



HAL
open science

Asymptotically optimal priority policies for indexable and non-indexable restless bandits

Ina Maria Maaïke Verloop

► **To cite this version:**

Ina Maria Maaïke Verloop. Asymptotically optimal priority policies for indexable and non-indexable restless bandits. 2014. hal-00743781v3

HAL Id: hal-00743781

<https://hal.science/hal-00743781v3>

Preprint submitted on 7 Jul 2014 (v3), last revised 29 Feb 2016 (v6)

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.

Asymptotic optimal control of multi-class restless bandits

I.M. Verloop

CNRS ; IRIT ; 2 rue C. Camichel, F-31071 Toulouse, France
Université de Toulouse ; INP ; *IRIT* ; F-31071 Toulouse, France

July 7, 2014

Abstract

We study the asymptotic optimal control of multi-class restless bandits. A restless bandit is a controllable stochastic process whose state evolution depends on whether or not the bandit is made active. The aim is to find a control that determines at each decision epoch which bandits to make active in order to minimize the overall average cost associated to the states the bandits are in. Since finding the optimal control is typically intractable, we study an asymptotic regime instead that is obtained by letting the number of bandits that can be simultaneously made active grow proportionally with the population of bandits. We consider both a fixed population of bandits as well as a dynamic population of bandits where bandits can depart and new bandits can arrive over time to the system. We propose a class of priority policies, obtained by solving a linear program, that are proved to be asymptotically optimal under a global attractor property and a technical condition. Indexability of the bandits is not required for the result to hold. For a fixed population of bandits, the technical condition reduces to checking a unichain property of the priority policy. For a dynamic population of bandits we present a large class of restless bandit problems for which the technical condition is always satisfied.

In case the bandits are indexable, we prove that Whittle's index policy is included in the class of asymptotically optimal policies. This generalizes the result of Weber and Weiss (1990), who showed asymptotic optimality of Whittle's index policy for a symmetric fixed population of bandits, to the setting of (i) several classes of bandits, (ii) multiple actions, and (iii) possible arrivals of new bandits. In order to prove the main results we combine fluid-scaling techniques with linear programming results. This is a different proof approach than that taken in Weber and Weiss, and, in contrary to the latter, allows to include arrivals of new bandits to the system.

We then present a multi-class $M/M/S+M$ queue, which is an example of a dynamic population of indexable restless bandits. This model satisfies the technical condition as well as the global attractor property, henceforth, asymptotic optimality of an index policy follows.

For non-indexable bandits we describe how to select priority policies from the class of asymptotically optimal policies and present numerical evidence that, outside the asymptotic regime, the performance of our proposed priority policies is nearly optimal.

Keywords: Restless bandits, asymptotic optimality, Whittle's index policy, arm-acquiring bandits.

1 Introduction

Multi-armed bandit problems are concerned with the optimal dynamic activation of several competing bandits taking into account that at each moment in time α bandits can be made active. The aim is to find a control that determines at each decision epoch which bandits to make active in order to minimize the overall cost associated to the states the bandits are in. In the by now classical multi-armed bandit model [19] it is assumed that only active bandits can change from state. In [51] Whittle introduced so-called restless bandits, where a bandit can also change its state while it is passive (that is, not active), possibly according to a different law from that which applies when it is active. The multi-armed restless bandit problem is a stochastic optimization problem that has gained popularity due to its multiple applications in for example sequential selection trials in medicine, sensor management, manufacturing systems, queueing and communication networks, control theory, economics, etc. We refer to [20, 32, 52] for further references, applications, and possible extensions studied in the literature.

In 1979, Gittins introduced index-based policies for the non-restless bandit problem. He associated to each bandit an index being a function of the state a bandit is in and defined the policy that makes those

α bandits active having currently the greatest indices (see [18]). This policy is known as the Gittins index policy. It was first proved by Gittins that this policy is optimal in case $\alpha = 1$ [18] for the time-average and discounted cost criteria. Extensions of the optimality result of Gittins index policy when new bandits may arrive over time (for example Poisson arrivals or Bernoulli arrivals) were obtained in [49, 50]. Note that for $\alpha > 1$ the optimality result does not necessarily go through. In [39] sufficient conditions on the reward processes are given though in order to guarantee optimality of the Gittins policy for the discounted cost criterion when $\alpha > 1$.

In the presence of restless bandits, finding an optimal control is typically intractable. Whittle [51] proposed therefore to solve a relaxed optimization problem where the constraint of having at most α bandits active at a time is relaxed to a time-average or discounted version of the constraint. He observed that in case the bandits satisfy a so-called indexability property, an optimal solution to the relaxed optimization problem can be described by index values. The latter, in their turn, provide a heuristic for the original restless bandit problem, which is in the literature referred to as Whittle's index policy. This policy reduces to Gittins index policy when passive bandits are static (the non-restless case). In [21] the authors extend Whittle's index heuristic to the setting where for each bandit from multiple actions may be chosen, i.e., representing a divisible resource to a collection of bandits. Over the years, Whittle's index policy has been extensively applied and numerically evaluated in various application areas such as e.g. wireless downlink scheduling [5, 38], systems with delayed state observation [14], broadcast systems [41], multi-channel access models [1, 31], stochastic scheduling problems [2, 23, 36] and scheduling in the presence of impatient customers [7, 22, 29, 37].

As opposed to Gittins policy, Whittle's index policy is in general not an optimal solution. Only in some particular cases optimality has been proved. See for example [1, 31] where this has been proved for a symmetric restless bandit problem modeling a multi-channel access system. For a general restless bandit model, in [27] Whittle's index policy has been shown to be optimal for $\alpha = 1$ in case (i) there is one dominant bandit or when (ii) all bandits immediately reinitialize when made passive. Other results on optimality of Whittle's index policy exist for *asymptotic regimes*: In [51] Whittle conjectured that Whittle's index policy is nearly optimal as the number of bandits that can be simultaneously made active grows proportionally with the total number of bandits in the system. In the case of *symmetric* bandits, i.e., all bandits are governed by the same transition rules, this conjecture was proved by Weber and Weiss [47] assuming that the differential equation describing the fluid approximation of the system has a global attractor. Not all bandit problems satisfy this condition though, as was demonstrated in [47]. A more recent result on asymptotic optimality is in [38] where the authors consider a particular model as studied in [31] with two classes of bandits. They prove asymptotic optimality of Whittle's index policy under a recurrence condition. The latter condition replaces the global attractor condition needed in [47] and is numerically verified to hold for their model.

For a dynamic population of restless bandits, that is, when new bandits can arrive to the system, there exist few papers on the performance of index policies. We refer to [5, 6] and [7] where this has been studied in the context of wireless downlink channels and queues with impatient customers. In particular, in [5] Whittle's index policy was obtained assuming no future arrivals of bandits and numerically shown to perform well when including arrivals of bandits (the dynamic setting). In [6] it was shown that this heuristic is in fact maximum stable and asymptotically fluid optimal under the dynamic setting. We note that the asymptotic regime studied in [6] is different than the one as proposed by Whittle [47]. More precisely, in [6] at most one bandit can be made active at a time (the fluid scaling is obtained by scaling both space and time), while in [47] (as well as in this paper) the number of active bandits scales. In fact, in [6] it was shown that any policy that gives strict priority to a bandit being currently in its best possible channel condition (breaking ties according to a size-based rule), is asymptotically fluid optimal.

Including arrivals of new "entities" to the system can also be modelled by a fixed population of restless bandits. In that case a bandit represents a certain type of entities, and the state of a bandit represents the number of this type of entities. Hence, a new arrival will change the state of the bandit. In the context of queueing systems this has been studied for example in [2, 22, 29]. A Whittle's index obtained from the relaxation of this problem formulation can be state-dependent, i.e., Whittle's index of a certain type of tasks depends on the number of tasks present in the system and on the arrival characteristics. This is contrary to for example [5, 7], where the index, obtained using the dynamic population formulation, does not depend on the number of bandits present or on the arrival characteristics. Asymptotic optimality results for a fixed population of bandits modeling arrivals of new "entities" has been obtained in for example [22] where Whittle's index is shown to be optimal both in the light-traffic and the heavy-traffic limit.

In this paper we study the asymptotic optimal control of a general multi-class restless bandit problem and consider both a fixed population of bandits as well as a dynamic scenario where bandits can arrive and depart from the system. The asymptotic regime is obtained by letting the number of bandits that can be simultaneously made active grow proportionally with the population of bandits. We derive a set of priority policies that are asymptotically optimal when certain conditions are satisfied. For a fixed population of bandits these conditions reduce to a differential equation having a global attractor, which coincides with the condition as needed by Weber and Weiss [47]. For a dynamic population of bandits additional technical conditions are needed due to the infinite state space. To illustrate the applicability of the results, we present a large class of restless bandit problems for which we show the technical condition is satisfied. This class is characterized by the fact that a bandit that is kept passive will eventually leave the system. This can represent many practical situations such as impatient customers, companies that go bankrupt, perishable items, etc. We then present a multi-class $M/M/S+M$ queue, which is a very particular example of the above described class for which the technical conditions are satisfied. We describe a priority policy that satisfies the global attractor property and hence asymptotic optimality follows.

In case the bandits are indexable we prove that Whittle’s index policy is contained in the class of asymptotically optimal policies, both for a fixed population as well as for a dynamic population of bandits. This generalizes the result of Weber and Weiss [47] from a symmetric fixed population of bandits to the setting of: (i) several classes of bandits, and (ii) possible arrivals of new bandits. Another extension presented in the paper is the possibility of choosing among multiple actions per bandit. This is referred to as “superprocess” in the literature [20]. We derive that under a strong form of indexability the asymptotic optimality results for the priority policies and Whittle’s index rule go through.

As opposed to Whittle’s index policy, we do not need the restless bandits to be indexable in order to define asymptotically optimal policies. In case the bandits are non-indexable, we describe how to select priority policies that are asymptotically optimal and present numerical evidence that, outside the asymptotic regime, the performance of our proposed priority policies is nearly optimal.

We will consider a generalization of the standard restless bandit formulation: Instead of having at each moment in time exactly α bandits active, we allow strictly less than α bandits to be active at a time. We handle this by introducing so-called dummy bandits. In particular, we derive that it is asymptotically optimal to make those bandits active having currently the greatest, *but positive*, Whittle’s indices. Hence, whenever a bandit is in a state having a negative Whittle’s index, this bandit will never be made active.

Our proof technique relies on a combination of fluid-scaling techniques and linear programming results: First we describe the fluid dynamics of the restless bandit problem taking only into account the average behavior of the original stochastic system. The optimal equilibrium points of the fluid dynamics are described by an LP problem. We prove that the optimal value of the LP provides a lower bound on the cost of the original stochastic system. The optimal fluid equilibrium point is then used to describe priority policies for the original system whose fluid-scaled cost coincides with the lower bound, and are hence referred to as asymptotically optimal policies. In order to prove that Whittle’s index policy is one of these asymptotically optimal policies we then reformulate the relaxed optimization problem by the corresponding LP problem whose optimal solution is proved to coincide with that of the LP problem corresponding to the fluid problem as described above. This is a different proof approach than that taken in [47] and allows to include arrivals of bandits to the system, whereas the approach of [47] does not. The proposed class of priority policies are based on the optimal solution(s) of the LP formulation. Other papers where LP solutions have been used to derive asymptotic optimal policies are for example [26] where in the context of loss networks an admission control policy is derived by targeting at each moment in time for the optimal solution of the LP using a penalty function, and [16, 30] in the context of revenue management where the LP solution determines the set of products to be offered over a finite-time horizon.

In the context of restless bandits, an LP-based proof approach was previously used in [9, 34, 35]. In [34, 35] it allowed to characterize and compute indexability of restless bandits. In [9] a set of LP relaxations was presented, providing performance bounds for the restless bandit problem under the discounted-cost criterion. In addition, in [9] the primal-dual index heuristic was proposed (its definition does not require indexability of the system), and proved to have a suboptimality guarantee. We will see that an adapted version of the primal-dual index heuristic is included in the set of priority policies for which we obtain asymptotic optimality results.

The main contributions of this paper are the following:

- For a multi-class restless bandit problem (possibly non-indexable) with either a fixed or dynamic

population of bandits and possibly multiple actions per state, we determine a class of priority policies that are asymptotically optimal (Proposition 3.11) if the corresponding ODE has a global attractor and certain technical conditions (Condition 3.10) are satisfied.

- We show that Condition 3.10 is satisfied for a large class of restless bandit problems. In particular, for a fixed population of bandits under a unichain assumption and for a dynamic population when passive bandits will eventually leave the system (Proposition 5.1).
- In case the bandits are indexable, we show that Whittle's index policy is inside the class of priority policies, both for a fixed population of bandits (Proposition 4.3) and for a dynamic population of bandits (Proposition 4.6). Hence, we obtain that if both the global attractor property and Condition 3.10 are satisfied, then Whittle's index is asymptotically optimal.
- For non-indexable bandits we describe how to select priority policies from the class of asymptotically optimal policies (Section 7.1) and we numerically show that outside the asymptotic regime their suboptimality gap is very small (Section 7.2).

The remainder of the paper is organized as follows. In Section 2 we define the multi-class restless bandit problem. The asymptotic optimality results for a class of priority policies are presented in Section 3. In Section 4 we define Whittle's index policy and prove its asymptotic optimality, both for a fixed population as well as for a dynamic population of bandits. In Section 5 we discuss the conditions required in order to prove the asymptotic optimality result. In Section 6 we present the $M/M/S+M$ queue as an indexable restless bandit example for which a robust priority policy can be derived that is asymptotically optimal. Section 7 focuses on the selection of asymptotically optimal priority policies for non-indexable bandits and numerically evaluates their performance.

2 Model description

We consider a multi-class restless bandit problem in continuous time. There are K classes of bandits. New class- k bandits arrive according to a Poisson process with arrival rate $\lambda_k \geq 0$, $k = 1, \dots, K$. At any moment in time, a class- k bandit is in a certain state $j \in \{1, 2, \dots, J_k\}$, with $J_k < \infty$. When a class- k bandit arrives, with probability $p_k(0, j)$ this bandit starts in state $j \in \{1, \dots, J_k\}$.

Decision epochs are defined as the moments when an event takes place, i.e., an arrival of a new bandit, a change in the state of a bandit, or a departure of a bandit. At each decision epoch, the controller can choose for each bandit between two actions: action $a = 0$, that is, making the bandit passive, or action $a = 1$, that is, making the bandit active. When control a is performed on a class- k bandit in state i , it goes to state j after an exponentially distributed amount of time with rate $q_k^a(i, j)$, $j = 0, 1, \dots, J_k$, $j \neq i$, where we interpret $j = 0$ as the fact that the bandit has departed from the system. Note that the evolution of one bandit (given its action) is independent of that of all the other bandits. We write $q_k^a(j, j) = -\sum_{i=0, i \neq j}^{J_k} q_k^a(j, i)$ and define $q_k^a(j) := -q_k^a(j, j)$ as the total transition rate out of state j for a class- k bandit under action a . At any moment in time at most α bandits can be made active. This is a generalization of the standard restless bandit formulation for a fixed population of bandits where at each moment in time *exactly* α bandits needs to be made active¹.

We note that this problem is referred to as the *restless* bandit problem since even though a bandit is passive, it can still change from state.

A policy decides at each decision epoch which bandits are made active. For a given policy π we define $X^\pi(t) := (X_{j,k}^{\pi,a}(t); k = 1, \dots, K, j = 1, \dots, J_k, a = 0, 1)$ with $X_{j,k}^{\pi,a}(t)$ the number of class- k bandits at time t that are in state j and see action a . We further denote by $X_{j,k}^\pi(t) := \sum_{a=0}^1 X_{j,k}^{\pi,a}(t)$ the total number of class- k bandits in state j and $X_k^\pi(t) := \sum_{j=1}^{J_k} X_{j,k}^\pi(t)$ the total number of class- k bandits.

The cost per unit of time of having a class- k customer in state j under action a is equal to $C_{j,k}^a \in \mathbb{R}$. We note that the cost $C_{j,k}^a$ can be negative, i.e., representing a reward. The objective is to minimize

¹The standard formulation can be retrieved by replacing $C_{j,k}^0$ (as defined in this section) with $C_{j,k}^0 + C$ for all j, k , where C represents an additional cost of having an passive bandit. The average additional cost for having passive bandits in the system is equal to $(N - A)C$, with N the total number of bandits in the system (we have a fixed population of bandits) and A the average number of active bandits in the system. When C is large enough, the cost of having inactive bandits becomes more important than the original cost function. Hence, an optimal policy will necessarily have $A = \alpha$. How to share the number of active places among the bandits in the system will be determined by Whittle's index.

the long-run average holding cost among all Markovian policies (policies that base their decisions on the current state and time), i.e., find a Markovian policy π^* that minimizes

$$\sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left(\int_{t=0}^T C_{j,k}^a X_{j,k}^{\pi,a}(t) dt \right), \quad (1)$$

under the constraint that at any moment in time at most α bandits can be made active, that is,

$$\sum_{k=1}^K \sum_{j=1}^{J_k} X_{j,k}^1(t) \leq \alpha, \text{ for all time } t. \quad (2)$$

Throughout this paper we will consider two different scenarios: a fixed population of bandits, and a dynamic population of bandits. Below we further specify the model for the different settings:

- *Fixed population:* In this case there are no new arrivals of bandits, i.e., $\lambda_k = 0$ for all $k = 1, \dots, K$, and there are no departures, i.e., $q_k^a(j, 0) = 0$ for all $k = 1, \dots, K$, $j = 1, \dots, J_k$, $a = 0, 1$. Hence, $X_k^\pi(t) = X_k(0)$ for any time t and any policy π .
- *Dynamic population:* In this case there are new arrivals of bandits, i.e., $\lambda_k > 0$ for all $k = 1, \dots, K$, and each bandit can depart from the system, i.e., for each class k there is at least one state j and one action a such that $q_k^a(j, 0) > 0$.

For a given policy π , we will call the system stable if the process $\vec{X}^\pi(t)$ has a unique invariant probability distribution. In the case of a fixed population of bandits, the state space is finite, hence the process $X^\pi(t)$ being unichain would be a sufficient condition for stability of the policy π . In the case of a dynamic population of bandits the stability condition is more involved. Whether or not the system is stable can depend strongly on the employed policy. In Section 3 we will state necessary stability conditions for the dynamic restless bandit problem.

