Journal of Machine Learning Research 20 (2019) 1-93 Submitted 3/17; Revised 2/19; Published 3/19

Optimal Policies for Observing Time Series and Related Restless Bandit Problems

Christopher R. Dance CHRIS.DANCE@NAVERLABS.COM
Tomi Silander TOMI.SILANDER@NAVERLABS.COM
NAVER LABS Europe
6 chemin de Maupertuis
Meylan, 38240, France

Editor: Nikos Vlassis

Abstract

The trade-off between the cost of acquiring and processing data, and uncertainty due to a lack of data is fundamental in machine learning. A basic instance of this trade-off is the problem of deciding when to make noisy and costly observations of a discrete-time Gaussian random walk, so as to minimise the posterior variance plus observation costs. We present the first proof that a simple policy, which observes when the posterior variance exceeds a threshold, is optimal for this problem. The proof generalises to a wide range of cost functions other than the posterior variance. It is based on a new verification theorem by Niño-Mora that guarantees threshold structure for Markov decision processes, and on the relation between binary sequences known as *Christoffel words* and the dynamics of discontinuous nonlinear maps, which frequently arise in physics, control and biology.

This result implies that optimal policies for *linear-quadratic-Gaussian control with costly observations* have a threshold structure. It also implies that the restless bandit problem of observing multiple such time series, has a well-defined *Whittle index policy*. We discuss computation of that index, give closed-form formulae for it, and compare the performance of the associated index policy with heuristic policies.

Keywords: restless bandits, Whittle index, Christoffel words, Sturmian words, Kalman filter, linear-quadratic-Gaussian control

1. Introduction

This paper answers three closely-related questions about discrete-time filtering of scalar time series with costly observations, where the nature of the observations is controlled through a query action. The first two questions concern the structure of optimal policies for observing a single time series so as to minimise either a function of the posterior variance (Theorem 1, Excerpt) or a quadratic function of the system state and control input (Corollary 2). The third question concerns the observation of several such time series with a constraint on the number of time series that can be observed simultaneously. This is an instance of a *restless bandit problem* and it is interesting to know that the problem has a well-defined *Whittle index policy* (Theorem 1).

This introduction begins with the time-series model (Section 1.1) that the three questions have in common. It then motivates, formulates and states the key results for each question in turn (Sections 1.2 to 1.4). It concludes with an intuitive guide to the main concepts involved in the proof (Section 1.5) and a description of the structure of the rest of the paper (Section 1.6).

Notation. In this paper, $\mathbb{Z}_+ = \{0, 1, 2, \dots\}$, $\mathbb{Z}_{++} = \{1, 2, 3, \dots\}$, $\mathbb{R}_+ = [0, \infty)$, $\mathbb{R}_{++} = (0, \infty)$ and $\overline{\mathbb{R}} = [-\infty, \infty] = \mathbb{R} \cup \{-\infty, \infty\}$ denotes the extended real numbers. Left limits are denoted by $f(y^-) = \lim_{x \uparrow y} f(x)$ and right limits by $f(y^+) = \lim_{x \downarrow y} f(y)$. The terms *increasing* and *decreasing* are used in the strict sense, with non-decreasing and non-increasing used otherwise.

1.1. Time-Series Model

We consider the classic discrete-time scalar normally-distributed state-space model. In this model, the state is partially observed through measurements as fully described by the conditional dependencies

$$\left.\begin{aligned} Z_0 &\sim \mathcal{N}(z_0, v_0) \\ Z_{t+1}|Z_t, u_t &\sim \mathcal{N}(\bar{A}Z_t + \bar{B}u_t, \Sigma_Z) \\ Y_{t+1}|Z_{t+1}, a_t &\sim \mathcal{N}(Z_{t+1}, \Sigma_Y(a_t)) \end{aligned}\right\} \quad \text{for } t \in \mathbb{Z}_+. \tag{1}$$

The state Z_t is a real-valued random variable with initial mean z_0 and variance v_0. The sequence of states depends on the control or exogenous input $u_t \in \mathbb{R}$. The measurement Y_{t+1} is a real-valued random variable which depends on a query action $a_t \in \{0, 1\}$. The variances $\Sigma_Z, \Sigma_Y(0), \Sigma_Y(1) > 0$ and real-valued parameters $\bar{A}, \bar{B}$ are known. Query action $a_t = 1$ is assumed to correspond to a higher-quality observation than query action $a_t = 0$, so that $\Sigma_Y(1) < \Sigma_Y(0)$ and it is possible that $\Sigma_Y(0) = \infty$ which represents a totally uninformative observation or no observation at all.

