



HAL
open science

Tabular and Deep Learning for the Whittle Index

Francisco Robledo Relaño, Vivek Borkar, Urtzi Ayesta, Konstantin Avrachenkov

► **To cite this version:**

Francisco Robledo Relaño, Vivek Borkar, Urtzi Ayesta, Konstantin Avrachenkov. Tabular and Deep Learning for the Whittle Index. ACM Transactions on Modeling and Performance Evaluation of Computing Systems, 2024, 10.1145/3670686 . hal-04599291

HAL Id: hal-04599291

<https://hal.science/hal-04599291>

Submitted on 4 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Tabular and Deep Learning for the Whittle Index

FRANCISCO ROBLEDO RELAÑO*, UPV/EHU, Univ. of the Basque Country, Spain

VIVEK BORKAR, Indian Institute of Technology Bombay, India

URTZI AYESTA, CNRS-IRIT, UPV/EHU and IKERBASQUE, France and Spain

KONSTANTIN AVRACHENKOV, INRIA Sophia Antipolis, France

The Whittle index policy is a heuristic that has shown remarkably good performance (with guaranteed asymptotic optimality) when applied to the class of problems known as Restless Multi-Armed Bandit Problems (RMABPs). In this paper we present QWI and QWINN, two reinforcement learning algorithms, respectively tabular and deep, to learn the Whittle index for the total discounted criterion. The key feature is the use of two time-scales, a faster one to update the state-action Q -values, and a relatively slower one to update the Whittle indices. In our main theoretical result we show that QWI, which is a tabular implementation, converges to the real Whittle indices. We then present QWINN, an adaptation of QWI algorithm using neural networks to compute the Q -values on the faster time-scale, which is able to extrapolate information from one state to another and scales naturally to large state-space environments. For QWINN, we show that all local minima of the Bellman error are locally stable equilibria, which is the first result of its kind for DQN-based schemes. Numerical computations show that QWI and QWINN converge faster than the standard Q -learning algorithm, neural-network based approximate Q -learning and other state of the art algorithms.

CCS Concepts: • **Computing methodologies** → **Sequential decision making**.

Additional Key Words and Phrases: Machine learning, Reinforcement Learning, Whittle Index, Markov Decision Problem, Multi-armed Restless Bandit

ACM Reference Format:

Francisco Robledo Relano*, Vivek Borkar, Urtzi Ayesta, and Konstantin Avrachenkov. 2022. Tabular and Deep Learning for the Whittle Index. 1, 1 (May 2022), 22 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 INTRODUCTION

Markov decision processes (MDPs) provide a mathematical framework to model sequential decision problems. Formally, an MDP is a stochastic control process where the objective is to maximise a long-run payoff. At each time step, depending on the state and action taken by the MDP, the decision maker receives a reward and reaches a new random state. Due to their broad applicability, MDPs are found in many areas, including artificial intelligence, economics, communications and operations research.

Solving an MDP typically involves dynamic programming, a method often computationally prohibitive for realistically sized models. Consequently, significant attention has been devoted to classes of MDPs that are either analytically

* Corresponding author. The authors are listed in reversed alphabetical order.

Authors' addresses: Francisco Robledo Relano*, frobledo96@gmail.com, UPV/EHU, Univ. of the Basque Country, Donostia, Gipuzkoa, Spain, 20018; Vivek Borkar, borkar.vs@gmail.com, Indian Institute of Technology Bombay, India; Urtzi Ayesta, urtzi.ayesta@irit.fr, CNRS-IRIT, UPV/EHU and IKERBASQUE, France and Spain; Konstantin Avrachenkov, kavrachenkov@inria.fr, INRIA Sophia Antipolis, France.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

Manuscript submitted to ACM

tractable or amenable to precise approximations. This paper focuses on one such class - Restless Multi-Armed Bandit Problems (RMABPs), initially introduced in [26]. In an RMABP there are N concurrent projects or bandit's arms. The decision maker knows the states of all arms and the reward in every state, and aims at maximizing the long-term reward. At every decision epoch, the decision maker activates $M < N$ arms, and the state of active and passive arms evolve stochastically. We note that RMABPs are a generalization of multi-arm bandit problems (MABP), in which only one arm can be activated at a time, and the state of passive arms does not evolve. It was shown by Gittins that the optimal solution to a MABP is an index policy, nowadays known as Gittins' index policy, see [7].

Over the years, RMABPs have gained significant traction, finding applications in diverse domains such as inventory routing, machine maintenance, healthcare systems, and networking. Solving RMABPs precisely is typically infeasible, except for simplified examples, and they are classified as PSPACE-hard problems [17]. In [26] Whittle developed a heuristic by solving a relaxed version of the RMABP in which $M < N$ arms are activated on average. The resulting policy, known as Whittle index policy, relies on calculating Whittle index for each of the arms, and activating in every decision epoch the M arms with the highest Whittle indices. It has been reported on numerous instances that Whittle index policy provides strikingly good performance, and it has been shown to be asymptotically optimal as the number of arms grows large [25]. As expected, Whittle index reduces to Gittins index when applied to an MABP.

In recent years there has been a surge in the interest in developing algorithms to learn the Whittle index in model-free setting. The first paper aimed at learning index policies is to the best of our knowledge, [4], which considered an MABP and developed an algorithm that learns Gittins' indices. Regarding Whittle index, one of the first papers was [5], which proposed an algorithm that did not converge to the correct values. More recently, the authors of [6] focused on the resolution of the restless multi-armed bandit by iteratively obtaining a policy similar to the Whittle index policy, but not necessarily a precise value of the Whittle index themselves. Furthermore, the algorithm in [6] is a tabular only algorithm, making it unsuitable for large state spaces where the exploration of all states becomes impossible.

Further contributions to the field include the work by [11], which established a foundation for the indexability concept in restless bandit problems and the associated optimality of Whittle indices for dynamic channel access. [8] showcased the adaptability of the Whittle index for stochastic scheduling within a restless multi-armed bandit framework. [22] offered insights into asymptotically optimal priority policies in settings involving both indexable and non-indexable bandits, illustrating the robustness of Whittle index-based heuristics. Moreover, [14] provided an examination of threshold-indexability and highlighted the practical utility of the Whittle index for single-project control and as a heuristic policy.

In their exploration of the time-average criterion, Avrachenkov and Borkar [2] devised a tabular algorithm that successfully converges to the Whittle index. The tabular setting with the time-average criterion was then extended to a Neural Network (NN) setting in [16]. A finite-time analysis of the scheme of [2] with a neural network approximate has been also recently performed in [27]. Killian et al. [9] explore the online learning landscape for multi-action RMABs, presenting a Q-learning Lagrange policy algorithm tailored for restless multi-armed bandits with multiple discrete actions. Another notable mention is the NeurWIN algorithm [13], a NN gradient-based reinforcement learning approach for solving RMABPs. NeurWIN, focused on maximizing discounted long-term rewards, does not, however, guarantee convergence to the true Whittle indices.

This paper introduces QWI and QWINN, two novel reinforcement learning algorithms, engineered to master the Whittle index in the realm of discounted rewards. Both algorithms leverage a dual time-scale strategy: a rapid cycle for updating state-action Q -values and a slower one for refining Whittle indices. Our principal theoretical finding asserts that QWI, with its tabular form, consistently aligns with the Whittle indices across various RMABPs. Additionally, we

unveil QWINN, an evolution of QWI that integrates neural networks, drawing inspiration from DQN [12]. QWINN excels in large state spaces due to its neural network-driven extrapolation capabilities. We establish, for the first time to our knowledge, that QWINN ensures that local minima of the Bellman error function form stable equilibria in the differential equation limit of the DQN model. Thus, when QWINN starts within the attraction domain of such equilibria, convergence is highly probable. This interesting finding for DQN-based methods marks a significant advance in the field. Our numerical comparisons, pitting QWI and QWINN against established algorithms like standard Q -learning, DQN, and NeurWIN, demonstrate superior performance in both convergence rate and discounted reward optimization, with QWINN exhibiting remarkable proficiency in deriving accurate Whittle indices from limited data samples.

We compare numerically the performance of QWI and QWINN with several other algorithms, including the standard Q -learning algorithm, a vanilla implementation of DQN and NeurWIN. Unlike [23] and [15], we do not analyse the regret of our algorithms. Instead, we will study the convergence to the theoretical values of the Whittle indices and assess the performance of the algorithms over time as a function of the policy they have learned so far. Our results show that both QWI and QWINN outperform the other algorithms both in terms of the rate of convergence as well as the discounted reward. QWINN stands out in its ability to obtain good Whittle indices with far fewer samples than QWI, and thanks to the extrapolation of the information performed by the neural networks, it is able to obtain good predictions of the indices of these states even with few or no samples. On the other hand, QWI, given enough samples in each state/action pair, is able to reliably converge to the correct values of the Whittle indices.

2 RESTLESS MARKOVIAN BANDITS

2.1 Problem formulation and relaxation

In our study of Restless Markovian Bandits, we examine a Restless Multi-Armed Bandit Problem (RMABP) comprising N arms. Each arm i , where $i = 1, \dots, N$, possesses a state space \mathcal{S}^i of cardinality $|\mathcal{S}^i|$. The combined state space for all arms is denoted as $\mathcal{S} = \mathcal{S}^1 \times \dots \times \mathcal{S}^N$ with cardinality $|\mathcal{S}|$. For each arm at time step n , s_n^i represents the current state and a_n^i the chosen action. The reward garnered by arm i at this step is given by $r^i(s_n^i, a_n^i)$. We adopt the total discounted reward criterion with a discount factor $0 < \gamma < 1$. Active arms transition between states with probability $p(s_{n+1}^i | s_n^i, 1)$, while passive arms use a different transition probability $p(s_{n+1}^i | s_n^i, 0)$.

At each time step n , the control policy π observes the state of all the arms, and activates M arms ($a_n^i = 1$), whereas the rest remain passive with $a_n^i = 0$. The objective is to determine the optimal control policy π^* that solves:

$$V_{\pi^*}(s_1, \dots, s_N) = \max_{\pi} E_s \left[\sum_{n=1}^{\infty} \sum_{i=1}^N \gamma^n r^i(s_n^i, a_n^i) \right], \quad (1)$$

subject to:

$$\sum_{i=1}^N a_n^i = M, \quad n \geq 0, \quad (2)$$

where V_{π^*} is known as the value function.

