


Markovian Restless Bandits and Index Policies: A Review

José Niño-Mora 

Department of Statistics, Carlos III University of Madrid, 28903 Getafe, Madrid, Spain; jose.nino@uc3m.es

Abstract: The restless multi-armed bandit problem is a paradigmatic modeling framework for optimal dynamic priority allocation in stochastic models of wide-ranging applications that has been widely investigated and applied since its inception in a seminal paper by Whittle in the late 1980s. The problem has generated a vast and fast-growing literature from which a significant sample is thematically organized and reviewed in this paper. While the main focus is on priority-index policies due to their intuitive appeal, tractability, asymptotic optimality properties, and often strong empirical performance, other lines of work are also reviewed. Theoretical and algorithmic developments are discussed, along with diverse applications. The main goals are to highlight the remarkable breadth of work that has been carried out on the topic and to stimulate further research in the field.

Keywords: Markov decision processes; bandit problems; restless bandits; dynamic and stochastic resource allocation; index policies; online learning; regret analysis

MSC: 90B36; 90C39; 90C4



Citation: Niño-Mora, J. Markovian Restless Bandits and Index Policies: A Review. *Mathematics* **2023**, *11*, 1639. <https://doi.org/10.3390/math11071639>

Academic Editor: Andreas C. Georgiou

Received: 31 January 2023

Revised: 10 March 2023

Accepted: 25 March 2023

Published: 28 March 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

In the much-studied *multi-armed bandit problem* (MABP), one is presented with a number of generic dynamic and stochastic *projects* modeled as binary-action (active/selected or passive/rested) *Markov decision processes* (MDPs) (see, e.g., Puterman [1]) with the aim of maximizing the expected reward accrued by dynamically selecting one project to engage at each time. The problem, named after a multi-armed slot machine (or “bandit”) —whence the projects are commonly referred to in the literature as *bandits* or *arms*— is widely regarded as a major modeling framework for addressing the *exploration vs. exploitation trade-off* in widely diverse settings. A key assumption is that project states remain frozen while passive, which renders the problem tractable in some relevant scenarios, most notably under the expected total infinite-horizon (geometrically) discounted-reward criterion. In this case, as first shown by Gittins and Jones [2], one can attach to each project a scalar function of its state called an *index* that is based solely on the project’s characteristics, such that the resulting *priority-index policy*, which selects at each time a project with highest index, is optimal.

In a seminal work, Whittle [3] first considered a model that dropped the assumption of frozen states by allowing rested projects to change states. Projects then become *restless*, giving rise to the *restless multi-armed bandit problem* (RMABP). Whittle outlined potential applications to demonstrate the expanded modeling power of the new modeling framework, and, to overcome its intractability due to the *curse of dimensionality*, he proposed a tractable heuristic priority-index policy based on Lagrangian relaxation and decomposition ideas. He further conjectured a form of asymptotic optimality for this *Whittle index policy*.

The impact of Whittle’s work took off to a slow start, but over the following decade researchers began to realize the potential of Whittle’s model and solution approach. The number of researchers interested in the RMABP and its variants, as well as the resulting number of published papers on the subject, has grown steadily since then and has picked up steam in recent years. The literature on the RMABP, whether on its theoretical, algorithmic, or ap-

plication aspects, is currently vast to the point where it is virtually infeasible for researchers to keep up to date with the latest advances in the field.

This is the motivation of the present paper, which reviews and highlights a hopefully representative albeit limited and incomplete sample of the enormous body of work published on the topic by a myriad of researchers over the thirty-plus years since its inception by Whittle, emphasizing breadth over depth. The main goal is to stimulate further research on this subject both by bringing it to the attention of researchers who may not have encountered it previously, and by providing an overview of its wide-ranging possibilities and open problems for those who may have missed important research directions. This review focuses on solution approaches that develop, test, and analyze index policies due to their intuitive appeal, tractability, asymptotic optimality properties, and often strong empirical performance. The author has strived to go well beyond his comfort zone to provide a balanced coverage of the subject, minimizing bias due to his familiarity with index policies for models with known parameters, although some inevitable bias may remain.

The review is organized as follows. Section 2 surveys the antecedents to the RMABP, in particular, the classic MABP and the Gittins index policy. Section 3 formulates the RMABP and outlines the Whittle index policy. Section 4 reviews works on the complexity, approximation, and relaxations of the RMABP. Section 5 focuses on indexability, that is, the existence of the Whittle index and extensions. Section 6 discusses works on means of computing the Whittle index. Section 7 considers works that establish the optimality of the myopic index policy for the RMABP, whereas Section 8 reviews the asymptotic optimality of index policies. Section 9 surveys multi-action restless bandits. Section 10 considers policies that are different from Whittle's and based on Lagrangian and fluid relaxations. Section 11 addresses reinforcement learning and, in particular, Q-learning solution approaches. Section 12 surveys works on the RMABP from the perspective of online learning. Sections 13 and 14 are devoted to works on applications of the RMABP in diverse settings. The former section focuses on MDP models and the latter on *partially observable MDP* (POMDP) models. Finally, Section 15 concludes this paper.

2. Antecedents: Multi-Armed Bandits and the Gittins Index Policy

The investigation of bandit-like problems has its early roots in the seminal work of Thompson [4], who outlined a research-planning problem that aimed to allocate ("apportion" in the words of Thompson) two Bernoulli treatments with unknown success probabilities to a set of individuals based on prior statistical evidence using Bayes' rule to reduce the expected number that received inferior treatment. Thompson [5] further clarified the use and implementation of his proposed approach for dynamic sequential allocation (the basis of the widely investigated and applied *Thompson sampling* method), sketched an extension to multiple Bernoulli treatments, and even reported the results of a pioneering simulation study *avant la lettre*, illustrating its practical effectiveness.

This precursor to the MABP was later refined and cast into the framework of the then-budding theory of the *sequential design of experiments* (see Wald [6] for its first systematic account), which addressed the design of sequential sampling procedures "in which the size and composition of the samples are not fixed in advance but are functions of the observations themselves", as stated by Robbins [7]. In that paper, Robbins considered the two-armed bandit problem of sampling from two Bernoulli populations to maximize the expected number of successes from a fixed number of observations under the practically relevant assumption that population parameters are unknown. He thus adopted a *learning* perspective and showed that one can design sampling policies that, as the number of observations grows, asymptotically attain the maximum expected reward per observation.

Bradt et al. [8] adopted a Bayesian perspective and considered a finite-horizon two-armed bandit problem to maximize the expected sum of a fixed number of observations through sequential sampling from two unknown statistical populations. Their approach assumed a known prior distribution. They identified conditions under which the myopic index policy is optimal but, in general, noted the difficulty of elucidating the structure of

optimal policies. Yet, in the case of two Bernoulli populations, one with known parameter, they showed that the optimal policy is characterized by a critical quantity, which we would call an *index*, a function of both the remaining number of observations and the current posterior distribution of the unknown population. At each time, it is optimal to sample from the unknown population if and only if its current index value is greater than or equal to the known population’s parameter. Bellman [9] considered the problem under the infinite-horizon discounted reward criterion, also establishing the optimality of an *index policy*. In this case, the index is a function only of the posterior distribution of the success probability of the unknown population.

MABPs were later cast into the more general framework of multi-stage decision processes, in particular, MDPs (see the pioneering monograph of Bellman ([10] Chapter 11) and the later work of Howard [11], and for more recent accounts of MDPs, see, e.g., the comprehensive textbooks of Puterman [1] and Bertsekas [12,13]).

A general formulation of a standard version of the MABP in the MDP setting is as follows. A decision maker is faced with a finite collection of N reward-yielding dynamic and stochastic *projects*. The *state* $X_{n,t}$ of project $n = 1, \dots, N$ evolves over discrete time periods $t = 0, 1, 2, \dots$ over an infinite horizon across the *state space* \mathcal{X}_n . The evolution and reward of project n in period t depend both on the current state $X_{n,t}$ and the chosen *action* $A_{n,t}$, which can take two values: 1 (*active*, i.e., *engaging* the project) and 0 (*passive*, i.e., *resting* it). When the project occupies state $X_{n,t} = i_n$ and action $A_{n,t} = a_n$ is selected, the project yields an expected reward $r_n^{a_n}(i_n)$, which is time-discounted with the factor $0 < \beta < 1$ and its state moves to $X_{n,t+1} = j_n$ with a transition probability of $p_n^{a_n}(i_n, j_n)$, where $p_n^0(i_n, j_n) = \delta_{i_n j_n}$ (Kronecker’s delta) so passive projects do not change state. It is usually assumed that $r_n^0(i_n) \equiv 0$ for all $i_n \in \mathcal{X}_n$, i.e., passive projects give no rewards but this assumption is nonessential. At each time t , the decision maker observes the *joint state* $\mathbf{X}_t = (X_{n,t})_{n=1}^N$ and then selects one project so the *joint action* $\mathbf{A}_t = (A_{n,t})_{n=1}^N \in \{0, 1\}^N$ must satisfy

$$\sum_{n=1}^N A_{n,t} = 1, \quad t = 0, 1, 2, \dots \tag{1}$$

The class of non-anticipative *policies* that prescribes these feasible action choices is denoted as Π and the expectation starting from the joint state \mathbf{i} under policy π is denoted as $\mathbb{E}_{\mathbf{i}}^{\pi}[\cdot]$. The goal of the *infinite-horizon discounted MABP* is to find an *optimal policy* $\pi^* \in \Pi$, which, for *any* initial joint state \mathbf{i} , maximizes the expected total discounted reward earned over an infinite horizon, which is given by

$$\mathbb{E}_{\mathbf{i}}^{\pi^*} \left[\sum_{t=0}^{\infty} \sum_{n=1}^N r_n^{A_{n,t}}(X_{n,t}) \beta^t \right], \tag{2}$$

Note that applying the conventional solution approach based on solving numerically Bellman’s optimality equations using *dynamic programming* (DP) is hindered by the *curse of dimensionality* as, even in the finite state case, the number of joint states and hence of equations increases exponentially with the number of projects.

After being considered intractable for a long time, Gittins and Jones [2] made a breakthrough by demonstrating that the optimal policies for the above MABP have a strikingly simple structure. Specifically, for each project n , there exists an *index* $\varphi_n(i_n)$, which is a function of the project state i_n , that is based solely on its characteristics (rewards and transition law) such that the resulting *index policy*, which selects a project with the highest index value at each time, is optimal. This index was introduced in the aforementioned work of Bellman [9] for a special Bernoulli bandit model (see also the paper by Gittins [14] and the monographs by Gittins [15] and Gittins et al. [16]).

Researchers have utilized various approaches to offer different proofs for these celebrated results on the optimality of what is widely known as the *Gittins index* policy (see, e.g., Whittle [17], Varaiya et al. [18], Weber [19], Tsitsiklis [20], and Bertsimas and Niño-Mora [21]). Furthermore, the results have been extended to incorporate variations

that preserve the optimality of index policies, e.g., random project arrivals (Whittle [22]), non-Markovian evolution (Varaiya et al. [18]), the minimization of the expected cost for one project to reach a target state (Dumitriu et al. [23]), and constrained arm switches (Bao et al. [24]), among others.

In independent ground-breaking work, Klimov [25] presented a landmark result analogous to that in [2] by first establishing the optimality of an index policy for scheduling a multi-class M/G/1 queue with Bernoulli job feedback and linear holding costs under the average criterion, which is a semi-Markov MABP with project arrivals. The results were extended to the more general *branching bandit* problem by Meilijson and Weiss [26], who analyzed it under the average criterion, and Weiss [27], who analyzed it under the discounted criterion. For more on Klimov’s results, the reader is referred to the review article by Niño-Mora [28].

3. Restless Multi-Armed Bandits and the Whittle Index Policy

Whittle [3] introduced a major extension to the classic MABP by allowing rested projects to change state, in which case they are called *restless bandits*, with state transitions being independent across projects. The model was also considered in independent work by O’Flaherty [29]. Whittle noted that the restless feature substantially expanded the modeling power, which, however, came at the expense of tractability, as index policies are generally not optimal for the resulting RMABP. Whittle further allowed the number of projects selected at each time, M , to not necessarily be one. The *joint action* $\mathbf{A}_t = (A_{n,t})_{n=1}^N \in \{0, 1\}^N$ must thus satisfy

$$\sum_{n=1}^N A_{n,t} = M, \quad t = 0, 1, 2, \dots \tag{3}$$

Whittle focused on the average-reward criterion, the goal being to find a policy satisfying (3) that maximizes the steady-state reward

$$\mathbb{E}^\pi \left[\sum_{n=1}^N r_n^{A_n}(X_n) \right], \tag{4}$$

where (X_n, A_n) has the steady-state distribution of state-action pairs for project n under policy π and $\mathbb{E}^\pi[\cdot]$ denotes the expectation under this policy, which does not depend on the initial joint state under suitable regularity conditions. He considered a *relaxation* of this problem, where the sample-path constraint (3) is *relaxed* by requiring only that M projects be active *on average* so that

$$\mathbb{E}^\pi \left[\sum_{n=1}^N A_n \right] = M. \tag{5}$$

Whittle showed that this relaxed problem, whose optimal value provides an upper bound on that of the original problem, was amenable to a Lagrangian solution approach by attaching a Lagrange multiplier λ to constraint (5). In the resulting *Lagrangian relaxation*, the constraint is brought into the objective, with the multiplier playing the economic role of a *subsidy for passivity*. He further noted that the Lagrangian relaxation allows a decomposition, which entails single-project subproblems of the form

$$\underset{\pi_n \in \Pi_n}{\text{maximize}} \mathbb{E}^\pi [r_n^{A_n}(X_n) + \lambda(1 - A_n)], \tag{6}$$

where Π_n is the class of admissible policies for operating project n *in isolation*. Now, on the one hand, one can consider the *Lagrangian dual problem*, which is to find the best multiplier λ for which the Lagrangian relaxation yields the best possible upper bound. On the other hand, one must consider the properties of optimal policies for the subproblems in (6).