The observed history H_t at time t is $z_0, v_0, a_0, a_1, \dots, a_{t-1}, u_0, u_1, \dots, u_{t-1}, Y_1, Y_2, \dots, Y_t$. Under the Bayesian filter, the information state is given by the posterior mean $z_t := \mathbb{E}[Z_t|H_t]$ and variance $v_t := \mathbb{E}[(Z_t - z_t)^2|H_t]$. In this case, the Bayesian filter is the Kalman filter (Thiele, 1880; Kalman, 1960) and it follows that the information state undergoes the following Markovian transitions:

$$\left.\begin{aligned} &z_{t+1}|z_t, v_t, a_t, u_t \sim \mathcal{N}(\bar{A}z_t + \bar{B}u_t, \bar{A}^2 v_t + \Sigma_Z - \Phi_{a_t}(v_t)) \\ &v_{t+1}|z_t, v_t, a_t, u_t = \Phi_{a_t}(v_t) \end{aligned}\right\} \quad \text{for } t \in \mathbb{Z}_+ \tag{2}$$

where $\Phi_a : \mathbb{R}_+ \to \mathbb{R}_+$ for $a \in \{0, 1\}$ is the Möbius transformation

$$\Phi_a(v) := \frac{(\bar{A}^2 v + \Sigma_Z) \times \Sigma_Y(a)}{(\bar{A}^2 v + \Sigma_Z) + \Sigma_Y(a)}. \tag{3}$$

The above model excludes the relatively-simple case of learning a *stationary* parameter Z_t, since we assume that $\Sigma_Z > 0$. Under this assumption, to facilitate analysis, we shall scale the units of variance such that $\Sigma_Z = 1$. Specifically, we consider the (relative) variance state at time t given by $x_t := v_t/\Sigma_Z$ and the (relative) precision of the observations given by $\theta_a := \Sigma_Z/\Sigma_Y(a)$ for $a \in \{0, 1\}$. We also define the parameter $r := |\bar{A}|$, noting that $\bar{A}$ can be negative. After this change of coordinates, the variance transitions are now given by the Möbius transformation

$$\phi_a(x) := \frac{\Phi_a(\Sigma_Z x)}{\Sigma_Z} = \frac{(r^2 \Sigma_Z x + \Sigma_Z)\Sigma_Y(a)}{\Sigma_Z(r^2 \Sigma_Z x + \Sigma_Z + \Sigma_Y(a))} = \frac{r^2 x + 1}{\theta_a(r^2 x + 1) + 1}. \tag{4}$$

The above model also excludes *totally informative* observations, since we assume that $\Sigma_Y(1) > 0$. Such models are nevertheless partly-addressed in the limit as the precision $\theta_1 \to \infty$ in Proposition 31, where it is possible to find closed-form formulae relating to optimal policies for the three problems discussed in the rest of this introduction. The rationale for excluding such cases from our analysis is that our characterisation of the behaviour of maps-with-gaps (Theorem 12) no longer applies.

1.2. Optimal Policies for Observing a Single Time Series

The simplest problem addressed here involves an uncertainty cost $C(x)$, where x is the (relative variance) state introduced just above, and a price $\lambda \in \mathbb{R}$ to be paid every time query action $a = 1$ is taken. Price λ might reflect costs of energy, labour, communication, computational processing, hardware or risks associated with each measurement. Recall that a policy is *non-anticipative* if it selects actions at time t based only on information available up-to and including time t. The objective is to find a non-anticipative policy π that selects query actions so as to minimise the β-discounted performance functional, for discount factor $\beta \in [0, 1)$, for all initial states x,

$$\mathbb{E}_x^\pi \left[\sum_{t=0}^{\infty} \beta^t (\lambda A_t + C(X_t)) \right] \tag{5}$$

where $\mathbb{E}_x^\pi$ denotes the expectation over sequences $(X_t, A_t)_{t=0}^\infty$ of states X_t and actions A_t with initial state $X_0 = x$, where actions A_t are taken according to policy π and transitions are according to (2). Note that this performance functional depends neither on the control or exogenous input u_t nor on the posterior mean z_0 since these appear neither in the costs nor in the transitions of the posterior variance, which are given by the deterministic mapping (4). So, while in some settings, such as that considered in Section 1.3, it is natural to require policies π to also select the control u_t, in the setting considered here it does not matter which value for u_t is chosen and we only require the policy to select the query action a_t. Thus, this problem corresponds to the following deterministic dynamic program for value function $V : \mathbb{R}_{++} \to \mathbb{R}$,

$$V(x) = \min_{a \in \{0,1\}} \left\{ \lambda a + C(x) + \beta V(\phi_a(x)) \right\}. \tag{6}$$