Whittle's approach [26] relaxes the constraint to be satisfied on average, leading to an unconstrained problem formulation: $E_s \left[\sum_{n=0}^{\infty} \sum_{i=1}^N \gamma^n a_n^i \right] \leq M/(1 - \gamma)$. The problem then has an equivalent *unconstrained* formulation:

$$\max_{\pi} E_s \left[\sum_{n=1}^{\infty} \sum_{i=1}^N \gamma^n \left(r(s_n^i, a_n^i) + \lambda (1 - a_n^i) \right) \right] \quad (3)$$

where λ is a Lagrange multiplier associated with the constraint. We note that as λ increases, we expect the passive action to become attractive in more states, and as a consequence, the multiplier λ can be seen as a subsidy for passivity.

The key observation by Whittle is that problem (3) can be decomposed, and its solution is obtained by combining the solution of N independent problems. In other words, for each arm i we need to solve the associated Bellman equation given by:

$$V_{\pi}^i(s) = \max_{a \in \{0,1\}} [Q^i(s, a)] \quad (4)$$

where

$$Q_{\pi}^i(s, a) = a \left(r^i(s, 1) + \gamma \sum_j p^i(j|s, 1) V_{\pi}^i(j) \right) + (1 - a) \left(r^i(s, 0) + \lambda + \gamma \sum_j p^i(j|s, 0) V_{\pi}^i(j) \right) \quad (5)$$

The functions $V^i(s)$ and $Q^i(s, a)$ are known as the value function and the state-action function resp. for arm i .

2.2 Whittle index

The optimal action in (4) will be to activate a given arm i if $r^i(k, 1) + \gamma \sum_j p^i(j|k, 1) V^i(j) > r^i(k, 0) + \lambda + \gamma \sum_j p^i(j|k, 0) V^i(j)$, while otherwise the optimal action will be to keep it passive. For a given λ , the optimal policy π^* is then characterized by $\mathcal{S}(\lambda)$, the set of states in which the optimal action is to activate.

The problem is deemed *indexable* if the subset of states where activation is optimal decreases monotonically with increasing λ . For a specific state k , the Whittle index, $\lambda(k)$, equates the expected rewards for active and passive actions:

$$r^i(k, 1) + \gamma \sum_j p^i(j|k, 1) V_{\pi^*}^i(j) = r^i(k, 0) + \lambda(k) + \gamma \sum_j p^i(j|k, 0) V_{\pi^*}^i(j). \quad (6)$$

It thus follows that the Whittle index characterizes the optimal solution to the relaxed problem (3), which will simply activate all arms whose Whittle index is larger than λ .

In his seminal paper, Whittle introduced the heuristic policy for the problem (1), known nowadays as Whittle index policy, which is defined as the policy that at every time step n activates the M arms with the M highest Whittle indices. As explained in the introduction, this heuristic has shown to have a close to optimal performance, and to be asymptotically optimal as the values of both N and M tend to infinity, see [25] and [22].

3 LEARNING THE WHITTLE INDICES

In this section we present QWI and QWINN, the tabular and neural network-based algorithms to learn Whittle indices. Throughout this section we drop the dependence on the arm from the notation.

3.1 Tabular QWI algorithm

In this subsection, we delve into the Tabular QWI algorithm, a method designed to learn Whittle indices in a structured, iterative manner. The core idea is to iteratively adjust the Whittle index, refining our estimate of the state-action function until convergence. The Whittle index $\lambda(x)$ for a state x can be implicitly derived from the equilibrium condition in Equation (6):

$$\lambda(x) : Q(x, 1) - Q(x, 0) = 0. \quad (7)$$

This is an implicit equation, as the state-action function $Q(\cdot, \cdot)$ depends on the value of the multiplier $\lambda(x)$. The Whittle index is thus defined by coupled equations (5) and (7).

To find this equilibrium, we propose a method that alternates between estimating the state-action functions and updating the Whittle index. This process is conducted online, ensuring convergence to the actual Whittle index. The QWI algorithm employs stochastic approximation with multiple time scales (Chapter 6 in [3]) to achieve this.

For each arm i , from the samples (s_n, a_n, r_n, s_{n+1}) :

$$Q_{n+1}^x(s_n, a_n) = (1 - \alpha(n))Q_n^x(s_n, a_n) + \alpha(n) \left((1 - a_n)(r_0(s_n) + \lambda_n(x)) + a_n r_1(s_n) + \gamma \max_{v \in \{0,1\}} Q_n^x(s_{n+1}, v) \right) \quad (8)$$

and

$$\lambda_{n+1}(x) = \lambda_n(x) + \beta(n) (Q_n^x(x, 1) - Q_n^x(x, 0)). \quad (9)$$

Here, s_n denotes the state visited at time step n , x is a reference state for which we wish to learn Whittle's index, $r_0(s_n)$ and $r_1(s_n)$ are the sampled rewards for actions passive and active, respectively, and $\alpha(n)$ and $\beta(n)$ are learning parameters. The superscripts in (8) stand for the parametric dependence of the Q-values on x, λ , where x is a fixed reference state and $\lambda = \lambda(x)$ is a slowly varying parameter updated by (9). As usual in the literature, the learning parameters need to satisfy $\sum_n \alpha(n) = \infty$, $\sum_n \alpha(n)^2 < \infty$, $\sum_n \beta(n) = \infty$, $\sum_n \beta(n)^2 < \infty$. In addition, we require $\beta(n) = o(\alpha(n))$ in order to implement the two distinct time scales, namely, the relatively faster time scale for the updates of the state-action function, and the slow one for the Whittle indices. In our case, we use the following stepsizes that meet the above-mentioned conditions:

$$\alpha(n) = \frac{1}{\lceil n/5000 \rceil} \quad (10)$$

$$\beta(n) = \frac{1}{1 + \lceil n \log(n)/5000 \rceil} I\{(n) \bmod(50) = 0\} \quad (11)$$

The pseudo-code implementation of QWI is present in Alg. 1. For a given value of the discount parameter γ and exploration parameter ϵ , we initialise the N arms with a random state. At every decision epoch n , with probability $1 - \epsilon$, we choose the greedy action, i.e, we activate ($a_n^i = 1$) the M arms with largest $\lambda_n(x)$, while with probability ϵ we choose M arms at random. Rewards r_n^i are collected, the new states s_{n+1}^i are observed, and the values of the state-action function and the Whittle indices are updated by (8) and (9). For a given arm i , updating the value functions of all states $s \in \mathcal{S}^i$ necessitates iterating over all states in the state space and utilizing all reference states x . Thus, the computational complexity, defined as the total number of computations required for these updates, scales as $O(|\mathcal{S}^i|^2 |\mathcal{A}|)$, where $\mathcal{A} = \{0, 1\}$. We note that QWI learns simultaneously the Whittle indices of all the states of every arm on-line.

Algorithm 1 Tabular QWI Algorithm

Input: Discount parameter $\gamma \in (0, 1)$, exploration parameter $\epsilon \in [0, 1]$,
Output: Whittle index matrix for all states in each arm i
Initialize s_0 for all arms
for $n = 1 : n_{end}$ **do**
 Define action a_n^i through ϵ -greedy policy for each arm i
 Get new states s_{n+1}^i and rewards r_n^i from states s_n^i and actions a_n^i
 Update learning rate $\alpha(n)$, $\beta(n)$ as (10) and (11)
 Update (s_n^i, a_n^i) Q -values as (8)
 Update Whittle estimates for all states x in each arm i as (9)
end for

The next result shows that QWI converges to the Whittle index for any RMABP. The proof, adapted from [2], is given in Appendix A.

THEOREM 3.1. (Convergence of QWI) *Given learning parameters $\alpha(n)$ and $\beta(n)$ such that $\sum_n \alpha(n) = \sum_n \beta(n) = \infty$, $\sum_n \alpha(n)^2 < \infty$, $\sum_n \beta(n)^2 < \infty$ and $\beta(n) = o(\alpha(n))$ and that the problem satisfies the indexability condition, iterations (8) and (9) converge respectively to the state-action function of Whittle index policy, denoted by $Q_W(s, a)$, and to the Whittle indices $\lambda(s)$, i.e., $\lambda_n(s) \rightarrow \lambda(s)$ and $Q_n(s, a) \rightarrow Q_W(s, a)$ a.s. $\forall s \in S, a \in A$ as $n \rightarrow \infty$.*

3.2 QWINN algorithm

The QWINN algorithm extends the Whittle index heuristic of the QWI algorithm through the integration of neural networks for approximating Q -values. Distinct from the tabular approach in section 3.1, QWINN (algorithm 2) utilizes a neural network with three hidden layers (100,200,100 neurons each) connected by ReLU activation functions. This network inputs two state variables, the visited state s and the reference state x and outputs Q -values for both possible actions $Q_\theta^x(s) = [Q_\theta^x(s, 0) \quad Q_\theta^x(s, 1)]$. Following QWI's heuristic, QWINN updates its Q -value neural network estimator as:

$$Q_{target}^x(s_n, a_n) = (1 - a_n)(r_0(s_n) + \lambda(x)) + a_n r_1(s_n) + \gamma \max_{v \in A} Q_{\theta'}^x(s_{n+1}, v). \quad (12)$$

QWINN utilizes a tabular approach to update the Whittle indices, similar to QWI's heuristic in Equation (9). The Whittle indices are defined as:

$$\lambda_{n+1}(x) = \lambda_n(x) + \beta(n) \left(Q_\theta^x(x, 1) - Q_\theta^x(x, 0) \right). \quad (13)$$

One of the main changes with respect to the tabular algorithm in section 3.1 is the use of Double Q-Learning [21] [20], where we employ a second neural network for the computation of the $\max_{v \in A} Q_{\theta'}^x(s_{n+1}, v)$ term in Equation (12). The reason for this decision is due to the maximization bias of Q-learning: overestimating the Q -values results in this error increase over time, due to the target being $r + \gamma \max_v Q^x(s, v)$. The use of a second neural network for the target helps control this bias. This second neural network copies the parameter values of the main network Q_θ^x every 50 iterations.

