof of Lemma 3

The proof of each part is as follows:

1. Since the model is indexable and $y \in \Gamma_{d+1}$, $w(d) = \lambda_{d+1}$. Therefore, the optimal policy is indifferent between choosing the active and the passive action at $\lambda = \lambda_{d+1}$.
2. By definition, for any $\lambda \in (\lambda_d, \lambda_{d+1}]$, h_d is an optimal policy. Therefore, we have $J_\lambda^{(h_d, y)}(x) \geq J_\lambda^{(h_d)}(x)$ with $y \in \mathcal{X} \setminus \mathcal{P}_d$ for all $x \in \mathcal{X}$, with equality if $y \in \Gamma_{d+1}$ and $\lambda = \lambda_{d+1}$.

Funding information

This research was funded in part by the Innovation for Defence Excellence and Security (IDEaS) Program of the Canadian Department of National Defence through grant CFPMN2-037, and the Fonds de Recherche du Quebec—Nature et Technologies (FRQNT).

Competing interests

There were no competing interests to declare which arose during the preparation or publication process for this article.

References

- [1] ABAD, C. AND IYENGAR, G. (2016). A near-optimal maintenance policy for automated DR devices. *IEEE Trans. Smart Grid* **7**, 1411–1419.
- [2] AKBARZADEH, N. AND MAHAJAN, A. (2019). Dynamic spectrum access under partial observations: a restless bandit approach. In *Canadian Workshop on Information Theory*, Institute of Electrical and Electronics Engineers, Piscataway, NJ, pp. 1–6.
- [3] AKBARZADEH, N. AND MAHAJAN, A. (2019). Restless bandits with controlled restarts: indexability and computation of Whittle index. In *58th IEEE Conference on Decision and Control*, Institute of Electrical and Electronics Engineers, Piscataway, NJ, pp. 7294–7300.
- [4] AKBARZADEH, N. AND MAHAJAN, A. (2020). An improved algorithm to compute Whittle index. Available at <https://codeocean.com/capsule/8680851/tree/v1>.
- [5] ANSELL, P. S., GLAZEBROOK, K. D., NIÑO-MORA, J. AND O'KEEFFE, M. (2003). Whittle's index policy for a multi-class queueing system with convex holding costs. *Math. Operat. Res.* **57**, 21–39.
- [6] ARCHIBALD, T. W., BLACK, D. P. AND GLAZEBROOK, K. D. (2009). Indexability and index heuristics for a simple class of inventory routing problems. *Operat. Res.* **57**, 314–326.
- [7] AVRACHENKOV, K., AYESTA, U., DONCEL, J. AND JACKO, P. (2013). Congestion control of TCP flows in internet routers by means of index policy. *Computer Networks* **57**, 3463–3478.
- [8] AYESTA, U., ERAUSQUIN, M. AND JACKO, P. (2010). A modeling framework for optimizing the flow-level scheduling with time-varying channels. *Performance Evaluation* **67**, 1014–1029.
- [9] BERTSIMAS, D. AND NIÑO-MORA, J. (1996). Conservation laws, extended polymatroids and multiarmed bandit problems; a polyhedral approach to indexable systems. *Math. Operat. Res.* **21**, 257–306.
- [10] CHAKRAVORTY, J. AND MAHAJAN, A. (2014). Multi-armed bandits, Gittins index, and its calculation. In *Methods and Applications of Statistics in Clinical Trials: Planning, Analysis, and Inferential Methods*, Vol. 2, John Wiley, Hoboken, NJ, pp. 416–435.
- [11] DEO, S. *et al.* (2013). Improving health outcomes through better capacity allocation in a community-based chronic care model. *Operat. Res.* **61**, 1277–1294.
- [12] EGIDI, N. AND MAPONI, P. (2006). A Sherman–Morrison approach to the solution of linear systems. *J. Comput. Appl. Math.* **189**, 703–718.
- [13] GITTINS, J., GLAZEBROOK, K. AND WEBER, R. (2011). *Multi-armed Bandit Allocation Indices*. John Wiley, Chichester.
- [14] GITTINS, J. C. (1979). Bandit processes and dynamic allocation indices. *J. R. Statist. Soc. B* **41**, 148–177.
- [15] GLAZEBROOK, K., HODGE, D. AND KIRKBRIDE, C. (2013). Monotone policies and indexability for bidirectional restless bandits. *Adv. Appl. Prob.* **45**, 51–85.
- [16] GLAZEBROOK, K. AND MITCHELL, H. (2002). An index policy for a stochastic scheduling model with improving/deteriorating jobs. *Naval Res. Logistics* **49**, 706–721.
- [17] GLAZEBROOK, K. D., KIRKBRIDE, C. AND OUENNICHE, J. (2009). Index policies for the admission control and routing of impatient customers to heterogeneous service stations. *Operat. Res.* **57**, 975–989.
- [18] GLAZEBROOK, K. D., MITCHELL, H. M. AND ANSELL, P. S. (2005). Index policies for the maintenance of a collection of machines by a set of repairmen. *Europ. J. Operat. Res.* **165**, 267–284.
- [19] GLAZEBROOK, K. D., RUIZ-HERNANDEZ, D. AND KIRKBRIDE, C. (2006). Some indexable families of restless bandit problems. *Adv. Appl. Prob.* **38**, 643–672.
- [20] JACKO, P. (2012). Optimal index rules for single resource allocation to stochastic dynamic competitors. In *VALUETOOLS '11: Proc. 5th International ICST Conference on Performance Evaluation Methodologies and Tools*, Association for Computing Machinery, New York, pp. 425–433.
- [21] LIU, K. AND ZHAO, Q. (2010). Indexability of restless bandit problems and optimality of Whittle index for dynamic multichannel access. *IEEE Trans. Inf. Theory* **56**, 5547–5567.