The first question addressed in this paper is: *for what cost functions is a threshold policy optimal for this problem?* For instance, one may intuitively guess that optimal policies for variance minimisation with $C(x_t) = x_t$, for entropy minimisation with $C(x_t) = \log(x_t)$, or for precision maximisation with $C(x_t) = -1/x_t$, might involve making expensive observations at time t when the variance state x_t exceeds a threshold. The following assumption on the model covers these examples.

Assumption A1

(i) The state space $\mathcal{I}$ is either $[0, \infty)$ *or* $(0, \infty)$.

(ii) The discount factor β is in $[0, 1)$.

(iii) The transition functions $\phi_a : \mathbb{R}_+ \to \mathbb{R}_+$ for $a \in \{0, 1\}$ are of the form

$$\phi_a(x) = \frac{r^2 x + 1}{\theta_a(r^2 x + 1) + 1}$$

for some $0 \le \theta_0 < \theta_1 < \infty$ and some $r \in (0, 1]$.

(iv) The uncertainty cost function $C : \mathcal{I} \to \mathbb{R}$ is of the form $C(x) = \sum_{i=1}^{n_C} C_i(x)$ for some $n_C \in \mathbb{Z}_{++}$, where each of the functions $C_i : \mathcal{I} \to \mathbb{R}$ satisfies one of the following conditions:

C1. For $x \in \mathcal{I}$, the derivatives $C_i'(x) := \frac{d}{dx} C_i(x)$ and $C_i''(x) := \frac{d^2}{dx^2} C_i(x)$ exist and
- *the function $C_i(x)$ is concave,*
- *the function $\frac{1}{x^3} C_i''\left(\frac{1}{x}\right)$ is non-decreasing,*
- *and the function $\frac{1}{x^2} C_i'\left(\frac{1}{x}\right)$ is non-increasing and convex.*

C2. For $x \in \mathcal{I}$, the function $C_i(x)$ is non-decreasing, convex and differentiable.

(v) There exists a measurable function $w : \mathcal{I} \to [1, \infty)$ and constants $M > 0$ and $\gamma \in [\beta, 1)$ such that for every state $x \in \mathcal{I}$:

(a) $|C(x)| \le M w(x)$; and

(b) $\beta \max_{a \in \{0,1\}} w(\phi_a(x)) \le \gamma w(x)$.

Regarding part (i) of the above assumption, note that we may work with the interval $\mathcal{I} = (0, \infty)$ in cases where the cost function $C(x)$ is not a real number for $x = 0$, in order to include cases like $\log(x)$ and $-1/x$, but the results of the paper continue to hold in cases where $C(0)$ is defined. In fact, the main results of the paper also apply to smaller intervals $\mathcal{I}$, as long as the initial state is contained in the interval $\mathcal{I}$ and the open interval (y_1, y_0) is in $\mathcal{I}$, where y_a is the unique solution to $\phi_a(y_a) = y_a$ (see Section 2 for further discussion of such fixed points.)

In part (iii) of the above assumption, the inequality $\theta_0 < \theta_1$ is taken to be strict as otherwise the problem is trivial. Specifically, if $\theta_0 = \theta_1$ then the state sequence does not depend on the policy, so the policy that always takes action 0 is optimal if $\lambda > 0$, whereas the policy that always takes action 1 is optimal if $\lambda < 0$.

Part (iv) of the above assumption is at the heart of our answer to our first question: it represents the most general class of cost functions $C(\cdot)$ that our approach to proving optimality of threshold policies naturally allows for. Delving into the proof, the specific form of part (iv) results from a requirement that certain *majorisation* inequalities (Marshall et al., 2010) hold in Lemma 25.

Uncertainty cost functions $C(\cdot)$ satisfying part (iv) may be neither convex nor concave. For instance in the case $C(x) = (x^3 - 1)/x$, we can write $C(x) = \sum_{i=1}^{n_C} C_i(x)$ for $n_C = 2$ with $C_1(x) = x^2$, a convex function satisfying part C1, and with $C_2(x) = -1/x$, a concave function satisfying part C2. In this instance, the cost $C(x) = (x^3 - 1)/x$ is *unbounded* from both below and above (that is, $\lim_{x \downarrow 0} C(x) = -\infty$ and $\lim_{x \uparrow \infty} C(x) = \infty$). However, it is possible that the cost is *bounded* both below and above, for instance in the case $C(x) = x/(x+1)$.