Throughout the paper we will need to make additional restrictions on the model and policies considered in order for the results to hold. When doing so, we will make a distinction between (i) assumptions made on the model parameters (referred to as ‘‘Assumption’’) and (ii) conditions that are posed on the policy under investigation (referred to as ‘‘Condition’’).

Remark 2.1 (Multi actions) *In the model description we assumed there are only two possible actions per bandit: $a = 0$ (passive bandit) and $a = 1$ (active bandit). A natural generalization is to consider the possibility of multiple actions per bandit, that is, for a class- k bandit in state j the scheduler can chose from any action $a \in \{0, \dots, A_{j,k}\}$ and at most α bandits can be made active at a time, i.e., $\sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=1}^{A_{j,k}} X_{j,k}^a(t) \leq \alpha$. This is referred to as ‘‘superprocess’’ in the literature [20]. For the non-restless bandit problem with $\alpha = 1$ an index policy is known to be optimal in case each state has a dominant action, that is, if an optimal policy selects a class- k bandit in state j to be made active, it always chooses the same control $a_{j,k}$, with $a_{j,k} \in \{1, \dots, A_{j,k}\}$. A less strict condition is given in [20, Condition D].*

In this paper we consider the restless bandit problem and present the asymptotic optimality results for $A_{j,k} = 1$. We note that all results obtained will go through in the multi-action context when modifying accordingly the definition for the policies, see Remark 3.6 and Remark 4.1.

3 Fluid analysis and asymptotic optimality

In this section we consider a fluid formulation of the multi-class restless bandit problem, which allows to derive a class of priority policies that asymptotically minimize the cost of the original stochastic model (as given in (1)). More precisely, in Section 3.1 we introduce a fluid control problem and show that its optimal fluid cost provides a lower bound on the cost in the original stochastic model. In Section 3.2 we then define a class of priority policies for the original stochastic model based on the optimal fluid solution, which are shown to be asymptotically optimal in Section 3.3.

3.1 Fluid control problem and lower bound

The fluid control problem arises from the original stochastic model by only taking into account the mean drifts. For a given control $u(t)$, let $x_{j,k}^{u,a}(t)$ denote the amount of class- k fluid in state j under action a at time t and let $x_{j,k}^u(t) = x_{j,k}^{u,0}(t) + x_{j,k}^{u,1}(t)$ be the amount of class- k fluid in state j . Its dynamics is then described by

$$\frac{dx_{j,k}^u(t)}{dt} = \lambda_k p_k(0, j) + \sum_{a=0}^1 \sum_{i=1, i \neq j}^{J_k} x_{i,k}^{u,a}(t) q_k^a(i, j) - \sum_{a=0}^1 x_{j,k}^{u,a}(t) q_k^a(j), \quad (3)$$

with the constraint on the total amount of active fluid given by $\sum_{k=1}^K \sum_{j=1}^{J_k} x_{j,k}^{u,1}(t) \leq \alpha$, for all $t \geq 0$. We are interested in finding an optimal equilibrium point of the fluid dynamics, as given in (3), with as goal to minimize its holding cost. Hence, we pose the following linear optimization problem:

$$(LP) \quad \min_{(x_{j,k}^a)} \sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 C_{j,k}^a x_{j,k}^a$$

$$\text{s.t. } 0 = \lambda_k p_k(0, j) + \sum_{a=0}^1 \sum_{i=1, i \neq j}^{J_k} x_{i,k}^a q_k^a(i, j) - \sum_{a=0}^1 x_{j,k}^a q_k^a(j), \quad \forall j, k, \quad (4)$$

$$\sum_{k=1}^K \sum_{j=1}^{J_k} x_{j,k}^1 \leq \alpha, \quad (5)$$

$$\sum_{j=1}^{J_k} \sum_{a=0}^1 x_{j,k}^a = x_k(0), \text{ if } \lambda_k = 0, \quad \forall k, \quad (6)$$

$$x_{j,k}^a \geq 0, \quad \forall j, k, a, \quad (7)$$

where the constraint (6) follows since by (3) we have $\sum_{j=1}^{J_k} \frac{d}{dt} x_{j,k}^u(t) = 0$ if $\lambda_k = 0$, for all $t \geq 0$.

We denote by x^* an optimal solution of the above problem (LP), assuming it exists, and denote the optimal value by

$$V^*(\alpha) := \sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 C_{j,k}^a x_{j,k}^{*,a}.$$

We now prove some results for the optimization problem (LP). The proof of the following lemma may be found in Appendix A.

Lemma 3.1 *If there exists a policy π such that the process $X^\pi(t)$ has a unique invariant probability distribution with finite first moments, then the feasible set of (LP) is non-empty and $V^*(\alpha) < \infty$.*

As a consequence of Lemma 3.1 we get necessary conditions under which there exists a policy that makes the system stable and has finite first moments.

Corollary 3.2 *Consider a dynamic population of bandits. A necessary condition for the existence of a policy that is stable with finite first moments is*

$$\sum_{k=1}^K \sum_{j=1}^{J_k} y_{j,k}^{*1} \leq \alpha,$$

with $(y_{j,k}^{*a}) = \arg \min \{ \sum_{k=1}^K \sum_{j=1}^{J_k} x_{j,k}^1 : (x_{j,k}^a) \text{ satisfies (4) and (7)} \}$.

Proof: Assume there exists system is stable. Hence, by Lemma 3.1 the feasible set of (LP) is non-empty. That is, there exists an $(x_{j,k}^a)$ such that (4) and (7) hold and $\sum_{k=1}^K \sum_{j=1}^{J_k} x_{j,k}^1 \leq \alpha$. Hence, by definition of the optimal solution y^* we obtain $\sum_{k=1}^K \sum_{j=1}^{J_k} y_{j,k}^{*1} \leq \sum_{k=1}^K \sum_{j=1}^{J_k} x_{j,k}^1 \leq \alpha$. This concludes the proof. \square

The optimal solution of the fluid control problem (LP) serves as a lower bound on the cost of the original stochastic optimization problem (1), see the following lemma. The proof can be found in Appendix B.

Lemma 3.3 *We have*

$$\sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left(\int_0^T C_{j,k}^a X_{j,k}^{\pi,a}(t) dt \right) \geq V^*(\alpha), \quad (8)$$

for either

- a fixed population of bandits and for any policy π , or,
- a dynamic population of bandits and for any stable policy π , or,
- a dynamic population of bandits with $C_{j,k}^a > 0$ for all j, k, a and for any policy π that is either rate-stable (i.e., $\lim_{t \rightarrow \infty} \sum_{j,k} \frac{X_{j,k}^{\pi}(t)}{t} = 0$ almost surely) or mean rate-stable (i.e., $\lim_{t \rightarrow \infty} \sum_{j,k} \frac{\mathbb{E}(X_{j,k}^{\pi}(t))}{t} = 0$).

3.2 Priority policies

A priority policy is defined as follows. There is a predefined priority ordering on the states each bandit can be in. At any moment in time a priority policy makes active a maximum number of bandits being in the states having the highest priority among all the bandits present. In addition, the policy can prescribe that certain states are never made active.

Below we define a set of priority policies, Π^* , that will play a key role in the remaining of the paper. In Section 3.3 we will show that a policy in this class is asymptotically optimal if it satisfies certain conditions. We emphasize that we do not require the bandits to be indexable in order to define the set of priority policies Π^* . This as opposed to the definition of Whittle's index policy, which is only well defined in case the system is indexable. We note that Whittle's index policy is included in Π^* (in case the bandits are indexable), as will be proved in Section 4.3.

Definition 3.4 (Set of priority policies) *We define the set of priority policies Π^* as follows. Let $X^* := \{x^* \text{ an optimal solution of (LP)}\}$. Then,*

$$\Pi^* := \cup_{x^* \in X^*} \Pi(x^*),$$

where $\Pi(x^*)$ is the set of priority policies that satisfy the following rules:

1. A bandit in state (j, k) with $x_{j,k}^{*,1} > 0$ and $x_{j,k}^{*,0} = 0$ is given higher priority than a bandit in state (\tilde{j}, \tilde{k}) with $x_{\tilde{j},\tilde{k}}^{*,0} > 0$.
2. A bandit in state (j, k) with $x_{j,k}^{*,0} > 0$ and $x_{j,k}^{*,1} > 0$ is given higher priority than a bandit in state (\tilde{j}, \tilde{k}) with $x_{\tilde{j},\tilde{k}}^{*,0} > 0$ and $x_{\tilde{j},\tilde{k}}^{*,1} = 0$.
3. If $\sum_{k=1}^K \sum_{j=1}^{J_k} x_{j,k}^{*,1} < \alpha$, then any class- k bandit in state j with $x_{j,k}^{*,1} = 0$ and $x_{j,k}^{*,0} > 0$ will **never** be made active.

If there exists a policy such that the system is stable and has finite first moments, then by Lemma 3.1 the feasible set of (LP) is non-empty, and hence, the set Π^* is non-empty. In general, the set Π^* can consist of more than one policy. When choosing a policy from the set, preferably, we aim for a *robust* policy, that is, a policy described by a priority ordering that does not depend on α . (We note that for the (LP) problem, increasing α is equivalent to decreasing $x_k(0)$ or λ_k , $k = 1, \dots, K$, proportionally for the fixed or dynamic population, respectively.) In case the system is indexable, Whittle's index policy is a robust element of Π^* , see Section 4, and hence this seems to be a good choice. For a non-indexable system, the set Π^* might no longer contain a robust policy. In Section 7 we explain how to select in that case priority policies as α changes.

Before continuing we first give an example of Definition 3.4.

Example 3.5 *Assume $K = 2$ and $J_k = 2$. Let x^* be such that for class 1 we have $x_{1,1}^{*,0} = 0$, $x_{2,1}^{*,0} = 4$, $x_{1,1}^{*,1} = 3$, $x_{2,1}^{*,1} = 1$ and for class 2 we have $x_{1,2}^{*,0} = 2$, $x_{2,2}^{*,0} = 0$, $x_{1,2}^{*,1} = 0$, $x_{2,2}^{*,1} = 5$ and $\alpha = 10$. The priority policies associated to x^* in the set $\Pi(x^*)$, as defined in Definition 3.4, satisfy the following rules: By point 1.): Class-1 bandits in state 1 and class-2 bandits in state 2 are given the highest priority. By point*

3.): Since $x_{1,1}^{*,1} + x_{2,1}^{*,1} + x_{1,2}^{*,1} + x_{2,2}^{*,1} = 9 < \alpha$, class-2 bandits in state 1 are never made active. Hence, the set $\Pi(x^*)$ contains two policies: either give priority according to $(1,1) \succ (2,2) \succ (2,1)$ or give priority according to $(2,2) \succ (1,1) \succ (2,1)$, where a state (j,k) denotes a class- k bandit in state j . In both cases state $(1,2)$ is never made active.

Remark 3.6 (Multi actions) In this remark we explain how to define the set of priority policies Π^* in case of multiple actions per bandit. Recall that for a class- k bandit in state j the scheduler can chose from any action $a \in \{0, 1, 2, \dots, A_{j,k}\}$.

Similar to the non-restless bandit problem (see Remark 2.1), we are interested in priority policies such that if a class- k bandit in state j is chosen to be active, it will always be chosen to be made active in a fixed mode $a_{j,k}$. We therefore need to restrict the set X^* to optimal solutions of (LP) that satisfy $x_{j,k}^{*,a} x_{j,k}^{*,\tilde{a}} = 0$ for all $a, \tilde{a} \in \{1, \dots, A_{j,k}\}$. The latter implies that for all activation modes $a = 1, \dots, A_{j,k}$ one has $x_{j,k}^{*,a} = 0$, with the exception of at most one active mode, denoted by $a_{j,k}$. Hence we have $x_{j,k}^{*,a_{j,k}} \geq 0$ (instead of being equal to 0).

The set $\Pi(x^*)$ is now defined as in Definition 3.4, replacing the action $a = 1$ by $a = a_{j,k}$. All results obtained in Section 3 remain valid, by simply replacing $a = 1$ by $a = a_{j,k}$.

Remark 3.7 In [9] a heuristic is proposed for the multi-class restless bandit problem for a fixed population of bandits: the so-called primal-dual heuristic. This is defined based on the optimal (primal and dual) solution of a LP problem corresponding to the discounted-cost criterion. In fact, if the primal-dual heuristic would have been defined based on the problem (LP), it can be checked that it satisfies the properties of Definition 3.4, and hence is included in the set of asymptotically optimal priority policies Π^* .

In order to prove asymptotic optimality of a policy $\pi^* \in \Pi^*$, as will be done in Section 3.3, we will investigate its fluid dynamics, that is, we study the process $x^u(t)$, as defined in (3), with as control u the strict priority rule as given by π^* . For ease of notation we will write $x^*(t)$ instead of $x^{\pi^*}(t)$. We further denote by $S_{j,k}^*$ the set of all states (i,l) , $i = 1, \dots, J_l$, $l = 1, \dots, K$, such that class- l bandits in state i have higher priority than class- k bandits in state j under policy π^* , and I^* is the set of all states that will never be made active under policy π^* . The process $x^*(t)$ is hence defined as follows.

$$\frac{dx_{j,k}^*(t)}{dt} = \lambda_k p_k(0, j) + \sum_{a=0}^1 \sum_{i=1, i \neq j}^{J_k} x_{i,k}^{*,a}(t) q_k^a(i, j) - \sum_{a=0}^1 x_{j,k}^{*,a}(t) q_k^a(j), \quad \forall j, k, \quad (9)$$

$$\text{with } x_{j,k}^{*,1}(t) = \min \left(\left(\alpha - \sum_{(i,l) \in S_{j,k}^*} x_{i,l}^*(t) \right)^+, x_{j,k}^*(t) \right), \quad \text{if } (j, k) \notin I^*, \quad \forall j, k,$$

$$x_{j,k}^{*,1}(t) = 0, \quad \text{if } (j, k) \in I^*, \quad \forall j, k,$$

$$x_{j,k}^{*,0}(t) = x_{j,k}^*(t) - x_{j,k}^{*,1}(t), \quad \forall j, k.$$

It follows directly that an optimal solution x^* of (LP) is in fact an equilibrium point of the process $x^*(t)$ under a policy $\pi^* \in \Pi(x^*)$.

Lemma 3.8 Let $\pi^* \in \Pi^*$ and let x^* be a point such that $\pi^* \in \Pi(x^*)$. Then x^* is an equilibrium point of the process $x^*(t)$ as defined in (9) for the policy π^* .

Proof: Since x^* is an optimal solution of (LP), it follows directly from the definition of $\Pi(x^*)$ that x^* is an equilibrium point of the process $x^*(t)$. \square

Besides x^* being an equilibrium point, in order to prove asymptotic optimality of π^* it will be needed that x^* is the unique equilibrium point of (9) and in addition is a global attractor, i.e., all trajectories converge to x^* . This is not true in general, which is why we state it as a condition that a policy needs to satisfy. In Section 5.1 we will further comment on this condition.

Condition 3.9 Given a policy $\pi^* \in \Pi(x^*) \subset \Pi^*$, for an $x^* \in X^*$. The point x^* is the unique equilibrium point and the global attractor of the process $x^*(t)$ as defined in (9) for the policy π^* . I.e., for any initial point (for a given value of $x_k(0)$, $k = 1, \dots, K$, in case $\lambda_k = 0$) the process $x^*(t)$ converges to the point x^* .

3.3 Asymptotic optimality of priority policies

In this section we present results showing the asymptotic optimality of priority policies in the set Π^* . In particular, we obtain that, after a *fluid scaling*, the long-run average holding cost is optimized under a priority policy in Π^* .

We will consider the multi-class restless bandit problem in the following fluid-scaling regime: we scale by r both the arrival rates and the number of bandits that can be made active per time unit. That is, class- k bandits arrive at rate $\lambda_k \cdot r$, $k = 1, \dots, K$, and $\alpha \cdot r$ bandits can be made active at any moment in time. We let $X_{j,k}^r(0) = r \cdot x_{j,k}(0)$, with $x_{j,k}(0) \geq 0$.

For a given policy π , we denote by $X_{j,k}^{r,\pi,a}(t)$ the number of class- k bandits in state j experiencing action a at time t under scaling parameter r . We will be interested in the process after the fluid scaling, i.e., space is scaled linearly with the parameter r :

$$x_{j,k}^{r,\pi,a}(t) := X_{j,k}^{r,\pi,a}(t)/r. \quad (10)$$

We consider the cost of the stochastic model after fluid scaling, that is, we are interested in

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \mathbb{E} \left(\int_0^T C_{j,k}^a x_{j,k}^{r,\pi,a}(t) dt \right). \quad (11)$$

Our goal is to find policies that minimize the above as $r \rightarrow \infty$. We directly obtain from Lemma 3.3 that $V^*(\alpha)$ provides a lower bound on this cost. We will therefore call a policy *asymptotically optimal* when the fluid-scaled cost (11) converges to this lower bound.

In order to prove asymptotic optimality of priority policies in the set Π^* , we need the following technical condition.

Condition 3.10 *Given a policy $\pi^* \in \Pi^*$.*

- a) *The process $x^{r,\pi^*}(t)$ has a unique invariant probability distribution p^{r,π^*} , which has a finite first moment, for all r .*
- b) *The family $\{p^{r,\pi^*}, r\}$ is tight.*
- c) *The family $\{p^{r,\pi^*}, r\}$ is uniform integrable.*

For a fixed population of bandits, b) and c) are always satisfied since the state space of $x^{r,\pi^*}(t)$ is finite. A sufficient condition for Condition 3.10 a) to be satisfied (for the fixed population of bandits) is the Markov process $x^{r,\pi^*}(t)$ to be unichain (there is a state x such that under policy π^* there is a path from any state to x , i.e., state x is recurrent), for any r [44]. For the dynamic population setting, we will show in Section 5.2 the Condition 3.10 to be always satisfied for a large class of restless bandit problems.

We can now state the asymptotic optimality result.