In the following analysis, we investigate the local convergence properties of the DQN algorithm. We strengthen a necessary condition for local minima in optimization theory to an assumption and apply stochastic approximation theory to elucidate DQN's behavior in a neighborhood of a local minimum under specific conditions.

3.2.1 Preliminary Considerations. Consider the sequence $\tilde{\theta}_m = \theta_{T_n}$, defined for each $T_n \leq m < T_{n+1}$, where $T_n \uparrow \infty$. We specify that for some n , the parameter $\theta_n^* := \theta_{T_n}$ lies in a bounded neighborhood around a local minimum θ^* of the Bellman error function $\mathcal{E}(\cdot, \theta^*)$, defined as:

$$\mathcal{E}(\theta, \tau) := \mathbb{E} \left[\left\| \left[Q_\theta^X(s, a) - (1-a)(r_0(s) + \lambda) - ar_1(s) - \gamma \max_{v \in A} Q_\tau^X(s', v) \right] \right\|^2 \right].$$

Let $\nabla_1 \mathcal{E}, \nabla_1^2 \mathcal{E}$ denote the gradient and the Hessian of \mathcal{E} with respect to the first argument alone. Assuming that the Hessian $\nabla_1^2 \mathcal{E}(\theta^*, \theta^*)$ is positive definite, the inverse function theorem ensures a locally defined, bijective mapping $F(\tau) = (\nabla_1 \mathcal{E}(\cdot, \tau))^{-1}(\mathbf{0})$, where $\mathbf{0}$ is the zero vector, in a neighborhood of θ^* . Let us introduce an additional assumption.

ASSUMPTION 3.1. $F(\tau)$ is locally a contraction around the equilibrium point (θ^*, θ^*) .

3.2.2 Local Convergence.

THEOREM 3.2. For a θ^* as above, under Assumption 3.1 and the condition that the Hessian $\nabla_1^2 \mathcal{E}(\theta^*, \theta^*)$ is positive definite, the DQN algorithm exhibits local convergence to an open ball of radius $\frac{2\epsilon}{1-\alpha}$ centered at θ^* , for $T_{n+1} - T_n$ sufficiently large.

PROOF. By Assumption 3.1, $F(\theta^*)$ is locally a contraction with some factor $0 < \alpha < 1$, and assuming that ϵ is within the specified bounds, consider $F_n(\theta_n^*) := (\nabla_1 \mathcal{E}(\cdot, \theta_n^*))^{-1}(\mathbf{0})$. For $T_{n+1} - T_n$ sufficiently large, $\|\theta_{n+1}^* - F_n(\theta_n^*)\| < \epsilon$. By continuity, for θ in the ϵ -neighbourhood of θ^* , $\|F_n(\theta) - F_n(\theta_n^*)\| < \epsilon'$, for some ϵ' that we take to equal ϵ without loss of generality. Then

$$\begin{aligned} \|\theta_{n+1}^* - \theta^*\| &\leq \epsilon + \|F_n(\theta_n^*) - F(\theta^*)\| \\ &\leq \epsilon + \|F_n(\theta_n^*) - F(\theta_n^*)\| + \|F(\theta_n^*) - F(\theta^*)\| \\ &\leq 2\epsilon + \alpha \|\theta_n^* - \theta^*\|. \end{aligned}$$

Iterating, it follows that θ_n^* approaches the δ -ball centred at θ^* . This establishes the local convergence of the DQN algorithm within a neighborhood B of θ^* to a δ -neighborhood thereof. \square

3.2.3 Observations and Practical Implications. One of the key assumptions we have introduced in Theorem 3.2 is that $F(\tau)$, defined as $F(\tau) = (\nabla_1 \mathcal{E}(\cdot, \tau))^{-1}(\mathbf{0})$, is locally a contraction around the point θ^* . This assumption is supported by numerical analysis, which shows that in our examples, $F(\tau)$ is indeed a contraction around θ^* . However, finding sufficient general conditions for this contraction to hold falls outside the scope of this paper.

We have performed numerical calculations to illustrate that this assumption holds in the problems considered in the paper. In order to do so, we have proceed as follows. Linearizing $\nabla_1 \mathcal{E}(\theta, \tau)$ in the proximity (θ^*, θ^*) we get

$$\nabla_1 \mathcal{E}(\theta, \tau) \approx \nabla_1 \mathcal{E}(\theta^*, \theta^*) + \nabla_1^2 \mathcal{E}(\theta^*, \theta^*)(\theta - \theta^*) + \nabla_2 \nabla_1 \mathcal{E}(\theta^*, \theta^*)(\tau - \theta^*) + o(\theta - \theta^*, \tau - \theta^*).$$

Since, $\nabla_1 \mathcal{E}(\theta^*, \theta^*) = 0$, we get

$$(\theta - \theta^*) \approx -(\nabla_1^2 \mathcal{E}(\theta^*, \theta^*))^{-1} \nabla_2 \nabla_1 \mathcal{E}(\theta^*, \theta^*)(\tau - \theta^*).$$

To have a contraction (locally) around the equilibrium point, we need the spectrum of $-(\nabla_1^2 \mathcal{E}(\theta^*, \theta^*))^{-1} \nabla_2 \nabla_1 \mathcal{E}(\theta^*, \theta^*)$ to be inside the unit disc in the complex plane around the origin.

The foregoing assumes that \mathcal{E} is differentiable in the second argument, i.e., the target. More generally, a version of Danskin's theorem ensures existence of directional derivatives and the foregoing can be modified accordingly.

To calculate numerically the eigenvalues we have considered a random batch of training samples and compute the loss function. We then compute numerically the gradient of the loss function with respect to the first argument ($\nabla_1 \mathcal{E}(\cdot, \cdot)$) and then again, the gradient thereof with respect to the first argument ($\nabla_1^2 \mathcal{E}(\cdot, \cdot)$) and the second argument ($\nabla_2 \nabla_1 \mathcal{E}(\cdot, \cdot)$).

We have numerically verified that the condition on the spectrum for all the environments considered in the paper is satisfied. In the following histogram, we represent the eigenvalues obtained using the homogeneous environments in our paper, using a neural network with size $((2, 50), (50, 100), (100, 50), (50, 2))$, that is, with 10402 parameters.

Furthermore, the practical application of this contraction depends on the intervals $T_{n+1} - T_n$ being sufficiently large. The larger these intervals are, the smaller the choice of δ can be. This observation may not be consistent with settings where $T_n = nT$ is used for some fixed $T > 0$ with decreasing step sizes. Nevertheless, it remains valid when a constant step size is employed, provided that T is sufficiently large.

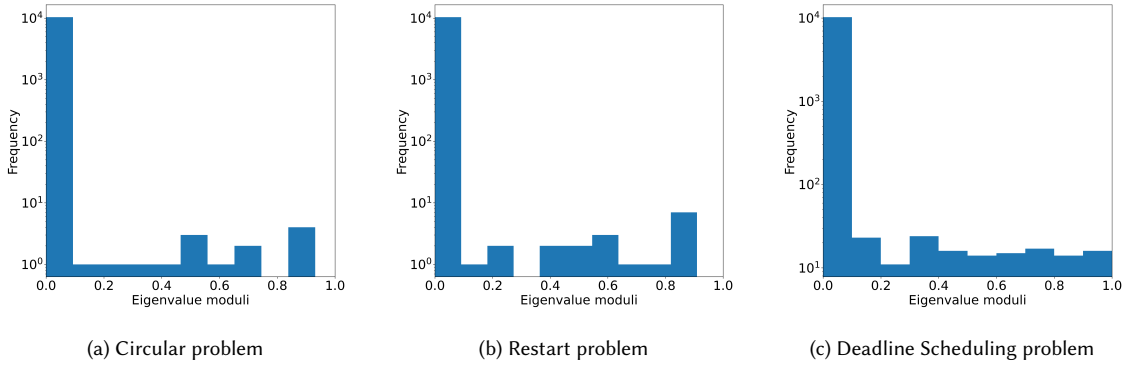


Fig. 1. Histogram of eigenvalue moduli of $-(\nabla_1^2 \mathcal{E}(\theta^*, \theta^*))^{-1} \nabla_2 \nabla_1 \mathcal{E}(\theta^*, \theta^*)$

3.3 Other algorithms

We describe now the other algorithms from the literature that we will consider in the numerical simulations.

3.3.1 NeurWIN algorithm. The NeurWIN algorithm, introduced in [13], is a neural network-based algorithm designed to directly compute the Whittle indices' estimates $\lambda_\theta(s_n)$ of a problem using as an input to the neural network only the state whose index is to be computed.

The NeurWIN algorithm considers that a policy of indices that achieves optimal rewards for RMABP is equivalent to Whittle's index policy. To obtain this policy, they define an activation function using a sigmoid function:

$$\sigma_m(f_\theta(s[t]) - \lambda) = \frac{1}{1 + e^{-m(f_\theta(s[t]) - \lambda)}}, \quad (14)$$

where m is a sensitivity parameter. This function selects action $a = 1$ with probability $\sigma_m(f_\theta(s[t]) - \lambda)$ and $a = 0$ with probability $1 - \sigma_m(f_\theta(s[t]) - \lambda)$. For each mini-batch of episodes, they randomly choose two states s_0 and s_1 , with s_0 serving as the fixed value $\lambda = f_\theta(s_0)$ and s_1 as the initial state. Multiple episodes are generated within each mini-batch, recording the actions and states visited based on the policy defined by Eq. (14). The gradients h_e are calculated as follows:

Algorithm 2 QWINN Algorithm

Input: Discount parameter $\gamma \in (0, 1)$, exploration parameter $\epsilon \in [0, 1]$,
Output: Whittle index vector for all states
Initialize s_0 for all arms
for $n = 1 : n_{end}$ **do**
 Define action a_n^i through ϵ -greedy policy for each arm i
 Get new states s_{n+1}^i and rewards r_n^i from states s_n^i and actions a_n^i
 Save each arm’s data into separated memories
 if number of samples in memory $>$ threshold **then**
 Update Whittle index learning rate $\beta(n)$ as (11)
 for arm $i \in N$ **do**
 for $x \in S$ **do**
 Predict Q -values of sample batch (s_n^i, a_n^i) using reference state x for Whittle index with Q_x^x
 Compute target Q -value for sample $(s_n^i, a_n^i, r_n^i, s_{n+1}^i, x)$ as (12), using secondary network $Q_{\theta'}^x$
 end for
 Compute the mean square error loss function between Q_{θ} and Q_{target}
 Update the θ parameters of the Q_{θ} regressor through backpropagation using an Adam optimizer
 Update the Whittle index for all states x as (13)
 if $n\%50 = 0$ **then**
 Copy the θ parameters from the main Q_{θ}^x neural network into the secondary $Q_{\theta'}^x$ neural network
 end if
 end for
 end if
end for

$$h_e \leftarrow \begin{cases} h_e + \nabla_{\theta} \ln(\sigma_m(f_{\theta}(s[t]) - \lambda)) & \text{if } a[t] = 1 \\ h_e + \nabla_{\theta} \ln(1 - \sigma_m(f_{\theta}(s[t]) - \lambda)) & \text{if } a[t] = 0 \end{cases} \quad (15)$$

for each sample at time t within the episode. After all episodes in the mini-batch are completed, the neural network parameters are updated as follows:

$$\theta \leftarrow \theta + L_b \sum_e (G_e - \bar{G}_b) h_e, \quad (16)$$

where G_e represents the discounted net rewards, G_b is the average of these rewards, and L_b is the learning rate.