In general part (iv) of the above assumption can always be satisfied with $n_C = 2$ terms. This is because finite sums of functions satisfying part C1 also satisfy part C1, since finite sums of functions satisfying part C2 also satisfy part C2, and since the zero-function ($C(x) = 0$ for all x) satisfies both cases.

Also part (iv) of the above assumption requires that functions C_i satisfying C2 have a derivative C_i'. This is simply for convenience in the proofs of Section 3.3. As such functions are real-valued convex functions, one can instead set C_i' equal to any subderivative at points where the derivative is not defined.

Part (v) is key to the *weighted supremum norm* approach to infinite-horizon discounted Markov decision problems, as is explained in detail by Wessels (1977) and Hernández-Lerma and Lasserre (1999, Chapter 8). It ensures that the infinite series in the objective function is well-defined for all policies π and initial states $x \in \mathcal{I}$, even though the cost function $C(\cdot)$ may be unbounded. Furthermore, although Bellman's equation may have multiple unbounded solutions $V(\cdot)$, part (v) guarantees that the optimal value function is the unique solution with $\sup_{x\in\mathcal{I}} |V(x)/w(x)| < \infty$. For $|C(x)| = x^q$ with $q \geq -1$, it is not hard to see that part (v) is satisfied for $w(x) := \max\{y_1/x, \alpha^{x-y_1}\}$ where y_1 is the unique positive solution of $\phi_1(x) = x$ and α is any number in the interval $(1, 1/\beta)$. Further, for this choice of $w(x)$, if $M(q)$ satisfies $x^q \leq M(q)w(x)$ and $C(x) = \sum_{k=1}^m a_k x^{q_k}$ for some $a_k \in \mathbb{R}$ and $q_k \geq -1$, then the triangle inequality gives $|\sum_{k=1}^m a_k x^{q_k}| \leq (\sum_{k=1}^m |a_k| M(q_k))\, w(x)$. Thus part (v) is satisfied for all the examples of the function $C(\cdot)$ considered above. Given the above assumption, we may now answer our first question. The full theorem (see page 14), of which the following is only an excerpt, also characterises the optimal choice of threshold and the associated Whittle index.

Theorem 1 (Excerpt) *Suppose A1 holds. Then for some threshold $s \in \bar{\mathbb{R}}$, the threshold policies $A_t = \mathbf{1}_{X_t > s}$ and $A_t = \mathbf{1}_{X_t \geq s}$ both minimise performance functional (5) for every initial state $x \in \mathcal{I}$.*

Theorem 1 only holds for cost functions that are functions of the posterior variance $v_t = \mathbb{E}[(Z_t - z_t)^2 | H_t]$ alone, although this includes non-quadratic expressions like $\mathbb{E}[|Z_t - z_t|^p | H_t]$ for $p > -1$, which is of the form $k(p)v_t^p$ for an appropriate function $k(\cdot)$. For cost functions that explicitly depend on the posterior mean, we have a Markov decision problem in a 2-dimensional real-valued state space and there are several possible extensions of the notion of a "threshold policy" that might be considered. Also, as transitions of the posterior mean depend on the values of the observations, the problem would no-longer be deterministic if dependence on the posterior mean were introduced.

From one perspective, this answer is a rare example of an explicit solution to a real-state partially-observed Markov decision process (POMDP). From another perspective, this answer is a rare example of an explicit solution to the problem of observation selection in sensor management (Hero and Cochran, 2011). Indeed, given a collection of variables which can (in principle) be observed and a single variable to predict, which are jointly Gaussian with known covariance, even the problem of deciding whether there exists a subset of k observations that reduces the prediction variance below a given threshold is NP-hard (Davis et al., 1997). Work has therefore focused on finding covariance structures for which the problem is tractable, for instance Das and Kempe (2008) show that selection of Gaussian observations with an exponential covariance can be solved by a simple discrete dynamic program, and on finding appropriate choices of cost functions for which there are guaranteed approximation algorithms (Krause et al., 2008, 2011; Badanidiyuru et al., 2014; Chen et al., 2014).