Proposition 3.11 *For a given policy $\pi^* \in \Pi(x^*) \subset \Pi^*$, assume Condition 3.9 and Condition 3.10 are satisfied. Then,*

$$\lim_{r \rightarrow \infty} \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \mathbb{E} \left(\int_0^T C_{j,k}^a x_{j,k}^{r,\pi^*,a}(t) dt \right) = V^*(\alpha).$$

In particular, we have

$$\begin{aligned} & \lim_{r \rightarrow \infty} \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \mathbb{E} \left(\int_0^T C_{j,k}^a x_{j,k}^{r,\pi^*,a}(t) dt \right) \\ & \leq \liminf_{r \rightarrow \infty} \liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \mathbb{E} \left(\int_0^T C_{j,k}^a x_{j,k}^{r,\pi,a}(t) dt \right), \end{aligned} \quad (12)$$

for either

- *a fixed population of bandits and for any policy π , or,*
- *a dynamic population of bandits and for any stable policy π , or,*

- a dynamic population of bandits with $C_{j,k}^a > 0$ for all j, k, a and for any policy π that is either rate-stable (i.e., $\lim_{t \rightarrow \infty} \sum_{j,k} \frac{X_{j,k}^\pi(t)}{t} = 0$ almost surely) or mean rate-stable (i.e., $\lim_{t \rightarrow \infty} \sum_{j,k} \frac{\mathbb{E}(X_{j,k}^\pi(t))}{t} = 0$).

The proof may be found in Appendix C and consists in the following steps: Given a policy $\pi^* \in \Pi(x^*)$, we show that the fluid-scaled steady-state queue length vector converges to x^* . Since x^* is an optimal solution of the fluid control problem (LP) and has cost value $V^*(\alpha)$, this implies that the fluid-scaled cost (11) under policy π^* converges to $V^*(\alpha)$. Since $V^*(\alpha)$ serves as a lower bound on the average cost, this allows us to conclude for asymptotic optimality of the priority policy π^* .

4 Whittle's index policy

In Section 3.3 we found that priority policies inside the set Π^* are asymptotically optimal. In this section we will show that Whittle's index policy is in fact included in Π^* and is hence asymptotically optimal.

In Section 4.1 we first define Whittle's index policy. In Section 4.2 and Section 4.3 we then give sufficient conditions under which Whittle's index policy is asymptotically optimal, both in the case of a fixed population of bandits, and in the case of a dynamic population of bandits, respectively.

4.1 Relaxed-constraint optimization problem and Whittle's indices

Whittle's index policy has been proposed by Whittle [51] as an efficient heuristic for the multi-class restless bandit problem ([20, Section 6] and [52, Section II.14]). Under this policy each bandit is associated a Whittle's index being a function of the state a bandit is in. Whittle's index policy makes those bandits active having currently the highest indices. In this section we will describe how these Whittle's indices are derived. In order to do so, we will need to introduce a slightly different optimization problem. In Section 4.2 and Section 4.3 we make the connection with the original problem (as stated in Section 2) as we show that Whittle's index policy is asymptotically optimal for the original problem.

In order to define Whittle's indices, we consider the following optimization problem: Assume at time 0 we have a population of bandits (at least one bandit from each class) and assume there are no arrivals in the future. At any moment in time at most α bandits can be made active and we restrict ourselves to the class of *stationary* and Markovian policies. As cost criterion we will consider one of the following: the average-cost criterion

$$\mathbb{C}^{av}(Y(\cdot)) := \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left(\int_0^T Y(t) dt \right), \quad (13)$$

or the discounted-cost criterion

$$\mathbb{C}^\beta(Y(\cdot)) := \mathbb{E} \left(\int_0^\infty e^{-\beta t} Y(t) dt \right),$$

for $\beta > 0$. The objective is to find a policy that minimizes

$$\sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \mathbb{C}^f(C_{j,k}^a X_{j,k}^{\pi,a}(\cdot)), \quad (14)$$

with $f \in \{av, \beta\}$. We note that the objective posed in (1) was that of the average-cost criterion. In Section 4.3 it will become clear why in this section we also need to introduce the discounted-cost criterion.

4.1.1 Relaxed-constraint optimization problem

The restless property of the bandits makes the above-described optimization problem often infeasible to solve. Instead, Whittle [51] proposed to study the so-called *relaxed-constraint optimization problem*, which is defined as follows: find a policy that minimizes (14) under the relaxed constraint

$$\sum_{k=1}^K \sum_{j=1}^{J_k} \mathbb{C}^f(X_{j,k}^{\pi,1}(\cdot)) \leq \alpha(f), \quad (15)$$

with $\alpha(av) = \alpha$ and $\alpha(\beta) = \int_0^\infty \alpha e^{-\beta t} dt = \alpha/\beta$ for $\beta > 0$. That is, the constraint that at most α bandits can be made active at any moment in time is replaced by its time-average or discounted version, (15). Hence, the cost under the optimal policy of the relaxed-constraint optimization problem provides a lower bound on the cost for any policy that satisfies the original constraint.

In standard restless bandit problems the constraint (15) needs to be satisfied in the strict sense, that is, with an “=” sign. In this paper we allowed however strictly less than α bandits to be active at a time. In order to be able to deal with this within the original framework of Whittle, we introduce so-called *dummy bandits*. That is, besides the initial population of bandits, we assume there are $\alpha(f)$ additional bandits that will never change state. We denote the state these bandits are in by B and the cost of having a dummy bandit in state B is $C_B^a = 0$, $a = 0, 1$. The introduction of these $\alpha(f)$ dummy bandits allows to reformulate the relaxed-constraint problem as follows: Minimize (14) under the relaxed constraint

$$\mathbb{C}^f(X_B^{\pi,1}(\cdot)) + \sum_{k=1}^K \sum_{j=1}^{J_k} \mathbb{C}^f(X_{j,k}^{\pi,1}(\cdot)) = \alpha(f). \quad (16)$$

This constraint is equivalent to (15) since, for a given set of active bandits, activating additional dummy bandits does not modify the behavior of the system.

Using the Lagrangian approach, we write the relaxed-constraint problem (minimize (14) under constraint (16)) as the problem of finding a policy π that minimizes

$$\sum_{k=1}^K \sum_{j=1}^{J_k} \mathbb{C}^f(C_{j,k}^0 X_{j,k}^{\pi,0}(\cdot) + C_{j,k}^1 X_{j,k}^{\pi,1}(\cdot) + \nu X_{j,k}^{\pi,1}(\cdot)) + \mathbb{C}^f(\nu X_B^{\pi,1}(\cdot)). \quad (17)$$

The Lagrange multiplier ν can be viewed as the cost to be paid per active bandit. From Lagrangian relaxation theory we have that there exists a value of the Lagrange multiplier ν such that the constraint (16) is satisfied.

Since there is no longer a common constraint for the bandits, problem (17) can be decomposed into several *subproblems*, one for each bandit: For each class- k bandit the subproblem is to minimize

$$\mathbb{C}^f(C_{J_k(\cdot),k}^{A_k(\cdot)} + \nu \mathbf{1}_{(A_k(\cdot)=1)}), \quad (18)$$

where $J_k(t)$ and $A_k(t)$ denote the state and action chosen for the class- k bandit at time t , respectively, and for each dummy bandit the problem is to minimize

$$\nu \mathbb{C}^f(\mathbf{1}_{(A_B(\cdot)=1)}), \quad (19)$$

with $A_B(t)$ the action chosen for the dummy bandit at time t .

For a given optimization criterion f , Whittle defines the index $\nu_{j,k}^f$ as the least value of ν for which it is optimal in (18) to make the class- k bandit in state j passive. We refer to $\nu_{j,k}^f$ as *Whittle's index* for the optimization criterion f . Similarly, we define the index ν_B^f as the least value of ν for which it could be optimal in (19) to make a dummy bandit passive.

A bandit is called *indexable* if the set of states where it is optimal in (18) to make a class- k bandit passive forms an increasing sequence of sets (from the empty set to the set $\{1, \dots, J_k\}$), as ν increases. We refer to the problem as being indexable if all bandits are indexable. Note that whether a problem is indexable or not can depend on the choice for f (and β). We refer to [36] for a survey on indexation results. In particular, [36] presents sufficient conditions for a restless bandit to be indexable and provides a method to calculate Whittle's indices. Sufficient conditions for indexability can also be found in [31, 46].

We note that the dynamics of a bandit in state B is independent of the action chosen. Since ν represents the cost to be paid when active, it will be optimal in (19) to make a bandit in state B passive if and only if $\nu \geq 0$. As a consequence, a dummy bandit is always indexable and $\nu_B^f = 0$.

If the bandit problem is indexable, an optimal policy for the subproblem (18) is then such that the class- k bandit in state j is made active if $\nu_{j,k}^f > \nu$, is made passive if $\nu_{j,k}^f < \nu$, and any action is optimal if $\nu_{j,k}^f = \nu$, [51].

An optimal solution to (14) under the relaxed constraint (16) is obtained by setting ν at the appropriate level ν^* such that (16) is satisfied and to make a class- k bandit in state j active if $\nu_{j,k}^f > \nu^*$, and passive if $\nu_{j,k}^f < \nu^*$. When a class- k bandit is in a state j such that $\nu_{j,k}^f = \nu^*$, one needs to appropriately

randomize the action in this state such that the relaxed constraint (16) is satisfied, [47, 51]. In case $\nu^* = 0$, we take the convention that the randomization is done among the bandits in state B (possible since there are exactly $\alpha(f)$ dummy bandits) while any class- k bandit in a state j with $\nu_{j,k}^f = 0$ is kept passive.

Since there are always $\alpha(f)$ bandits in state B , we necessarily have that an optimal action for class- k bandits in state j with $\nu_{j,k}^f \leq 0 = \nu_B^f$ (hence, having lower priority than bandits in state B) is to be passive. In particular, this implies that $\nu^* \geq 0$.

4.1.2 Whittle's index policy as heuristic

Obviously the above optimal control for the relaxed problem is not feasible for the original optimization problem (as posed in the beginning of Section 4.1) having as constraint that at most α bandits can be made active *at any moment in time*. Whittle [51] therefore proposed the following heuristic: At any moment in time, make α bandits active having currently the highest Whittle indices. Hence, if $\nu_{j,k}^f < \nu_{i,l}^f$, then a class- l bandit in state i is given higher priority than a class- k bandit in state j . Recall that a class- k bandit in state j with $\nu_{j,k}^f \leq 0$ will never be made active in the optimal solution of the relaxed problem. Analogously, we define that under Whittle's index policy ($\nu_{j,k}^f$) a class- k bandit in state j will never be made active if $\nu_{j,k}^f \leq 0$. Further, in case different states have the same Whittle index, an arbitrary fixed priority rule is used.

The above described heuristic is referred to in the literature as *Whittle's index policy* for the average cost ($f = av$) or discounted cost ($f = \beta$) criterion. In the next two sections we will prove that this heuristic is asymptotically optimal, both for the static and dynamic population.

Remark 4.1 (Multi actions) *In this remark we define Whittle's index policy in case of multiple actions. For that we need to assume a stronger form of indexability that says that for a class- k bandit in state j there is an index $\nu_{j,k}^f$ and an optimal activation mode $a_{j,k}^f$ such that an optimal solution for (18) is to be active in mode $a_{j,k}^f$ if $\nu < \nu_{j,k}^f$ and to be passive if $\nu > \nu_{j,k}^f$. Whittle's index rule is then defined as above, simply by replacing the action $a = 1$ by $a = a_{j,k}^f$.*

If the restless bandit problem satisfies this stronger form of indexability, then one can reduce the multi-action problem to the single-action problem and hence all asymptotic optimality results as obtained in Section 4.2 and Section 4.3 will be valid (simply replace action $a = 1$ by $a = a_{j,k}$).

4.2 Asymptotic optimality for a fixed population of bandits

In this section we will prove asymptotic optimality of Whittle's index policy in the setting of a fixed population of bandits. More precisely, we show that Whittle's index policy (defined for the time-average cost criterion $f = av$) is included in the set of asymptotically optimal policies Π^* , as obtained in Section 3.3, in case the bandit problem is indexable. We note that Whittles index policy is independent on the value for α and $x_k(0)$, which makes it a *robust* policy in the sense that the priority ordering does not depend on the number of bandits that can be made active or on the number of bandits present in each class.

We will need the following assumption (the same assumption was made in [47]).

Assumption 4.2 *The transition rates of the state of one class- k bandit are such that they form a unichain (there is a state $j \in \{1, \dots, J_k\}$ such that there is a path from any state $i \in \{1, \dots, J_k\}$ to state j), for all k , regardless of the policy employed.*

The next proposition shows that Whittle's index policy is included in the set of asymptotically optimal policies. The proof can be found in Appendix D.

Proposition 4.3 *Consider a fixed population of bandits. If Assumption 4.2 holds and if the restless bandit problem is indexable for the average-cost criterion, then there is an $x^* \in X^*$ such that Whittle's index policy ($\nu_{j,k}^{av}$) is included in the set $\Pi(x^*) \subset \Pi^*$.*

We can now conclude that Whittle's index policy is asymptotically optimal.

Corollary 4.4 Consider a fixed population of bandits. If the assumptions of Proposition 4.3 are satisfied and if Condition 3.9 holds for Whittle's index policy $(\nu_{j,k}^{av})$, then policy $(\nu_{j,k}^{av})$ is asymptotically optimal. That is, for any policy π ,

$$\begin{aligned} & \lim_{r \rightarrow \infty} \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \mathbb{E} \left(\int_0^T C_{j,k}^a x_{j,k}^{r, \nu_{j,k}^{av}, a}(t) dt \right) \\ & \leq \lim_{r \rightarrow \infty} \inf_{r \rightarrow \infty} \lim_{T \rightarrow \infty} \inf_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \mathbb{E} \left(\int_0^T C_{j,k}^a x_{j,k}^{r, \pi, a}(t) dt \right). \end{aligned}$$

Proof: From Proposition 3.11 and Proposition 4.3 we obtain the desired result. \square

The above corollary was previously proved by Weber and Weiss in [47] for the case of symmetric bandits, i.e., $K = 1$. We note that the assumptions made in [47] in order to prove the asymptotic optimality result are the same as the ones in Corollary 4.4.

The proof technique used in Weber and Weiss is different from the one used here. In [47] the cost under an optimal policy is lower bounded by the optimal cost in the relaxed problem and upper bounded by the cost under Whittle's index policy. By showing that both bounds converge to the same value, the fluid approximation, the asymptotic optimality of Whittle's index policy is concluded. In this paper we have used another approach in order to prove the result. In particular, our approach allowed to include the case of a dynamic population of bandits, see Section 4.3, whereas the approach of [47] did not (note that for a dynamic population, the optimal cost cannot be lower bounded by any relaxed optimization problem).

4.3 Asymptotic optimality for a dynamic population of bandits

In this section we will introduce an index policy based on Whittle's indices for the dynamic population of bandits and show it to be asymptotically optimal. More precisely, we show the index policy to be included in the set of asymptotically optimal policies Π^* , as obtained in Section 3.3.

Recall that our objective is to find a policy that asymptotically minimizes the average-cost criterion (13). However, we cannot make use of Whittle's index policy with indices $\nu_{j,k}^{av}$ as we did in the previous section, since any policy that makes sure that the class- k bandit leaves after a finite amount of time, which is necessary for the system to be stable, has an average cost equal to zero and is hence an optimal solution of (18) for $f = av$. In order to derive a non-trivial index rule, the authors of [5, 7] consider instead the Whittle indices corresponding to the discounted-cost criterion ($f = \beta$, $\beta > 0$). An index rule for the average-cost criterion is then obtained by considering the limiting values as $\beta \downarrow 0$. We propose here the same: for each class k , let $\beta_l \rightarrow 0$ be some subsequence such that a class- k bandit is indexable under the β_l -discounted criterion and the limit

$$\nu_{j,k}^{lim} := \lim_{l \rightarrow \infty} \nu_{j,k}^{\beta_l}$$

is well-defined for all $j = 1, \dots, J_k$. We will refer to the index policy $(\nu_{j,k}^{lim})$ as Whittle's index policy for the average-cost criterion in the dynamic population setting².

We will need the following assumption on the model parameters.

Assumption 4.5 For all $k = 1, \dots, K$, the set of optimal solutions of the linear program

$$\begin{aligned} & \min_x \sum_{j=1}^{J_k} (C_{j,k}^0 x_{j,k}^0 + C_{j,k}^1 x_{j,k}^1 + \nu x_{j,k}^1) \\ & \text{s.t. } 0 = \lambda_k p_k(0, j) + \sum_{a=0}^1 \sum_{i=1, i \neq j}^{J_k} x_{i,k}^a q_k^a(i, j) - \sum_{a=0}^1 x_{j,k}^a q_k^a(j), \quad \forall j, \\ & x_{j,k}^a \geq 0, \quad \forall j, a, \end{aligned}$$

is bounded when $\nu > 0$.

²In the case of multiple actions per bandit, as defined in Remark 2.1, we need to assume the stronger form of indexability as defined in Remark 4.1, where the optimal activation mode for a class- k bandit in state j , $a_{j,k}^{\beta_l}$, cannot depend on β_l , for β_l small enough (i.e., $a_{j,k}^{\beta_l} = a_{j,k}$ for l large enough).

We note that this assumption is always satisfied in case $C_{j,k}^0 > 0$ for all j, k , since $x_{j,k}^{*,1}$ and $x_{j,k}^{*,0}$ are upperbounded by the cost value of a feasible solution divided by $\nu + C_{j,k}^1 > 0$ and $C_{j,k}^0 > 0$, respectively.

The next proposition shows that Whittle's index policy ($\nu_{j,k}^{lim}$) is included in the set of asymptotically optimal policies. The proof can be found in Appendix E.

Proposition 4.6 *Consider a dynamic population of bandits. If Assumption 4.5 holds and if the discounted restless bandit problem is indexable for $\beta_1 \leq \bar{\beta}$, with $0 < \bar{\beta} < 1$, then there is an $x^* \in X^*$ such that Whittle's index policy ($\nu_{j,k}^{lim}$) is included in the set $\Pi(x^*) \subset \Pi^*$.*

We can now conclude for asymptotic optimality of Whittle's index policy.