In order to compare the convergence speed of all these algorithms to the desired policy, we have decided to represent each NeurWIN update as a unique transition, so that at each iteration all algorithms are updated simultaneously. It is also important to highlight that, unlike in QWI and QWINN, the training and execution phases are separate in NeurWIN, i.e., we first need to learn the indices of each arm separately.

3.3.2 Q -learning and DQN algorithms. We also consider the vanilla implementations of Q -learning and DQN [19]. Q -learning is known to converge to the optimal policy, [24], but it suffers from the “curse of dimensionality”. DQN does not have performance guarantees, but in practice performs better thanks to the interpolation of the state information.

As with QWINN, DQN training is performed through batches of samples stored in memory, where the input to the neural network is the combination of arm states, while the outputs are the Q -values for each possible combination of

arm activations. In the following results, Q -learning uses $\alpha = \frac{1}{\#N(s,a)}$ as the learning step size, where $\#N(s,a)$ is the number of times that said state (s,a) has been visited, while DQN has been trained using a fixed learning rate $lr = 0.001$.

4 NUMERICAL RESULTS

In this section we compare the performance of our algorithms QWI and QWINN with respect to Q -learning, DQN and NeurWIN. We consider three different RMABPs used in previous literature: the "restart problem" proposed in [2], the "deadline scheduling problem" studied in [28] and the "circular environment problem" in [5]. An interesting feature of these examples is that the Whittle index can be calculated in closed-form, which allows us to assess the accuracy of the estimates of the Whittle indices obtained with the various algorithms.

We recall that Whittle index policy is a heuristic, i.e., it does not characterize the optimal solution to an RMABP, even though it has often been reported that its sub-optimality gap is very small. For this reason, we consider as a baseline the performance of the optimal policy π^* wherever the latter can be computed. We recall that the optimal policy π^* solves Bellman's Optimality Equation, see for example [18],

$$\begin{aligned} V_{\pi^*}(s) &= \max_{a \in \{0,1\}^N} \left(r(s,a) + \gamma \sum_{s'} p(s'|s,a) V_{\pi^*}(s') \right), \\ \text{s.t. } \sum_{i=1}^N a^i(t) &= m, \forall t, \end{aligned} \tag{17}$$

where $s = (s^1, \dots, s^N)$ denotes the states of all the arms, $r(s,a) = \sum_{i=1}^N r^i(s^i, a^i)$, and $p(s'|s,a) = \prod_{i=1}^N p^i(s'_i | s_i, a_i)$. We solve equation (17) by value Iteration.

We also use Bellman's equation in order to assess the performance of each algorithm during its training for small enough problems. Let π denote the current policy of each one of the algorithms at a given time. Then, the value function $V_{\pi}(s)$ characterizing its performance is the solution of Bellman's equation

$$V_{\pi}(s) = \sum_a \pi(a|s) \left(r(s,a) + \gamma \sum_{s'} p(s'|s,a) V_{\pi}(s') \right), \tag{18}$$

where $\pi(a|s)$ is the probability of taking action $a \in \{0,1\}^N$ in state s . We also solve Equation (18) using value Iteration.

In order to compare the performance of a policy π with respect to the optimal policy π^* we use the Bellman Relative Error defined as

$$BRE(\pi, \pi^*) = \frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} \frac{|V_{\pi}(s) - V_{\pi^*}(s)|}{V_{\pi^*}(s)}. \tag{19}$$

This metric will be particularly useful in subsections 4.1, 4.2 and 4.4, where the reduced state space allows us to easily compute it, and to study the convergence of QWI, QWINN, NeurWIN, Q -learning and DQN to the optimal policy. Throughout this section, we will denote by π_n^P , $P \in \{QWI, QWINN, DQN, Q\}$, the Whittle index heuristic estimated by algorithm P at time n . We consider a discount factor $\gamma = 0.9$ and an exploration parameter $\epsilon = 1$ (i.e., pure off-policy mode).

4.1 Restart problem (homogeneous arms)

In this section we consider a "restart problem" with state space $S = \{0, 1, 2, 3, 4\}$. From any state, the active action ($a = 1$) brings the arm to the initial state with probability 1. i.e., $p(0|s, 1) = 1$, for all s . The transitions probability matrix in the

case of passive action is:

$$P_0 = \begin{pmatrix} 1-x & x & 0 & 0 & 0 \\ 1-x & 0 & x & 0 & 0 \\ 1-x & 0 & 0 & x & 0 \\ 1-x & 0 & 0 & 0 & x \\ 1-x & 0 & 0 & 0 & x \end{pmatrix}.$$

The expected conditional reward is given by $r(s, 0) = y^{s+1}$ for passive actions, while $r(s, 1) = 0$ for active actions ($a = 1$). We consider the following setting: $N = 5$ arms; one arm $M = 1$ is activated every time epoch; and $x = y = 0.9$. The Whittle index can be analytically calculated in this problem, yielding $\lambda(0) = -0.9$, $\lambda(1) = -0.7371$, $\lambda(2) = -0.5373$, $\lambda(3) = -0.3188$ and $\lambda(4) = -0.0939$ while using a discount factor $\gamma = 0.9$. We note that the Whittle index is increasing as the state increases, which is consistent with the reward structure. We note that from the description of the problem, a good policy will tend to keep all arms around state 0, as the rewards are larger there.

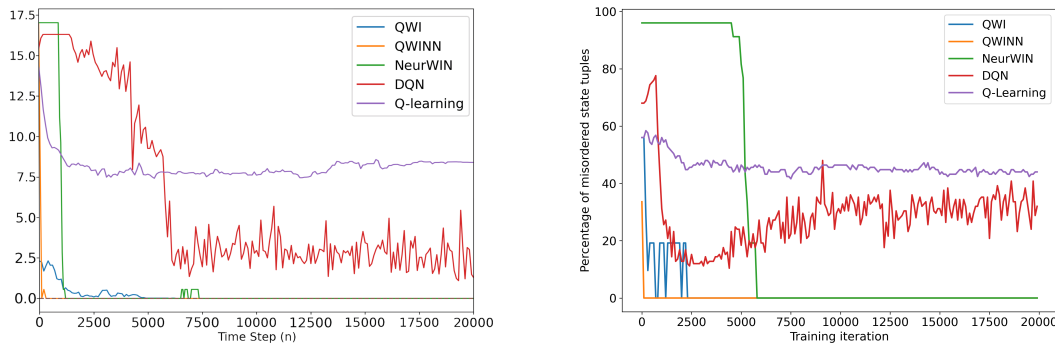


Fig. 2. Performance graphs for the “homogeneous restart” problem. (left) Bellman Relative Error $BRE(\pi_n^P)$, $P \in \{QWI, QWINN, NeurWIN, DQN, Q-learning\}$ during training for the “homogeneous restart” problem, $N = 5, M = 1, |S| = 5$; and (right) Percentage of states in which an optimal action is not performed in the “restart” problem for homogeneous arms

In Figure 2 (left) we depict the Bellman Relative Error for all the algorithms, namely, we depict $BRE(\pi_n^P, \pi^*)$, $P \in \{QWI, QWINN, NeurWIN, Q-learning, DQN\}$, for all the time steps n , while Figure 2 (right) represents the percentage of states in which an optimal action is not taken. We observe that QWI, NeurWIN and QWINN learn the optimal policy, QWINN being the fastest. We note that even though Whittle index policy need not be optimal, it might happen in such a simple problem. On the other hand, Q-learning does not learn the optimal policy, even though we know from theory that as $n \rightarrow \infty$, it is the only algorithm among the approaches considered here for which convergence to the optimal policy is guaranteed [24]. The vanilla DQN implementation does better, but it requires longer to converge to the optimal policy.

An important observation is that the performance of an index policy, in particular the values plotted in Figure 2 (left), only depend on the relative ordering between the states, and not on the precise value of the Whittle index. We thus explore the accuracy of the estimates of QWI/QWINN (see Figure 3a) and NeurWIN (see Figure 3b) during their training. More specifically, in Figure 3a we plot the averaged (across the 5 arms) estimates of Whittle indices. In the case of

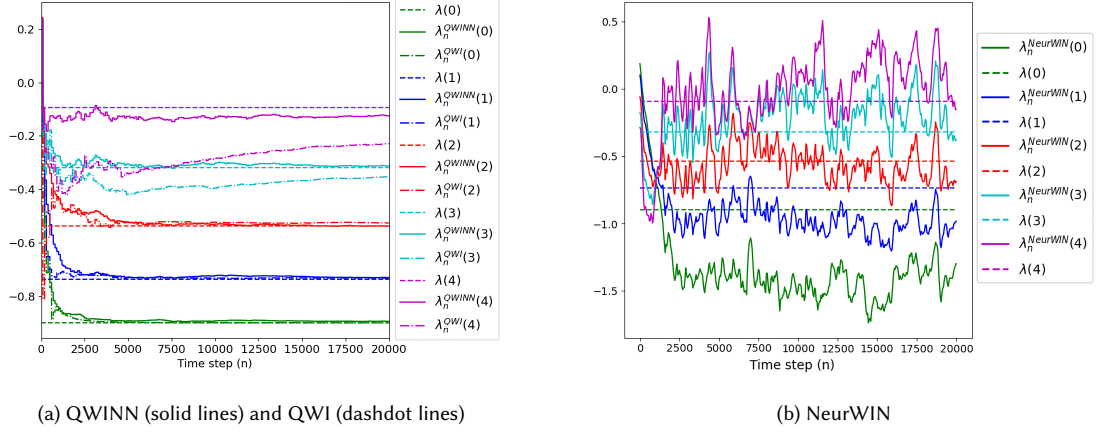


Fig. 3. Evolution of the Whittle index estimates for the restart problem

NeurWIN, there is a single agent cloned for the 5 arms, and therefore, in Figure 3b we depict the estimate of this single agent.