1.3. The Linear Quadratic Gaussian Problem with Costly Observations

The second question addressed by this paper is: *when are threshold policies optimal for making observations in a generalisation of the linear-quadratic-Gaussian control problem in which observations are costly but controlled through a query action?* Specifically, suppose the states and observations are as in (1) but the objective is to find a non-anticipative policy π that selects a feedback-control action $u_t \in \mathbb{R}$ and a sensor-query action $a_t \in \{0, 1\}$ so as to minimise the β-discounted performance functional

$$\mathbb{E}^{\pi}_{z_0,v_0}\left[\sum_{t=0}^{\infty} \beta^t (DZ_t^2 + Fu_t^2 + \lambda A_t)\right],$$

where $D, F \in \mathbb{R}_+$ and the expectation $\mathbb{E}^{\pi}_{z_0,v_0}$ is over the Markovian transitions (2) under policy π when the initial state Z_0 is normally distributed with mean z_0 and variance v_0. In this expression $\lambda \in \mathbb{R}$ denotes a cost to be paid each time action $A_t = 1$ is taken, just as in (5), and if $\lambda = 0$ then the problem reduces to the classical linear-quadratic-Gaussian control problem.

The problem with price $\lambda \neq 0$ is an old but unsolved problem, although it has been known for a long time that in finite-horizon versions of the problem the optimal timing of observations can in principle be determined *a priori*, and does not depend on the values of the observations (Kushner, 1964; Meier et al., 1967). Thus, the subproblems of state-estimation, observation scheduling and control can be decoupled, and this is often known as the *separation principle*. Meier et al. (1967) addressed the observation scheduling sub-problem using a tree search for forward dynamic programming. Other numerical attacks on the problem have included formulating it as a two-point boundary value problem, whether in continuous time (Athans, 1972) or discrete time (Kerr, 1981). Wu and Arapostathis (2008) study the long-term average and infinite-horizon discounted version of the problem and prove that a separation principle still holds over the infinite horizon, while Molin and Hirche (2009) extended these results to a setting where measurements are subject to random packet losses. Zhao et al. (2014) studied the infinite-horizon average-cost version of the problem, showing that optimal policies can be approximated arbitrarily closely by periodic schedules, but without giving any insight into the structure of such periodic schedules, or results about the discounted case.

Studies of linear quadratic control with costly observations were initially motivated by aerospace applications in which telemetry data from space vehicles was transmitted over band-limited links to ground stations (Kushner, 1964; Meier et al., 1967). More recently, continued research has been motivated by applications to networked control systems for unmanned aerial vehicles (Seiler, 2001), vehicle control (Daoud et al., 2006) and teleoperation (Hirche et al., 2007).

An immediate corollary of Theorem 1 is the following answer to the above question.

Corollary 2 *Suppose* $\bar{A} \in [-1, 1]$, $\bar{B} \in \mathbb{R}\backslash\{0\}$, $D \in \mathbb{R}_{++}$, $F \in \mathbb{R}_+$, $\beta \in (0, 1)$, $\Sigma_Y(a) \in [0, \infty]$ *for* $a \in \{0, 1\}$ *with* $\Sigma_Y(0) \geq \Sigma_Y(1)$, *and that* $\lambda \in \mathbb{R}$. *Then, there exists a threshold* $s \in \overline{\mathbb{R}}$ *such that an optimal policy for the problem of linear-quadratic-Gaussian control with costly observations is to set*

$$a_t = \begin{cases} 1 & \text{if } v_t \geq s \\ 0 & \text{if } v_t < s \end{cases} \qquad \text{and} \qquad u_t = -\left(\frac{\bar{A}}{\bar{B} + \frac{F}{\beta \bar{B} R}}\right) z_t$$

whatever the initial state, where R *is the unique positive root of the quadratic equation*

$$-\beta \bar{B}^2 R^2 + (\beta \bar{B}^2 D + \beta \bar{A}^2 F - F)R + DF = 0.$$

A proof of Corollary 2 is presented in Section 3.6.

1.4. Multi-Target Tracking and Restless Bandits

This paper also addresses the problem of monitoring *multiple* time series so as to maintain a precise belief while imposing a *constraint* on the number of time series that may be sensed at each time. This problem is often called the *multi-target tracking* problem. Multiple heuristics have been proposed for this problem. The simplest heuristics include round-robin schedules, and greedy or myopic schedules (Oshman, 1994). More sophisticated heuristics exploit probabilistic sensor allocations based on "steady-state" covariance matrices in continuous time (Mourikis and Roumeliotis, 2006; Le Ny et al., 2011) or in discrete time (Gupta et al., 2006). Also, Whittle (1988) proposed a *restless bandit* heuristic, and one of the applications of that heuristic is to the multi-target tracking problem. As that heuristic is a major focus of the current paper, we review the literature on the restless bandit approaches to multi-target tracking straight after explaining what restless bandits are below.