- [22] LOTT, C. AND TENEKETZIS, D. (2000). On the optimality of an index rule in multichannel allocation for single-hop mobile networks with multiple service classes. *Prob. Eng. Inf. Sci.* **14**, 259–297.
- [23] NIÑO-MORA, J. (2001). Restless bandits, partial conservation laws and indexability. *Adv. Appl. Prob.* **33**, 76–98.
- [24] NIÑO-MORA, J. (2002). Dynamic allocation indices for restless projects and queueing admission control: a polyhedral approach. *Math. Program.* **93**, 361–413.
- [25] NIÑO-MORA, J. (2007). Dynamic priority allocation via restless bandit marginal productivity indices. *TOP* **15**, 161–198.
- [26] NIÑO-MORA, J. (2006). Restless bandit marginal productivity indices, diminishing returns, and optimal control of make-to-order/make-to-stock M/G/1 queues. *Math. Operat. Res.* **31**, 50–84.
- [27] PAPADIMITRIOU, C. H. AND TSITSIKLIS, J. N. (1999). The complexity of optimal queueing network control. *Math. Operat. Res.* **24**, 293–305.
- [28] PUTERMAN, M. L. (2014). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley, Hoboken, NJ.
- [29] QIAN, Y., ZHANG, C., KRISHNAMACHARI, B. AND TAMBE, M. (2016). Restless poachers: handling exploration–exploitation tradeoffs in security domains. In *AAMAS '16: Proc. 2016 International Conference on Autonomous Agents and Multiagent Systems*, Association for Computing Machinery, New York, pp. 123–131.
- [30] WANG, J., REN, X., MO, Y. AND SHI, L. (2020). Whittle index policy for dynamic multichannel allocation in remote state estimation. *IEEE Trans. Automatic Control* **65**, 591–603.
- [31] WEBER, R. R. AND WEISS, G. (1990). On an index policy for restless bandits. *J. Appl. Prob.* **27**, 637–648.
- [32] WHITTLE, P. (1988). Restless bandits: activity allocation in a changing world. *J. Appl. Prob.* **25**, 287–298.
- [33] YU, Z., XU, Y. AND TONG, L. (2018). Deadline scheduling as restless bandits. *IEEE Trans. Automatic Control* **63**, 2343–2358.