In Figure 3a we note that the QWI's estimates for the indices of states 0,1 and 2 converge, whereas the estimates of states 3 and 4 need longer simulations due to the fact that these states are visited less frequently, while QWINN converges rapidly to the theoretical values of all the indices. It is not surprising that the estimates for states 3 and 4 require longer as they are visited less frequently. We also observe that the estimates are stable, i.e, there are no fluctuations. In Figure 3b we note that the ordering with NeurWIN is correct, but that the fluctuations do not vanish. The latter might suggest that in the case of heterogeneous arms, NeurWIN might swap the ordering of the states, which might yield a sub optimality gap larger than with QWI and QWINN. In Section 4.2 we show that this is indeed the case.

4.2 Restart problem (heterogeneous arms)

In this section we want to investigate the impact that errors in the estimation of the Whittle index has on the performance of the arms. In order to do so, we consider a restart problem with $N = 3$ arms, each one having a different set of parameters, and as before we consider $M = 1$. As in the previous section, we assume that for every arm $x = y$, and that this value is taken uniformly at random from the set $\{0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$. We simulate 20 such instances, and to assess the sub-optimality gap we consider the Bellman Averaged Relative Error of algorithm P defined as:

$$BRE(P) = \frac{1}{20} \sum_{k=1}^{20} BRE_k(\pi_k^P, \pi_k^*), \quad (20)$$

where $BRE_k(\pi_k^P, \pi_k^*)$ denotes the Bellman Relative Error for instance k , π_k^* is the optimal policy for instance k , and π_k^P is the Whittle Index policy for instance k obtained with algorithm P .

In Figure 4 (right) we plot the evolution of $BRE(P)$ for each of the algorithms during the training. We note that the state space is smaller than in the previous section, and as a consequence Q -learning converges faster, and that DQN essentially learns the optimal policy. Regarding QWI and QWINN, we observe that whereas the relative errors of QWI and QWINN are negligible and indistinguishable from DQN, NeurWIN incurs a tangible error of around 3%. The latter

is due to lack of accuracy in the estimation of Whittle index, which can induce a wrong ordering among the states. One way to show convergence to the correct index policy is to plot, for all possible combinations of states, what percentage of these do not perform the same action that Whittle index policy would render. In other words, during each instant of training, and for each possible combination of states among the different arms of the problem, we evaluate what action each algorithm’s policy would perform against the Whittle index policy corresponding to that problem. In Figure 4 (left) we plot the percentage of state combinations in which the theoretical Whittle index action is not taken. For the restart problem with heterogeneous arms, using $N = 3$, the total number of state combinations is 5^3 . As we can see, while QWINN reaches the theoretical Whittle index policy for most of the states in the first iterations, the error of QWI is around 10% whereas that of NeurWIN is around 20%. The former is due to the difficulty of visiting the higher states of the problem while the latter is due to inaccurate index predictions, leading to sub-optimal policies especially in heterogeneous cases.

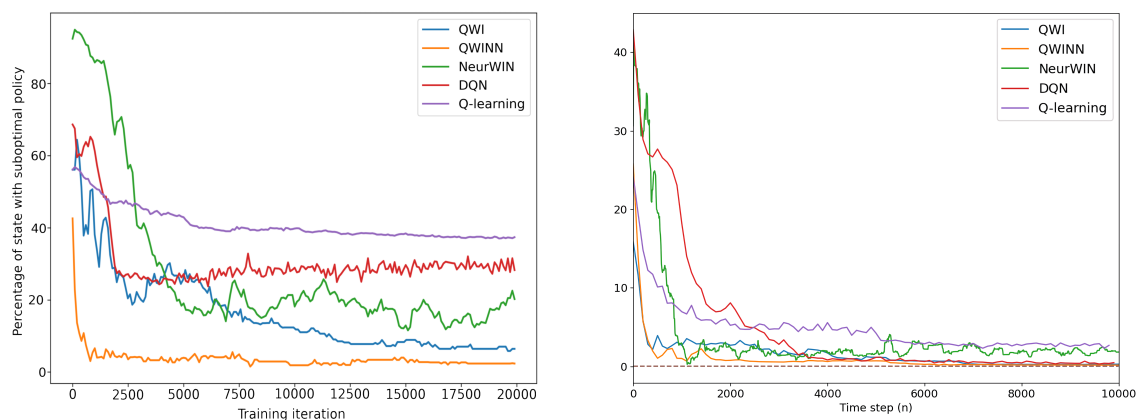


Fig. 4. Performance graphs for the QWI, QWINN, NeurWIN, DQN and Q-learning algorithms: assignment of optimal policies in the heterogeneous “restart” problem (left) and discounted rewards for the homogeneous “deadline scheduling” problem (right).

4.3 Deadline scheduling problem

In this section we consider the deadline scheduling problem studied in [13], and for which the source code is publicly available. This problem, proposed in [28], has states formed by two different variables: The service time $B \in [0, 9]$ and the deadline $T \in [0, 12]$, leading to a total of $|S| = 130$ states (although several of these states are not accessible by the Markov chain). Therefore, the variable B represents the amount of workload pending to complete a certain job while the variable T represents the remaining time available to perform it. When there is no job at the i th position, the state is $(0, 0)$, while otherwise it is (T_n^i, B_n^i) . In each iteration, the state transition $s_n^i = (T_n^i, B_n^i)$ depends on the action a_n^i performed:

$$s_{n+1}^i = \begin{cases} (T_n^i - 1, (B_n^i - a_n^i)^+) & \text{if } T_n^i > 1, \\ (T, B) \text{ with prob. } Q(T, B) & \text{if } T_n^i \leq 1, \end{cases}$$

where $b^+ = \max(b, 0)$. When $T = 1$, the deadline for completing the job ends, and the new state (including the empty state $(0, 0)$) is chosen uniformly at random. The value of B , the workload to be completed, is only reduced if the action

is active. If the scheduler reaches the state $(T = 1, B > 0)$, the job could not be finished on time, and an extra penalty $F(B_n^i - a_n^i) = 0.2(B_n^i - a_n^i)^2$ is incurred. In addition, activating the arms involves a fixed cost $c = 0.8$, resulting in the following rewards:

$$r_n^i(s_n^i, a_n^i, c) = \begin{cases} (1-c)a_n^i & \text{if } B_n^i > 0, T_n^i > 1, \\ (1-c)a_n^i - F(B_n^i - a_n^i) & \text{if } B_n^i > 0, T_n^i = 1, \\ 0 & \text{otherwise} \end{cases}$$

In [28] the authors provide an expression for Whittle index, given by:

$$\lambda(T, B, c) = \begin{cases} 0 & \text{if } B = 0, \\ 1 - c & \text{if } 1 \leq B \leq T - 1, \\ \gamma^{T-1}F(B - T + 1) - \gamma^{T-1}F(B - T) + 1 - c & \text{if } T \leq B \end{cases}$$

In the following results we will consider the homogeneous case using $N = 5$ and $M = 2$, with a processing cost $c = 0.8$ for all arms and a discount factor $\gamma = 0.9$.

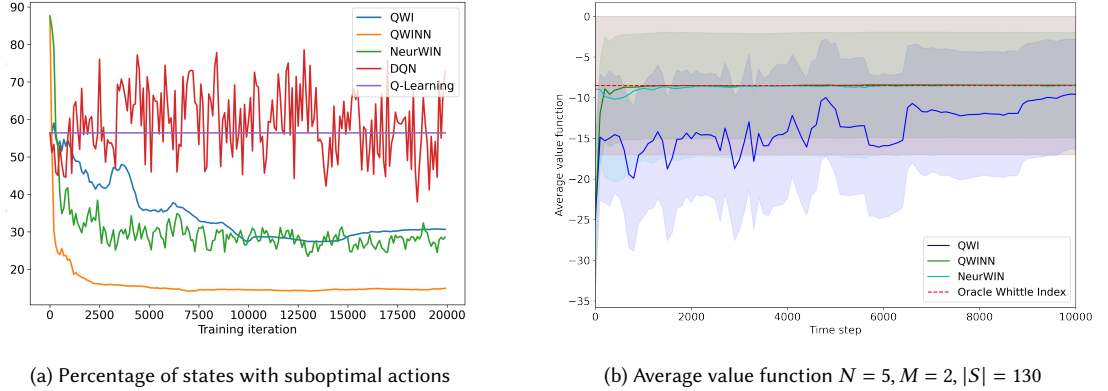


Fig. 5. Performance graphs for “deadline scheduling problem” using $N = 5, M = 2, |S| = 130$.

In Figure 5a we plot the percentage of state combinations in which the action of an algorithm would diverge from that prescribed by Whittle index policy. In other words, this happens when the algorithm fails to identify the best arm. We observe that the QWINN is the algorithm with the smallest number of misorderings, and that both QWI and NeurWIN fail to identify the best arm in around 30% of the states. By inspection, we have observed that QWI’s estimates are poor for states with deadline $T > 9$, which are rarely visited. This is a clear example of the limitations of a tabular algorithm when solving a problem with a large state space: by not having enough samples of these states, the tabular QWI algorithm is not able to obtain good predictions of the Whittle indices. On the other hand, we have not identified a particular set of states in which NeurWIN and QWINN fail.