Whittle proposed the following *indexability* property for a single project: as the passivity subsidy λ increases from $-\infty$ to $+\infty$, the set of project states where it is optimal in (6) to rest the project increases monotonically from the empty set to the full project state

space. The index $\varphi_n(i_n)$ for project n in state i_n is then defined as the value of λ at which both actions are optimal in i_n , which is unique under indexability.

If each constituent project in an RMABP is indexable, each single-project subproblem in (6) is solved optimally by an index policy with index $\varphi_n(i_n)$. It is optimal to select (resp. rest) project n in state i_n if and only if $\varphi_n(i_n) \geq \lambda$ (resp. $\varphi_n(i_n) \leq \lambda$). In such a case, Whittle proposed to use the resulting priority-index policy as a heuristic rule for the RMABP. At each time, select M projects with the highest index values. He further conjectured that, although the *Whittle index policy* will generally be suboptimal, it should have a form of *asymptotic optimality*. In the case of a population of projects with a fixed number of distinct project types in fixed proportions, as M and N grow to infinity with a fixed ratio M/N , the average reward per project should converge to that under an optimal policy.

Whittle also proposed to use this approach to define an index policy for the discounted RMABP, noting that for non-restless projects his index reduces to the Gittins index.

4. Complexity, Approximation, and Relaxations

In contrast to the classic MABP for which the Gittins index policy allows breaking the curse of dimensionality under the infinite-horizon discounted-reward criterion, the RMAB is, in general, computationally intractable, as first shown by Papadimitriou and Tsitsiklis [30]. Among a host of ground-breaking complexity results, they showed the RMABP to be PSPACE-hard, even for projects with deterministic state transitions. As they stated, this “is considered a much more convincing evidence of intractability than NP-hardness”.

Regarding theoretical performance guarantees for heuristic policies, Guha et al. [31] presented the first *constant-factor approximation algorithms* based on *linear programming* (LP) duality for several classes of RMABPs under the average-reward criterion, including a $(2 + \epsilon)$ -approximation algorithm for the *feedback RMABP*, which includes the *partially observable MDP* (POMDP) RMABP model for opportunistic spectrum access considered in Liu and Zhao [32,33] and in Le Ny et al. [34]. They considered an index policy that is closely related to Whittle’s, and extended their results to the wider class of *monotone bandits*.

Other researchers have followed in the footsteps of [31] to design constant-factor approximation algorithms for related RMABP classes that incorporate additional relevant features (see, e.g., Wan and Xu [35], Xu and Song [36], and Xu and Wang [37], which incorporate, respectively, project weights, channel interferences, and multiple constraints).

As outlined above, Whittle proposed a problem relaxation to approximate the RMABP. This idea was extended by Bertsimas and Niño-Mora [38] who considered a hierarchy of N increasingly stronger *linear programming* (LP) relaxations (where N is the number of projects), with increasing size and hence complexity. The first relaxation in the hierarchy is that in [3], whereas the last is the standard exact formulation of exponential size in N . They further proposed and tested an alternative index policy based on the optimal solution to the first-order relaxation, which does not require projects to be indexable.

Whittle’s relaxation approach to the RMABP was further extended by Hawkins [39] and then by Adelman and Mersereau [40] into the broader setting of so-called *weakly coupled* MDPs, in which a set of sample-path linking constraints couples a collection of otherwise independent constituent subproblems. They compared bounds and policies obtained from two different problem relaxations, which were based, respectively, on LP-based *approximate DP* (see, e.g., the monographs by Powell [41] and Bertsekas [13]) and Lagrangian relaxation. Among their results, they showed that both relaxations entail fitting an additively separable value function approximation to the Bellman optimality equations and established that the approximate DP bound was tighter than the one based on Lagrangian relaxation.

The analyses of relaxations in [39,40] were further refined by Brown and Zhang [42], who provided theoretical justification in the form of sufficient conditions for the empirically observed fact that the gap between the two different upper bounds on the optimal value considered in [40] was typically small.

5. Indexability

Whittle [3] pointed out that “One would very much like to have simple sufficient conditions for indexability; at the moment, none are known”. Most works where Whittle’s index policy has been applied to particular RMABP models (see, e.g., the references in Sections 13 and 14) have followed the following scheme to argue that individual projects are indexable: (i) show that the subproblems in (6) (or their discounted counterparts) can be solved optimally using a certain family of structured policies for any value of the multiplier λ , typically threshold policies for linearly-ordered project state spaces; (ii) obtain an optimal policy within such a family of policies as a function of λ for these subproblems, e.g., an optimal threshold policy; and (iii) show that the function derived in the previous step satisfies a required form of monotonicity in λ , under which the Whittle index is then obtained by inverting this function. In some works, step (i) is carried out rigorously, e.g., using threshold policies in Le Ny et al. [34], Liu and Zhao [32,33], and Liu et al. [43]; using stopping rules in Fryer and Harms [44]; and using another family of policies in Caro and Yoo [45]. Yet, often, the required optimality of the family of structured policies under consideration is merely postulated or conjectured as researchers either bypass it, as, e.g., in Whittle ([46] Chapter 14.6) and in Veatch and Wein [47], or find that the proof is elusive, as in Dance and Silander [48].

A different approach based on *partial conservation laws* (PCLs) that provides sufficient indexability conditions for general restless bandits has been developed by the author in a series of papers [49–53]. It proceeds by (i) positing a family of structured policies that one guesses might be optimal for the single-project subproblems in (6); (ii) proving that a certain *marginal work* project metric is positive for the postulated family of policies; and (iii) showing that the *marginal productivity index* computed by an adaptive-greedy algorithm given in [49] is generated in a monotonic fashion. Steps (ii) and (iii) can be carried out either numerically or analytically. The main result is that under the conditions in (ii) and (iii), it follows simultaneously in one fell swoop that both the postulated family of policies is optimal for the single-project subproblems *and* the obtained marginal productivity index is the project’s Whittle index. Furthermore, under the condition in (ii), the model is indexable consistently with the postulated family of policies if and only if the condition in (iii) holds. This approach allowed to overcome the long-standing problem of establishing the optimality of threshold policies and proving indexability for the scalar Kalman filter restless bandit model, as demonstrated in the groundbreaking work of Dance and Silander [54].

6. Whittle Index Computation

In some restless bandit models, the Whittle index can be derived in closed form, e.g., in Whittle [3], Whittle ([46] Chapter 14.6), Veatch and Wein [47], and Liu et al. [43]. For other models where the Whittle index needs to be computed, an exact algorithm applied to a general n -state restless bandit satisfying PCLs was provided in [49,50] and a fast-pivoting implementation with complexity $O(n^3)$ was presented in [55,56], extending the author’s work in [57] on efficient computation of the Gittins index. Ref. [55] further provided an index algorithm, also with $O(n^3)$ complexity but with a larger leading constant, for checking indexability and computing the index without the need to satisfy the PCLs. Note that these algorithms *both* check for indexability *and* compute the Whittle index when it exists. In special cases, the complexity is reduced, e.g., in the birth–death queueing model in [50], where it is shown that threshold policies are optimal and the Whittle index is computed in $O(n)$ time. For models where projects have a continuous, real state space, the PCL approach also provides a means of computing the Whittle index, as shown in [53].

Another approach to computation of the Whittle index was provided by Qian et al. [58], who introduced sufficient conditions for indexability in their model of optimal patrol policies and provided an algorithm to test for indexability. Given indexability, they proposed to use binary search to compute the Whittle index.

Akbarzadehand and Mahajan [59] extended the adaptive-greedy algorithm in [49,50] to a version that can compute the Whittle index for an arbitrary indexable restless bandit.

They further presented an efficient implementation of their algorithm with $O(n^3)$ complexity for an n -state project and provided alternative sufficient conditions for indexability.

7. Optimality of the Myopic Policy

Although index policies are generally suboptimal for the RMABP, a substantial amount of research has been devoted to identifying conditions under which the *myopic policy*, which prioritizes projects based on their current expected reward when selected, is optimal. However, these results typically require strong assumptions about the constituent projects, most notably that they are homogenous.

The first work in this area was by O'Flaherty [29], who provided sufficient conditions for the myopic policy to be optimal in two-project RMABPs. The paper outlined ideas supporting the validity of these conditions but did not provide formal proofs.

Ehsan and Liu [60] considered an RMABP model for optimal dynamic server allocation in a multi-class single-server discrete-time queue with delayed backlog information and convex nondecreasing holding costs to minimize the expected total discounted cost over a finite or infinite horizon. They derived sufficient conditions for a myopic policy to be optimal in certain regions of the state space, in the cases of batch and Poisson job arrivals, which take the form of sufficient separation among indices.

Ahmad et al. [61] considered a POMDP RMABP model with $M = 1$ active projects at each time for dynamic multi-channel access with identical channels and demonstrated the optimality of the myopic policy under both discounted- and average-reward criteria, provided that state transitions exhibit a positive correlation over time. Liu et al. [62] incorporated error-prone channel state detection and demonstrated the optimality of the myopic policy for the two-channel case.

In a series of papers, Wang and coworkers investigated conditions under which the myopic policy is optimal for several POMDP RMABP models of optimal dynamic multi-channel access. In [63], Wang et al. considered the dynamic multi-channel access model with identical channels and the class of so-called standard reward functions, for which a closed-form condition on the discount factor ensures optimality of the myopic policy under the discounted criterion. The results were also extended to the average-reward criterion. Wang and Chen [64] considered a general setting and introduced three axioms that characterized the so-called regular functions, which yielded closed-form conditions under which the myopic policy is optimal. The results were demonstrated in RMAB problems arising in multi-channel opportunistic access with heterogeneous channels. These results were extended by Wang et al. [65] to a framework based on the so-called g-regular functions. Wang et al. [66] considered a multi-channel access problem where each of N channels is modeled as a multi-state Markov chain rather than as a two-state Markov chain as in prior work. They demonstrated that the myopic policy reduces channel selection to a round-robin policy whose optimality is established for accessing $M = N - 1$ out of N channels. Wang et al. [67] further extended this work by identifying two sets of sufficient conditions on both the eigenvalues and eigenmatrices resulting from channel-state transition matrices that guarantee the optimality of the myopic policy (see also the monograph by Wang and Chen [68]).

Ouyang and Teneketzis [69] considered a POMDP dynamic multi-channel access RMABP model, where the underlying channel-state Markov chain has an arbitrary finite number of states. They presented sufficient conditions for the optimality of a myopic sensing policy over a finite-time horizon under discounted and undiscounted-reward criteria.

Blasco and Gündüz [70] proposed a POMDP RMABP model for a multi-access wireless network with transmitting nodes, each with an energy-harvesting (EH) device and a finite-capacity rechargeable battery, with the goal of maximizing throughput. Under certain conditions on the EH processes and battery sizes, the optimality of the myopic policy is shown.

The Age of Information (AoI) is a popular metric for capturing information freshness, based on the time elapsed since the most recently delivered packet in a communication

node. Kadota et al. [71] proposed a model for minimizing the AoI and demonstrated that, in symmetric networks, the myopic policy, which prioritizes older packets for transmission, is optimal.

8. Asymptotic Optimality of Index Policies

Although index policies are generally suboptimal for an RMABP where M projects out of N are selected at each time, a substantial amount of research has sought to establish the asymptotic optimality of index policies, most notably Whittle's, in a regime where both M and N increase to infinity in a fixed ratio. Assuming a homogeneous population of projects, Weber and Weiss [72] demonstrated that Whittle's relaxation of the RMABP is asymptotically tight, in that the optimal average reward per project is the same for the relaxed and original problems in the aforementioned limit. They further showed that, although Whittle's conjecture on the asymptotic optimality of his index policy does not hold generally, it does hold under a certain condition: when "the differential equation describing the fluid approximation to the index policy has a globally stable equilibrium point." Although that condition is typically hard to check, they demonstrated in [73] that it is satisfied in the case of three-state projects.

Bagheri and Scaglione [74] introduced a significant extension to the dynamic multi-channel access RMABP model with two-state channels, the so-called cognitive compressive sensing problem, where the maximum number of channels to be sensed at each time can vary. They provided conditions under which the myopic index policy is asymptotically optimal.

Larrañaga et al. [75] considered an RMABP model for the optimal scheduling of a multi-class queue with convex holding costs and user impatience under the average cost criterion. They derived index policies and demonstrated that Whittle's index policy is asymptotically optimal in both light- and heavy-traffic regimes.

Ouyang et al. [76] considered a downlink scheduling problem modeled as a POMDP RMABP problem. They established the asymptotic optimality of Whittle's index policy for two classes of positively correlated channels under the two-state channel model, assuming a recurrence condition that can be verified numerically.

Verloop [77] considered a set of priority policies for RMABPs with possibly non-indexable projects that can incorporate project arrivals and departures, as well as multi-action projects, which are asymptotically optimal when the differential equation for the system's fluid approximation has a global attractor, similarly as in [72]. She further demonstrated that these results can be applied to Whittle's index policy in the case of indexable projects.

Fu et al. [78] considered an RMABP dynamic job assignment model for a server farm with multiple heterogeneous servers to optimize energy efficiency. They showed that under certain conditions, Whittle's index policy is asymptotically optimal, which requires a significant extension to the approach in [72]. Motivated by geographically distributed server farms where available servers for a given job are job-dependent, Fu and Moran [79] extended the model in [78]. They substantially improved on previous work by establishing not only the asymptotic optimality of an index policy but also its exponential convergence rate.

Hu and Frazier [80] considered a finite-horizon RMABP, derived an index policy based on optimal solutions of single-project subproblems, and argued its asymptotic optimality, showing an optimality gap of $o(N)$ for an N -project model.