Corollary 4.7 *Consider a dynamic population of bandits. If the assumptions of Proposition 4.6 are satisfied and if Condition 3.9 and Condition 3.10 hold for Whittle's index policy ($\nu_{j,k}^{lim}$), then policy ($\nu_{j,k}^{lim}$) is asymptotically optimal, that is,*

$$\begin{aligned} & \lim_{r \rightarrow \infty} \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \mathbb{E} \left(\int_0^T C_{j,k}^a x_{j,k}^{r, \nu^{lim}, a}(t) dt \right) \\ & \leq \lim_{r \rightarrow \infty} \inf \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \mathbb{E} \left(\int_0^T C_{j,k}^a x_{j,k}^{r, \pi, a}(t) dt \right), \end{aligned} \quad (20)$$

for any policy π that is stable or for any policy π that is rate-stable or mean rate-stable with $C_{j,k}^a > 0$ for all j, k, a .

Proof: The result follows directly from Proposition 3.11 and Proposition 4.6. □

The above result for the dynamic population shows that the heuristic ($\nu_{j,k}^{lim}$), which is based on a model without arrivals, is in fact nearly optimal in the presence of arrivals. In addition, as for the fixed population, Whittle's indices defined for the dynamic setting are *robust* in the sense that they do not depend on the arrival characteristics of new bandits or on the exact number of bandits that can be made active, α . The Whittle's indices do depend on f (and β) though.

5 On the technical conditions

For a class of priority policies Π^* we showed in Proposition 3.11 asymptotic optimality under the global attractor property (Condition 3.9) and a technical condition (Condition 3.10). In this section we further discuss these two conditions: In Section 5.1 we comment on the global attractor property and in Section 5.2 we describe a general class of restless bandit problems for which Condition 3.10 is always satisfied.

5.1 Condition 3.9

Condition 3.9 is concerned with the process $x^*(t)$, defined by the ODE (9), to have a global attractor. We note that in [47] the same global attractor property was required in order to prove asymptotic optimality of Whittle's index policy for a fixed population of symmetric bandits ($K = 1$).

For a fixed population of symmetric bandits ($K=1$) with $J = 2$ or $J = 3$, the global attractor property is proved to always hold, see [48]. For $K > 1$ and/or more than three states, no sufficient conditions are available in order for x^* to be a global attractor of $x^*(t)$. A necessary condition in order to have a global attractor was provided in [47, Lemma 2]. There it is proved that for a fixed population with symmetric bandits, indexability is necessary in order for a robust priority policy (that is, a priority ordering that does not depend on α or on $x_k(0)$, $k = 1, \dots, K$) to satisfy the global attractor property (for any value of α and $x_k(0)$). In addition, this robust priority policy will coincide with Whittle's index policy. In case there does not exist a robust policy in Π^* , this implies that the problem is non-indexable. However, we emphasize that in that case the global attractor property might still hold for policies in Π^* .

The asymptotic optimality result under Whittle's index policy for the case $K = 1$ [47] has been cited extensively, where the global attractor property needed for the asymptotic optimality result has to be

verified often numerically. Note that in the context of mean field interaction models, convergence of the stationary measure also relies on a global attractor assumption of the corresponding ODE, see for example [8]. The authors of [38] proved asymptotic optimality of Whittle’s index policy for a very specific model with only two classes of bandits (fixed population of bandits) under a recurrence condition. The latter condition replaced the global attractor condition, however, the authors need as well to resort to numerical experiments in order to verify this recurrence condition.

In the remainder of this section we describe the necessity of the global attractor property and the technical challenges in case this condition is not satisfied.

Optimal fluid control problems have been widely studied in the literature in order to obtain asymptotically optimal policies for the stochastic model. In the context of this paper, the fluid control problem related to our results would be to find the optimal control $u^*(t)$ that minimizes

$$\lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 C_{j,k}^a x_{j,k}^{u,a}(t) dt, \quad (21)$$

where the dynamics of $x_{j,k}^{u,a}(t)$ is described by (3). The optimal control $u^*(t)$ needs to be translated back to the stochastic model in such a way that it is asymptotically optimal. When stating the global attractor property, the above is exactly what we have in mind. In fact, instead of solving this transient fluid control problem, we directly consider an optimal equilibrium point of the fluid model and propose a priority policy based on this equilibrium point. When the global attractor property is satisfied, this implies that the optimal equilibrium point is indeed reached by the associated strict priority control, and hence this priority control solves (21).

When, for any $\pi^* \in \cup_{x^* \in X^*} \Pi(x^*) = \Pi^*$, the global attractor property is not satisfied, this means that there does not exist a priority control $u(t) = \pi^* \in \Pi(x^*)$ such that the fluid process $x^{\pi^*}(t)$ converges to x^* . In that case, we can be in either one of the following two situations: 1) There exists a control $u^*(t)$ for which the process $x^{u^*}(t)$ does have as global attractor $x^* \in X^*$, where X^* was defined as the set of optimal equilibrium points. This control $u^*(t)$ is most likely not of priority type. 2) There does not exist any control that has a global attractor $x^* \in X^*$. In the latter case, the optimal control $u^*(t)$ can be such that the process $x^{u^*}(t)$ behaves cyclic or converges to a non-optimal equilibrium point. Hence, in case Condition 3.9 is not satisfied, in both Situation 1) and Situation 2), one needs to determine the exact transient behaviour of the optimal control of (21), $u^*(t)$, which in its turn needs to be translated back to the stochastic model. This is a challenging research problem and is subject for future research.

5.2 Condition 3.10

First assume we have a fixed population of bandits. As noted in Section 3.3, points b) and c) of Condition 3.10 are always satisfied. A sufficient condition for a) to be satisfied, i.e., the stochastic processes of the number of bandits under policy π^* to have a unique invariant probability distribution, is the Markov process under policy π^* to be unichain.

For a dynamic population of bandits we present a very general class of restless bandit problems for which Condition 3.10 is satisfied. More precisely, we consider problems in which bandits that are kept passive will eventually leave the system. For many real-life situations this assumption arises naturally, think for example of customers that become impatient and abandon the queue/system, companies that go bankrupt, perishable items, etc. The proof of the proposition may be found in Appendix F.

Proposition 5.1 *Assume for each class k the state 0 is positive recurrent under the policy that always keeps the class- k bandit passive. For any priority policy π for which $X^\pi(t)$ is irreducible, Condition 3.10 is satisfied.*

We note that a very specific example inside the above described class of problems is a bandit problem in which only active bandits are allowed in the system, i.e., $q_k^0(i, 0) = \infty$, for all $k = 1, \dots, K$, $i = 1, \dots, J_k$. This could describe for example the hiring process where new candidates are modeled by new arriving bandits, room occupation in a casualty departments where patients require direct attention, or a loss network. When $q_k^0(i, 0) = \infty$ for all k, i , the state space is finite (there can be at most α bandits present in the system), hence, for this particular example, Condition 3.10 would also follow directly from a unichain assumption.

Condition 3.10 a) concerns stability of the system under a strict priority policy resulting from the fluid analysis. In general, care has to be taken when applying a fluid optimal control directly to the stochastic

system. For example, in the context of scheduling in networks, there exist examples where priority policies resulting from the fluid model do not succeed in making the system stable, see for example [43, 45]. We believe that for any restless bandit problem the priority policies in Π^* will be stable though, however, this is a subject for future research. As an example, we refer to [6] where a particular restless bandit problem was studied that modelled a system with state-dependent capacity. In that problem, certain priority policies (e.g. the myopic $c\mu$ rule) yield an unstable system, while other priority policies, including Whittle's index policy, keep the system stable.

6 Case study: a multi-server queue with abandonments

In this section we study a multi-class multi-server system with impatient customers, the multi-class $M/M/S+M$ system. This is an example of a restless bandit problem with a dynamic population. We will derive a robust priority policy that is in the class Π^* and show that it satisfies the two conditions needed in order to conclude for asymptotic optimality.

The impact of abandonments has attracted considerable interest from the research community, with a surge in recent years. To illustrate the latter, we can mention the recent Special Issue in Queueing Systems on abandonments in queueing systems [25] and the survey paper [13] on abandonments in a many-server setting

We consider a multi-class system with S servers working in parallel. At any moment in time, each server can serve at most one customer. Class- k customers arrive according to a Poisson process with rate $\lambda_k > 0$ and require an exponentially distributed service with mean $1/\mu_k < \infty$. Server s , $s = 1, \dots, S$ works at speed 1. Customers waiting (being served) abandon the queue after an exponentially distributed amount of time with mean $1/\theta_k$ ($1/\tilde{\theta}_k$), with $\theta_k > 0$, $\tilde{\theta}_k \geq 0$, for all k . Having one class- k customers waiting in the queue (in service) costs c_k (\tilde{c}_k) per unit of time. Each abandonment of a waiting class- k customer (class- k customer being served) costs d_k (\tilde{d}_k). We are interested in finding a policy that minimizes the total long-run average cost

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^K \mathbb{E} \left(\int_0^T (c_k X_k^{\pi,0}(t) + \tilde{c}_k X_k^{\pi,1}(t)) dt + d_k R_k^\pi(T) + \tilde{d}_k \tilde{R}_k^\pi(T) \right),$$

where $X_k^{\pi,0}(t)$ ($X_k^{\pi,1}(t)$) denotes the number of class- k customers in the queue (in service) at time t and $R_k^\pi(t)$ ($\tilde{R}_k^\pi(t)$) denotes the number of abandonments of waiting class- k customers (class- k customers being served) in the interval $[0, t]$, under policy π .

Representing each customer in the queue (in service) by a passive (active) bandit, the problem can be addressed within the framework of a restless bandit model with the following parameters: $J_k = 1$, $q_k^0(1, 0) = \theta_k > 0$, $q_k^1(1, 0) = \mu_k + \tilde{\theta}_k$, $C_{1,k}^0 = c_k + d_k \theta_k$, $C_{1,k}^1 = \tilde{c}_k + \tilde{d}_k \tilde{\theta}_k$, $k = 1, \dots, K$, and $\alpha = S$, where we used that $\mathbb{E}(R_k^\pi(T)) = \theta_k \mathbb{E}(\int_0^T X_k^{\pi,0}(t) dt)$ and $\mathbb{E}(\tilde{R}_k^\pi(T)) = \tilde{\theta}_k \mathbb{E}(\int_0^T X_k^{\pi,1}(t) dt)$. A bandit can only be in two states (state 0 or state 1), hence indexability follows directly (for any choice of β).

We can now define an index policy that we will prove to be included in the set Π^* . For each class k we set:

$$\iota_k := q_k^1(1, 0) \left(\frac{C_{1,k}^0}{q_k^0(1, 0)} - \frac{C_{1,k}^1}{q_k^1(1, 0)} \right).$$

The index policy (ι_k) is then defined as follows: At any moment in time serve (at most) S customers present in the system that have the highest, strictly positive, index values, ι_k . In case a customer belongs to a class that has a negative index value, this customer will never be served.

Before continuing, we first give an interpretation of the index ι_k . The term $1/q_k^a(1, 0)$ is the time it takes until a bandit under action a leaves the system. Hence, $C_{1,k}^a/q_k^a(1, 0)$ is the cost for applying action a on a class- k bandit until it leaves the system. The difference $\frac{C_{1,k}^0}{q_k^0(1, 0)} - \frac{C_{1,k}^1}{q_k^1(1, 0)}$ is the reduction in cost when making a class- k bandit active (instead of keeping him passive), so that finally the index ι_k represents the reduction in cost per time unit in case class k is made active. Also note that the index rule (ι_k) is robust, i.e., it does not depend on the arrival rate of the customers or the number of servers present in the system.

Proposition 6.1 *Policy (ι_k) is contained in the set Π^* . In addition, policy (ι_k) is the unique robust policy in the set Π^* , that is, it is independent on S and on λ_k , $k = 1, \dots, K$. Hence, priority policy (ι_k) coincides with Whittle's index policy $(\nu_{j,k}^{lim})$.*

Proof: For the multi-class multi-server system with abandonments the linear program (LP) is given by:

$$\begin{aligned} \min_x \quad & \sum_k (c_k x_k^0 + \tilde{c}_k x_k^1 + d_k \theta_k x_k^0 + \tilde{d}_k \tilde{\theta}_k x_k^1), \\ \text{s.t.} \quad & 0 = \lambda_k - \mu_k x_k^1 - \theta_k x_k^0 - \tilde{\theta}_k x_k^1, \\ & \sum_{k=1}^K x_k^1 \leq S, \text{ and } x_k^0, x_k^1 \geq 0. \end{aligned} \quad (22)$$

Equation (22) implies $x_k^0 = \frac{\lambda_k - (\mu_k + \tilde{\theta}_k)x_k^1}{\theta_k}$. Hence, the above linear program is equivalent to solving

$$\begin{aligned} \max_x \quad & \sum_k ((c_k + d_k \theta_k) \frac{\mu_k + \tilde{\theta}_k}{\theta_k} - \tilde{c}_k - \tilde{d}_k \tilde{\theta}_k) x_k^1, \\ \text{s.t.} \quad & \sum_{k=1}^K x_k^1 \leq S, \text{ and } 0 \leq x_k^1 \leq \frac{\lambda_k}{\mu_k + \tilde{\theta}_k}. \end{aligned}$$

An optimal solution is to assign maximum values to x_k^1 for those classes having the highest values for $(c_k + d_k \theta_k) \frac{\mu_k + \tilde{\theta}_k}{\theta_k} - \tilde{c}_k - \tilde{d}_k \tilde{\theta}_k = \iota_k$, with $\iota_k > 0$, until the constraint $\sum_k x_k^1 \leq S$ is saturated. Denote this optimal solution by x^* . Assume the classes are ordered such that $\iota_1 \geq \iota_2 \geq \dots \geq \iota_K$. Hence, there is an l such that: for all $k < l$ it holds that $x_k^{*,1} = \frac{\lambda_k}{\mu_k + \tilde{\theta}_k}$, and hence $x_k^{*,0} = 0$, for $k = l$ it holds that $0 \leq x_l^{*,1} \leq \frac{\lambda_l}{\mu_l + \tilde{\theta}_l}$ and hence $x_l^{*,0} \geq 0$, and for all $k > l$ it holds that $x_k^{*,1} = 0$. This gives that the index policy (ι_k) is included in the set $\Pi(x^*) \subset \Pi^*$, see Definition 3.4.

It follows directly that (ι_k) is the unique robust policy. For any value of S or λ_k , Whittle's index policy is in the set Π^* (Proposition 4.6). Hence, we obtain that (ι_k) and Whittle's index policy have the same priority ordering. \square

Note that the $M/M/S+M$ system belongs to the class of problems as described in Proposition 5.1. Hence, Condition 3.10 is satisfied. The global attractor property follows in the same way as in [3]. We therefore have the following optimality result for the index policy (ι_k) .

Proposition 6.2 *Consider a system with Sr servers working in parallel and arrival rates $\lambda_k r$, $k = 1, \dots, K$. The index policy (ι_k) is asymptotically optimal as $r \rightarrow \infty$, i.e., for any policy π ,*

$$\begin{aligned} & \lim_{r \rightarrow \infty} \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^K \mathbb{E} \left(\int_0^T ((c_k + d_k \theta_k) x_k^{r, \iota, 0}(t) + (\tilde{c}_k + \tilde{d}_k \tilde{\theta}_k) x_k^{r, \iota, 1}(t)) dt \right) \\ & \leq \lim_{r \rightarrow \infty} \inf \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^K \mathbb{E} \left(\int_0^T ((c_k + d_k \theta_k) x_k^{r, \pi, 0}(t) + (\tilde{c}_k + \tilde{d}_k \tilde{\theta}_k) x_k^{r, \pi, 1}(t)) dt \right). \end{aligned}$$

Proof: In Proposition 6.1 we showed that (ι_k) is included in $\Pi(x^*)$, with x^* as given in the proof of Proposition 6.1. In Appendix H we prove that the process $x^*(t)$ corresponding to policy (ι_k) , as defined in (9), has the point x^* as a unique global attractor, i.e., Condition 3.9 is satisfied. Together with Proposition 5.1 we then obtain that the index policy (ι_k) is asymptotically optimal. \square

Remark 6.3 (Existing results in literature) *In [29] a single-server queue with abandonments has been studied. Whittle's index policy is derived by modeling the system as a fixed population of restless bandits, where each bandit represents a queue, hence, the state of a bandit is the number of customers in the queue. (The latter implies that $J_k = \infty$ for all k , hence it does not fall inside the framework of this paper). The results obtained in [29] hold for general holding cost functions. In the case of linear holding costs, as considered in this section, the index rule as derived in [29] coincides with (ι_k) . In particular, this implies that, even though in [29] the arrival characteristics are taken into account, the there obtained*

index policy is independent on λ_k or on the number of customers present in the system when considering linear holding cost.

For the case $\tilde{c}_k = 0, \tilde{\theta}_k = 0$ and $c_k + d_k \theta_k > 0$, the asymptotic optimality of the policy (ι_k) has been previously proved in [3, 4]. Note that in this setting the performance criterion is the weighted number of customers present in the queue. If $\sum \lambda_k / \mu_k > S$, that is, the overload situation, the fluid-scaled cost $V^*(S)$ will be non-zero, and hence the optimality result is useful. This is not the case when $\sum \lambda_k / \mu_k < S$, the underloaded system, as was also observed in [3, 4]: In underload we have for any non-idling policy $x_k^{*,0} = 0, \forall k$, see Equation (49), which implies $V^*(S) = 0$, that is, in equilibrium the cost is zero for any non-idling policy. Instead, in [28] the transient behavior of the fluid model has been studied for the underload setting. It is shown that the optimal transient fluid control is in fact a state-dependent strategy and hence no longer a strict priority policy.

We further note that for a discrete-time model with one server and for the particular case of $\tilde{\theta}_k = 0, \tilde{c}_k = c_k > 0$, Whittle's index ν_k^{lim} has been derived in [7]. This index coincides with the Whittle's index ι_k for the continuous-time model.³

7 Non-indexable restless bandits

The set of priority policies Π^* , which are asymptotically optimal, consists of more than one policy, and hence it is not direct which priority policy to choose. For an indexable restless bandit problem, Whittle's index policy is inside the set Π^* and is robust, that is, it is independent on the exact values for $\alpha, \lambda_k, x_k(0)$, $k = 1, \dots, K$. This is therefore an obvious choice. In particular, in the literature Whittle's index policy has been extensively tested numerically for different applications and shown to perform well, see for example [2, 1, 6, 7, 14, 22, 23, 29, 31, 37, 41] and examples in the book [20]. In this section we therefore focus our attention on non-indexable restless bandits. In Section 7.1 we describe how to select a priority policy from the (possibly large) set of asymptotically optimal priority policies Π^* and in Section 7.2 its performance is numerically evaluated outside the asymptotic regime.