One example of a real-world application of the discrete-time problem, which was our original motivation for studying the problems in this paper, is the measurement of on-street parking occupancy (Dey, 2014), in a setting where cheap-but-low-quality observations are available through payment data (at parking meters or through mobile phones), expensive-but-high-quality observations are available through portable cameras, which are moved daily or weekly (and thus in discrete time), and there are a limited number of portable cameras with which to observe many streets.

To formulate the problem, suppose there are $n \in \mathbb{Z}_{++}$ independent time series of the form (1), indexed by $i \in \{1, 2, \ldots, n\}$, and time series i has state $Z_{t,i}$ at time $t \in \mathbb{Z}_+$. Each time series may have its own parameters $z_{i,0}, v_{i,0}, \bar{A}_i, \bar{B}_i, \Sigma_{Z_i}$, and its own input $u_{i,t}$. As in Section 1.1, we scale the posterior variance of each time series to get a variance state $x_{i,t}$ on a state space $\mathcal{I}_i$. As in Section 1.2, each time series has its own uncertainty cost $C_i : \mathcal{I}_i \to \mathbb{R}$. Corresponding to these time series there are n query actions $a_{i,t} \in \{0, 1\}$ at each time t which specify the nature of the observation $Y_{i,t}$ of time series i. These observations have their own parameters $\Sigma_{Y_i} : \{0, 1\} \to (0, \infty]$. However, these actions are subject to the constraint that only $m \in \mathbb{Z}_{++}$ with $m < n$ expensive observations can be made at each time. As in Section 1.1, the transitions of the posterior variance are given by the Möbius transformation (4), which does not involve the exogenous inputs $u_{i,t}$, and which is deterministic.

The problem is then to find a history-dependent randomised policy π that minimises the total β-discounted uncertainty cost

$$\mathbb{E}_x^\pi \left[\sum_{i=1}^{n} \sum_{t=0}^{\infty} \beta^t C_i(X_{t,i}) \right]$$

for any initial state $x \in \mathcal{I}_1 \times \cdots \times \mathcal{I}_n$, subject to the constraint that policy π makes m observations at each time, so that

$$\sum_{i=1}^{n} A_{t,i} = m \qquad \text{for } t \in \mathbb{Z}_+,$$

where $\mathbb{E}_x^\pi$ denotes the expectation over sequences $(X_t, A_t)_{t=0}^{\infty}$ with initial state $X_{0,1} = x_1, \ldots, X_{0,n} = x_n$, where actions are taken according to the potentially non-deterministic policy π and transitions are according to (4). It is equally possible to work with the constraint $\sum_{i=1}^{n} A_{t,i} \leq m$ as discussed below.

Restless Bandits. The multi-target tracking problem is an instance of a *restless bandit problem* (Whittle, 1988). Typically, such problems are defined in terms of a set of $n \in \mathbb{Z}_{++}$ two-action Markov decision processes (MDPs), although generalisations to a time-varying number of MDPs (Verloop, 2016) and to more than two actions per MDP (Glazebrook et al., 2011) have been explored. The two actions are usually referred to as *active* or *play* versus *inactive* or *passive* and each of the MDPs is referred to as an *project* or *arm.*

In a restless bandit problem, these n MDPs are coupled into a single MDP as follows. The state space is the Cartesian product of the state spaces of the projects, and the state of each project transitions independently of the other projects given the actions taken on that project. Thus the transitions of a project depend only on the actions taken on that project and on that project's current state. The objective is to find a non-anticipative policy that minimises the sum of the projects' individual performance functionals if those functionals all represent costs (or that maximises the sum of the projects' individual performance functionals if those functionals all represent rewards), for all initial states. The precise notion of the performance functional for an individual project depends on the setting: infinite horizon average cost and infinite horizon discounted cost settings are both commonly considered.

However, the action space is only a subset of the Cartesian product of the action spaces of the projects, as there is a constraint on the number m of projects that are simultaneously active at each time, where $m \in \mathbb{Z}_{++}$ with $m < n$. Typically, the constraint is that *exactly* m projects are active at each time, but this is readily relaxed to a constraint that *at most* m projects are active by including "dummy projects", whose cost is always zero, in the population of n projects. More general constraints have been explored (Niño-Mora, 2015), in which each project consumes resources as a function of both its state and the action taken, and the total cost of the resources consumed at each time is constrained. In the absence of any such action constraint, the problem would be solved by applying an optimal policy for each project independently. Moreover, it turns out that if the constraint were only on the (discounted) time-average number of projects that are simultaneously active, rather than a constraint at each time, the problem could again be separated into n smaller problems after introducing a Lagrange multiplier. Indeed, this observation was one motivation for the Whittle-index approach first proposed in Whittle (1988), as discussed below.