Zayas-Cabán et al. [81] considered a finite-horizon RMABP with multi-action projects and time-dependent upper bounds on the number of projects that can be selected at each time. They derived a heuristic policy based on the optimal solution to an LP relaxation and proved its asymptotic optimality. Their analysis does not rely on indexability or stability conditions and applies to the model variant with project arrivals. The proposed policy was shown to have an optimality gap of $O(\sqrt{N} \log N)$.

Maatouk et al. [82] considered an RMABP model for the optimal scheduling of transmissions over unreliable channels to minimize the average age of an information metric and establish the asymptotic optimality of the Whittle index policy under a recurrence condition that can be verified numerically. Kriouile et al. [83] considered a model extension and presented a novel approach to show the asymptotic optimality under less stringent conditions.

Brown and Smith [84] presented index policies based on Lagrangian relaxation with an $O(\sqrt{N})$ optimality gap.

In a promising recent work, Zhang and Frazier [85] considered the finite-horizon RMABP and identified a non-degeneracy condition and a class of so-called fluid-priority policies for which the asymptotic optimality gap was $O(1)$. When the condition failed to hold, they showed that fluid-priority policies still had an optimality gap of $O(\sqrt{N})$.

In a more recent work, Gast et al. [86] presented a framework for analyzing policies for both finite- and infinite-horizon RMABPs. Most notably, they provided conditions for a policy to be asymptotically optimal with an exponential convergence rate and presented a so-called LP-index policy, with provably strong asymptotic optimality properties.

9. Multi-Action Bandits

In the standard RMABP introduced by Whittle [3], individual projects only allow two modes of operation, active and passive. However, a more general model was proposed much earlier, where individual projects were modeled as multi-action MDPs, which contained a passive action. The resulting model, in which only one project can be active at each time (i.e., operated with an action different from the passive one) is known in the literature as a *bandit superprocess* and its inception was credited to Nash [87] by Gittins [14].

Whittle [17] provided an optimality proof for the Gittins index policy that was extended to bandit superprocesses, subject to a condition on dominating policies. The proof used a key construct later termed the Whittle integral. Brown and Smith [88] showed that this integral gives an upper bound on the value of a bandit superprocess, which is tighter than that obtained through Lagrangian relaxation. They further showed how to efficiently compute the integral. Hadfield-Menell and Russell [89] also considered bandit superprocesses, providing a constructive definition of the Whittle integral and providing an alternate computation method.

The extension of the Whittle index to multi-action projects was first outlined by Weber [90], who illustrated it in a particular model and further outlined a means of computing the resulting index. In [91], the author formalized the index extension in a general setting and demonstrated its applicability to a multi-armed multi-mode restless bandit problem, concerning the optimal dynamic allocation of a shared resource to a set of projects that can be operated in multiple modes while respecting a peak resource consumption limit. Sufficient PCL-based conditions were proposed to ensure both the existence of the index and the validity of an adaptive-greedy algorithm for its computation. These indices and further examples of their applications were also considered in the posterior work of Glazebrook et al. [92].

The approach in [91] was extended in [93] to a real-state project setting and demonstrated in a model for optimal dynamic energy management in a wireless sensor network. The approach outlined in [91] was developed and established rigorously by the author in [94], which extended the sufficient indexability conditions in [49] for binary-action restless bandits to *multi-gear bandits*.

In recent years there has been an increased interest in RMABPs for multi-action projects, as they allow modeling of more realistic situations with projects that allow multiple operating modes. However, the heuristic policies that have been proposed are based on different approaches than the aforementioned indices. Zayas-Cabán et al. [81] considered a policy for multi-action RMABPs obtained from an LP relaxation for which they established asymptotic optimality. The priority policies proposed by Verloop [77] can also be applied to multi-action RMABPs. More recently, Killian et al. [95,96] argued the relevance and

importance of investigating multi-action RMABPs and presented powerful new methods for obtaining asymptotically optimal policies. Xiong et al. [97,98] also considered multi-action RMABPs in a learning setting and developed efficient heuristic policies.

10. Lagrangian Index and Fluid Relaxation Policies

Recent works have extended the Lagrangian relaxation approach used by Whittle [3] to consider tighter Lagrangian relaxations with several multipliers, one per time period (for finite-horizon problems), as well as relaxations based on different ideas, in particular, fluid relaxations. These relaxations yield new policies with promising performance gains.

Brown and Smith [84] considered a finite-horizon RMABP model with time-dependent rewards for dynamic item selection, including, for example, the problem of dynamic product assortment with demand learning by a retailer considered by Caro and Gallien [99]. They proposed index policies and bounds based on a Lagrangian relaxation of the standard formulation that uses one Lagrange multiplier per period, and, under certain conditions, established their asymptotic optimality.

Brown and Zhang [100] studied an RMABP model, where projects can consume different amounts of a shared resource, and there is exogenous information, modeled as a finite-state Markov chain, which can affect the shared resource limit, as well as each project's rewards, transitions, and resource consumption. They developed a Lagrangian relaxation and a "dynamic fluid relaxation" that provided upper bounds on the optimal value, as well as heuristic policies. The dynamic fluid relaxation bound and policy were shown to be asymptotically optimal, whereas those obtained from the Lagrangian relaxation were shown not to be asymptotically optimal.

Hao et al. [101] investigated a deadline scheduling problem with randomly arriving jobs that need to be processed before their deadlines expire, motivated by the problem of scheduling multiple electric vehicles in a charging station. They proposed a Lagrangian-based index policy, established its asymptotic optimality, and demonstrated through numerical experiments that the policy substantially outperforms Whittle's index policy.

11. Reinforcement Learning and Q-Learning Approaches

A difficulty that arises when addressing a real-world problem via a Markovian RMABP model is that its parameters (rewards and transition probabilities) are unknown, which has motivated research on this issue, among which reinforcement learning and, in particular, Q-learning (see Watkins and Dayan [102]) approaches stand out (see, e.g., Chapters 6 and 7 in Bertsekas [13] and the monograph by Powell [103]).

Fu et al. [104] considered the infinite-horizon average-cost RMABP and developed a tractable reinforcement learning algorithm based on parallel Q-learning recursions, which learns an approximation to the Whittle index. The approach was tested numerically and it was shown that its performance was close to that of Whittle's index policy in the model with all parameters known.

Wu et al. [105] addressed the RMABP using a deep learning state-aware value function approximation approach, where the joint value function was approximated by a linear combination of individual project value functions.

Li et al. [106] investigated a particular application of the RMABP to resource scheduling in autonomous driving. They considered a Whittle index policy solution approach and approximated the index through a deep reinforcement learning method.

Biswas et al. [107] considered RMABP models in public health settings and proposed Whittle-index-based Q-learning schemes that were shown to converge to the performance of the Whittle index policy with all parameters known.

Avrachenkov and Borkar [108] presented a reinforcement learning method for RMABPs under the average-reward criterion and leveraged the properties of Whittle's index policy to reduce the search space of Q-learning and improve computational efficiency. They provided a convergence analysis and reported on numerical experiments that supported their findings.

Killian et al. [96] developed learning algorithms for multi-action RMABPs and combined Lagrangian relaxation and Q-learning. They showed that under certain conditions, the proposed learning scheme converged to the asymptotically optimal multi-action RMAB policy. They further proposed another scheme that attains the asymptotic optimality properties of a Lagrange policy for multi-action RMABs through Q-learning.

Nakhleh et al. [109] proposed a neural Whittle index network that can learn the Whittle indices for a given model by exploiting their properties. This motivated the use of deep reinforcement learning for training the neural network. The approach was demonstrated in computational experiments.

Nakhleh and Hou [110] considered problems with optimal threshold policies to learn the optimal threshold. They developed an online policy that was shown to outperform other reinforcement learning algorithms by exploiting special structure. They further applied the results to efficiently learn the Whittle index.

12. Regret-Based Online Learning

In a significant area of research that considers finite-horizon bandit models with unknown parameters from an online learning perspective, the performance of the proposed policies is assessed using a measure of *regret*, which compares the performance of a policy against that of an oracle, the offline optimum, which implements the optimal policy based on known parameter values. The pioneering work in this vein was the paper by Robbins [7], who showed the existence of an asymptotically optimal policy, namely having a regret that grows sublinearly in the horizon T for a two-armed Bernoulli bandit model with unknown parameters. In the MABP setting, Lai and Robbins [111] proved a landmark result, establishing a logarithmic lower bound on the regret for projects with i.i.d. rewards, as well as a policy asymptotically attaining it. Anantharam et al. [112] extended the results to MABPs with project transition probabilities parameterized by a scalar.

More recent work has sought to extend these results to online learning in RMABP models with unknown parameters (see, e.g., Chapters 4 and 5 in the monograph by Zhao [113]). As pointed out by Zhao, a key issue hindering regret analysis in such a setting is that the offline optimum when all parameters are known is typically unavailable, which is typically handled by considering instead a proxy oracle that corresponds to a particular policy, leading to the notion of *weak regret*.

Major research efforts have been devoted to devising policies for these RMABPs, which are often not index-based, along with corresponding regret analyses. It has been shown that, under suitable assumptions, some policies attain a logarithmic regret for particular RMABP models, as demonstrated in the work of Filippi et al. [114] and Tekin and Liu [115]. A more general regret bound of order $O(\sqrt{T})$ was established by Ortner et al. [116]. A different but related problem was studied by Garivier and Moulines [117], where project reward distributions were shown to abruptly change at unknown times.

Gupta et al. [118] considered online RMABPs where unknown reward probabilities are assumed to drift over time and proposed a dynamic extension of Thompson sampling, which was shown to outperform alternative methods.

Dai et al. [119] addressed the online RMABP using an approach based on the assumption that the optimal policy belongs to a finite set of policies. They proposed a corresponding learning policy and illustrated it in an opportunistic spectrum access model. They established that the proposed policy attains near-logarithmic regret. Liu et al. [120] proposed a different policy with logarithmic-order regret.

Modi et al. [121] focused on a particular online model of opportunistic spectrum access for which they provided an online learning policy with logarithmic-order regret.

Grünwälder and Khaleghi [122] considered a more general setting where project dynamics have long-range dependence, including Markov chain dynamics as a special case. They proposed policies for this setting and provided corresponding regret analyses.

A policy with logarithmic-order regret was also obtained by Agrawal and Asawa [123] in an opportunistic spectrum access model. Gafni and Cohen [124] considered a different

class of RMAB problems motivated by communication networks and financial investment applications. They proposed a policy for which they established a logarithmic-order regret, as well as a finite-sample regret bound.

Xu et al. [125] considered an online risk-averse RMABP model. They proposed an index policy for the problem and showed that it attains a regret of order $O(\log T/T)$.

Gafni et al. [126] provided an extension to the online RMABP that incorporated an exogenous global Markov process governing the rewards distribution of each project. They proposed a policy that achieved a logarithmic-order regret over time.

Gafni and Cohen [127] considered a dynamic multi-channel access model for wireless networks, in which each channel has a different rate for each user. They proposed a policy for this model that attains a logarithmic-order regret.

Xiong et al. [97] considered a finite-horizon online multi-action RMABP model. They proposed an index policy that is defined for non-indexable projects and showed that it achieves a sublinear regret with low computational complexity. The results were extended by Xiong et al. [98] to the average-reward criterion.

13. Applications: MDP Models

In Ref. [3], Whittle illustrated the RMABP by sketching several potential applications, including the allocation of alternative medical treatments to patients for a condition caused by a continuously mutating virus, tracking surveillance of enemy submarines by aircraft, and dynamic activation and deactivation of a pool of workers subject to tiring and recovery. In Chapter 14.6 of his monograph [46], he further outlined a machine maintenance model and derived its Whittle index. Numerous researchers have since demonstrated that a wide range of problems encountered in real-world applications can be modeled as RMABPs or variants thereof. This allows for the use of general solution methods developed for this framework. Below is a selection of such applications, focusing on RMABP models that can be formulated as MDPs with known model parameters.

13.1. Variants of the MABP

Several variants of the classic MABP can be reformulated as RMABPs. Consider, e.g., the MABP with switching penalties (costs or delays) considered by Banks and Sundaram [128], who showed that index policies are generally suboptimal in such a case, and Asawa and Teneketzis [129], who proposed a heuristic index policy. In [130–132], the author demonstrated that the Asawa and Teneketzis index is the Whittle index of the problem in its restless reformulation and developed Whittle index algorithms for its efficient computation exploiting special structure. The results were extended to restless bandits with switching penalties, as outlined in [133]. Different policies for bandits with switching penalties were considered by Le Ny and Feron [134], Caro and Gallien [99], and Arlotto et al. [135].

In [136,137], the author considered the finite-horizon MABP and its extensions and showed that such problems can be reformulated as infinite-horizon RMABP models. This yielded a new efficient algorithm for computing a classic finite-horizon priority index by deploying a Whittle index algorithm in such a setting. Alternative finite-horizon index policies for the MABP were considered by Caro and Gallien [99].

Dayanik et al. [138] considered a MABP where projects are not always available for selection. They showed that index policies are not optimal for this problem variant. They reformulated it as an RMABP and derived and analyzed its Whittle index policy.

Caro and Yoo [45] proposed a MABP model with random response delays. They showed that the resulting RMABP reformulation is indexable and computed the Whittle index for the special Beta-Bernoulli Bayesian learning model. Computational experiments showed that the index policy achieved near-optimal performance.

13.2. Queueing Models

Veatch and Wein [47] formulated the problem of optimal scheduling in a multi-class M/M/1 make-to-stock queue as an RMABP model. They showed that the lost sales case is indexable and evaluated and tested the resulting Whittle index policy. They further argued that the backorder case is non-indexable, which was also pointed out by Whittle in ([46] Chapter 14.7). In [51], the author considered a model extension to a multi-class make-to-order/make-to-stock M/G/1 queue with possibly nonlinear holding costs in the backorder case and showed that, in contrast to the aforementioned non-indexability results, the resulting model was, in fact, indexable under an extended notion of indexability, relative to a mixed average-bias criterion (see also Ansell et al. [139] for an application of Whittle's index policy to the optimal scheduling of a make-to-order multi-class M/M/1 queue with nonlinear costs).