However, it is important to highlight that these misorderings do not necessarily imply a significant loss of performance. In Figure 5b, we present the average value function for the QWI, QWINN, and NeurWIN algorithms during training, comparing them to the performance of the Whittle index Oracle policy for $N = 5$ and $M = 2$. As observed, while QWI takes longer to reach performance comparable to QWINN and NeurWIN, the latter two achieve close to identical

performance to the Whittle index Oracle policy within the first few hundred iterations, indicating that the suboptimal states of NeurWIN do not significantly affect the evaluation of its policy. The performance of DQN and Q-learning has not been included in this analysis as neither is able to learn an efficient policy due to the size of the coupled state space.

In Figure 6 we extend our analysis to a larger number of arms, $N = 100$ while using $T \in [0, 17]$ and $B \in [0, 15]$, resulting in a total of $|S| = 255$ states, with different activation levels $M = \{20, 50\}$. Due to the increased state space, we limit our analysis to only QWINN and NeurWIN algorithms. In both Figure 6a for $M = 20$ and Figure 6b for $M = 50$, QWINN has a higher average performance with respect to NeurWIN, especially in the case of $M = 20$, where QWINN’s results are more stable. The difference between the two algorithms is smaller for the $M = 50$ case as the agent is saturated by having to choose more arms, reducing the difference between the policies of the two algorithms.

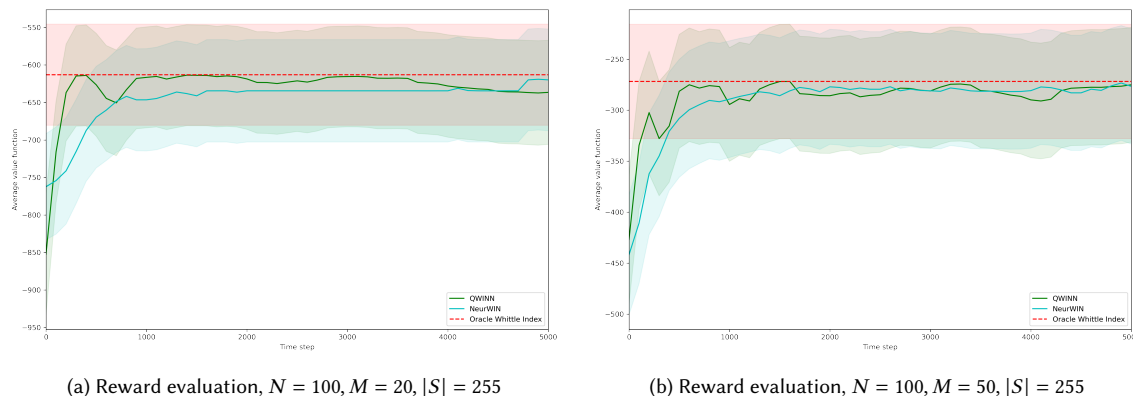


Fig. 6. Performance graphs for “deadline scheduling” problem with $|S| = 255$ states and $N = 100$ arms using $M = \{20, 50\}$ active arms

Lastly, let us examine the scenario of heterogeneous arms, where each arm has a state space of size $|S| = 130$ and $N = 100$. To differentiate the arms, we assign different activation cost values $c = \{0.1, 0.3, 0.6, 0.8\}$ for every 25 arms. Figure 7 illustrates the convergence of the algorithms. QWI falls short of reaching an optimal policy, while both QWINN and NeurWIN converge to Whittle’s index Oracle policy. However, NeurWIN experiences performance perturbations caused by inaccurate index predictions, which become more pronounced with an increasing number of heterogeneous arms. In contrast, QWINN converges smoothly to the optimal policy without such performance fluctuations. It is worthy to note that while NeurWIN learns and then commits, QWINN operates in pure on-line mode.

4.4 Circular problem

Finally, we consider here the “circular” problem, with a state space $S = \{0, 1, 2, 3\}$ [5]. In this problem, with an active action the process remains in its current state with probability 0.6, or increments positively with probability 0.4. Similarly, with a passive action the process remains in its current state with probability 0.6, or decrements negatively with probability 0.4. It is called circular, as an increment (decrement) from state 3 (0) brings the process to state 0 (3). The reward function does not depend on the action performed, but only on the state, being $R(0) = -1, R(1) = R(2) = 0, R(3) = 1$. Using a discount parameter of $\gamma = 0.9$, the theoretical values of Whittle index for each state are $\lambda(0) = -0.4390, \lambda(1) = 0.4390, \lambda(2) = 0.8652, \lambda(3) = -0.8652$, and we note that the best state is 1, and the second best 2.

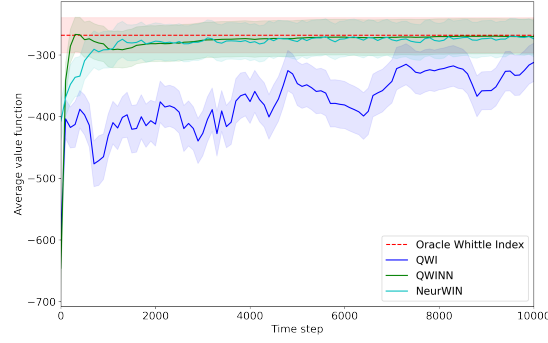


Fig. 7. Reward evaluation of the NeurWIN and QWINN algorithms for a heterogeneous deadline scheduling problem with a state space of size $|S| = 130$, $N = 100$ arms, and $M = 25$. Each of the four groups of 25 arms has a distinct cost (0.1, 0.3, 0.6, 0.8)

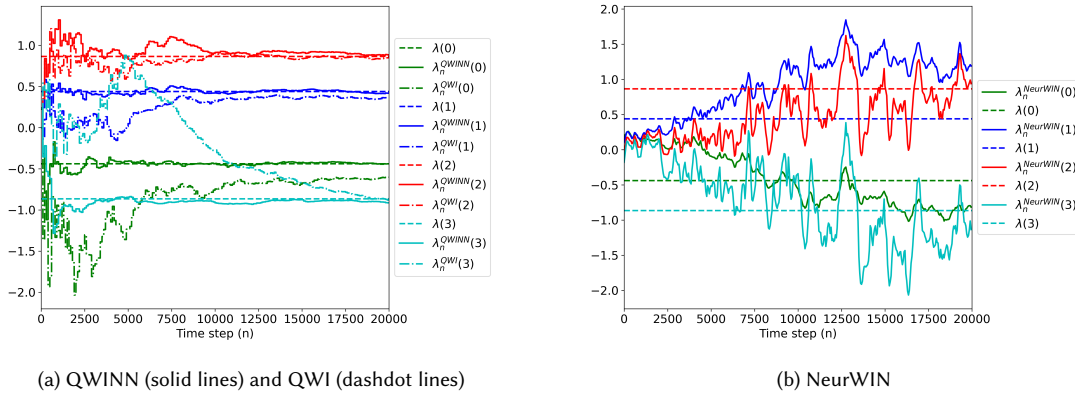
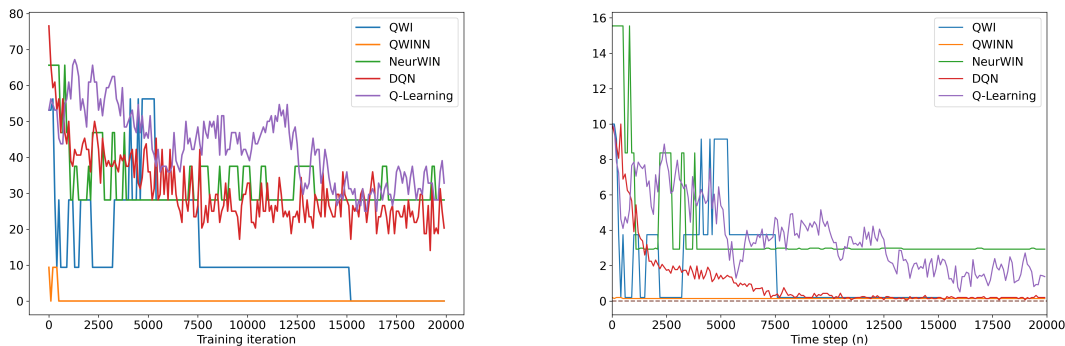


Fig. 8. Evolution of the Whittle index estimates for the circular problem

In Figure 8, we plot the Whittle index estimates for the QWI, QWINN and NeurWIN algorithms for this problem. We note that QWI and QWINN order correctly all states, but that with NeurWIN the estimate of the index of state 2 is consistently larger than that of state 1. We also note that QWINN's estimates converge the fastest. In Figure 9a we plot the percentage of states in which the optimal action for the problem is not realized employing $N = 3$ and $M = 1$ with $\gamma = 0.9$ in the QWI, QWINN, NeurWIN and classical DQN and Q-learning algorithms. Of all these algorithms, only QWINN and later QWI, are able to obtain an optimal action for all possible combinations of problem states. The rest of the algorithms remain with an error rate of 30%, due in part to the difficulty of obtaining sufficient samples for all possible states in Q-learning, a poor regression in the case of DQN or a poor assignment of the indices in the case of NeurWIN, as can be seen in the Figure 8b. The effect of these policies can be seen in Figure 9b, where we plot the Bellman Relative Error $BRE(P)$ for QWI, QWINN, NeurWIN, DQN and Q-learning. As can be seen, none of the index heuristic algorithms achieves an optimal policy, since for this problem with this number of arms the optimal policy does not coincide with Whittle's index policy. Nevertheless, QWI and QWINN are able to achieve very good performance with respect to the optimal policy, especially QWINN which achieves this policy in the first few hundred iterations.



(a) Percentage of states in which an optimal action is not performed in the “circular” problem with homogeneous arms
 (b) Bellman Relative Error $BRE(\pi_n^P)$, $P \in \{QWI, QWINN, NeurWIN, DQN, Q-Learning\}$

Fig. 9. Performance graphs for “circular problem” with $N = 3, M = 1, |S| = 4$.

On the other hand, due to NeurWIN’s misordering of the indices, it shows a larger sub-optimality gap throughout the training. Q-learning and especially DQN perform very well, despite their poor state assignments, because the regret of these assignments is relatively small.