7.1 Policy selection

In order to define how to select priority policies from the set Π^* , we need the following technical lemma that gives a characterization of an optimal solution of the (LP) problem. Note that this lemma is valid for both indexable and non-indexable examples. We refer to Appendix G for the proof.

Lemma 7.1 *In case of a dynamic population of bandits, assume that the set of optimal solutions of (LP) is bounded and either $p_k(0, j) > 0$, for all j, k , or $C_{j,k}^0 > 0$, for all j, k .*

For either a fixed or dynamic population of bandits, there exists at least one optimal solution of (LP), x^ , such that $x_{j,k}^{*,0} x_{j,k}^{*,1} > 0$ for at most one state (j, k) .*

The assumption (for the dynamic population) on the boundedness of optimal solutions of (LP) is always satisfied in case $C_{j,k}^0 > 0$ for all j, k , since $x_{j,k}^1 \leq \alpha$ and $x_{j,k}^{*,0}$ is upperbounded by the cost value of a feasible solution divided by $C_{j,k}^0 > 0$.

We can now describe a method to select priority policies from the set Π^* as α changes. We note that the set Π^* and an optimal solution x^* depend on α . In the remainder of this section we will therefore write $\Pi^*(\alpha)$ and $x^*(\alpha)$ instead of Π^* and x^* to emphasize the dependence on the α .

We have from Lemma 7.1 that for a fixed α there exists at least one optimal solution of (LP), $x^*(\alpha)$, such that $x_{j,k}^{*,0}(\alpha) x_{j,k}^{*,1}(\alpha) > 0$ for at most one state (j, k) . When choosing a policy from $\Pi^*(\alpha)$, we focus on such a solution $x^*(\alpha)$ and the corresponding set $\Pi(x^*(\alpha))$. Now define $0 = \alpha_0 < \alpha_1 < \alpha_2 < \dots$, such that for $\alpha \in [\alpha_i, \alpha_{i+1})$ the components that have positive $x_{j,k}^{*,0}(\alpha)$ remain unchanged and the same for the components that have positive $x_{j,k}^{*,1}(\alpha)$ (or in other words, the basis of an optimal solution of (LP) does not change in an interval $[\alpha_i, \alpha_{i+1})$). For any $\alpha \in [\alpha_i, \alpha_{i+1})$ we then choose the same priority ordering. We do this in the following way: All states $(j_1, k_1), \dots, (j_H, k_H)$ having $x_{j_h, k_h}^{*,1}(\alpha) > 0$ and $x_{j_h, k_h}^{*,0}(\alpha) = 0$, $h = 1, \dots, H$, receive highest priority, see Definition 3.4. The priority ordering within these states can be chosen such that the ordering changes minimally as α moves to other intervals. The state (j_m, k_m)

³Note that for the model of [7], the fluid scaled cost is always strictly positive ($V^*(S) = 0$ would imply that $x_k = 0$. However this contradicts with Equation (22) which would read $0 = \lambda_k$). Hence the asymptotic optimality result is useful for both an underload as well as for an overload regime.

$\alpha = 1$	Optimal basic solution			
$1 \leq x(0) \leq 2.4$	$x_1^{*,0} = 0, x_1^{*,1} > 0$	$x_2^{*,0} > 0, x_2^{*,1} = 0$	$x_3^{*,0} > 0, x_3^{*,1} = 0$	(5) not binding
$2.4 \leq x(0) \leq 3.6$	$x_1^{*,0} = 0, x_1^{*,1} > 0$	$x_2^{*,0} > 0, x_2^{*,1} > 0$	$x_3^{*,0} > 0, x_3^{*,1} = 0$	(5) binding
$3.6 \leq x(0) \leq 7.36$	$x_1^{*,0} > 0, x_1^{*,1} > 0$	$x_2^{*,0} = 0, x_2^{*,1} > 0$	$x_3^{*,0} > 0, x_3^{*,1} = 0$	(5) binding
$7.36 \leq x(0)$	$x_1^{*,0} > 0, x_1^{*,1} = 0$	$x_2^{*,0} = 0, x_2^{*,1} > 0$	$x_3^{*,0} > 0, x_3^{*,1} = 0$	(5) binding

Table 1: Optimal basic solutions for the (LP) problem

	priority ordering	always passive	name of policy
$x(0) = 1, 2$	1	2,3	prio1
$x(0) = 3$	$1 \succ 2$	3	prio12
$x(0) \geq 4$	$2 \succ 1$	3	prio21

Table 2: Selected priority policies

with $x_{j_m, k_m}^{*,1}(\alpha) > 0$ and $x_{j_m, k_m}^{*,0}(\alpha) > 0$ (there is only one such state) has lower priority than the states associated to $h = 1, \dots, H$, by Definition 3.4. For states with $x_{j_l, k_l}^{*,1}(\alpha) = 0$ and $x_{j_l, k_l}^{*,0}(\alpha) > 0$, we have to distinguish between two situations: (i) in case $\sum_{k=1}^K \sum_{j=1}^{J_k} x_{j,k}^{*,1}(\alpha) < \alpha$, that is, there is capacity left unused, then any bandit with $x_{j_l, k_l}^{*,1}(\alpha) = 0$ and $x_{j_l, k_l}^{*,0}(\alpha) > 0$ will never be made active (by Definition 3.4), (ii) when $\sum_{k=1}^K \sum_{j=1}^{J_k} x_{j,k}^{*,1}(\alpha) = \alpha$, we choose to allow to make (some of) those states (j_l, k_l) active only if the so-resulting final priority ordering appears as well as a priority ordering for values of $\tilde{\alpha}$ that are in an higher interval (only taking into account the states having $x_{j,k}^{*,1}(\tilde{\alpha}) > 0$). That is, such states (j_l, k_l) are made active, only in case this would be an optimal solution for the (LP) problem in case there is more capacity α available.

In the next section the selection method is explained for a particular example.

7.2 Performance evaluation

We now turn our attention to a particular non-indexable example. We took the continuous-time version of the example as given in [36, Section 2.2]. There is a fixed population of bandits and each bandit can be in three states. The cost structure is given by

$$(C_1^0, C_2^0, C_3^0) = (-0.458, -0.5308, -0.6873) \quad \text{and} \quad (C_1^1, C_2^1, C_3^1) = (-0.9631, -0.7963, -0.1057),$$

and $Q^0 = (q^0(i, j))_{i,j}$ and $Q^1 = (q^1(i, j))_{i,j}$ are given by

$$Q^0 = \begin{pmatrix} -0.8098 & 0.4156 & 0.3942 \\ 0.5676 & -0.5809 & 0.0133 \\ 0.0191 & 0.1097 & -0.1288 \end{pmatrix} \quad \text{and} \quad Q^1 = \begin{pmatrix} -0.2204 & 0.0903 & 0.1301 \\ 0.1903 & -0.8137 & 0.6234 \\ 0.2901 & 0.3901 & -0.6802 \end{pmatrix}. \quad (23)$$

We note that the global attractor property is satisfied since a bandit can only be in three states, [48]. In addition, any policy gives a unichain Markov chain, hence Condition 3.10 is satisfied. We therefore have that any priority policy in Π^* is asymptotically optimal when both α and $x(0)$ grow large. Our aim in this section is to numerically evaluate the performance of priority policies in Π^* outside the asymptotic regime. In particular, we evaluate the performance when $\alpha = 1$, that is, at most one bandit can be made active at a time, and we let the number of bandits, $x(0)$, vary.

For a given value of α and $x(0)$, the set Π^* consists of more than one policy. Before presenting the numerical results, we will therefore first describe the priority policies we considered using the selection method as given in the previous section. In Table 1 one can find the structure of an optimal basic solution of (LP), obtained numerically, when fixing $\alpha = 1$ and letting the number of bandits present in the system, $x(0)$, increase. We note that equivalently we could have taken $x(0) = \bar{x}$ fixed and let α decrease, simply by a change of variable in the (LP) problem ($y_{j,k}^\alpha = x_{j,k}^\alpha \bar{x} / x(0)$ and $\alpha = \bar{x} / x(0)$).

We can now characterize the priority policies, see also Table 2. Consider $x(0) = 1$ or $x(0) = 2$. In that case we derive from Table 1 that a bandit in state 1 receives priority (by Definition 3.4). Since the constraint (5) is not binding, a bandit in state 2 or 3 will never be made active (by Definition 3.4). Hence, $\Pi(x^*)$ consists of the policy prio1. Now consider $x(0) = 3$. Then, Definition 3.4 prescribes that, for any

policy in $\Pi(x^*)$, state 1 has strict priority over state 2, and state 3 is either never made active, or has lowest priority. Note that for smaller values of $x(0)$ (equivalent to considering higher values of α), state 3 is not made active either, hence we chose to have state 3 passive, that is, we focus on policy prio12. Now consider $4 \leq x(0) \leq 7$. Then, Definition 3.4 prescribes that, for any policy in $\Pi(x^*)$, state 2 has strict priority over state 1, and state 3 is either never made active, or has lowest priority. Since state 3 is never made active for $x(0) < 4$, we chose to do the same for $4 \leq x(0) \leq 7$, that is, we focus on policy prio21. Now consider $x(0) \geq 8$. Then Definition 3.4 prescribes that, for any policy in $\Pi(x^*)$, state 2 has strict priority and that states 1 and 3 are either never made active or have lowest priority. For smaller values of $x(0)$, class 1 is made active, while class 3 is never made active. Hence, we chose to do the same for $x(0) \geq 8$, that is, we focus on policy prio21.

We have numerically evaluated the performance of the priority policies as described in Table 2 against both the optimal policy (obtained numerically by value iteration) and against other priority policies. In Figure 1 (left) we plot the relative suboptimality gap (in %) for the different policies when $\alpha = 1$ and let the number of bandits, $x(0)$, vary on the horizontal axis. The line referred to as “LP-based” plots for each given $x(0)$ the selected priority policy as given in Table 2. We observe that these selected policies always have the smallest sub-optimality gap.

The labels prio123 and prio213 refer to the priority policies $1 \succ 2 \succ 3$ and $2 \succ 1 \succ 3$. Prio123 and prio213 are inside the class of asymptotically optimal policies, Π^* , for $x(0) = 3$ and $x(0) \geq 4$, respectively, however, the selection process as described in Section 7.1 does not select these policies. In fact, we observe that prio123 and prio213 are outperformed by our selected priority policies. Below we will see that this suboptimality gap can be made arbitrarily large.

The difference in performance between different priority policies is not that large in this example. For other instances of non-indexable bandits, including dynamic populations, the differences can be larger though. For this particular example, we note however that the suboptimality gap can be made arbitrarily large by adequately changing the values for $C_3^1, q_{3,1}^1$ and $q_{3,2}^1$. These parameters do not affect the performance of policies that never make state 3 active (including the priority policies in Table 2), but do influence the performance of prio123 and prio213. By making the cost of being active in state 3, C_3^1 , larger, and the transition rates when being active in state 3 smaller, the performance of these policies degrades. As an example, in Figure 1 (right) we plot the suboptimality gaps when taking $C^1(3) = 5$ and $q_{3,1}^1 = q_{3,2}^1 = 0.001$ and we observe larger optimality gaps. Furthermore, note that for a fixed $x(0)$ the gap will grow linearly in C_3^1 and hence can be made arbitrarily large.




Figure 1: Suboptimality gap of priority policies for non-indexable example

8 Conclusion and further research

We have characterized a class of priority policies that are asymptotically optimal for the general multi-class restless-bandit problem. We studied both the setting of a fixed population of bandits as well as

a dynamic population of bandits. In both cases we showed that Whittle’s index policy is included in the class of asymptotically optimal policies. Our result extends that of Weber and Weiss [47] in several directions: (i) we allow various classes of restless bandits and multiple actions, (ii) we allow arrivals of new restless bandits to the system, and (iii) we do not need the restless bandits to be indexable in order to define asymptotically optimal policies. In order to obtain the asymptotic optimality results we combined fluid-scaling techniques with linear programming results. This is a different approach than that taken in [47] and in particular allows us to include arrivals of bandits to the system.

We considered a slight variation of the standard restless bandit formulation: Instead of having at each moment in time exactly α bandits active, we allowed strictly less than α bandits to be active at a time. By introducing so-called dummy bandits we showed that restless bandits having a negative Whittle’s index are never made active under an optimal policy of the relaxed optimization problem or under Whittle’s index policy.

As future work it would be interesting to further understand the necessity of the global attractor property in order for asymptotic optimality results to hold. In addition, it would be interesting to investigate whether Condition 3.10 would hold in greater generality for restless bandit problems. So far, we showed it to hold in case any passive bandit will eventually leave. A further interesting thread would be to estimate the suboptimality gap of the proposed priority policies outside the fluid-scaling regime.

Acknowledgements

The author is grateful to Urtzi Ayesta, Nicolas Gast, Peter Jacko, José Niño-Mora, and Philippe Robert for valuable discussions and comments.

References

- [1] S.H.A. Ahmad, M. Liu, T. Javidi, Q. Zhao, and B. Krishnamachari. Optimality of myopic sensing in multichannel opportunistic access. *IEEE Transactions on Information Theory*, 55:4040–4050, 2009.
- [2] P.S. Ansell, K.D. Glazebrook, J. Niño-Mora, and M. O’Keeffe. Whittle’s index policy for a multi-class queueing system with convex holding costs. *Mathematical Methods of Operations Research*, 57:21–39, 2003.
- [3] R. Atar, C. Giat, and N. Shimkin. The $c\mu/\theta$ rule for many-server queues with abandonment. *Operations Research*, 58(5):1427–1439, 2010.
- [4] R. Atar, C. Giat, and N. Shimkin. On the asymptotic optimality of the $c\mu/\theta$ rule under ergodic cost. *Queueing Systems*, 67(2):127–144, 2011.
- [5] U. Ayesta, M. Erausquin, and P. Jacko. A modeling framework for optimizing the flow-level scheduling with time-varying channels. *Performance Evaluation*, 67:1014–1029, 2010.
- [6] U. Ayesta, M. Erausquin, M. Jonckheere, and I.M. Verloop. Scheduling in a random environment: stability and asymptotic optimality. *IEEE/ACM Transactions on Networking*, 21(1):258–271, 2013.
- [7] U. Ayesta, P. Jacko, and V. Novak. A nearly-optimal index rule for scheduling of users with abandonment. In *Proceedings of IEEE INFOCOM*, Hong Kong, 2011.
- [8] M. Benaïm and J-Y Le Boudec. A class of mean field interaction models for computer and communication systems. *Performance Evaluation*, 65:823–838, 2008.
- [9] D. Bertsimas and J. Niño-Mora. Restless bandits, linear programming relaxations, and a primal-dual index heuristic. *Operations Research*, 48(1):80–90, 2000.
- [10] P. Billingsley. *Convergence of Probability Measures*. Wiley, New York, 1999.
- [11] M.J. Cánovas, M.A. López, and J. Parra. On the continuity of the optimal value in parametric linear optimization: Stable discretization of the Lagrangian dual of nonlinear problems. *Set-Valued Analysis*, 13:69–84, 2005.
- [12] E. Çinlar. *Introduction to Stochastic Processes*. Prentice-Hall, New Jersey, 1975.

- [13] J.G. Dai and S. He. Many-server queues with customer abandonment: A survey of diffusion and fluid approximations. *Journal of Systems Science and Systems Engineering*, 21:1–36, 2012.
- [14] N. Ehsan and M. Liu. On the optimality of an index policy for bandwidth allocation with delayed state observation and differentiated services. In *Proceedings of IEEE INFOCOM*, Hong Kong, 2004.
- [15] S.N. Ethier and T.G. Kurtz. *Markov Processes: Characterization and Convergence*. Wiley, New York, 1986.
- [16] G. Gallego, G. Iyengar, R. Phillips, and A. Dubey. Managing flexible products on a network. CORC technical report TR-2004-01, Columbia University, New York. 2004.
- [17] N. Gast and B. Gaujal. A mean field model of work stealing in large-scale systems. In *Proceedings of ACM SIGMETRICS*, pages 13–24, New York NY, USA, 2010.
- [18] J.C. Gittins. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society. Series B*, 41(2):148–177, 1979.
- [19] J.C. Gittins. *Multi-Armed Bandit Allocation Indices*. Wiley, Chichester, 1989.
- [20] J.C. Gittins, K.D. Glazebrook, and R.R. Weber. *Multi-Armed Bandit Allocation Indices*. Wiley, Chichester, 2011.
- [21] K.D. Glazebrook, D.J. Hodge, and C. Kirkbride. General notions of indexability for queueing control and asset management. *Annals of Applied Probability*, 21:876–907, 2011.
- [22] K.D. Glazebrook, C. Kirkbride, and J. Ouenniche. Index policies for the admission control and routing of impatient customers to heterogeneous service stations. *Operations Research*, 57:975–989, 2009.
- [23] K.D. Glazebrook and H.M. Mitchell. An index policy for a stochastic scheduling model with improving/deteriorating jobs. *Naval Research Logistics*, 49:706–721, 2002.
- [24] X. Guo, O. Hernández-Lerma, and T. Prieto-Rumeau. A survey of recent results on continuous-time Markov decision processes. *TOP*, 14:177–261, 2006.
- [25] J. Hasenbein and D. Perry (Eds). Special issue on queueing systems with abandonments. *Queueing Systems*, 75:111–384, 2013.
- [26] G. Iyengar and K. Sigman. Exponential penalty function control of loss networks. *Annals of Applied Probability*, 14(4):1698–1740, 2004.
- [27] P. Jacko. Optimal index rules for single resource allocation to stochastic dynamic competitors. In *Proceedings of ValueTools*, 2011.
- [28] M. Larrañaga, U. Ayesta, and I.M. Verloop. Dynamic fluid-based scheduling in a multi-class abandonment queue. In *Proceedings of Performance*, Vienna, Austria, 2013.
- [29] M. Larrañaga, U. Ayesta, and I.M. Verloop. Index policies for multi-class queues with convex holding cost and abandonments. In *Proceedings of ACM SIGMETRICS*, Austin TX, USA, 2014.
- [30] G. Liu and G. van Ryzin. On the choice-based linear programming model for network revenue management. *Manufacturing & Service Operations Management*, 10(2):288–310, 2008.
- [31] K. Liu and Q. Zhao. Indexability of restless bandit problems and optimality of Whittle index for dynamic multichannel access. *IEEE Transactions on Information Theory*, 56:5547–5567, 2010.
- [32] A. Mahajan and D. Teneketzis. Multi-armed bandit problems. In *Foundations and Application of Sensor Management*, eds. A.O. Hero III, D.A. Castanon, D. Cochran and K. Kastella., pages 121–308, Springer-Verlag, 2007.
- [33] S.P. Meyn and R.L. Tweedie. *Markov Chains and Stochastic Stability*. Cambridge University Press, 2009.