In [50], the author illustrated the applicaiton of the PCL-based theoretical and algorithmic framework presented in that paper in the context of a general model for the optimal control of admission to a birth–death queue, which was solved by an extended Whittle index policy. The model was shown to be a building block for an RMABP formulation of a comprehensive model for the optimal dynamic control of admission and routing to parallel queues. This formulation included features such as job abandonments and finite buffers and was further developed in [140–143].

Raissi-Dehkordi and Baras [144] considered a model for optimal pull broadcast scheduling in information delivery systems, which they modeled through discrete-time bulk service queues. They derived and tested the Whittle index policy for an RMABP formulation of the model.

Dusonchet and Hongler [145] derived the Whittle index for a make-to-stock queue with backorders under the discounted cost criterion. They further outlined a theory of Whittle indexability for continuous time and state restless bandits and demonstrated their approach using several random dynamic models, including diffusion processes.

Goyal et al. [146] considered a discrete-time queueing model for the optimal scheduling of multimedia transmissions over a polled multiaccess fading channel, and investigated the Whittle index policy for the resulting RMABP formulation.

In Ref. [147], the author formulated the problem of the optimal scheduling of a multi-class queue with finite buffers as an RMABP, established its indexability relative to a bias optimality criterion, and developed the corresponding Whittle index policy, which yielded nontrivial insights.

Cao and Nyberg [148] proposed a Markovian model for the optimal dynamic admission control of multi-class traffic to a finite shared buffer. They established the model's indexability, evaluated the Whittle index, and numerically tested the resulting index policy.

Borkar and Pattathil [149] investigated the so-called egalitarian processor sharing queueing model, which they reformulated as an RMABP and established its indexability. They showed how to compute the index and demonstrated its near-optimal performance through numerical experiments.

13.3. Web Crawling

O'Meara and Patel [150] considered the problem of optimal scheduling of a web robot to construct and maintain topic-specific web indexes. They realized that this problem could be modeled as an RMABP and developed a reinforcement learning algorithm for its approximate solution via an index policy similar to Whittle's.

In Ref. [151], the author proposed a Markovian model for the optimal dynamic scheduling of page refreshes in a local repository of copies of randomly changing remote web pages. The model was reformulated as an RMABP and Whittle's index policy was derived and tested. Avrachenkov and Borkar [152] considered a model for the optimal scheduling of a web crawler to retrieve ephemeral content from various sites. They showed that this model fits into the RMABP framework and developed a Whittle index policy, which they tested numerically.

13.4. Public Health Interventions

Deo et al. [153] developed a model for optimal community-based healthcare delivery for a chronic disease, which was formulated as a variant of the finite-horizon RMABP. They designed a myopic index heuristic policy and tested its performance using real data, demonstrating significant performance gains over the benchmark policy.

Ayer et al. [154] proposed an RMABP model to support the prioritization of hepatitis C treatment decisions in U.S. prisons. They established indexability and derived the Whittle index in closed form, deriving insights from it. They further proposed an adjusted closed-form index policy that was designed to overcome the limitations of Whittle's index policy. The model was validated using real-world data against benchmark policies.

Mate et al. [155,156] discussed the use of the RMABP framework in public health and reported on a field study aimed at assisting local health delivery agents to improve maternal and child health (see also the related work of Biswas et al. [107]).

13.5. Communication Networks

Restless bandits have been widely deployed as a modeling framework in communication networks, which is currently one of its major application areas. In this type of setting, Wei et al. [157] considered the problem of optimal relay selection in wireless cooperative networks, incorporating finite-state Markov channels, adaptive modulation and coding, and residual relay energy. They showed that this problem can be formulated as an indexable RMABP. Simulation results demonstrated the effectiveness of the Whittle index rule.

Wei and Neely [158] proposed a model of power-aware throughput maximization in a multi-user file-downloading system. The model was formulated as a variant of the RMABP. An index policy that was different from Whittle's and based on a Lyapunov indexing approach was proposed and tested.

Aalto et al. [159] investigated optimal opportunistic scheduling for downlink data traffic in a wireless cell with time-varying channels to minimize flow-level holding costs. They developed a size-aware index policy by deploying the Whittle index in a novel way. A numerical study demonstrated the improved performance achieved by the proposed policy.

Borkar et al. [160] proposed a multi-user energy-efficient scheduling model in which each user has a separate queue and a cost is incurred for holding packets in each queue. Packets are transmitted through a shared channel with time-varying quality that can differ across users. Additionally, the cost incurred, i.e., energy consumed, for packet transmissions is a function of the channel quality. Indexability was proven for the average-cost criterion, Whittle's index was evaluated, and the resulting policy was tested. The Whittle index policy was shown to outperform previously considered policies such as max-weight scheduling and weighted fair scheduling.

Sun et al. [161] considered a model for heterogeneous cellular networks with a macro base station and multiple small base stations. User equipment cell association was investigated to maximize the long-run average system throughput. The model was formulated as an RMABP for which index policies were derived and tested.

Aalto et al. [162] investigated the opportunistic scheduling of downlink data traffic with partial channel information to minimize flow-level holding costs. The paper extended earlier work and, in particular, showed the indexability of the flow-level opportunistic scheduling problem with partial channel information in part of the parameter space. The authors derived a formula for the Whittle index, established optimality of threshold policies, and tested numerically Whittle's index policy by comparing it against alternative policies.

Wang et al. [163] considered a scheduling problem where a server opportunistically serves multiple user classes under time-varying multiple-state Markov channels with the goal of minimizing the average waiting cost. They reformulated the problem as an RMABP but noted that indexability was still open in this setting. They presented sufficient conditions on a channel-state transition matrix implying indexability and obtained

the Whittle index in closed form. For the general case, they proposed an approximate Whittle index.

Sun et al. [164] addressed the Age of Information (AoI) minimization scheduling problem over a star-topology wireless network, which was formulated as an RMABP. Indexability was shown and the Whittle index was obtained in closed form. The index was extended to incorporate more realistic features. Numerical studies demonstrated the effectiveness of the proposed index policies.

Chen et al. [165] proposed to use the uncertainty of information, as measured by Shannon's entropy, as an information freshness metric. They considered a model where a central monitor observes N binary Markov chains through $M < N$ communication channels and developed scheduling policies to minimize the long-run average uncertainty of information. The problem was cast as an RMAB and an index policy was developed and tested achieving excellent empirical performance.

Singh et al. [166] considered the problem of determining which base station a user should associate with in a dense millimeter-wave network. The objective was to design an association policy to minimize the weighted average sojourn time for users in the system. The problem was formulated as an RMABP, which was shown to be indexable. Whittle's index policy was derived and tested in a simulation study, in which it outperformed alternative user association policies.

13.6. Miscellaneous Applications

Given the huge variety of applications in which the RMAB framework has been deployed, this section discusses several papers that do not neatly fit into the other categories discussed previously.

Caro and Gallien [99] considered a model for the optimal dynamic product assortment of a retailer to maximize the overall profit for a selling season. The problem was formulated as a finite-horizon MABP with Bayesian learning in which multiple projects can be selected at each time. A closed-form index policy, which approximates Whittle's index policy, was derived, analyzed, and tested.

Huberman and Wu [167] considered a model for the automatic generation of a ranking of information sources to be presented to users with limited attention. The objective was to maximize the total expected utility. The problem was formulated as an RMABP with dual-speed projects, which were shown to be PCL-indexable by Glazebrook et al. [168], and hence their Whittle index can be computed using the adaptive-greedy algorithm in [49].

Kumar and Saranga [169] developed near-optimal obsolescence mitigation policies based on an RMABP model and Whittle's index policy.

Temple and Frazzoli [170] considered a well-studied online search problem known as the Cow Path Problem (CPP), which is typically addressed via competitive analysis. They instead considered an MDP formulation for which they proved that the relaxed RMABP version was indexable, which allowed them to derive a Whittle index policy.

He et al. [171] considered a model for the opportunistic scheduling of low-priority jobs onto under-utilized cloud resources left by high-priority jobs. Assuming that the availability of servers to low-priority jobs can be modeled as on/off Markov chains, they formulated the problem as an RMAB. They established indexability and derived closed-form Whittle index formulae. A numerical study using real data center traces demonstrated the effectiveness of the policy compared with alternative rules.

Taylor and Mathieu [172] proposed a dynamic demand response model formulated as an RMABP. Whittle index policies were derived and used to rank loads. Numerical experiments showed that the resulting policy significantly outperforms the naïve greedy policy.

Sun and Ma [173] considered a crowd-sensing model in mobile social networks (MSNs) to maximize social welfare under a coverage constraint. The model was formulated as an RMABP, its indexability was proven, and its Whittle index was used to design novel incentive schemes, which were shown to outperform previously proposed policies.

Lin et al. [174] proposed a model for optimal forward-looking experiential learning problems, which was formulated as an RMABP. Under certain assumptions, indexability was established and the Whittle index and other index policies were derived and tested through numerical experiments with real data, which demonstrated its near-optimal utility and showed that they outperformed alternative rules.

Guo et al. [175] developed scheduling policies for networked cyber-physical systems that satisfy inter-delivery time requirements of clients connected through wireless channels. The problem was formulated as an infinite-state risk-sensitive MDP model. Among the results, when channels were not relatively reliable, they presented a Whittle-like index policy for the model. Simulation results demonstrated the effectiveness of the proposed index policy.

Yu et al. [176] addressed the stochastic deadline scheduling problem, which they formulated as an RMABP model. This was shown to be indexable and the Whittle index was derived in closed form. The Whittle index policy was shown to be asymptotically optimal.

Qian et al. [58] considered a model for the optimal scheduling of patrol policies motivated by security domains with frequent interactions between defenders and attackers such as wildlife protection, which they formulated as an RMAB. They provided sufficient conditions for indexability and an algorithm to test it.

Avrachenkov et al. [177] investigated a model for Generalized Additive Increase Multiplicative Decrease (G-AIMD) dynamics for resource allocation under a fairness-based utility function. The model was formulated as an RMABP, indexability was established in special cases, and a means of computing the Whittle index was presented. The index policy was numerically tested and it was shown to achieve near-optimal performance.

Borkar et al. [178] considered an MDP model for optimal resource allocation in cloud computing based on dynamic pricing. The model was cast as an RMABP, indexability was proven, and an iterative scheme for computing the Whittle index was provided.

Motivated by the challenge of incorporating user feedback to tailor system operation for improved individual user satisfaction, Menner and Zeilinger [179] proposed a model to optimize the collection and processing of user feedback to maximize a user comfort measure. The model was formulated as an RMABP, which was shown to be PCL-indexable. The Whittle index for this model was computed using the adaptive-greedy algorithm in [49]. Furthermore, the authors considered a learning model where transition probabilities are unknown and developed and tested an approach that combined restless bandit indices with upper confidence bound algorithms.

Motivated by recommendation systems, Jhunjhunwala et al. [180] considered an RMABP where, in each time period, the reward for selecting a project depends on the time that has elapsed since the project was last selected. They characterized the optimal policy with respect to the long-run average-reward criterion.

Abbou and Makis [181] considered a maintenance planning model, where available repairmen are dynamically allocated to a set of unreliable production facilities with machines that incur losses due to degradation. The aim was to find a scheduling policy for maintenance interventions that minimizes production losses per period. The model was formulated as an RMABP and was shown to be indexable. The Whittle index policy was shown to achieve near-optimal performance and to outperform previously proposed policies.

Gerum et al. [182] developed an RMABP model for optimal inspection and maintenance scheduling policies in railway systems. Indexability was demonstrated and Whittle indices were obtained, providing a novel index policy for such a system, which was shown to be highly effective in a data-driven setting.

Li et al. [183] considered an RMABP model for scheduling the allocation of limited resources to a large number of jobs such as medical treatments with random lifetimes and service times after a mass-casualty event, where jobs were initially subject to triage and classified. Whittle indices were derived, and another solution approach was considered

through a nonstandard Lagrangian relaxation. Numerical experiments demonstrated that the second approach achieved better performance than the first one.

Fu and Moran [79] studied a job-assignment model in a large-scale server farm system with geographically distributed heterogeneous servers. The goal was to maximize the system's energy efficiency through dynamic load control on the networked servers. A scalable job-assignment policy was presented. Drawing on the asymptotic optimality analysis of Weber and Weiss [72], it was shown that the proposed policy quickly (exponentially) approaches asymptotic optimality, which was verified through numerical experiments. Fu et al. [184] extended the model and analysis to a setting with varying requests and limited-capacity resources that are shared by requests.

Dahiya et al. [185] considered the dynamic allocation of human operators in a system with semi-autonomous robots that are required to perform independent task sequences, but are subject to the risk of getting stuck or failing. A human operator can assist the robot. The model was formulated as an RMABP and its indexability was proven under certain conditions. A simulation study demonstrated the near optimality of Whittle's index policy.

Motivated by mobile intervention problems, Ou et al. [186] utilized RMABs with network effects that were not amenable to computational solution using standard methods. They proposed a new solution approach for the networked RMABs, provided sufficient conditions for the optimality of their approach, and demonstrated its strong empirical performance in real-world scenarios.

14. Applications: POMDP Models

A growing area of research considers RMABPs with continuous-state projects that arise from POMDP models (see, e.g., the monograph by Krishnamurthy [187]), particularly in multi-target tracking and sensor scheduling applications (see also the monograph of Wang and Chen [68]).

La Scala and Moran [188] considered an RMABP model for multi-target tracking with project states following scalar Kalman filter dynamics in a POMDP framework. Le Ny et al. [34] and Liu and Zhao [32] proposed equivalent POMDP RMABP models that arise in different applications, the scheduling of unmanned aerial vehicles in the former paper and multi-channel opportunistic access in the latter, with imperfectly observed Gilbert–Elliott (two-state Markov chain) channels. Both papers established indexability of the individual projects, evaluated Whittle's index policy, and tested it numerically, reporting excellent performance.