Figure 10 showcases the discounted value function achieved by the QWI, QWINN, and NeurWIN algorithms for different scenarios. In Figure 10a, we focus on the circular dynamics problem with a state space of $|S|=4, N=100$, and $M=20$. Figures 10b and 10c explore cases where we expand the state space to $|S|=50$, with only the first and last states having non-zero rewards (-1 and +1, respectively), along with $N=100, M=50$ (Figure 10b) and $N=100, M=70$ (Figure 10c).

In all these cases, NeurWIN struggles to generate an effective policy due to the incorrect ordering of its indices. Notably, in Figure 10a, where the state space is small, QWI achieves a Whittle index Oracle policy before QWINN. However, for the larger and sparser state spaces in Figures 10b and 10c, QWINN outperforms QWI by a significant margin.

5 CONCLUSIONS

In this paper we have developed two algorithms, QWI and QWINN, capable of learning the Whittle index for the total discounted criterion. We have presented analytical results about the asymptotic behavior of QWI as well as DQN and QWINN, the latter being the first of its kind. For QWINN, our analysis rigorously establishes the local convergence of the DQN algorithm under specific conditions. By strengthening assumptions on the Hessian’s positive definiteness and applying stochastic approximation theory, we demonstrate that the DQN algorithm reliably converges to a neighborhood of a local minimum, particularly when iterates start close to this minimum and the algorithm incorporates a reset mechanism. Through numerical simulations, we have compared the performance of QWI, QWINN with respect to other relevant algorithms, namely Q-learning, DQN and NeurWIN. Our results show that QWI, QWINN, NeurWIN are all more efficient than Q-learning in obtaining policies with near-optimal performance. QWI estimates very accurately the Whittle index in problems of small to moderate size. QWINN is able to accurately learn Whittle indices for problems with larger state spaces, and it is typically the fastest algorithm to converge. NeurWIN’s estimates

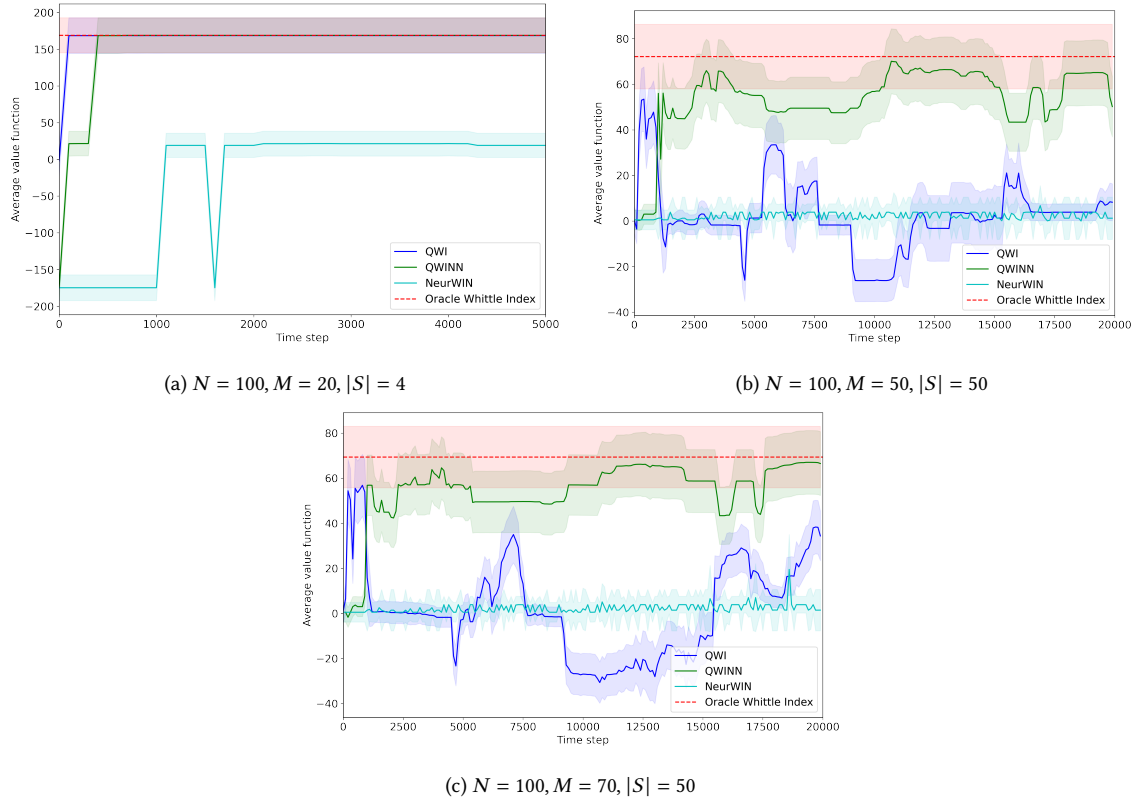


Fig. 10. Evaluation of the rewards obtained from the index policy for the circular dynamics problem for the QWI, QWINN and NeurWIN algorithms.

are not always accurate, but the latter does not necessarily imply a loss in performance. On the other hand, a simple problem like the circular problem shows that a wrong ordering can lead to a substantial performance degradation.

An interesting and challenging direction for future research is the case of restless bandits with partial observations. The neural network based approach will be particularly useful there in order to mitigate the complexity of the belief state. Regret analysis in this framework is another research direction that needs to be addressed.

ACKNOWLEDGEMENTS

We would like to express our gratitude to the reviewers and editors, who provided invaluable guidance to improve the paper.

F. Robledo and U. Ayesta have received funding from the Department of Education of the Basque Government through the Consolidated Research Group MATHMODE (IT1456-22). Research partially supported by the French "Agence Nationale de la Recherche (ANR)" through the project ANR-22-CE25-0013-02 (ANR EPLER) and DST-Inria Cefipra project LION. V. Borkar has received funding from the S. S. Bhatnagar Fellowship from Council of Scientific and Industrial Research, Government of India and Google Research.

REFERENCES

- [1] Jinane Abounadi, Dimitrib Bertsekas, and Vivek S Borkar. 2001. Learning algorithms for Markov decision processes with average cost. *SIAM Journal on Control and Optimization* 40, 3 (2001), 681–698. https://epubs.siam.org/doi/abs/10.1137/S0363012999361974?casa_token=1DzrZBygwIAAAAA:1R6PB1DKopfx50oH5D9xAql38TAL8XbUL1TY3k8FJkc5NrkemLARZoiEd1paX3bA6zMAiHikLZWiw
- [2] Konstantin E. Avrachenkov and Vivek S. Borkar. 2022. Whittle index based Q-learning for restless bandits with average reward. *Automatica* 139 (May 2022), 110186. <https://doi.org/10.1016/j.automatica.2022.110186>
- [3] Vivek S Borkar et al. 2008. Stochastic approximation: a dynamical systems viewpoint. *Cambridge Books* (2008).
- [4] Michael O Duff. 1995. Q-learning for bandit problems. In *Machine Learning Proceedings 1995*. Elsevier, 209–217. <https://www.sciencedirect.com/science/article/pii/B9781558603776500347>
- [5] Jing Fu, Yoni Nazarathy, Sarat Moka, and Peter G Taylor. 2019. Towards Q-learning the Whittle Index for Restless Bandits. In *2019 Australian & New Zealand Control Conference (ANZCC)*. IEEE, 249–254. <https://doi.org/10.1109/anzcc47194.2019.8945748>
- [6] Lachlan J. Gibson, Peter Jacko, and Yoni Nazarathy. 2021. A Novel Implementation of Q-Learning for the Whittle Index. In *Performance Evaluation Methodologies and Tools (Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering)*, Qianchuan Zhao and Li Xia (Eds.). Springer International Publishing, Cham, 154–170. https://doi.org/10.1007/978-3-030-92511-6_10
- [7] J.C. Gittins, K. Glazebrook, and R. Weber. 2011. *Multi-armed Bandit Allocation Indices*. Wiley.
- [8] Yu-Pin Hsu. 2018. Age of Information: Whittle Index for Scheduling Stochastic Arrivals. In *2018 IEEE International Symposium on Information Theory (ISIT)*. 2634–2638. <https://doi.org/10.1109/ISIT.2018.8437712> ISSN: 2157-8117.
- [9] Jackson A. Killian, Arpita Biswas, Sanket Shah, and Milind Tambe. 2021. Q-Learning Lagrange Policies for Multi-Action Restless Bandits. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining (KDD '21)*. Association for Computing Machinery, New York, NY, USA, 871–881. <https://doi.org/10.1145/3447548.3467370>
- [10] Chandrashekar Lakshminarayanan and Shalabh Bhatnagar. 2017. A stability criterion for two timescale stochastic approximation schemes. *Automatica* 79 (2017), 108–114. <https://doi.org/10.1016/j.automatica.2016.12.014>
- [11] Keqin Liu and Qing Zhao. 2010. Indexability of Restless Bandit Problems and Optimality of Whittle Index for Dynamic Multichannel Access. *IEEE Transactions on Information Theory* 56, 11 (Nov. 2010), 5547–5567. <https://doi.org/10.1109/TIT.2010.2068950>
- [12] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *nature* 518, 7540 (2015), 529–533.
- [13] Khaled Nakhleh, Santosh Ganji, Ping-Chun Hsieh, I-Hong Hou, and Srinivas Shakkottai. 2021. NeurWIN: Neural Whittle Index Network For Restless Bandits Via Deep RL. In *Advances in Neural Information Processing Systems*, Vol. 34. Curran Associates, Inc., 828–839. <https://proceedings.neurips.cc/paper/2021/hash/0768281a05da9f27df178b5c39a51263-Abstract.html>
- [14] José Niño-Mora. 2020. A verification theorem for threshold-indexability of real-state discounted restless bandits. *Mathematics of Operations Research* 45, 2 (2020), 465–496. <https://doi.org/10.1287/moor.2019.0998>
- [15] Ronald Ortner, Daniil Ryabko, Peter Auer, and Rémi Munos. 2012. Regret bounds for restless markov bandits. In *International conference on algorithmic learning theory*. Springer, 214–228. https://link.springer.com/chapter/10.1007/978-3-642-34106-9_19
- [16] Tejas Pagare, Vivek Borkar, and Konstantin Avrachenkov. 2023. Full Gradient Deep Reinforcement Learning for Average-Reward Criterion. In *Learning for Dynamics and Control Conference*. PMLR, 235–247.
- [17] C.H. Papadimitriou and J.N. Tsitsiklis. 1999. The Complexity of Optimal Queueing Network. *Mathematics of Operations Research* 24, 2 (1999), 293–305.
- [18] Martin L. Puterman. 2014. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons. Google-Books-ID: VvB-jBAAQBAJ.
- [19] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- [20] Hado Van Hasselt. 2010. Double Q-learning. In *Advances in Neural Information Processing Systems*, Vol. 23. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2010/hash/091d584fcd301b442654dd8c23b3fc9-Abstract.html>
- [21] Hado Van Hasselt, Arthur Guez, and David Silver. 2016. Deep Reinforcement Learning with Double Q-Learning. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 30. <https://ojs.aaai.org/index.php/AAAI/article/view/10295>
- [22] I. M. Verloop. 2016. Asymptotically optimal priority policies for indexable and nonindexable restless bandits. *The Annals of Applied Probability* 26, 4 (Aug. 2016), 1947–1995. <https://doi.org/10.1214/15-AAP1137>
- [23] Siwei Wang, Longbo Huang, and John C. S. Lui. 2020. Restless-UCB, an Efficient and Low-complexity Algorithm for Online Restless Bandits. In *Advances in Neural Information Processing Systems*, Vol. 33. Curran Associates, Inc., 11878–11889. <https://proceedings.neurips.cc/paper/2020/hash/89ae0fe22c47d374bc9350ef99e01685-Abstract.html>
- [24] Christopher JCH Watkins and Peter Dayan. 1992. Q-learning. *Machine learning* 8, 3-4 (1992), 279–292. <https://link.springer.com/content/pdf/10.1007/BF00992698.pdf>
- [25] Richard R Weber and Gideon Weiss. 1990. On an index policy for restless bandits. *Journal of applied probability* (1990), 637–648. <https://doi.org/10.2307/3214547>
- [26] Peter Whittle. 1988. Restless bandits: Activity allocation in a changing world. *Journal of Applied Probability* 25, A (Jan. 1988), 287–298. <https://doi.org/10.2307/3214163>