- [34] J. Niño-Mora. Restless bandits, partial conservation laws and indexability. *Advances in Applied Probability*, 33(1):76–98, 2001.
- [35] J. Niño-Mora. Characterization and computation of restless bandit marginal productivity indices. In *ACM SMCTools*, New York, USA, 2007.
- [36] J. Niño-Mora. Dynamic priority allocation via restless bandit marginal productivity indices. *TOP*, 15:161–198, 2007.
- [37] J. Niño-Mora. Marginal productivity index policies for admission control and routing to parallel multi-server loss queues with reneging. *Lecture Notes in Computer Science*, 4465:138–149, 2007.
- [38] W. Ouyang, A. Eryilmaz, and N.B. Shroff. Asymptotically optimal downlink scheduling over Markovian fading channels. In *Proceedings of IEEE INFOCOM*, Orlando FL, USA, 2012.
- [39] D.G. Pandalis and D. Teneketzis. On the optimality of the Gittins index rule for multi-armed bandits with multiple plays. *Mathematical Methods of Operations Research*, 50:449–461, 1999.
- [40] M.L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, New York, 1994.
- [41] V. Raghunathan, V. Borkar, M. Cao, and P.R. Kumar. Index policies for real-time multicast scheduling for wireless broadcast systems. In *Proceedings of IEEE Conf. Comput. Commun.*, pages 1570–1578, 2008.
- [42] P. Robert. *Stochastic Networks and Queues*. Springer-Verlag, New York, 2003.
- [43] A.N. Rybko and A.L. Stolyar. Ergodicity of stochastic processes describing the operation of open queueing networks. *Problems of Information Transmission*, 28:199–220, 1992.
- [44] H.C. Tijms. *A First Course in Stochastic Models*. Wiley, England, 2003.
- [45] I.M. Verloop and R. Núñez-Queija. Assessing the efficiency of resource allocations in bandwidth-sharing networks. *Performance Evaluation*, 66:59–77, 2009.
- [46] R.R. Weber. Comments on: Dynamic priority allocation via restless bandit marginal productivity indices. *TOP*, 15:211–216, 2007.
- [47] R.R. Weber and G. Weiss. On an index policy for restless bandits. *Journal of Applied Probability*, 27:637–648, 1990.
- [48] R.R. Weber and G. Weiss. Addendum to “On an index policy for restless bandits”. *Journal of Applied Probability*, 23:429–430, 1991.
- [49] G. Weiss. Branching bandit processes. *Probability in the Engineering and Informational Sciences*, 2:269–278, 1988.
- [50] P. Whittle. Arm-acquiring bandits. *The Annals of Probability*, 9(2):284–292, 1981.
- [51] P. Whittle. Restless bandits: Activity allocation in a changing world. *Journal of Applied Probability*, 25:287–298, 1988.
- [52] P. Whittle. *Optimal Control, Basics and Beyond*. John Wiley & Sons, 1996.

Appendix A: Proof of Lemma 3.1

Let π be a policy for which a unique invariant distribution exists having finite first moment. Let $X_k(0) = x_k(0)$.

For a fixed population, we have that $\lim_{t \rightarrow \infty} \frac{X_{j,k}^\pi(t)}{t} \leq \lim_{t \rightarrow \infty} \frac{X_k(0)}{t} = 0$, for all j, k . For a dynamic population, stability of policy π implies rate-stability, that is, $\lim_{t \rightarrow \infty} \frac{X_{j,k}^\pi(t)}{t} = 0$, for all j, k . Hence, in both cases we conclude that

$$\lim_{t \rightarrow \infty} \frac{X_{j,k}^\pi(t)}{t} = 0, \text{ for all } j, k. \quad (24)$$

Note that $\int_0^t X_{j,k}^{\pi,a}(s)ds$ is the total aggregated amount of time spent on action a on class- k bandits in state j during the interval $(0, t]$. Hence, we can write the following sample-path construction of the process $X_{j,k}^\pi(t)$:

$$X_{j,k}^\pi(t) = X_{j,k}^\pi(0) + N^{\lambda_k p_k(0,j)}(t) + \sum_{a=0}^1 \sum_{i=1, i \neq j}^{J_k} N^{q_k^a(i,j)} \left(\int_0^t X_{i,k}^{\pi,a}(s)ds \right) - \sum_{a=0}^1 N^{q_k^a(j)} \left(\int_0^t X_{j,k}^{\pi,a}(s)ds \right), \quad (25)$$

where $N^{\lambda_k p_k(0,j)}(t)$, $N^{q_k^a(i,j)}(t)$ and $N^{q_k^a(j)}(t)$ are independent Poisson processes having as rates $\lambda_k p_k(0, j)$, $q_k^a(i, j)$ and $q_k^a(j)$, respectively, $i, j = 1, \dots, J_k$, $k = 1, \dots, K$, $a = 0, 1$. By the ergodic theorem [12], we obtain that $\frac{1}{t} \int_0^t X_{j,k}^{\pi,a}(s)ds$ converges to the mean, denoted by $\bar{X}_{j,k}^{\pi,a} < \infty$, for all j, k, a . Hence, when dividing both sides in (25) by t , using that $N^\theta(at)/t \rightarrow a\theta$ as $t \rightarrow \infty$, and together with (24), we obtain that

$$0 = \lambda_k p_k(0, j) + \sum_{a=0}^1 \sum_{i=1, i \neq j}^{J_k} q_k^a(i, j) \bar{X}_{i,k}^{\pi,a} - \sum_{a=0}^1 \bar{X}_{j,k}^{\pi,a} q_k^a(j), \text{ a.s.},$$

that is, \bar{X}^π satisfies Equation (4). By definition, \bar{X}^π satisfies $\sum_{k,j} \bar{X}_{j,k}^{\pi,1} \leq \alpha$, $\bar{X}_{j,k}^{\pi,a} \geq 0$ and if $\lambda_k = 0$, then $\sum_{j=1}^{J_k} \sum_{a=0}^1 \bar{X}_{j,k}^{\pi,a} = x_k(0)$. Hence, \bar{X}^π is a feasible solution of (LP).

Since the feasible set is non-empty and the objective is to minimize the cost, the optimal value satisfies $V^*(\alpha) < \infty$. \square

Appendix B: Proof of Lemma 3.3

By Fatou's lemma we have

$$\sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left(\int_0^T C_{j,k}^a X_{j,k}^{\pi,a}(t) dt \right) \geq \sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \mathbb{E} \left(\liminf_{T \rightarrow \infty} \frac{1}{T} \int_0^T C_{j,k}^a X_{j,k}^{\pi,a}(t) dt \right).$$

Hence, it is sufficient to prove that

$$\sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \liminf_{T \rightarrow \infty} \frac{1}{T} \int_0^T C_{j,k}^a X_{j,k}^{\pi,a}(t) dt \geq V^*(\alpha), \text{ almost surely.} \quad (26)$$

We consider a fixed realization ω of the process. If $\sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \liminf_{T \rightarrow \infty} \frac{1}{T} \int_0^T C_{j,k}^a X_{j,k}^{\pi,a}(t) dt = \infty$, then (26) is trivially true, since $V^*(\alpha) < \infty$ (see Lemma 3.1). Hence, assume this is not the case, and consider the subsequence t_n corresponding to the liminf sequence in (26). Note that for a fixed population we have that since $X_{j,k}^{\pi,a}(t) \leq X_k(0)$, for all t , $\frac{1}{T} \int_0^T X_{j,k}^{\pi,a}(t) dt$ are bounded sequences. Hence, there is a subsequence t_{n_l} of t_n such that $\frac{1}{t_{n_l}} \int_0^{t_{n_l}} C_{j,k}^a X_{j,k}^{\pi,a}(t) dt$ converges to a constant $\bar{X}_{j,k}^{\pi,a}$, for all j, k, a . In the case of a dynamic population, given the policy π is stable, we have by the ergodicity theorem [12] that $\frac{1}{T} \int_0^T X_{j,k}^{\pi,a}(t) dt$ converges to the mean, here denoted by $\bar{X}_{j,k}^{\pi,a}$.

In addition, it holds that $\lim_{t \rightarrow \infty} X_{j,k}^{\pi,a}(t)/t \leq 0$, for all j, k, a . For the fixed population this follows since $\lim_{t \rightarrow \infty} X_{j,k}^{\pi,a}(t)/t \leq \lim_{t \rightarrow \infty} X_k(0)/t = 0$, and for the dynamic population this follows since any stable policy is rate stable.

Hence, when studying (25) in the point t_{n_l} , dividing both sides by t_{n_l} and using that $N^\theta(t)/t \rightarrow \theta$ as $t \rightarrow \infty$, we obtain

$$0 = \lambda_k p_k(0, j) + \sum_{a=0}^1 \sum_{i=1, i \neq j}^{J_k} q_k^a(i, j) \bar{X}_{i,k}^{\pi,a} - \sum_{a=0}^1 \bar{X}_{j,k}^{\pi,a} q_k^a(j).$$

By (2) we also have that $\sum_{k=1}^K \sum_{j=1}^{J_k} \bar{X}_{j,k}^{\pi,1} \leq \alpha$, hence \bar{X}^π is a feasible solution of (LP). We can now conclude that

$$\begin{aligned} \sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \liminf_{T \rightarrow \infty} \frac{1}{T} \int_0^T C_{j,k}^a X_{j,k}^{\pi,a}(t) dt &= \sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \lim_{l \rightarrow \infty} \frac{1}{t_{n_l}} \int_0^{t_{n_l}} C_{j,k}^a X_{j,k}^{\pi,a}(t) dt \\ &= \sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 C_{j,k}^a \bar{X}_{j,k}^{\pi,a} \geq V^*(\alpha), \end{aligned}$$

which proves (8).

We now assume that π is rate-stable and that $C_{j,k}^a > 0$, for all j, k, a . Again we consider a fixed realization ω of the process. If $\sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \liminf_{T \rightarrow \infty} \frac{1}{T} \int_0^T C_{j,k}^a X_{j,k}^{\pi,a}(t) dt = \infty$, then the result is trivially true, since $V^*(\alpha) < \infty$ (see Lemma 3.1). Hence, assume this is not the case, and consider the subsequence t_n corresponding to the liminf sequence. So

$$\sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \lim_{n \rightarrow \infty} \frac{1}{t_n} \int_0^{t_n} C_{j,k}^a X_{j,k}^{\pi,a}(t) dt < \infty. \quad (27)$$

Since $C_{j,k}^a > 0$, this implies that the sequence $\frac{1}{t_n} \int_0^{t_n} X_{j,k}^{\pi,a}(t) dt$ is bounded, for all j, k, a . By the Bolzano-Weierstrass theorem, there exists a subsubsequence t_{n_l} of t_n and values $\bar{X}_{j,k}^{\pi,a}$'s such that $\lim_{l \rightarrow \infty} \frac{1}{t_{n_l}} \int_0^{t_{n_l}} X_{j,k}^{\pi,a}(t) dt = \bar{X}_{j,k}^{\pi,a}$, for all j, k, a . In addition, by rate stability we have that $\lim_{t \rightarrow \infty} X_{j,k}^{\pi,a}(t)/t = 0$, a.s., for all j, k, a . The proof follows now again as in the fixed population case.

The proof in the case of mean-rate stability goes along similar lines as that for rate stability and is therefore not included here. \square

Appendix C: Proof of Proposition 3.11

We denote by $S_{j,k}^*$ the set of all states (i, l) , $i = 1, \dots, J_l$, $l = 1, \dots, K$, such that class- l bandits in state i have higher priority than class- k bandits in state j under policy π^* , and I^* is the set of all states that will never be made active under policy π^* . The transition rates of the process $x^{r,\pi^*}(t)$ are then defined as follows:

$$x \rightarrow x + \frac{e_{j,k}}{r} \quad \text{at rate } r \lambda_k p_k(0, j), \quad k = 1, \dots, K, \quad j = 1, \dots, J_k, \quad (28)$$

$$x \rightarrow x - \frac{e_{j,k}}{r} \quad \text{at rate } r \sum_{a=0}^1 x_{j,k}^a q_k^a(j, 0), \quad k = 1, \dots, K, \quad j = 1, \dots, J_k, \quad (29)$$

$$x \rightarrow x - \frac{e_{j,k}}{r} + \frac{e_{i,k}}{r} \quad \text{at rate } r \sum_{a=0}^1 x_{j,k}^a q_k^a(j, i), \quad k = 1, \dots, K, \quad i, j = 1, \dots, J_k, \quad i \neq j, \quad (30)$$

where $x_{j,k}^1 = \min\left((\alpha - \sum_{(i,l) \in S_{j,k}^*} x_{i,l})^+, x_{j,k}\right)$, if $(j, k) \notin I^*$, and $x_{j,k}^1 = 0$ otherwise, $x_{j,k}^0 = x_{j,k} - x_{j,k}^1$, and $e_{j,k}$ is a vector composed of all zeros except for component (j, k) which is one.

From (28)–(30) it follows that there exists a continuous function $b_l(x)$, with $l \in \mathcal{L}$ and \mathcal{L} composed of a finite number of vectors in $\mathbb{N}^{\sum_k J_k}$, such that the transition rates of the process $x^{r,\pi^*}(t)$ from x to $x + l/r$ have the form $r b_l(x)$. Hence, the process $x^{r,\pi^*}(t)$ belongs to the family of density dependent population processes as defined in [15, Chapter 11].

Note that the process $x^*(t)$ as defined in (9) can equivalently be written as $\frac{dx^*(t)}{dt} = F(x^*(t))$, with $F(x^*) = \sum_{l \in \mathcal{L}} l b_l(x^*)$. We note that $F(\cdot)$ is Lipschitz continuous. From Condition 3.9 we have that x^* is the unique global attractor of $x^*(t)$.

The family $\{p^{r,\pi^*}\}$ being tight, together with Condition 3.9 and [17, Theorem 4], it follows that $p^{r,\pi^*}(x)$ converges to the Dirac measure in x^* , the global attractor of $x^*(t)$. Hence, we can write

$$\begin{aligned} & \sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \lim_{r \rightarrow \infty} \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left(\int_0^T C_{j,k}^a x_{j,k}^{r,\pi^*,a}(t) dt \right) \\ &= \sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \lim_{r \rightarrow \infty} \sum_x p^{r,\pi^*}(x) C_{j,k}^a x_{j,k}^a = \sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 C_{j,k}^a x_{j,k}^{*,a} = V^*(\alpha), \end{aligned}$$

where the first step follows from the ergodicity theorem [44, 12] (applicable since the first moment of p^{r,π^*} is finite), the second step (interchange of limit and summation) follows from uniform integrability of $\{p^{r,\pi^*}\}$ and the fact that p^{r,π^*} converges to the Dirac measure in x^* , and the last step follows since x^* is an optimal solution of (LP).

We conclude the proof by noting that $V^*(\alpha)$ is a lower bound on the steady-state cost, as shown in Lemma 3.3. \square

Appendix D: Proof of Proposition 4.3

Recall that the relaxed optimization problem for $f = av$ consists in finding a stationary and Markovian policy that minimizes

$$\sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^1 \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left(\int_0^T C_{j,k}^a X_{j,k}^{\pi,a}(t) dt \right), \quad (31)$$

under the relaxed constraint

$$\sum_{k=1}^K \sum_{j=1}^{J_k} \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left(\int_0^T X_{j,k}^{\pi,1}(t) dt \right) \leq \alpha. \quad (32)$$

For a given policy π , we denote by $x_{j,k}^{\pi,a}$ the (stationary) state-action frequencies, that is, the average fraction of time the class- k bandit is in state j and action a is chosen. Assumption 4.2 implies that these frequencies exist and satisfy the balance equations, that is, they satisfy

$$0 = \sum_{a=0}^1 q_k^a(j) x_{j,k}^{\pi,a} - \sum_{a=0}^1 \sum_{i=1, i \neq j}^{J_k} q_k^a(i, j) x_{i,k}^{\pi,a}, \quad \forall j.$$

We restrict ourselves to the class of policies that are symmetric for bandits in the same class⁴. Having $x_k(0)$ bandits in class k , Equations (31) and (32) can then equivalently be written as

$$\sum_{k=1}^K x_k(0) \sum_{j=1}^{J_k} \left(C_{j,k}^0 x_{j,k}^{\pi,0} + C_{j,k}^1 x_{j,k}^{\pi,1} \right) \quad \text{and} \quad \sum_{k=1}^K x_k(0) \sum_{j=1}^{J_k} x_{j,k}^{\pi,1} \leq \alpha,$$

respectively. We can now formulate the relaxed optimization problem as the following linear program (D):

$$\begin{aligned} (D) \quad & \min_x \sum_{k=1}^K x_k(0) \sum_{j=1}^{J_k} \left(C_{j,k}^0 x_{j,k}^0 + C_{j,k}^1 x_{j,k}^1 \right) \\ \text{s.t.} \quad & 0 = \sum_{a=0}^1 q_k^a(j) x_{j,k}^a - \sum_{a=0}^1 \sum_{i=1, i \neq j}^{J_k} q_k^a(i, j) x_{i,k}^a, \quad \forall j, k, \\ & \sum_{k=1}^K x_k(0) \sum_{j=1}^{J_k} x_{j,k}^1 \leq \alpha, \\ & \sum_{j=1}^{J_k} \sum_{a=0}^1 x_{j,k}^a = 1, \quad \forall k, \quad x_{j,k}^a \geq 0, \quad \forall k, j, a. \end{aligned} \quad (33)$$

We have that for any feasible solution $(x_{j,k}^a)$ of (D) there is a stationary policy π such that the state-action frequencies $x_{j,k}^{\pi,a}$ coincide with the value of the feasible solution $x_{j,k}^a$ [40, Theorem 8.8.2 b)]. Hence, for any optimal (symmetric) policy π^* of the relaxed optimization problem, the state-action frequencies $x_{j,k}^{\pi^*,a}$ provide an optimal solution of (D). We further note that $(x_{j,k}^{\pi^*} x_k(0))$ is an optimal solution of (LP).