Liu and Zhao [33] considered a class of RMABPs that arise in dynamic multi-channel access applications. They proved the indexability and derived the Whittle index in closed form for the discounted- and average-reward criteria. The Whittle index policy was shown to be optimal in the case of identical projects under certain conditions. A numerical study demonstrated the near-optimal performance of Whittle's index policy and the related Lagrangian performance bound.

Le Ny et al. [189] considered the problem of scheduling observations using a number of sensors on a higher number of targets whose states obey continuous-time Kalman filter dynamics. Whittle's approach was deployed in this RMABP, indexability was argued, and in the scalar case with identical sensors, the Whittle index was derived in closed form.

Liu et al. [43] investigated a class of RMABPs where the active action resets the system's evolution. Indexability was established and the Whittle index was derived in closed form. The results were applied to opportunistic spectrum access and supervisory control systems. These results were extended by Akbarzadehand and Mahajan [190].

Gan and Chen [191] introduced a new sensing policy for dynamic multi-channel access with primary and secondary users based on an RMABP formulation. Whittle's index policy was derived and applied and various results on the related Lagrangian relaxation were obtained.

He et al. [192] considered a POMDP model for topology tracking in a dynamic network with limited monitoring resources where links are modeled as independent on-off Markov

chains. The goal was to maximize the overall tracking accuracy of link states. A version of the model based on link sampling was formulated as an RMABP and its indexability was proven under certain conditions, allowing the deployment of Whittle's index policy for link sampling. Numerical results attested to the strong performance of the proposed approaches.

Meshram et al. [193] considered a restless bandit model for a recommendation system. They characterized the optimal discounted policy for the single-project case with two underlying states and argued that in a certain special case the optimal policy is of threshold type, which they view as a relevant step towards establishing indexability in future work.

Ouyang et al. [194] considered a model of opportunistic multi-user scheduling in downlink networks with Markovian outage channels, which was cast as an RMABP. They showed that the model is indexable and obtained the Whittle index policy in closed form. Numerical experiments demonstrated the policy's near-optimal performance.

Taboada et al. [195] addressed the problem of scheduling traffic flows in wireless downlink systems under limited channel-state feedback to minimize the mean flow delay. The problem was cast into the RMABP framework with POMDP projects, indexability was argued, and Whittle's index was evaluated and tested in numerical experiments.

Meshram et al. [196] considered an RMABP where each project can be in one of two states, which are not observable and need to be inferred from the current belief state and a possible binary signal. Single-project subproblems were shown to admit an approximate threshold-type optimal policy in certain cases in which they satisfy an approximate indexability property. In cases where the optimality of threshold policies can be established, indexability was argued and the Whittle index was calculated.

Elmaghraby et al. [197] considered the problem of dynamic channel allocation for femtocells sharing the use of a regular macrocell spectrum. The channel state is not observable and macrocell user feedback is utilized. The problem was cast as an RMABP with POMDP projects and an approximate Whittle index policy was derived. A numerical study demonstrated the effectiveness of the proposed policy compared to the myopic rule.

Mehta et al. [198] proposed an RMABP with constrained availability of projects and Markovian state evolution, with the true states being hidden. The optimality of a threshold policy was argued, and based on this, indexability was shown. A formula for the Whittle index was derived for the rested case and an index algorithm was provided for the restless case.

Kaza et al. [199] considered a class of RMABPs with hidden states that allow cumulative feedback. They showed that individual project subproblems, which they called lazy restless bandits (LRBs), have optimal policies of threshold type. Indexability was argued and the Whittle index was derived in closed form for two sets of special cases. An index-computing algorithm was provided and an extensive numerical study was reported.

Yang and Luo [200] considered a massive multiple-input multiple-output (MIMO) system with a smaller number of available orthogonal pilot sequences than users. The resulting pilot allocation problem was modeled as a POMDP with Gauss–Markov fading channels. The problem was cast as an RMABP and the Whittle index was approximately derived. Numerical results were reported that demonstrated a strong performance for the index policy.

Hsu et al. [201] considered a wireless broadcast network in which a base station updates users about random information arrivals subject to a transmission capacity constraint. The problem was cast into the RMABP framework with POMDP projects. Structural results on optimal policies were identified, allowing the deployment of Whittle's indexability and index policy. An online version was further considered. The results were validated in a numerical study.

Wang et al. [202] studied dynamic channel allocation for the estimation of remote states in multi-agent systems. Whittle's approach was deployed in an RMABP formulation of the problem, and its strong performance was demonstrated through numerical experiments.

Chen and Ephremides [203] investigated a discrete-time model in which a base station simultaneously updates multiple users. The goal was to design a scheduling policy that minimizes an Age of Incorrect Information (AoII) metric for imperfect channel-state information. Whittle's index policy was derived under a simple condition. To avoid dealing with indexability, an alternative index policy was presented. A numerical study tested the performance of the policies considered.

Kang and Joo [204] considered a model for minimizing information mismatch under limited communication capabilities for a system in which a base station collects time-varying state information from multiple sources and makes decisions based on the collected information. An RMABP formulation with POMDP projects was shown to be indexable and its Whittle index was obtained in closed form in certain scenarios.

Motivated by public health interventions, Li and Varakantham [205] incorporated fairness constraints into an RMABP model. A modified Whittle index was derived and, for the case where transition probabilities are unknown, a model-free learning method was provided. The results were validated in a numerical study.

Tong et al. [206] considered age-of-information minimization in Internet-of-Things networks with correlated sources. The problem was formulated as a correlated RMAB. A generalized Whittle index and a generalized partial Whittle index were derived for the identical and non-identical channel settings, respectively, and corresponding index policies were proposed. A numerical study demonstrated that the policies are nearly-optimal, as they approach given Lagrangian-based bounds, and outperform state-of-the-art alternative policies.

15. Conclusions

The work reviewed herein demonstrates without a doubt that after more than three decades since their inception, restless bandits continue to be a vibrant research area full of interesting problems to address in theoretical, algorithmic, and application realms. It appears that the main avenues for further research should include the following: (1) deepening the understanding of indexability and developing easier-to-apply sufficient indexability conditions, both for models with binary-action and multi-action projects; (2) developing a full understanding of methods for designing asymptotically optimal index policies as the size of the model scales; (3) devising methods suitable for very large-scale models arising from applications; (4) developing and refining methods for designing and analyzing effective policies for models with unknown parameters that need to be learned online; and (5) further extending the scope and modeling power of the RMABP framework.

Funding: This work was funded in part by the Spanish State Research Agency (*Agencia Estatal de Investigación*, AEI) under grant PID2019-109196GB-I00/AEI/10.13039/501100011033 and by the *Comunidad de Madrid* in the context of a multi-year agreement with Carlos III University of Madrid within the activity "*Excelencia para el Profesorado Universitario*" in the framework of the V Regional Plan of Scientific Research and Technological Innovation 2016–2020.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Conflicts of Interest: The author declares no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

DP	Dynamic programming
LP	Linear programming
MDP	Markov decision process
POMDP	Partially observable Markov decision process
MABP	Multi-armed bandit problem
RMABP	Restless multi-armed bandit problem

References

1. Puterman, M.L. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*; Wiley: New York, NY, USA, 1994.
2. Gittins, J.C.; Jones, D.M. A dynamic allocation index for the sequential design of experiments. In *Progress in Statistics, Proceedings of the European Meeting of Statisticians, Budapest, Hungary, 31 August–5 September 1972*; Colloquia Mathematica Societatis János Bolyai; Gani, J.; Sarkadi, K.; Vincze, I., Eds.; North-Holland: Amsterdam, The Netherlands, 1974; Volume 9, pp. 241–266.
3. Whittle, P. Restless bandits: Activity allocation in a changing world. *J. Appl. Probab.* **1988**, *25A*, 287–298. [[CrossRef](#)]
4. Thompson, W.R. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* **1933**, *25*, 275–294. [[CrossRef](#)]
5. Thompson, W.R. On the theory of apportionment. *Am. J. Math.* **1935**, *57*, 450–456. [[CrossRef](#)]
6. Wald, A. *Sequential Analysis*; Wiley: New York, NY, USA, 1947.
7. Robbins, H. Some aspects of the sequential design of experiments. *Bull. Amer. Math. Soc.* **1952**, *58*, 527–535. [[CrossRef](#)]
8. Bradt, R.N.; Johnson, S.M.; Karlin, S. On sequential designs for maximizing the sum of n observations. *Ann. Math. Statist.* **1956**, *27*, 1060–1074. [[CrossRef](#)]
9. Bellman, R. A problem in the sequential design of experiments. *Sankhyā Indian J. Stat.* **1956**, *16*, 221–229.
10. Bellman, R. *Dynamic Programming*; Princeton University Press: Princeton, NJ, USA, 1957.
11. Howard, R.A. *Dynamic Programming and Markov Processes*; Wiley: New York, NY, USA, 1960.
12. Bertsekas, D.P. *Dynamic Programming and Optimal Control*, 4th ed.; Athena Scientific: Belmont, MA, USA, 2017; Volume I.
13. Bertsekas, D.P. *Dynamic Programming and Optimal Control—Approximate Dynamic Programming*, 4th ed.; Athena Scientific: Nashua, NH, USA, 2012; Volume II.
14. Gittins, J.C. Bandit processes and dynamic allocation indices (with discussion). *J. Roy. Statist. Soc. Ser. B* **1979**, *41*, 148–177.
15. Gittins, J.C. *Multi-Armed Bandit Allocation Indices*; Wiley: Chichester, UK, 1989.
16. Gittins, J.C.; Glazebrook, K.; Weber, R. *Multi-Armed Bandit Allocation Indices*, 2nd ed.; Wiley: Chichester, UK, 2011.
17. Whittle, P. Multi-armed bandits and the Gittins index. *J. Roy. Statist. Soc. Ser. B* **1980**, *42*, 143–149. [[CrossRef](#)]
18. Varaiya, P.P.; Walrand, J.C.; Buyukkoc, C. Extensions of the multiarmed bandit problem: The discounted case. *IEEE Trans. Automat. Control* **1985**, *30*, 426–439. [[CrossRef](#)]
19. Weber, R. On the Gittins index for multiarmed bandits. *Ann. Appl. Probab.* **1992**, *2*, 1024–1033. [[CrossRef](#)]
20. Tsitsiklis, J.N. A short proof of the Gittins index theorem. *Ann. Appl. Probab.* **1994**, *4*, 194–199. [[CrossRef](#)]
21. Bertsimas, D.; Niño-Mora, J. Conservation laws, extended polymatroids and multiarmed bandit problems; a polyhedral approach to indexable systems. *Math. Oper. Res.* **1996**, *21*, 257–306. [[CrossRef](#)]
22. Whittle, P. Arm-acquiring bandits. *Ann. Probab.* **1981**, *9*, 284–292. [[CrossRef](#)]
23. Dumitriu, I.; Tetali, P.; Winkler, P. On playing golf with two balls. *SIAM J. Discret. Math.* **2003**, *16*, 604–615. [[CrossRef](#)]
24. Bao, W.; Cai, X.; Wu, X. A general theory of multiarmed bandit processes with constrained arm switches. *SIAM J. Control Optim.* **2021**, *59*, 4666–4688. [[CrossRef](#)]
25. Klimov, G.P. Time-sharing service systems. I. *Theory Probab. Appl.* **1974**, *19*, 532–551. [[CrossRef](#)]
26. Meilijson, I.; Weiss, G. Multiple feedback at a single-server station. *Stoch. Process. Appl.* **1977**, *5*, 195–205. [[CrossRef](#)]
27. Weiss, G. Branching bandit processes. *Probab. Eng. Inform. Sci.* **1988**, *2*, 269–278. [[CrossRef](#)]
28. Niño-Mora, J. Klimov’s model. In *Wiley Encyclopedia of Operations Research and Management Science*; Cochran, J.J., Cox, L.A., Jr., Keskinocak, P., Kharoufeh, J.P., Smith, J.C., Eds.; Wiley: New York, NY, USA, 2011. [[CrossRef](#)]
29. O’Flaherty, B. Some results on two-armed bandits when both projects vary. *J. Appl. Probab.* **1989**, *26*, 655–658. [[CrossRef](#)]
30. Papadimitriou, C.H.; Tsitsiklis, J.N. The complexity of optimal queuing network control. *Math. Oper. Res.* **1999**, *24*, 293–305. [[CrossRef](#)]
31. Guha, S.; Munagala, K.; Shi, P. Approximation algorithms for restless bandit problems. *J. ACM* **2010**, *58*, 3. [[CrossRef](#)]
32. Liu, K.; Zhao, Q. A restless bandit formulation of opportunistic access: Indexability and index policy. In Proceedings of the 5th IEEE Annual Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks Workshops, San Francisco, CA, USA, 16–20 June 2008; pp. 1–5.
33. Liu, K.; Zhao, Q. Indexability of restless bandit problems and optimality of Whittle index for dynamic multichannel access. *IEEE Trans. Inform. Theory* **2010**, *56*, 5547–5567. [[CrossRef](#)]
34. Le Ny, J.; Dahleh, M.; Feron, E. Multi-UAV dynamic routing with partial observations using restless bandit allocation indices. In Proceedings of the American Control Conference (ACC), Seattle, WA, USA, 11–13 June 2008; pp. 4220–4225.
35. Wan, P.J.; Xu, X.H. Weighted restless bandit and its applications. In Proceedings of the 35th IEEE International Conference on Distributed Computing Systems (ICDCS), Columbus, OH, USA, 29 June–2 July 2015; pp. 507–516.
36. Xu, X.H.; Song, M. Approximation algorithms for wireless opportunistic spectrum scheduling in cognitive radio networks. In Proceedings of the 35th Annual IEEE International Conference on Computer Communications (INFOCOM), San Francisco, CA, USA, 10–14 April 2016; pp. 1–7.
37. Xu, X.H.; Wang, L.X. Efficient algorithm for multi-constrained opportunistic wireless scheduling. In Proceedings of the 16th IEEE International Conference on Mobility, Sensing and Networking (MSN), Tokyo, Japan, 17–19 December 2020; pp. 169–173.
38. Bertsimas, D.; Niño-Mora, J. Restless bandits, linear programming relaxations, and a primal-dual index heuristic. *Oper. Res.* **2000**, *48*, 80–90. [[CrossRef](#)]