Let us relate the above definition to the typical usage of the term *bandit* in the machine-learning literature. In that context, multi-armed bandits are reinforcement-learning problems involving a set of projects whose reward distributions are unknown. At each time, the learner must select which project to play. Such bandits involve a trade-off between exploring projects to acquire information about their expected payoffs and exploiting projects with the highest expected payoffs. In the simplest versions of such problems, where the prior on the reward distributions is independent over projects, each project can be viewed as an MDP whose state corresponds to the belief about that project's payoff distribution. Each time the project is played, its reward is observed and this belief is updated. Such updates correspond to state transitions. Each time the project is inactive, its state does not change.

If we allow projects to make general Markovian state transitions, not just transitions corresponding to belief updates, while preserving the requirement that a project only changes state when it is played, then we arrive at a more general class of problems known as *ordinary* or *classical bandits* (Gittins et al., 2011). In turn, restless bandits generalise ordinary bandits in two ways. Firstly, restless bandits allow more than one project to be simultaneously active (if $m > 1$). Secondly, restless bandits allow the state of a project to change even when the project is not active, which is why they are called *restless*.

While this additional generality is important in modelling real-world problems, it comes at a price. On the one hand, the *Gittins index policy* is optimal for ordinary bandit problems and can be computed in polynomial time for problems with finite state spaces (Niño-Mora, 2007). On the other hand, it is in general PSPACE-hard (Papadimitriou and Tsitsiklis, 1999; Guha et al., 2010) to find policies that approximate optimal policies for restless bandit problems with finite state spaces to any non-trivial factor. At first glance, this might suggest that the multi-target tracking problem addressed here, with uncountable state-space $\mathbb{R}_+$ or $\mathbb{R}_{++}$, is impossibly difficult. At second glance, this poses an interesting question: for which restless bandit problems can we find approximately-optimal policies efficiently?

Whittle Index Policy. Whittle (1988) proposed a policy which generalises the Gittins index policy to restless bandit problems. This policy associates a real (or in some definitions an extended-real) number $\lambda_i^*(x_i)$ called the *Whittle index* with the state x_i of each project i. The policy then plays the m projects with the largest Whittle indices at each time, or for restless bandits with more general constraints, it selects a subset of projects with the largest Whittle indices such that the constraint is met. Ties are usually broken uniformly at random or according to a predefined priority ordering.

Whittle's index policy has been the subject of great interest for computational, empirical and theoretical reasons. The policy is potentially attractive in terms of computational cost as it reduces the original restless bandit problem, whose state space is the Cartesian product of the state spaces of the projects, to the computation of n Whittle indexes for individual projects. The policy is also attractive from a systems-architecture point-of-view, as it allows one to mix-and-match different types of projects, and it naturally accommodates the arrival or departure of projects in the sense that the Whittle index does not depend on the number of projects n. Additionally, extensive numerical tests of Whittle's policy in different applications repeatedly demonstrate that it performs remarkably well when the projects are all indexable. Indeed, 12 references to such empirical work are cited in Section 8 of Verloop (2016).

The literature contains several definitions of the Whittle index $\lambda_i^*(x_i)$ of project i, which are not equivalent in general, although they turn out to be equivalent for the problem addressed in this paper. The definition in Whittle (1988) is slightly informal and does not clearly distinguish between these definitions. All the definitions involve a modified version of project i's MDP, called the *λ-price problem*. For restless bandits with a constraint on the number m of projects that can be simultaneously active, the λ-price problem involves replacing the reward $r_i(x_i, a_i)$ for taking action a_i in state x_i by $r_i(x_i, a_i) - \lambda a_i$ where $\lambda \in \mathbb{R}$ represents a price for taking the active action $a_i = 1$. Verloop (2016) then defines $\lambda_i^*(x_i)$ as the least price λ for which action $a_i = 0$ is optimal for the λ-price problem in state x_i. Meanwhile, Guha et al. (2010) define $\lambda_i^*(x_i)$ as the largest price λ for which the actions $a_i = 0$ and $a_i = 1$ are both optimal for the λ-price problem in state x_i.