We assume the restless bandit problem is indexable. Hence, an optimal policy of the relaxed optimization problem is described in Section 4.1, and will be denoted here by $\tilde{\pi}^*$. We recall that policy $\tilde{\pi}^*$ is described by a value $\nu^* \geq 0$ and is such that a class- k bandit in state j is served if $\nu_{j,k}^{av} > \nu^*$ and is not served if $\nu_{j,k}^{av} < \nu^*$. Hence, the state-action frequencies under $\tilde{\pi}^*$ satisfy

$$\begin{aligned} x_{j,k}^{\tilde{\pi}^*,0} &= 0 \quad \text{when } \nu_{j,k}^{av} > \nu^*, \\ x_{j,k}^{\tilde{\pi}^*,1} &= 0 \quad \text{when } \nu_{j,k}^{av} < \nu^*. \end{aligned} \quad (34)$$

By definition of policy $\tilde{\pi}^*$, for states (\hat{j}, \hat{k}) with $\nu_{\hat{j},\hat{k}}^{av} = \nu^*$ a class- \hat{k} bandit in state \hat{j} is made active with a certain probability, hence $x_{\hat{j},\hat{k}}^{\tilde{\pi}^*,0} \geq 0$ and $x_{\hat{j},\hat{k}}^{\tilde{\pi}^*,1} \geq 0$.

⁴We can do this without loss of generality, since this is the case for the optimal solution of the relaxed problem as given by Whittle.

Since Whittle's index policy gives priority to bandits having highest index value, we directly obtain that Whittle's index policy ($\nu_{j,k}^{av}$) satisfies points 1 and 2 of Definition 3.4 when setting $x^* = (x_{j,k}^{\tilde{\pi}^*} x_k(0))$. We now treat point 3 of Definition 3.4: Assume $\sum_{k=1}^K \sum_{j=1}^{J_k} x_{j,k}^{\tilde{\pi}^*,1} x_k(0) < \alpha$. Hence, under the optimal policy, on average, strictly less than α bandits are made active. This implies that the remaining fraction of the time the policy makes dummy bandits in state B active. Hence, $\nu_B^{av} \geq \nu^*$. Since $\nu^* \geq 0$ and $\nu_B^{av} = 0$ we necessarily have $\nu^* = 0$. A policy satisfies point 3 of Definition 3.4 if it never makes a class- k bandit in state j active that satisfies

$$x_{j,k}^{\tilde{\pi}^*,1} = 0 \text{ and } x_{j,k}^{\tilde{\pi}^*,0} > 0. \quad (35)$$

From (34) (with $\nu^* = 0$), we obtain that (35) implies $\nu_{j,k}^{av} \leq 0$. By definition of Whittle's index policy, a bandit in a state such that $\nu_{j,k}^{av} \leq 0$ will never be made active, hence point 3 is satisfied. We therefore conclude that Whittle's index policy ($\nu_{j,k}^{av}$) is included in the set of priority policies $\Pi(x^*) \subset \Pi^*$, with $x^* = (x_{j,k}^{\tilde{\pi}^*} x_k(0))$. \square

Appendix E: Proof of Proposition 4.6

Let $\beta \leq \bar{\beta}$ and $\beta > 0$. Whittle's index $\nu_{j,k}^\beta$ results from solving the following problem for a class- k bandit:

$$\min_{A_k(\cdot)} \mathbb{E} \left(\int_0^\infty e^{-\beta t} (C_{J_k(t),k}^{A_k(t)} + \nu \mathbf{1}_{(A_k(t)=1)}) dt \right), \quad (36)$$

see (18), where $A_k(t) \in \{0, 1\}$ and $J_k(t)$ denotes the state of the class- k bandit. This is a continuous-time discounted Markov decision problem in a finite state space. After uniformization ([24, Remark 3.1], [40, Section 11.5.2]) this is equivalent to a *discrete-time* discounted Markov decision problem with discount factor $\tilde{\beta} = \frac{\bar{q}}{\beta + \bar{q}}$, cost function $\tilde{C}_{j,k}^a = \frac{C_{j,k}^a + \nu \mathbf{1}_{(a=1)}}{\beta + \bar{q}}$, and transition probabilities $\tilde{p}_k^a(i, j) = \frac{q_k^a(i, j)}{\bar{q}} + \mathbf{1}_{(i=j)}$ (recall that $q_k^a(i, i) = -q_k^a(i) = -\sum_{j=0, i \neq j}^{J_k} q_k^a(i, j)$), where $\bar{q} := \max_{i,k,a} q_k^a(i) < \infty$. In LP formulation the discrete-time MDP for the class- k bandit is then as follows (see [40, Section 6.9]):

$$\begin{aligned} & \max_v \sum_{j=1}^{J_k} \gamma_{j,k} v(j) \\ \text{s.t. } & v(i) - \tilde{\beta} \sum_{j=0}^{J_k} \tilde{p}_k^a(i, j) v(j) \leq \tilde{C}_{i,k}^a, \quad \forall i = 1, \dots, J_k, \quad a = 0, 1, \end{aligned}$$

with $\gamma_{j,k} > 0$ arbitrary. In fact, we will make the choice $\gamma_{j,k} = \lambda_k(p_{0j}^k + \epsilon)$, with $\epsilon > 0$. The dual of the above LP is

$$\begin{aligned} (D_k(\beta, \epsilon)) \quad & \min_x \sum_{j=1}^{J_k} \frac{C_{j,k}^0 x_{j,k}^0 + C_{j,k}^1 x_{j,k}^1 + \nu x_{j,k}^1}{\beta + \bar{q}} \\ \text{s.t. } & 0 = \lambda_k(p_k(0, j) + \epsilon) + \sum_{a=0}^1 \sum_{i=1, i \neq j}^{J_k} \frac{q_k^a(i, j)}{\beta + \bar{q}} x_{i,k}^a - \frac{\beta}{\beta + \bar{q}} \sum_{a=0}^1 x_{j,k}^a - \sum_{a=0}^1 \frac{q_k^a(j)}{\beta + \bar{q}} x_{j,k}^a, \quad \forall j \\ & x_{j,k}^a \geq 0, \quad \forall j, a. \end{aligned} \quad (37)$$

As stated in Section 4.1 ("Optimal solution per bandit"), indexability implies that an optimal policy for the subproblem (36) is described by a priority ordering according to the indices $\nu_{j,k}^\beta$: an optimal action in state j is $a = 1$ if $\nu_{j,k}^\beta > \nu$ and $a = 0$ if $\nu_{j,k}^\beta < \nu$. By [40, Theorem 6.9.4], this implies that there exists an optimal solution to $(D_k(\beta, \epsilon))$, denoted by $x_k^*(\beta, \epsilon)$, such that

$$\begin{aligned} x_{j,k}^{*,0}(\beta, \epsilon) &= 0 \quad \text{when } \nu_{j,k}^\beta > \nu, \\ x_{j,k}^{*,1}(\beta, \epsilon) &= 0 \quad \text{when } \nu_{j,k}^\beta < \nu. \end{aligned}$$

Since $\lim_{l \rightarrow \infty} \nu_{j,k}^{\beta_l} = \nu_{j,k}^{lim}$, we obtain that there exists an $L(\nu)$ such that for all $l > L(\nu)$ it holds that

$$x_{j,k}^{*,0}(\beta_l, \epsilon) = 0 \quad \text{when } \nu_{j,k}^{lim} > \nu, \quad (38)$$

$$x_{j,k}^{*,1}(\beta_l, \epsilon) = 0 \quad \text{when } \nu_{j,k}^{lim} < \nu. \quad (39)$$

By change of variable $\tilde{x}_{j,k}^a = x_{j,k}^a / (\beta + \bar{q})$ we obtain that $\tilde{x}_k^*(\beta_l, \epsilon)$ satisfies (38) and (39) and is an optimal solution of $(\tilde{D}_k(\beta_l, \epsilon))$ defined as:

$$\begin{aligned}
(\tilde{D}_k(\beta, \epsilon)) \quad & \min_{\tilde{x}} \sum_{j=1}^{J_k} (C_{j,k}^0 \tilde{x}_{j,k}^0 + C_{j,k}^1 \tilde{x}_{j,k}^1 + \nu \tilde{x}_{j,k}^1) \\
\text{s.t.} \quad & 0 = \lambda_k(p_k(0, j) + \epsilon) + \sum_{a=0}^1 \sum_{i=1, i \neq j}^{J_k} q_k^a(i, j) \tilde{x}_{i,k}^a - \beta \sum_{a=0}^1 \tilde{x}_{j,k}^a - \sum_{a=0}^1 q_k^a(j) \tilde{x}_{j,k}^a, \quad \forall j, \quad (40) \\
& \tilde{x}_{j,k}^a \geq 0, \quad \forall j, a.
\end{aligned}$$

By Assumption 4.5 we have that the set of optimal solutions of $(\tilde{D}_k(0, 0))$ is bounded and non-empty when $\nu > 0$. Hence, from [11, Corollary 1] we obtain that the correspondence that gives for each (β, ϵ) the set of optimal solutions of $(\tilde{D}_k(\beta, \epsilon))$ is upper semicontinuous in the point $(\beta, \epsilon) = (0, 0)$. It is a compact-valued correspondence (after summing (40) over all j , we have that $\tilde{x}_k = \lambda_k(1 + \epsilon J_k) / \beta, \beta > 0$). Hence, it follows that there exists a sequence $(\beta_{l_n}, \epsilon_n)$ (with β_{l_n} a subsequence of β_l and $\epsilon_n \rightarrow 0$) such that $\tilde{x}_{j,k}^{*,a}(\beta_{l_n}, \epsilon_n) \rightarrow \tilde{x}_{j,k}^{*,a}$, as $n \rightarrow \infty$, and with \tilde{x}_k^* an optimal solution of $(\tilde{D}_k(0, 0))$. For a fixed ν , the components of $\tilde{x}^*(\beta_l, \epsilon)$ that are zero are independent of the exact values for $\epsilon > 0$, and $l > L(\nu)$, see (38) and (39). Hence, the limit \tilde{x}^* , which is an optimal solution of $(\tilde{D}_k(0, 0))$, has the same components equal to zero, i.e., (38) and (39) are satisfied for \tilde{x}^* .

Below we will show that there exists a value ν^* such that there is a vector \tilde{y}^* that satisfies the following: (i) \tilde{y}_k^* is an optimal solution of $(\tilde{D}_k(0, 0))$, for all k , with $\nu = \nu^*$, (ii) \tilde{y}^* is an optimal solution of (LP), and (iii) the Whittle index policy $(\nu_{j,k}^{lim})$ is included in the set $\Pi(\tilde{y}^*) \in \Pi^*$. The latter then concludes the proof.

In the remainder of the proof we denote by $\tilde{x}_k^*(\nu)$ the above-described optimal solution \tilde{x}_k^* of $(\tilde{D}_k(0, 0))$ for a given value ν . We have the following properties:

- **Property 1:**

$$\sum_{k=1}^K \sum_{j=1}^{J_k} \tilde{x}_{j,k}^{*,1}(\infty) \leq \alpha. \quad (41)$$

This can be seen as follows. As $\nu \rightarrow \infty$, the objective of $(\tilde{D}_k(0, 0))$ is to minimize $\sum_{j=1}^{J_k} \tilde{x}_{j,k}^1$. For any feasible solution x of (LP), x_k is in the feasible set of $\tilde{D}_k(0, 0)$. Hence, $\sum_{j=1}^{J_k} \tilde{x}_{j,k}^{*,1}(\infty) \leq \sum_{j=1}^{J_k} x_{j,k}^1$ with x a feasible solution of (LP). In addition, we have that $\sum_{k=1}^K \sum_{j=1}^{J_k} x_{j,k}^1 \leq \alpha$ with x a feasible solution of (LP). This proves (41).

- **Property 2:**

$$\sum_{j=1}^{J_k} \tilde{x}_{j,k}^{*,1}(\nu) \geq \sum_{j=1}^{J_k} \tilde{x}_{j,k}^{*,1}(\tilde{\nu}), \quad \text{for } \nu < \tilde{\nu}. \quad (42)$$

This can be seen as follows: By definition we have $\sum_{j=1}^{J_k} \sum_{a=0}^1 C_{j,k}^a \tilde{x}_{j,k}^{*,a}(\nu) + \nu \sum_{j=1}^{J_k} \tilde{x}_{j,k}^{*,1}(\nu) \leq \sum_{j=1}^{J_k} \sum_{a=0}^1 C_{j,k}^a \tilde{x}_{j,k}^{*,a}(\tilde{\nu}) + \nu \sum_{j=1}^{J_k} \tilde{x}_{j,k}^{*,1}(\tilde{\nu})$ and $\sum_{j=1}^{J_k} \sum_{a=0}^1 C_{j,k}^a \tilde{x}_{j,k}^{*,a}(\tilde{\nu}) + \tilde{\nu} \sum_{j=1}^{J_k} \tilde{x}_{j,k}^{*,1}(\tilde{\nu}) \leq \sum_{j=1}^{J_k} \sum_{a=0}^1 C_{j,k}^a \tilde{x}_{j,k}^{*,a}(\nu) + \tilde{\nu} \sum_{j=1}^{J_k} \tilde{x}_{j,k}^{*,1}(\nu)$. Subtracting the latter inequality from the first, we obtain Equation (42).

- **Property 3:**

$$\sum_{j=1}^{J_k} \tilde{x}_{j,k}^{*,1}(\nu) < \infty \quad \text{for } \nu > 0. \quad (43)$$

This follows since by Assumption 4.5 the set of optimal solutions of $(\tilde{D}_k(0, 0))$ is bounded for $\nu > 0$.

We define $\bar{\alpha} := \sum_{k=1}^K \sum_{j=1}^{J_k} \tilde{x}_{j,k}^{*,1}(0)$. Equations (41)–(43) imply that there exists a $\nu^* \geq 0$ such that

$$\sum_{k=1}^K \sum_{j=1}^{J_k} \tilde{x}_{j,k}^{*,1}((\nu^*)^-) \geq \min(\alpha, \bar{\alpha}) \quad \text{and} \quad \sum_{k=1}^K \sum_{j=1}^{J_k} \tilde{x}_{j,k}^{*,1}((\nu^*)^+) \leq \min(\alpha, \bar{\alpha}). \quad (44)$$

From standard LP theory we know that there exists a $\bar{\nu} < \infty$ such that $\tilde{x}_k^*(\bar{\nu})$ is an optimal solution of $(D_k(0, 0))$ for all $\nu \geq \bar{\nu}$, that is $\tilde{x}_k^*(\nu) = \tilde{x}_k^*(\bar{\nu})$ for $\nu \geq \bar{\nu}$. Hence, we can take $\nu^* < \infty$.

From (42) and (44) we obtain that there exists a $\tilde{y}^* = (\tilde{y}_{j,k}^{*,a})$ with $\tilde{y}_{j,k}^*$ being a convex combination of $\tilde{x}_{\tilde{k}}^*((\nu^*)^-)$ and $\tilde{x}_{\tilde{k}}^*((\nu^*)^+)$ and for $k \neq \tilde{k}$, $\tilde{y}_{j,k}^*$ being equal to either $\tilde{x}_{\tilde{k}}^*((\nu^*)^-)$ or $\tilde{x}_{\tilde{k}}^*((\nu^*)^+)$, such that $\sum_{k=1}^K \sum_{j=1}^{J_k} \tilde{y}_{j,k}^{*,1} = \min(\alpha, \bar{\alpha})$. Note that \tilde{y}^* is still a solution of $(\tilde{D}_k(0, 0))$, for all k . Now, if $\alpha = \min(\bar{\alpha}, \alpha)$, it follows directly that \tilde{y}^* is also an optimal solution of (LP). If instead $\bar{\alpha} = \min(\bar{\alpha}, \alpha)$, then $\nu^* = 0$ and hence \tilde{y}^* is an optimal solution of $(\tilde{D}_k(0, 0))$ with $\nu = 0$. After summing over k , the latter has the same objective function as (LP). Together with $\sum_{k=1}^K \sum_{j=1}^{J_k} \tilde{y}_{j,k}^{*,1} = \bar{\alpha} \leq \alpha$, it follows that \tilde{y}^* is also an optimal solution of (LP).

It remains to be proved that the Whittle index policy is included in the set $\Pi(\tilde{y}^*) \subset \Pi^*$. Assume for class \tilde{k} the states are ordered such that $\nu_{j_1, \tilde{k}}^{lim} \leq \nu_{j_2, \tilde{k}}^{lim} < \dots \leq \dots \leq \nu_{j_{J_{\tilde{k}}, \tilde{k}}}^{lim}$. From $\nu^* < \infty$ and Properties (38)–(39) (which hold for $\tilde{x}^*(\nu)$) we have that there are n^* and \tilde{n}^* , $n^* \leq \tilde{n}^*$, such that $\nu_{j_{n^*, \tilde{k}}} = \dots = \nu_{j_{\tilde{n}^*, \tilde{k}}} = \nu^*$ and

$$\begin{aligned} \tilde{x}_{j_m, \tilde{k}}^{*,1}((\nu^*)^-) &= 0, \text{ for all } m = 1, \dots, n^*, \\ \tilde{x}_{j_m, \tilde{k}}^{*,0}((\nu^*)^-) &= 0, \text{ for all } m = n^* + 1, \dots, J, \end{aligned}$$

and

$$\begin{aligned} \tilde{x}_{j_m, \tilde{k}}^{*,1}((\nu^*)^+) &= 0, \text{ for all } m = 1, \dots, \tilde{n}^*, \\ \tilde{x}_{j_m, \tilde{k}}^{*,0}((\nu^*)^+) &= 0, \text{ for all } m = \tilde{n}^* + 1, \dots, J. \end{aligned}$$

The vector $\tilde{y}_{\tilde{k}}^*$ is a convex combination of $\tilde{x}_{\tilde{k}}^*((\nu^*)^-)$ and $\tilde{x}_{\tilde{k}}^*((\nu^*)^+)$, hence $\tilde{y}_{j_m, \tilde{k}}^{*,1} = 0$ for all $m \leq n^*$ and $\tilde{y}_{j_m, \tilde{k}}^{*,0} = 0$ for all $m \geq \tilde{n}^* + 1$. Hence, Whittle's index policy $(\nu_{j,k}^{lim})$ satisfies items 1 and 2 of Definition 3.4 with $x^* = \tilde{y}^*$.