39. Hawkins, J.T. A Lagrangian Decomposition Approach to Weakly Coupled Dynamic Optimization Problems and Its Applications. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA, USA, 2003.
40. Adelman, D.; Mersereau, A.J. Relaxations of weakly coupled stochastic dynamic programs. *Oper. Res.* **2008**, *56*, 712–727. [[CrossRef](#)]
41. Powell, W.B. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, 2nd ed.; Wiley: Hoboken, NJ, USA, 2011.
42. Brown, D.B.; Zhang, J.W. On the strength of relaxations of weakly coupled stochastic dynamic programs. *Oper. Res.* **2022**. [[CrossRef](#)]
43. Liu, K.; Weber, R.; Zhao, Q. Indexability and Whittle index for restless bandit problems involving reset processes. In Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC), Orlando, FL, USA, 12–15 December 2011; pp. 7690–7696.
44. Fryer, R.; Harms, P. Two-armed restless bandits with imperfect information: Stochastic control and indexability. *Math. Oper. Res.* **2018**, *43*, 399–427. [[CrossRef](#)]
45. Caro, F.; Yoo, O.S. Indexability of bandit problems with response delays. *Probab. Eng. Inform. Sci.* **2010**, *24*, 349–374. [[CrossRef](#)]
46. Whittle, P. *Optimal Control: Basics and Beyond*; Wiley: Chichester, UK, 1996.
47. Veatch, M.H.; Wein, L.M. Scheduling a multiclass make-to-stock queue: Index policies and hedging points. *Oper. Res.* **1996**, *44*, 634–647. [[CrossRef](#)]
48. Dance, C.R.; Silander, T. When are Kalman-filter restless bandits indexable? In Proceedings of the Conference on Neural Information Processing Systems (NIPS), Montreal, QB, Canada, 7–12 December 2015; Cortes, C., Lee, D.D., Sugiyama, M., Garnett, R., Eds.; MIT Press: Cambridge, MA, USA, 2015; pp. 1711–1719.
49. Niño-Mora, J. Restless bandits, partial conservation laws and indexability. *Adv. Appl. Probab.* **2001**, *33*, 76–98. [[CrossRef](#)]
50. Niño-Mora, J. Dynamic allocation indices for restless projects and queueing admission control: A polyhedral approach. *Math. Program.* **2002**, *93*, 361–413. [[CrossRef](#)]
51. Niño-Mora, J. Restless bandit marginal productivity indices, diminishing returns and optimal control of make-to-order/make-to-stock M/G/1 queues. *Math. Oper. Res.* **2006**, *31*, 50–84. [[CrossRef](#)]
52. Niño-Mora, J. Dynamic priority allocation via restless bandit marginal productivity indices (with discussion). *TOP* **2007**, *15*, 161–198. [[CrossRef](#)]
53. Niño-Mora, J. A verification theorem for threshold-indexability of real-state discounted restless bandits. *Math. Oper. Res.* **2020**, *45*, 465–496. [[CrossRef](#)]
54. Dance, C.R.; Silander, T. Optimal policies for observing time series and related restless bandit problems. *J. Mach. Learn. Res.* **2019**, *20*, 35.
55. Niño-Mora, J. Characterization and computation of restless bandit marginal productivity indices. In Proceedings of the 1st International ICST Workshop on Tools for solving Structured Markov Chains (SMCTools), Nantes, France, 26 October 2007; Buchholz, P., Dayar, T., Eds.; ICST: Brussels, Belgium, 2007. ACM International Conference Proceeding Series. [[CrossRef](#)]
56. Niño-Mora, J. A fast-pivoting algorithm for Whittle’s restless bandit index. *Mathematics* **2020**, *8*, 2226. [[CrossRef](#)]
57. Niño-Mora, J. A $(2/3)n^3$ fast-pivoting algorithm for the Gittins index and optimal stopping of a Markov chain. *INFORMS J. Comput.* **2007**, *19*, 596–606. [[CrossRef](#)]
58. Qian, Y.; Zhang, C.; Krishnamachari, B.; Tambe, M. Restless poachers: Handling exploration-exploitation tradeoffs in security domains. In Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), Singapore, Singapore, 9–13 May 2016; pp. 123–131.
59. Akbarzadeh, N.; Mahajan, A. Conditions for indexability of restless bandits and an $\mathcal{O}(K^3)$ algorithm to compute Whittle index. *Adv. Appl. Probab.* **2022**, *54*, 1164–1192. [[CrossRef](#)]
60. Ehsan, N.; Liu, M. Server allocation with delayed state observation: Sufficient conditions for the optimality of an index policy. *IEEE Trans. Wirel. Comm.* **2009**, *8*, 1693–1705. [[CrossRef](#)]
61. Ahmad, S.H.A.; Liu, M.Y.; Javidi, T.; Zhao, Q. Optimality of myopic sensing in multichannel opportunistic access. *IEEE Trans. Inform. Theory* **2009**, *55*, 4040–4050. [[CrossRef](#)]
62. Liu, K.; Zhao, Q.; Krishnamachari, B. Dynamic multichannel access with imperfect channel state detection. *IEEE Trans. Signal Process.* **2010**, *58*, 2795–2808. [[CrossRef](#)]
63. Wang, K.H.; Liu, Q.; Chen, L. Optimality of greedy policy for a class of standard reward function of restless multi-armed bandit problem. *IET Signal Process.* **2012**, *6*, 584–593. [[CrossRef](#)]
64. Wang, K.H.; Chen, L. On optimality of myopic policy for restless multi-armed bandit problem: An axiomatic approach. *IEEE Trans. Signal Proc.* **2012**, *60*, 300–309. [[CrossRef](#)]
65. Wang, K.H.; Chen, L.; Liu, Q. On optimality of myopic policy for opportunistic access with nonidentical channels and imperfect sensing. *IEEE Trans. Veh. Tech.* **2014**, *63*, 2478–2483. [[CrossRef](#)]
66. Wang, K.H.; Chen, L.; Yu, J.H.; Zhang, D.Z. Optimality of myopic policy for multistate channel access. *IEEE Comm. Lett.* **2016**, *20*, 300–303. [[CrossRef](#)]
67. Wang, K.H.; Yu, J.H.; Chen, L.; Zhou, P.; Win, M.Z. Optimal myopic policy for restless bandit: A perspective of eigendecomposition. *IEEE J. Sel. Top. Signal Process.* **2022**, *16*, 420–433. [[CrossRef](#)]
68. Wang, K.H.; Chen, L. *Restless Multi-Armed Bandit in Opportunistic Scheduling*; Springer: Cham, Switzerland, 2021.

69. Ouyang, W.Z.; Teneketzis, D. On the optimality of myopic sensing in multi-state channels. *IEEE Trans. Inform. Theory* **2014**, *60*, 681–696. [[CrossRef](#)]
70. Blasco, P.; Gündüz, D. Multi-access communications with energy harvesting: A multi-armed bandit model and the optimality of the myopic policy. *IEEE J. Sel. Areas Commun.* **2015**, *33*, 585–597. [[CrossRef](#)]
71. Kadota, I.; Sinha, A.; Uysal-Biyikoglu, E.; Singh, R.; Modiano, E. Scheduling policies for minimizing age of information in broadcast wireless networks. *IEEE/ACM Trans. Netw.* **2018**, *26*, 2637–2650. [[CrossRef](#)]
72. Weber, R.R.; Weiss, G. On an index policy for restless bandits. *J. Appl. Probab.* **1990**, *27*, 637–648. [[CrossRef](#)]
73. Weber, R.R.; Weiss, G. Addendum to: “On an index policy for restless bandits”. *Adv. Appl. Probab.* **1991**, *23*, 429–430. [[CrossRef](#)]
74. Bagheri, S.; Scaglione, A. The restless multi-armed bandit formulation of the cognitive compressive sensing problem. *IEEE Trans. Signal Process.* **2015**, *63*, 1183–1198. [[CrossRef](#)]
75. Larrañaga, M.; Ayesta, U.; Verloop, I.M. Asymptotically optimal index policies for an abandonment queue with convex holding cost. *Queueing Syst.* **2015**, *81*, 99–169. [[CrossRef](#)]
76. Ouyang, W.Z.; Eryilmaz, A.; Shroff, N.B. Downlink scheduling over Markovian fading channels. *IEEE/ACM Trans. Netw.* **2016**, *24*, 1801–1812. [[CrossRef](#)]
77. Verloop, I.M. Asymptotically optimal priority policies for indexable and nonindexable restless bandits. *Ann. Appl. Probab.* **2016**, *26*, 1947–1995. [[CrossRef](#)]
78. Fu, J.; Moran, B.; Guo, J.; Wong, E.W.M.; Zukerman, M. Asymptotically optimal job assignment for energy-efficient processor-sharing server farms. *IEEE J. Sel. Areas Commun.* **2016**, *34*, 4008–4023. [[CrossRef](#)]
79. Fu, J.; Moran, B. Energy-efficient job-assignment policy with asymptotically guaranteed performance deviation. *IEEE/ACM Trans. Netw.* **2020**, *28*, 1325–1338. [[CrossRef](#)]
80. Hu, W.; Frazier, P.I. An asymptotically optimal index policy for finite-horizon restless bandits. *arXiv* **2017**, arXiv:1707.00205.
81. Zayas-Cabán, G.; Jasin, S.; Wang, G. An asymptotically optimal heuristic for general nonstationary finite-horizon restless multi-armed, multi-action bandits. *Adv. Appl. Probab.* **2019**, *51*, 745–772. [[CrossRef](#)]
82. Maatouk, A.; Kriouile, S.; Assaad, M.; Ephremides, A. On the optimality of the Whittle’s index policy for minimizing the age of information. *IEEE Trans. Wirel. Comm.* **2021**, *20*, 1263–12770. [[CrossRef](#)]
83. Kriouile, S.; Assaad, M.; Maatouk, A. On the global optimality of Whittle’s index policy for minimizing the age of information. *IEEE Trans. Inf. Theory* **2022**, *68*, 572–600. [[CrossRef](#)]
84. Brown, D.B.; Smith, J.E. Index policies and performance bounds for dynamic selection problems. *Manag. Sci.* **2020**, *66*, 3029–3050. [[CrossRef](#)]
85. Zhang, X.Y.; Frazier, P.I. Restless bandits with many arms: Beating the Central Limit Theorem. *arXiv* **2021**, arXiv:2107.11911.
86. Gast, N.; Gaujal, B.; Yan, Y. LP-based policies for restless bandits: Necessary and sufficient conditions for (exponentially fast) asymptotic optimality. *arXiv* **2022**, arXiv:2106.10067.
87. Nash, P. Optimal Allocation of Resources between Research Projects. Ph.D. Thesis, University of Cambridge, Cambridge, UK, 1973.
88. Brown, D.B.; Smith, J.E. Optimal sequential exploration: Bandits, clairvoyants, and wildcats. *Oper. Res.* **2013**, *61*, 644–665. [[CrossRef](#)]
89. Hadfield-Menell, D.; Russell, S. Multitasking: Efficient optimal planning for bandit superprocesses. In Proceedings of the 31st Conference on Uncertainty in Artificial Intelligence (UAI), Amsterdam, The Netherlands, 12–16 July 2015; Meila, M., Heskes, T., Eds.; AUAI Press: Corvallis, OR, USA, 2015; pp. 345–354.
90. Weber, R. Comments on: “Dynamic priority allocation via restless bandit marginal productivity indices” [TOP 15 (2007),161–198] by J. Niño-Mora. *TOP* **2007**, *15*, 211–216. [[CrossRef](#)]
91. Niño-Mora, J. An index policy for multiarmed multimode restless bandits. In Proceedings of the 3rd International Conference on Performance Evaluation Methodologies and Tools (ValueTools), Athens, Greece, 20–24 October 2008; Baras, J., Courcoubetis, C., Eds.; ICST: Brussels, Belgium, 2008. ACM International Conference Proceedings Series. [[CrossRef](#)]
92. Glazebrook, K.D.; Hodge, D.J.; Kirkbride, C. General notions of indexability for queueing control and asset management. *Ann. Appl. Probab.* **2011**, *21*, 876–907. [[CrossRef](#)]
93. Niño-Mora, J. Index-based dynamic energy management in a multimode sensor network. In Proceedings of the 6th International Conference on Network Games, Control and Optimization (NetGCoop), Avignon, France, 28–30 November 2012; pp. 92–95. Available online: <https://ieeexplore.ieee.org/document/6486131> (accessed on 1 January 2023).
94. Niño-Mora, J. Multi-gear bandits, partial conservation laws, and indexability. *Mathematics* **2022**, *10*, 2497. [[CrossRef](#)]
95. Killian, J.A.; Perrault, A.; Tambe, M. Beyond “to act or not to act”: Fast Lagrangian approaches to general multi-action restless bandits. In Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS), Online, 3–7 May 2021; Endriss, U., Nowé, A., Dignum, F., Lomuscio, A., Eds.; International Foundation for Autonomous Agents and Multiagent Systems: Richland, SC, USA, 2021; pp. 710–718.
96. Killian, J.A.; Biswas, A.; Shah, S.; Tambe, M. Q-learning Lagrange policies for multi-action restless bandits. In Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, Singapore, Singapore, 14–18 August 2021; pp. 871–881.
97. Xiong, G.J.; Li, J.; Singh, R. Reinforcement learning augmented asymptotically optimal index policy for finite-horizon restless bandits. In Proceedings of the 36th AAAI Conference on Artificial Intelligence, Seoul, Republic of Korea, 22 February–1 March 2022; pp. 8726–8734.