- [27] Guojun Xiong and Jian Li. 2023. Finite-Time Analysis of Whittle Index based Q-Learning for Restless Multi-Armed Bandits with Neural Network Function Approximation. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- [28] Zhe Yu, Yunjian Xu, and Lang Tong. 2018. Deadline scheduling as restless bandits. *IEEE Trans. Automat. Control* 63, 8 (2018), 2343–2358. <https://ieeexplore.ieee.org/abstract/document/8295041/>

A PROOF OF THEOREM 3.1

The following assumptions are required:

(C1) Step sizes $\{\alpha(n)\}$ satisfy, for $x \in (0, 1)$,

$$\sup_n \frac{\alpha(\lfloor xn \rfloor)}{\alpha(n)} < \infty,$$

$$\sup_{y \in [x, 1]} \left| \frac{\sum_{m=0}^{\lfloor yn \rfloor} \alpha(m)}{\sum_{m=0}^n \alpha(m)} - 1 \right| \rightarrow 0$$

(C2) The problem is Whittle indexable, that is, given the subset of states $\Pi(\lambda)$ where the optimal action is the passive action due to the extra reward λ of this action, the number of states within this subset $\Pi(\lambda)$ must grow monotonically as λ increases.

We will follow the general scheme proposed in [10] and [3]. Consider a generalised version of equations (8) and (9) such that:

$$x_{n+1} = x_n + a(n) \left[h(x_n, y_n) + M_{n+1}^{(1)} \right] \quad (21)$$

$$y_{n+1} = y_n + b(n) \left[g(x_n, y_n) + M_{n+1}^{(2)} \right] \quad (22)$$

where (21) represents Equation (8), the fast time-scale where we compute the Q -values, and (22) represents Equation (9), the slower time-scale in which the Whittle indices are updated. In these equations, the functions h and g are continuous Lipschitz functions, the M_n are martingale difference sequences representing noise terms, and $a(n)$ and $b(n)$ are step-size terms satisfying $\frac{b(n)}{a(n)} \rightarrow 0$ as $n \rightarrow \infty$, in addition to the usual conditions $\sum_n a(n) = \sum_n b(n) = \infty$ and $\sum_n a(n)^2, \sum_n b(n)^2 < \infty$.

First, let us define $F_{su}^\lambda(\Psi(j, b))$ and $M_{n+1}(s, u)$ such that:

$$F_{su}^\lambda(\Psi(j, b)) = (1 - u)(R_0(s) + \lambda) + uR_1(s) + \gamma \sum_j p(j|i, u) \max_{v \in \{0,1\}} \Psi(j, v) \quad (23)$$

$$M_{n+1}(s, u) = (1 - u)(R_0(s) + \lambda_n(x)) + uR_1(s) + \max_{v \in \{0,1\}} Q_n(x_{n+1}, v) - F_{su}^{\lambda_n(x)}(Q_n) \quad (24)$$

We can now rewrite the Equation (8) as:

$$Q_{n+1}^x(s, u) = Q_n^x(s, u) + \alpha(n) \left[F_{su}^{\lambda_n(x)}(Q_n) - Q_n + M_{n+1}(s, u) \right] \quad (25)$$

Comparing equations (8) and (25) we make the correspondence $a(n) = \alpha(n)$, $h(x_n, y_n) = F_{su}^{\lambda_n(x)}(Q_n) - Q_n$, where $x_n = Q_n$ and $y_n = \lambda_n$ are the Q -value and Whittle index estimate respectively and $M_{n+1}(s, u)$ is the martingale difference sequence $M_{n+1}^{(1)}$. On the other hand, equations (9) and (22) correspond to $b(n) = \beta(n)$, $g(x_n, y_n) = Q_n^x(x, 1) - Q_n^x(x, 0)$ and a martingale difference sequence $M_{n+1}^{(2)} = 0$.

As in [3], we will consider three necessary conditions for equations (21) and (22) to be stable and converge to their optimal values $x_n \rightarrow x^*$ and $y_n \rightarrow y^*$.

A1 h and g must be Lipschitz continuous.

A2 $\{M_n^{(1)}\}$ and $\{M_n^{(2)}\}$ are martingale difference sequences.

A3 $\{a(n)\}$ and $\{b(n)\}$ satisfy:

- $a(n) > 0, b(n) > 0$

- $\sum_n a(n) = \sum_n b(n) = \infty, \sum_n (a(n)^2 + b(n)^2) < \infty$
- $\frac{b(n)}{a(n)} \rightarrow \infty$

The full proof of A1 can be easily verified on page 687 of [1]. In our notation, $M_{n+1}(s, u)$ corresponds to $M_{n+1}^{(1)}$ and $M_{n+1}^{(2)} = 0$, thus satisfying condition A2. Finally, A3 is also verified given that $\beta(n) = o(\alpha(n))$.

Let us assume that the equations (8) and (9) are bounded. We will prove this condition later. First we will rewrite the equation for the calculation of the indices (9) as:

$$\lambda_{n+1}(x) = \lambda_n(x) + \alpha(n) \left(\frac{\beta(n)}{\alpha(n)} \right) (Q_n^x(x, 1) - Q_n^x(x, 0)) \quad (26)$$

Let $\tau(n) = \sum_{m=0}^n \alpha(m)$, $m \geq 0$. We define $\bar{Q}(t)$, $\bar{\lambda}(t)$ as the interpolation of the trajectories of Q_n^x and $\lambda_n(x)$ on each interval $[\tau(n), \tau(n+1)]$, $n \geq 0$ as:

$$\bar{Q}(t) = Q(n) + \left(\frac{t - \tau(n)}{\tau(n+1) - \tau(n)} \right) (Q(n+1) - Q(n)) \quad (27)$$

$$\begin{aligned} \bar{\lambda}(t) &= \lambda(n) + \left(\frac{t - \tau(n)}{\tau(n+1) - \tau(n)} \right) (\lambda(n+1) + \lambda(n)) \\ t &\in [\tau(n), \tau(n+1)] \end{aligned} \quad (28)$$

which track the asymptotic behavior of the coupled o.d.e.s

$$\dot{Q}(t) = h(Q(t), \lambda(t)), \dot{\lambda} = 0$$

where the latter is a consequence of $\frac{\beta(n)}{\alpha(n)} \rightarrow 0$ in (26). From the reference frame of $Q(t)$, $\lambda(\cdot)$ is a constant of value λ' . Because of this, the first o.d.e. becomes $\dot{Q} = h(Q(t), \lambda')$, which is well posed and bounded, and has an asymptotically stable equilibrium at Q_λ^* (Theorem 3.4, p. 689 in [1]). This implies that $Q_n^x - Q_\lambda^* \rightarrow 0$ as $n \rightarrow \infty$.

On the other hand, for $\lambda(t)$, let us consider a second trajectory on another time scale, such that:

$$\begin{aligned} \tilde{\lambda}(t) &= \lambda(n) + \left(\frac{t - \tau'(n)}{\tau'(n+1) - \tau'(n)} \right) (g(n+1) - g(n)) \\ t &\in [\tau'(n), \tau'(n+1)], \tau'(n) = \sum_{m=0}^n \beta(m), n \geq 0 \end{aligned} \quad (29)$$

which tracks the o.d.e.:

$$\dot{\Lambda}(t) = Q_{\Lambda(t)}^*(x, 1) - Q_{\Lambda(t)}^*(x, 0)$$

If $\Lambda(t) > \lambda(x)$ (excess subsidy), the passive mode is preferred, i.e., $Q_{\Lambda(t)}^*(x, 0) > Q_{\Lambda(t)}^*(x, 1)$, making the r.h.s of the previous equation < 0 and $\Lambda(t)$ decreases. Likewise, if the opposite strict inequality holds, the r.h.s. is > 0 and $\lambda(t)$ increases. Thus, the trajectory of $\Lambda(\cdot)$ remains bounded. As in the previous case, being a well-defined bounded scalar o.d.e., it converges to an asymptotically stable equilibrium where Λ satisfies $Q_\Lambda^*(x, 1) = Q_\Lambda^*(x, 0)$. This point is where both policies are equally desirable, i.e. Λ is the Whittle index.