If $\sum_{k=1}^K \sum_{j=1}^{J_k} \tilde{y}_{j,k}^{*,1} < \alpha$, then since $\sum_{k=1}^K \sum_{j=1}^{J_k} \tilde{y}_{j,k}^{*,1} = \min(\alpha, \bar{\alpha})$ we have $\bar{\alpha} < \alpha$, so $\nu^* = 0$. This implies that for any state (j, k) with $\tilde{y}_{j,k}^{*,1} = 0$ and $\tilde{y}_{j,k}^{*,0} > 0$ it follows from Property (38) that $\nu_{j,k}^{lim} < (\nu^*)^+ = 0^+$. Hence, by definition of Whittle's index policy $(\nu_{j,k}^{lim})$, a bandit in this state will never be made active, which implies that item 3 in Definition 3.4 is satisfied for $x^* = \tilde{y}^*$. It hence follows that Whittle's index policy $(\nu_{j,k}^{lim})$ is included in the set of priority policies $\Pi(\tilde{y}^*) \subset \Pi^*$. \square

Appendix F: Proof of Proposition 5.1

Consider an arbitrary priority policy π . By assumption we have that $X^{r,\pi}(t)$ is irreducible. The Markov process $X^{r,\pi}(t)$ has unbounded transition rates, however, it follows that it does not die in finite time (upward jumps are of the order 1). Hence, once we prove the multi-step drift criterion [33, 42], we can conclude that there is a unique invariant distribution measure.

Hence, we need to prove that the multi-step drift criterion is satisfied: We will show that there is a $\delta > 0$, a $T < \infty$ and a stopping time τ such that $\mathbb{E}(\sum_{k=1}^K \sum_{j=1}^{J_k} X_k^{r,\pi,0}(\tau) | X(0) = x) - \sum_{k=1}^K \sum_{j=1}^{J_k} x_{j,k}^0 \leq -\delta$ for any $x \in C^c$, with $C = \{x : \sum_{k=1}^K \sum_{j=1}^{J_k} x_{j,k}^0 \leq c\}$ and such that $\mathbb{E}(\tau | X(0) = x) \leq T$.

We define the stopping time as the first moment that an active bandit is made passive. Hence, during the interval $[0, \tau]$ the collection of passive bandits does not change.

First assume there exists an x such that $\mathbb{E}(\tau | X(0) = x) = \infty$. This implies that when starting in state x , the collection of passive and active bandits remains fixed. Hence, each passive bandit evolves according to the transition rates $q_{j,k}^0$ and hence, the number of passive class- k bandits is equivalent to that in an $M/G/\infty$ queue (with arrival rate $\lambda_k r$ and phase-type distributed service requirements as described by the transitions of an inactive class- k bandit), which is stable. By irreducibility, if there is an x , such that $\mathbb{E}(\tau | X(0) = x) = \infty$, for any starting point, the process will be in a finite expected amount of time in state x , hence, stability follows.

We now assume $\mathbb{E}(\tau | X(0) = x) < \infty$ for all x . Since there is a finite number of states $J_k < \infty$ and the state transitions are exponential, it follows directly that there exists a $T < \infty$ such that $\mathbb{E}(\tau | X(0) = x) < T$ for all x . Note that the passive bandits behave independently during the interval $[0, \tau]$. The probability that an inactive bandit departs in the interval $[0, \tau]$ can be lower bounded by p_0 with $p_0 > 0$.

This follows from the assumption that state 0 is positive recurrent under the policy that always keeps the class- k bandit passive. Hence, the mean number of passive bandits that leave during the interval $[0, \tau]$ is given by $p_0 \sum_{k=1}^K \sum_{j=1}^{J_k} x_{j,k}^0$. We therefore have as mean drift

$$\begin{aligned} & \mathbb{E}\left(\sum_{k=1}^K \sum_{j=1}^{J_k} X_k^{r,\pi,0}(\tau) | X(0) = x\right) - \sum_{k=1}^K \sum_{j=1}^{J_k} x_{j,k}^0 \\ &= \lambda r \mathbb{E}_x(\tau) + 1 - p_0 \sum_k \sum_{j=1}^{J_k} x_{j,k}^0 < \lambda r T + 1 - p_0 c = -\delta, \end{aligned}$$

for $x \in C^c$ and $c := (\lambda r T + 1 + \delta)/p_0$. The $+1$ in the mean drift is due to the active bandit that becomes passive at time τ . By the multistep drift criterion we obtain that there is a unique invariant probability distribution for the process $X^{r,\pi}(t)$, for any r . Recall that we denote this distribution by $p^{r,\pi}$.

The number of passive class- k bandits in state j , $X_{j,k}^{r,\pi,0}(t)$, can be stochastically upper bounded by the number of customers in an $M/G/\infty$ queue with arrival rate $\bar{\lambda}_{j,k} r$, $\bar{\lambda}_{j,k} := \lambda_k p_{0,j} + \max_i q_k^1(i, j) \alpha$, and service requirement $B_{j,k}$. Here $B_{j,k}$ is described by the state transitions of a passive class- k bandit starting in state j until it exists the system (while kept passive). We denote the total number of customers in this $M/G/\infty$ queue by $Y_{j,k}^r(t)$. The latter serves as an upper bound on $X_{j,k}^{r,\pi,0}(t)$ since $\bar{\lambda}_{j,k}$ is an upper bound for the arrival rate of passive class- k bandits (in state j): (i) passive bandits that are made active in the original process, are considered to stay forever passive in the upper bound process, (ii) $\max_i q_k^1(i, j) \alpha$ represents the maximum at which rate active bandits can be made passive.

The stationary distribution of the process $\{Y_{j,k}^r(t)\}$, is given by a Poisson distribution with parameter $\bar{\lambda}_{j,k} r \mathbb{E}(B_{j,k})$ [42]. It can be checked that this distribution converges to the Dirac measure in the point $\bar{\lambda}_{j,k} \mathbb{E}(B_{j,k})$, as $r \rightarrow \infty$. By Prohorov's theorem it then follows that the family $\{Y_{j,k}^r/r\}$ is tight [42]. Furthermore, since $\mathbb{E}(Y_{j,k}^r/r) = \bar{\lambda}_{j,k} \mathbb{E}(B_{j,k})$ and $\mathbb{E}(\lim_{r \rightarrow \infty} Y_{j,k}^r/r) = \bar{\lambda}_{j,k} \mathbb{E}(B_{j,k})$, a.s., we obtain from [10, Theorem 3.6] that the family $\{Y_{j,k}^r/r\}$ is uniform integrable.

Since $\sum_{j,k} Y_{j,k}^r(t)/r$ represents a stochastic upper bound on the queue length process $\sum_{k=1}^K X_k^{r,\pi}(t)/r$, we obtain that the family $\{p^{r,\pi}\}$ is tight and uniform integrable as well. \square

Appendix G: Proof of Lemma 7.1

For the fixed population, the total number of constraints in (LP) is $\sum_{k=1}^K J_k + 1 + K$. However, since $\sum_{k=1}^K \lambda_k = 0$, one of the constraints in (4) is redundant for each k . Hence, the number of independent constraints in (LP) is $\sum_{k=1}^K J_k + 1$.

Since the feasible set of (LP) is bounded, from standard LP theory, see [40, Theorem D.1a], we obtain that there exists an optimal basic feasible solution x^* to (LP). Hence, x^* has $\sum_{k=1}^K J_k + 1$ basic terms and all other terms are equal to zero. If $x_{j,k}^* > 0$ for all (j, k) , then for any (j, k) there is an action a such that $x_{j,k}^{*,a} = 0$, and in at most one combination (j, k) the components $x_{j,k}^{*,a}$ can be positive in both actions. Hence, x^* satisfies the property in Definition 3.4.

Otherwise, let S denote the set of states (i, l) such that $x_{i,l}^* = 0$. By (4), if $(j, k) \in S$, then $\sum_{a=0}^1 \sum_{i \neq j} x_{i,k}^{*,a} q_k^a(i, j) = 0$. That is, $x_{i,k}^{*,a} q_k^a(i, j) = 0$ for all $i = 1, \dots, J_k$, $a = 0, 1$, if $(j, k) \in S$. Hence, for $(j, k) \notin S$, Equation (4) in the point x^* can be rewritten as

$$0 = \sum_{a=0}^1 \sum_{i=1, i \neq j, (i,k) \in S^c}^{J_k} x_{i,k}^{*,a} q_k^a(i, j) - \sum_{a=0}^1 x_{j,k}^{*,a} \bar{q}_k^a(j), \quad \forall j, k,$$

with $\bar{q}_k^a(j) := \sum_{i=0, i \neq j, (i,k) \in S^c}^{J_k} q_k^a(j, i)$. Hence, x^* (restricted to the states $(j, k) \in S^c$) is an optimal solution of (LP) restricted to the set of states S^c . Similar as above, the latter has an optimal basic solution with $|S^c| + 1$ basic terms (and all other terms equal to zero). Let y^* denote such an optimal basic solution. Note that y^* is also an optimal solution of (LP) when setting $y_{j,k}^* = 0$ for all states $(j, k) \in S$.

If $y_{j,k}^* > 0$ for all $(j, k) \notin S$, then, since it has $|S^c| + 1$ basic terms, it satisfies that for any (j, k) there is an action a such that $y_{j,k}^{*,a} = 0$, and in at most one combination (j, k) the components $y_{j,k}^{*,a}$ can be positive in both actions. Hence, y^* satisfies the property in Definition 3.4.

If $y_{j,k}^* = 0$ for some $(j, k) \notin S$, the above procedure can be repeated until one ends up with an optimal basic solution that satisfies the properties as given in Definition 3.4. LP-

Now assume a dynamic population of bandits. First assume $p_k(0, j) > 0$ for all k, j . By (4) we have that any feasible solution of (LP) has $x_{j,k} > 0$. Hence, for each (j, k) there exists at least one action a such that $x_{j,k}^a > 0$. Since the set of optimal solutions of (LP) is non-empty and bounded, from standard LP theory, see [40, Theorem D.1a], we obtain that there exists a bounded optimal basic feasible solution x^* to (LP). We know that x^* has $\sum_{k=1}^K J_k + 1$ basic terms (the number of constraints), and all other terms are equal to zero. Since $x_{j,k}^* > 0$ for all j, k , this implies that for any (j, k) there is one action a such that $x_{j,k}^{*,a} = 0$, and in at most one combination (j, k) the components $x_{j,k}^{*,a}$ can be positive in both actions $a = 0$ and $a = 1$.

Now assume $C_{j,k}^0 > 0$ for all j, k . This implies that for all $\epsilon > 0$ small enough the set of optimal solutions of the $(LP(\epsilon))$ problem is bounded and non-empty, where $LP(\epsilon)$ is defined by

$$\begin{aligned}
(LP(\epsilon)) \quad & \min_x \sum_{k=1}^K \sum_{j=1}^{J_k} \sum_{a=0}^{A_{j,k}} C_{j,k}^a x_{j,k}^a \\
\text{s.t.} \quad & 0 = \lambda_k(p_k(0, j) + \epsilon) + \sum_{a=0}^1 \sum_{i=1, i \neq j}^{J_k} x_{i,k}^a q_k^a(i, j) - \sum_{a=0}^1 x_{j,k}^a q_k^a(j), \quad \forall j, k, \\
& \sum_{k=1}^K \sum_{j=1}^{J_k} x_{j,k}^1 \leq \alpha, \\
& x_{j,k}^a \geq 0, \quad \forall j, k, a.
\end{aligned} \tag{45}$$

(Note that the assumption $C_{j,k}^0 > 0$ for all j, k as stated in Lemma 7.1 could have been replaced by the weaker assumption that the set of optimal solutions of $(LP(\epsilon))$ is bounded and non-empty.) By sensitivity results of linear programming theory we have that the same basis provides an optimal solution for $(LP(\epsilon))$, for all $0 \leq \epsilon < \bar{\epsilon}$ and $\bar{\epsilon} \geq 0$ small enough. We denote the corresponding optimal solution by $x^*(\epsilon)$. By (45) we have that $x_{j,k}^{*,a}(\epsilon) > 0$ for all $\epsilon > 0$. Since for any $0 < \epsilon < \bar{\epsilon}$ the basis of $x^*(\epsilon)$ is the same, we conclude that for any state (j, k) there is one action a (independent on ϵ) such that $x_{j,k}^{*,a}(\epsilon) = 0$ and for at most one state (j, k) (independent of ϵ) the components $x_{j,k}^{*,a}(\epsilon)$ can be strictly positive for both actions $a = 0$ and $a = 1$.

Note that $(LP(0)) = (LP)$. Hence, using [11, Corollary 1] we obtain that the correspondence that gives for each ϵ the set of optimal solutions of $(LP(\epsilon))$ is upper semicontinuous in the point $\epsilon = 0$. Being a compact-valued correspondence, it follows that there exists a sequence ϵ_l such that $\epsilon_l \rightarrow 0$ and $x^*(\epsilon_l) \rightarrow x^*$, with x^* being an optimal solution of (LP). Being the limit, x^* has the same components equal to zero (and maybe even more) as $x^*(\epsilon)$ (with $\epsilon < \bar{\epsilon}$). Hence, x^* has the property as stated in the lemma. \square

Appendix H: Condition 3.9 for an $M/M/S+M$ queue

Assume the classes are reordered such that $\iota_1 \geq \iota_2 \geq \dots \geq \iota_K$. We further denote by $I^* := \{l : \iota_l \leq 0\}$, the set of classes that will never be served. This set is of the form $I^* = \{\hat{l}, \dots, K\}$. Under policy (ι_k) , the ODE as defined in (9) is given by

$$\frac{dx_k^*(t)}{dt} = \lambda_k - x_k^{*,0}(t)\theta_k - x_k^{*,1}(t)(\mu_k + \tilde{\theta}_k), \quad \forall k, \tag{46}$$

$$\text{with } x_k^{*,1}(t) = \min \left((S - \sum_{l=1}^{k-1} x_l^*(t))^+, x_k^*(t) \right), \quad \text{if } k < \hat{l}, \quad \forall k, \tag{47}$$

$$x_k^{*,1}(t) = 0, \quad \text{if } k \geq \hat{l}, \quad \forall k, \tag{48}$$

$$x_k^{*,0}(t) = x_k^*(t) - x_k^{*,1}(t), \quad \forall k.$$

This ODE has a unique equilibrium point, which is given by

$$x_k^{*,0} = 0, \quad x_k^{*,1} = \frac{\lambda_k}{\mu_k + \tilde{\theta}_k}, \quad \text{for } k = 1, \dots, \hat{k}, \quad (49)$$

$$x_{\hat{k}+1}^{*,0} = \frac{\lambda_k - (\mu_k + \tilde{\theta}_k)(S - \sum_{l=1}^{\hat{k}} \frac{\lambda_l}{\mu_l + \tilde{\theta}_l})}{\theta_k}, \quad x_{\hat{k}+1}^{*,1} = S - \sum_{l=1}^{\hat{k}} \frac{\lambda_l}{\mu_l + \tilde{\theta}_l}, \quad \text{if } \hat{k} + 1 < \hat{l}, \quad (50)$$

$$x_k^{*,0} = \frac{\lambda_k}{\theta_k}, \quad x_k^{*,1} = 0, \quad \text{for } k \geq \min(\hat{k} + 2, \hat{l}), \quad (51)$$

where $\hat{k} = \arg \max\{k = 0, 1, \dots, \hat{l} - 1 : \sum_{l=1}^k \frac{\lambda_l}{\mu_l + \tilde{\theta}_l} \leq S\}$. This can be seen as follows. If x^* is an equilibrium point, it follows from (46) that

$$\frac{\lambda_k}{\mu_k + \tilde{\theta}_k} = x_k^{*,1} + x_k^{*,0} \frac{\theta_k}{\mu_k + \tilde{\theta}_k}. \quad (52)$$

We first prove (49). Let $k = 1$ and assume $1 \leq \hat{k}$. Hence, we have $\frac{\lambda_1}{\mu_1 + \tilde{\theta}_1} < S$. By (52) we obtain $x_1^{*,1} < S$. Together with (47), that is, $x_1^{*,1} = \min(S, x_1^*)$, we obtain $x_1^{*,1} = x_1^*$ and hence $x_1^{*,0} = 0$. From (52) we obtain that $x_1^{*,1} = \frac{\lambda_1}{\mu_1 + \tilde{\theta}_1}$. The proof of (49) continues by induction. Assume (49) holds for $k \leq l - 1$, and let $l \leq \hat{k}$. For $k \leq l - 1$ we have that $x_k^{*,1} = \frac{\lambda_k}{\mu_k + \tilde{\theta}_k}$. Since $\sum_{k=1}^l \frac{\lambda_k}{\mu_k + \tilde{\theta}_k} \leq S$, by (47) we obtain that $x_l^{*,1} = x_l^*$ and hence $x_l^{*,0} = 0$. From (52) we then obtain that (49) holds for $k = l$ as well.

We now prove (50). Let $\hat{k} + 1 < \hat{l}$. From (49) and (50) we obtain that $S - \sum_{l=1}^{\hat{k}} x_l^* < x_{\hat{k}+1}^*$. So by (47) we obtain $x_{\hat{k}+1}^{*,1} = S - \sum_{l=1}^{\hat{k}} \frac{\lambda_l}{\mu_l + \tilde{\theta}_l}$ as stated in (50).

We now prove (51). From (49) and (50) we obtain that $S \leq \sum_{l=1}^{\hat{k}+1} x_l^*$, hence $x_k^{*,1} = 0$ for k such that $\hat{k} + 1 < k < \hat{l}$. Equation (51) for $k \geq \hat{l}$ follows directly from (48).

In addition, x^* is a global attractor, as was shown in [3, Appendix] (this can be seen by replacing the μ_i in [3] by $\mu_i + \tilde{\theta}_i$, making the ODE in [3] coincide with our ODE (46)).