98. Xiong, G.J.; Wang, S.; Li, J. Learning infinite-horizon average-reward restless multi-action bandits via index awareness. In Proceedings of the 36th Conference on Neural Information Processing Systems (NeurIPS), New Orleans, LA, USA, 28 November–9 December 2022; Oh, A.H., Agarwal, A., Belgrave, D., Cho, K., Eds.; Neural Information Processing Systems (NIPS): La Jolla, CA, USA, 2022; Advances in Neural Information Processing Systems.
99. Caro, F.; Gallien, J. Dynamic assortment with demand learning for seasonal consumer goods. *Manag. Sci.* **2007**, *53*, 276–292. [[CrossRef](#)]
100. Brown, D.B.; Zhang, J.W. Dynamic programs with shared resources and signals: Dynamic fluid policies and asymptotic optimality. *Oper. Res.* **2022**, *70*, 3015–3033. [[CrossRef](#)]
101. Hao, L.L.; Xu, Y.J.; Tong, L. Asymptotically optimal Lagrangian priority policy for deadline scheduling with processing rate limits. *IEEE Trans. Automat. Control* **2022**, *67*, 236–250. [[CrossRef](#)]
102. Watkins, C.; Dayan, P. Q-learning. *Mach. Learn.* **1992**, *8*, 279–292. [[CrossRef](#)]
103. Powell, W.B. *Reinforcement Learning and Stochastic Optimization: A Unified Framework for Sequential Decisions*; Wiley: Hoboken, NJ, USA, 2022.
104. Fu, J.; Nazarathy, Y.; Moka, S.; Taylor, P.G. Towards Q-learning the Whittle index for restless bandits. In Proceedings of the Australian & New Zealand Control Conference (ANZCC), Auckland, New Zealand, 27–29 November 2019; pp. 249–254. [[CrossRef](#)]
105. Wu, S.; Zhao, J.; Tian, G.; Wang, J. State-aware value function approximation with attention mechanism for restless multi-armed bandits. In Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI), Online, 19–26 August 2021; Zhou, Z.H., Ed.; IJCAI: Cape Town, South Africa, 2021; pp. 458–464. [[CrossRef](#)]
106. Li, M.S.; Gao, J.; Zhao, L.; Shen, X.M. Adaptive computing scheduling for edge-assisted autonomous driving. *IEEE Trans. Veh. Tech.* **2021**, *70*, 5318–5331. [[CrossRef](#)]
107. Biswas, A.; Aggarwal, G.; Varakantham, P.; Tambe, M. Learn to intervene: An adaptive learning policy for restless bandits in application to preventive healthcare. In Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI), Online, 19–26 August 2021; Zhou, Z.H., Ed.; pp. 4039–4046. [[CrossRef](#)]
108. Avrachenkov, K.E.; Borkar, V.S. Whittle index based Q-learning for restless bandits with average reward. *Automatica* **2022**, *139*, 110186. [[CrossRef](#)]
109. Nakhleh, K.; Ganji, S.; Hsieh, P.C.; Hou, I.H.; Shakkottai, S. NeurWIN: Neural Whittle index network for restless bandits via deep RL. In Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS), Online, 6–14 December 2021; Advances in Neural Information Processing Systems; Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P.S., Wortman Vaughan, J., Eds.; Neural Information Processing Systems (NIPS): La Jolla, CA, USA, 2021; Volume 34.
110. Nakhleh, K.; Hou, I.H. DeepTOP: Deep threshold-optimal policy for MDPs and RMABs. In Proceedings of the 36th Conference on Neural Information Processing Systems (NeurIPS), New Orleans, LA, USA, 28 November–9 December 2022; Advances in Neural Information Processing Systems; Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., Oh, A., Eds.; Neural Information Processing Systems (NIPS): La Jolla, CA, USA, 2022; Volume 35.
111. Lai, T.L.; Robbins, H. Asymptotically efficient adaptive allocation rules. *Adv. Appl. Math.* **1985**, *6*, 4–22. [[CrossRef](#)]
112. Anantharam, V.; Varaiya, P.; Walrand, J. Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays. II. Markovian rewards. *IEEE Trans. Automat. Control* **1987**, *32*, 977–982. [[CrossRef](#)]
113. Zhao, Q. *Multi-Armed Bandits: Theory and Applications to Online Learning in Networks*; Morgan & Claypool: San Rafael, CA, USA, 2020.
114. Filippi, S.; Cappé, O.; Garivier, A. Optimally sensing a single channel without prior information: The tiling algorithm and regret bounds. *IEEE J. Sel. Top. Signal Process.* **2011**, *5*, 68–76. [[CrossRef](#)]
115. Tekin, C.; Liu, M.Y. Online learning of rested and restless bandits. *IEEE Trans. Inf. Theory* **2012**, *58*, 5588–5611. [[CrossRef](#)]
116. Ortner, R.; Ryabko, D.; Auer, P.; Munos, R. Regret bounds for restless Markov bandits. *Theoret. Comput. Sci.* **2014**, *558*, 62–76. [[CrossRef](#)]
117. Garivier, A.; Moulines, E. On upper-confidence bound policies for switching bandit problems. In Proceedings of the 22nd International Conference on Algorithmic Learning Theory (ALT), Espoo, Finland, 5–7 October 2011; Lecture Notes in Artificial Intelligence; Kivinen, J., Szepesvári, C., Ukkonen, E., Zeugmann, T., Eds.; Springer: Berlin, Germany, 2011; Volume 6925, pp. 174–188.
118. Gupta, N.; Granmo, O.C.; Agrawala, A. Thompson sampling for dynamic multi-armed bandits. In Proceedings of the 10th International Conference on Machine Learning and Applications and Workshops, Honolulu, HI, USA, 18–21 December 2011; pp. 484–489. [[CrossRef](#)]
119. Dai, W.H.R.; Gai, Y.; Krishnamachari, B.; Zhao, Q. The non-Bayesian restless multi-armed bandit: A case of near-logarithmic regret. In Proceedings of the 36th IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Prague, Czech Republic, 22–27 May 2011; pp. 2940–2943.
120. Liu, H.Y.; Liu, K.Q.; Zhao, Q. Learning in a changing world: Restless multiarmed bandit with unknown dynamics. *IEEE Trans. Inf. Theory* **2013**, *59*, 1902–1916. [[CrossRef](#)]
121. Modi, N.; Mary, P.; Moy, C. QoS driven channel selection algorithm for cognitive radio network: Multi-user multi-armed bandit approach. *IEEE Trans. Cogn. Commun. Netw.* **2017**, *3*, 49–66. [[CrossRef](#)]
122. Grünewälder, S.; Khaleghi, A. Approximations of the restless bandit problem. *J. Mach. Learn. Res.* **2019**, *20*, 14.

123. Agrawal, H.; Asawa, K. Decentralized learning for opportunistic spectrum access: Multiuser restless multiarmed bandit formulation. *IEEE Syst. J.* **2020**, *14*, 2485–2496. [[CrossRef](#)]
124. Gafni, T.; Cohen, K. Learning in restless multiarmed bandits via adaptive arm sequencing rules. *IEEE Trans. Automat. Control* **2021**, *66*, 5029–5036. [[CrossRef](#)]
125. Xu, J.Y.; Chen, L.J.; Tang, O. An online algorithm for the risk-aware restless bandit. *Eur. J. Oper. Res.* **2021**, *290*, 622–639. [[CrossRef](#)]
126. Gafni, T.; Yemini, M.; Cohen, K. Learning in restless bandits under exogenous global Markov process. *IEEE Trans. Signal Process.* **2022**, *70*, 5679–5693. [[CrossRef](#)]
127. Gafni, T.; Cohen, K. Distributed learning over Markovian fading channels for stable spectrum access. *IEEE Access* **2022**, *10*, 46652–46669. [[CrossRef](#)]
128. Banks, J.S.; Sundaram, R.K. Switching costs and the Gittins index. *Econometrica* **1994**, *62*, 687–694. [[CrossRef](#)]
129. Asawa, M.; Teneketzis, D. Multi-armed bandits with switching penalties. *IEEE Trans. Automat. Control* **1996**, *41*, 328–348. [[CrossRef](#)]
130. Niño-Mora, J. Computing an index policy for bandits with switching penalties. In Proceedings of the 1st International ICST Workshop on Tools for solving Structured Markov Chains (SMCTools), Nantes, France, 26 October 2007; Buchholz, P., Dayar, T., Eds.; ICST: Brussels, Belgium, 2007; ACM International Conference Proceedings Series. [[CrossRef](#)]
131. Niño-Mora, J. A faster index algorithm and a computational study for bandits with switching costs. *INFORMS J. Comput.* **2008**, *20*, 255–269. [[CrossRef](#)]
132. Niño-Mora, J. Fast two-stage computation of an index policy for multi-armed bandits with setup delays. *Mathematics* **2021**, *9*, 52. [[CrossRef](#)]
133. Niño-Mora, J. Marginal productivity index policies for scheduling restless bandits with switching penalties. In *Algorithms for Optimization with Incomplete Information*; Dagstuhl Seminar Proceedings, 16–21 January 2005; Albers, S., Möhring, R.H., Pflug, G.C., Schultz, R., Eds.; Schloss Dagstuhl—Leibniz-Zentrum für Informatik: Dagstuhl, Germany, 2005; Volume 05031. [[CrossRef](#)]
134. Le Ny, J.; Feron, E. Restless bandits with switching costs: Linear programming relaxations, performance bounds and limited lookahead policies. In Proceedings of the 25th American Control Conference (ACC), Minneapolis, MN, USA, 14–16 June 2006; pp. 1587–1592.
135. Arlotto, A.; Chick, S.E.; Gans, N. Optimal hiring and retention policies for heterogeneous workers who learn. *Manag. Sci.* **2014**, *60*, 110–129. [[CrossRef](#)]
136. Niño-Mora, J. A marginal productivity index policy for the finite-horizon multiarmed bandit problem. In Proceedings of the Joint 44th IEEE Conference on Decision and Control, and European Control Conference (CDC-ECC), Seville, Spain, 12–15 December 2005; pp. 1718–1722.
137. Niño-Mora, J. Computing a classic index for finite-horizon bandits. *INFORMS J. Comput.* **2011**, *23*, 173–330. [[CrossRef](#)]
138. Dayanik, S.; Powell, W.; Yamazaki, K. Index policies for discounted bandit problems with availability constraints. *Adv. Appl. Probab.* **2008**, *40*, 377–400. [[CrossRef](#)]
139. Ansell, P.S.; Glazebrook, K.D.; Niño Mora, J.; O’Keeffe, M. Whittle’s index policy for a multi-class queueing system with convex holding costs. *Math. Meth. Oper. Res.* **2003**, *57*, 21–39.
140. Niño-Mora, J. Marginal productivity index policies for admission control and routing to parallel multi-server loss queues with reneging. In Proceedings of the 1st EuroFGI Conference on Network Control and Optimization (NETCOOP), Avignon, France, 5–7 June 2007; Lecture Notes in Computer Science; Chahed, T., Tuffin, B., Eds.; Springer: Berlin/Heidelberg, Germany, 2007; Volume 4465, pp. 138–149.
141. Niño-Mora, J. Admission and routing of soft real-time jobs to multiclustres: Design and comparison of index policies. *Comput. Oper. Res.* **2012**, *39*, 3431–3444. [[CrossRef](#)]
142. Niño-Mora, J. Towards minimum loss job routing to parallel heterogeneous multiserver queues via index policies. *Eur. J. Oper. Res.* **2012**, *220*, 705–715. [[CrossRef](#)]
143. Niño-Mora, J. Resource allocation and routing in parallel multi-server queues with abandonments for cloud profit maximization. *Comput. Oper. Res.* **2019**, *103*, 221–236. [[CrossRef](#)]
144. Raissi-Dehkordi, M.; Baras, J.S. Broadcast scheduling in information delivery systems. In Proceedings of the IEEE Global Telecommunications Conference (GLOBECOM), Taipei, Taiwan, 17–21 November 2002; pp. 2935–2939.
145. Dusonchet, F.; Hongler, M.O. Continuous-time restless bandit and dynamic scheduling for make-to-stock production. *IEEE Trans. Robot. Automat.* **2003**, *19*, 977–990. [[CrossRef](#)]
146. Goyal, M.; Kumar, A.; Sharma, V. A stochastic control approach for scheduling multimedia transmissions over a polled multiaccess fading channel. *Wirel. Netw.* **2006**, *12*, 605–621. [[CrossRef](#)]
147. Niño-Mora, J. Marginal productivity index policies for scheduling a multiclass delay-/loss-sensitive queue. *Queueing Syst.* **2006**, *54*, 281–312. [[CrossRef](#)]
148. Cao, J.H.; Nyberg, C. Linear programming relaxations and marginal productivity index policies for the buffer sharing problem. *Queueing Syst.* **2008**, *60*, 247–269. [[CrossRef](#)]
149. Borkar, V.S.; Pattathil, S. Whittle indexability in egalitarian processor sharing systems. *Ann. Oper. Res.* **2022**, *317*, 417–437. [[CrossRef](#)]

150. O'Meara, T.; Patel, A. A topic-specific web robot model based on restless bandits. *IEEE Internet Comput.* **2001**, *5*, 27–35. [[CrossRef](#)]
151. Niño-Mora, J. A dynamic page-refresh index policy for web crawlers. In Proceedings of the 21st International Conference on Analytical and Stochastic Modelling Techniques and Applications (ASMTA), Budapest, Hungary, 30 June–2 July 2014; Lecture Notes in Computer Science; Sericola, B., Telek, M., Horváth, G., Eds.; Springer: Berlin/Heidelberg, Germany, 2014; Volume 8499, pp. 46–60.
152. Avrachenkov, K.E.; Borkar, V.S. Whittle index policy for crawling ephemeral content. *IEEE Trans. Control Netw. Syst.* **2018**, *5*, 446–455. [[CrossRef](#)]
153. Deo, S.; Iravani, S.; Jiang, T.T.; Smilowitz, K.; Samuelson, S. Improving health outcomes through better capacity allocation in a community-based chronic care model. *Oper. Res.* **2013**, *61*, 1277–1294. [[CrossRef](#)]
154. Ayer, T.; Zhang, C.; Bonifonte, A.; Spaulding, A.C.; Chhatwal, J. Prioritizing hepatitis C treatment in US prisons. *Oper. Res.* **2019**, *67*, 853–873. [[CrossRef](#)]
155. Mate, A.; Perrault, A.; Tambe, M. Risk-aware interventions in public health: Planning with restless multi-armed bandits. In Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), Online, 3–7 May 2021; Endriss, U., Nowé, A., Dignum, F., Lomuscio, A., Eds.; IFAAMAS: Richland, CA, USA, 2021; pp. 12017–12025.
156. Mate, A.; Madaan, L.; Taneja, A.; Madhiwalla, N.; Verma, S.; Singh, G.; Hegde, A.; Varakantham, P.; Tambe, M. A field study in deploying restless multi-armed bandits: Assisting non-profits in improving maternal and child health. In Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI), Online, 22 February–1 March 2022; pp. 12017–12025.
157. Wei, Y.; Yu, F.R.; Song, M. Distributed optimal relay selection in wireless cooperative networks with finite-state Markov channels. *IEEE Trans. Veh. Tech.* **2010**, *59*, 2149–2158.
158. Wei, X.H.; Neely, M.J. Power-aware wireless file downloading: A Lyapunov indexing approach to a constrained restless bandit problem. *IEEE/ACM Trans. Netw.* **2016**, *24*, 2264–2277. [[CrossRef](#)]
159. Aalto, S.; Lassila, P.; Osti, P. Whittle index approach to size-aware scheduling for time-varying channels with multiple states. *Queueing Syst.* **2016**, *83*, 195–225. [[CrossRef](#)]
160. Borkar, V.S.; Kasbekar, G.S.; Pattathil, S.; Shetty, P.Y. Opportunistic scheduling as restless bandits. *IEEE Trans. Control Netw. Syst.* **2018**, *5*, 1952–1961. [[CrossRef](#)]
161. Sun, Y.; Feng, G.; Qin, S.; Sun, S.S. Cell association with user behavior awareness in heterogeneous cellular networks. *IEEE Trans. Veh. Tech.* **2018**, *67*, 4589–4601. [[CrossRef](#)]
162. Aalto, S.; Lassila, P.; Taboada, I. Whittle index approach to opportunistic scheduling with partial channel information. *Perform. Eval.* **2019**, *136*, 102052. [[CrossRef](#)]
163. Wang, K.H.; Yu, J.H.; Chen, L.; Zhou, P.; Ge, X.H.; Win, M.Z. Opportunistic scheduling revisited using restless bandits: Indexability and index policy. *IEEE Trans. Wirel. Comm.* **2019**, *18*, 4997–5010. [[CrossRef](#)]
164. Sun, J.Z.; Jiang, Z.Y.; Krishnamachari, B.; Zhou, S.; Niu, Z.S. Closed-form Whittle's index-enabled random access for timely status update. *IEEE Trans. Comm.* **2020**, *68*, 1538–1551. [[CrossRef](#)]
165. Chen, G.P.; Liew, S.C.; Shao, Y.L. Uncertainty-of-information scheduling: A restless multiarmed bandit framework. *IEEE Trans. Inform. Theory* **2022**, *68*, 6151–6173. [[CrossRef](#)]
166. Singh, S.K.; Borkar, V.S.; Kasbekar, G.S. User association in dense mmWave networks as restless bandits. *IEEE Trans. Veh. Tech.* **2022**, *71*, 7919–7929. [[CrossRef](#)]
167. Huberman, B.A.; Wu, F. The economics of attention: Maximizing user value in information-rich environments. *Adv. Complex Syst.* **2008**, *11*, 487–496. [[CrossRef](#)]
168. Glazebrook, K.D.; Niño Mora, J.; Ansell, P.S. Index policies for a class of discounted restless bandits. *Adv. Appl. Probab.* **2002**, *34*, 754–774. [[CrossRef](#)]
169. Kumar, U.D.; Saranga, H. Optimal selection of obsolescence mitigation strategies using a restless bandit model. *Eur. J. Oper. Res.* **2010**, *200*, 170–180. [[CrossRef](#)]
170. Temple, T.; Frazzoli, E. Whittle-indexability of the cow path problem. In Proceedings of the American Control Conference (ACC), Baltimore, MD, USA, 30 June–2 July 2010; pp. 4152–4158.
171. He, T.; Chen, S.Y.; Kim, H.; Tong, L.; Lee, K.W. Scheduling parallel tasks onto opportunistically available cloud resources. In Proceedings of the IEEE 5th International Conference on Cloud Computing (CLOUD), Honolulu, HI, USA, 24–29 June 2012; pp. 180–187.
172. Taylor, J.A.; Mathieu, J.L. Index policies for demand response. *IEEE Trans. Power Syst.* **2014**, *29*, 1287–1295. [[CrossRef](#)]
173. Sun, J.; Ma, H. Heterogeneous-belief based incentive schemes for crowd sensing in mobile social networks. *J. Netw. Comput. Appl.* **2014**, *42*, 189–196. [[CrossRef](#)]
174. Lin, S.; Zhang, J.J.; Hauser, J.R. Learning from experience, simply. *Mark. Sci.* **2015**, *34*, 1–19. [[CrossRef](#)]
175. Guo, X.Y.; Singh, R.; Kumar, P.R.; Niu, Z.S. A risk-sensitive approach for packet inter-delivery time optimization in networked cyber-physical systems. *IEEE/ACM Trans. Netw.* **2018**, *26*, 1976–1989. [[CrossRef](#)]
176. Yu, Z.; Xu, Y.J.; Tong, L. Deadline scheduling as restless bandits. *IEEE Trans. Automat. Control* **2018**, *63*, 2343–2358. [[CrossRef](#)]
177. Avrachenkov, K.E.; Borkar, V.S.; Pattathil, S. Controlling G-AIMD by index policy. In Proceedings of the 56th IEEE Conference on Decision and Control (CDC), Melbourne, Australia, 12–15 December 2017; pp. 120–125.

178. Borkar, V.S.; Ravikumar, K.; Saboo, K. An index policy for dynamic pricing in cloud computing under price commitments. *Appl. Math.* **2017**, *44*, 215–245. [[CrossRef](#)]
179. Menner, M.; Zeilinger, M.N. A user comfort model and index policy for personalizing discrete controller decisions. In Proceedings of the 16th European Control Conference (ECC), Limassol, Cyprus, 12–15 June 2018; pp. 1759–1765.
180. Jhunjhunwala, P.R.; Moharir, S.; Manjunath, D.; Gopalan, A. On a class of restless multi-armed bandits with deterministic policies. In Proceedings of the International Conference on Signal Processing and Communications (SPCOM), Bangalore, India, 16–19 July 2018; pp. 487–491. [[CrossRef](#)]
181. Abbou, A.; Makis, V. Group maintenance: A restless bandits approach. *INFORMS J. Comput.* **2019**, *31*, 719–731. [[CrossRef](#)]
182. Gerum, P.C.L.; Altay, A.; Baykal-Gursoy, M. Data-driven predictive maintenance scheduling policies for railways. *Transport. Res. Part C Emerg. Technol.* **2019**, *107*, 137–154. [[CrossRef](#)]
183. Li, D.; Ding, L.; Connor, S. When to switch? Index policies for resource scheduling in emergency response. *Prod. Oper. Manag.* **2020**, *29*, 241–262. [[CrossRef](#)]
184. Fu, J.; Moran, B.; Taylor, P.G. A restless bandit model for resource allocation, competition, and reservation. *Oper. Res.* **2022**, *70*, 416–431. [[CrossRef](#)]
185. Dahiya, A.; Akbarzadeh, N.; Mahajan, A.; Smith, S.L. Scalable operator allocation for multi-robot assistance: A restless bandit approach. *IEEE Trans. Control Netw. Syst.* **2022**, *9*, 1397–1408. [[CrossRef](#)]
186. Ou, H.C.; Siebenbrunner, C.; Killian, J.; Brooks, M.B.; Kempe, D.; Vorobeychik, Y.; Tambe, M. Networked restless multi-armed bandits for mobile interventions. In Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS), Online, 9–13 May 2022; Faliszewski, P., Mascardi, V., Pelachaud, C., Taylor, M.E., Eds.; International Foundation for Autonomous Agents and Multiagent Systems: London, UK, 2022.
187. Krishnamurthy, V. *Partially Observed Markov Decision Processes: From Filtering to Controlled Sensing*; Cambridge University Press: Cambridge, UK, 2016.
188. La Scala, B.F.; Moran, B. Optimal target tracking with restless bandits. *Digit. Signal Process.* **2006**, *16*, 479–487. [[CrossRef](#)]
189. Le Ny, J.; Feron, E.; Dahleh, M. Scheduling continuous-time Kalman filters. *IEEE Trans. Automat. Control* **2011**, *56*, 1381–1394. [[CrossRef](#)]
190. Akbarzadeh, N.; Mahajan, A. Partially observable restless bandits with restarts: Indexability and computation of Whittle index. In Proceedings of the 61st IEEE Conference on Decision and Control (CDC), Cancún, Mexico, 6–9 December 2022; pp. 4898–4904.
191. Gan, X.; Chen, B. A novel sensing scheme for dynamic multichannel access. *IEEE Trans. Veh. Tech.* **2012**, *61*, 208–221. [[CrossRef](#)]
192. He, T.; Anandkumar, A.; Agrawal, D. Index-based sampling policies for tracking dynamic networks under sampling constraints. In Proceedings of the 30th IEEE International Conference on Computer Communications (INFOCOM), Shanghai, China, 10–15 April 2011; pp. 1233–1241.
193. Meshram, R.; Manjunath, D.; Gopalan, A. A restless bandit with no observable states for recommendation systems and communication link scheduling. In Proceedings of the 54th IEEE Conference on Decision and Control (CDC), Osaka, Japan, 15–18 December 2015; pp. 7820–7825.
194. Ouyang, W.Z.; Murugesan, S.; Eryilmaz, A.; Shroff, N.B. Exploiting channel memory for joint estimation and scheduling in downlink networks—A Whittle’s indexability analysis. *IEEE Trans. Inform. Theory* **2015**, *61*, 1702–1719. [[CrossRef](#)]
195. Taboada, I.; Liberal, F.; Fajardo, J.O.; Blanco, B. An index rule proposal for scheduling in mobile broadband networks with limited channel feedback. *Perform. Eval.* **2017**, *117*, 130–142. [[CrossRef](#)]
196. Meshram, R.; Manjunath, D.; Gopalan, A. On the Whittle index for restless multiarmed hidden Markov bandits. *IEEE Trans. Automat. Control* **2018**, *63*, 3046–3053. [[CrossRef](#)]
197. Elmaghraby, H.M.; Liu, K.Q.; Ding, Z. Femtocell scheduling as a restless multiarmed bandit problem using partial channel state observation. In Proceedings of the IEEE International Conference on Communications (ICC), Kansas City, MO, USA, 20–24 May 2018.
198. Mehta, V.; Meshram, R.; Kaza, K.; Merchant, S.N.; Desai, U.B. Rested and restless bandits with constrained arms and hidden states: Applications in social networks and 5G networks. *IEEE Access* **2018**, *6*, 56782–56799. [[CrossRef](#)]
199. Kaza, K.; Meshram, R.; Mehta, V.; Merchant, S.N. Sequential decision making with limited observation capability: Application to wireless networks. *IEEE Trans. Cogn. Commun. Netw.* **2019**, *5*, 237–251. [[CrossRef](#)]
200. Yang, F.; Luo, X. A restless MAB-based index policy for UL pilot allocation in massive MIMO over Gauss–Markov fading channels. *IEEE Trans. Veh. Technol.* **2020**, *69*, 3034–3047. [[CrossRef](#)]
201. Hsu, Y.P.; Modiano, E.; Duan, L.J. Scheduling algorithms for minimizing age of information in wireless broadcast networks with random arrivals. *IEEE Trans. Mob. Comput.* **2020**, *19*, 2903–2915. [[CrossRef](#)]
202. Wang, J.Z.; Ren, X.Q.; Mo, Y.L.; Shi, L. Whittle index policy for dynamic multichannel allocation in remote state estimation. *IEEE Trans. Automat. Control* **2020**, *65*, 591–603. [[CrossRef](#)]
203. Chen, Y.; Ephremides, A. Scheduling to minimize age of incorrect information with imperfect channel state information. *Entropy* **2021**, *23*, 1572. [[CrossRef](#)] [[PubMed](#)]
204. Kang, S.; Joo, C. Index-based update policy for minimizing information mismatch with Markovian sources. *J. Commun. Netw.* **2021**, *23*, 488–498. [[CrossRef](#)]

205. Li, D.; Varakantham, P. Efficient resource allocation with fairness constraints in restless multi-armed bandits. In Proceedings of the 38th Conference on Uncertainty in Artificial Intelligence (UAI), Eindhoven, The Netherlands, 1–5 August 2022; Cussens, J., Zhang, K., Eds.; PMLR: Birmingham, UK, 2022; pp. 1158–1167.
206. Tong, J.W.; Fu, L.Q.; Han, Z. Age-of-information oriented scheduling for multichannel IoT systems with correlated sources. *IEEE Trans. Wirel. Comm.* **2022**, *21*, 9775–9790. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.