In this paper, we use the following definition (Niño-Mora, 2014, 2015) which applies to restless bandits with general resource-consumption constraints. Let us drop the subscript i and consider a single project $\mathcal{P} = \langle \mathcal{X}, c, r, \mathscr{Q}, \beta \rangle$ with state space $\mathcal{X}$, resource function $c : \mathcal{X} \times \{0, 1\} \to \mathbb{R}$, reward function $r : \mathcal{X} \times \{0, 1\} \to \mathbb{R}$, transition law $\mathscr{Q}$ and discount factor $\beta \in [0, 1)$, with the following interpretation. At the start of period $t \in \mathbb{Z}_+$ the state $X_t \in \mathcal{X}$ is observed and an action $A_t \in \{0, 1\}$ is chosen. If $X_t = x$ and $A_t = a$ then $c(x, a)$ units of resource are consumed (the projects of the previous paragraph would have $c(x, a) = a$), the project yields a reward $r(x, a)$ and the state transitions to X_{t+1} which has distribution $\mathscr{Q}(\cdot|x, a)$. Let Π denote the class of all history-dependent randomised policies for project $\mathcal{P}$ and let $\mathbb{E}_x^\pi$ denote the expectation over sequences $(X_t, A_t)_{t=0}^\infty$ with initial state $X_0 = x$, where actions are taken according to policy π and transitions are according to $\mathscr{Q}$. Then, the *λ-price problem* for project $\mathcal{P}$ and price $\lambda \in \mathbb{R}$ is to find a policy $\pi_\lambda^* \in \Pi$ that maximises the performance functional

$$\mathbb{E}_x^\pi \left[\sum_{t=0}^\infty \beta^t \left(r(X_t, A_t) - \lambda c(X_t, A_t) \right) \right]$$

for all initial states $x \in \mathcal{X}$.

Definition 3 ***(Niño-Mora, 2014, 2015)***

*The **Whittle index** of project $\mathcal{P}$ in state x is a price $\lambda^*(x)$ for which*

1. *Action $a = 1$ is optimal in state x of the λ-price problem if and only if $\lambda \leq \lambda^*(x)$,*
2. *Action $a = 0$ is optimal in state x of the λ-price problem if and only if $\lambda \geq \lambda^*(x)$.*

*Project $\mathcal{P}$ is **indexable** if it has a Whittle index $\lambda^*(x)$ for all states x in its state space.*

For all of the above definitions, it is immediate that the Whittle index is unique if it exists. Verloop's definition has the advantage that the Whittle index, and hence the Whittle index policy, exist for a wider range of projects. On the other hand, if we know project i is indexable, the definition used in this paper has the advantage that we know we have found the Whittle index when we find a price $\lambda \in \mathbb{R}$ for which actions $a_i = 0$ and $a_i = 1$ are both optimal in state x_i of the λ-price problem.

Although Whittle's policy is not an optimal policy for general restless bandits, under certain sufficient conditions and for a certain limit, it is an *asymptotically* optimal policy. Specifically, in the limit as the number of projects n tends to infinity, while the number of projects that can be simultaneously active m varies in such a way that m/n is as constant as possible, the ratio of the cost-rate of Whittle's policy to the cost-rate of an optimal policy for the given n, m tends to one. Assuming an average-cost setting, for collections of identical projects whose size n does not vary with time, where each project has a finite state space, Whittle (1988) originally conjectured that it was sufficient that the identical project was indexable for such an asymptotic optimality result to hold. However, Weber and Weiss (1990) found counterexamples to this conjecture. Nevertheless, Weber and Weiss also found sufficient conditions for asymptotic optimality to hold, and those sufficient conditions imply that the projects are indexable (Lemma 2 of that paper). Under similar sufficient conditions, this asymptotic optimality result has recently been generalised by Verloop (2016) to restless bandits with dynamic populations of non-identical projects. Both the results of Weber and Weiss and the results of Verloop assume an average-cost setting and projects with finite state spaces. So new theoretical work may be required to understand asymptotic optimality for projects with uncountable state spaces, as studied here.

The Whittle index for a project is often written as the ratio of a *marginal reward* metric to a *marginal resource* metric, where:

- The *marginal reward* is the (discounted) reward-to-go by taking the active action ($a = 1$) then following an optimal policy minus the (discounted) reward-to-go by taking the passive action ($a = 0$) then following an optimal policy;
- and the *marginal resource* is the (discounted) resource-to-go by taking the active action ($a = 1$) then following an optimal policy minus the (discounted) resource-to-go by taking the passive action ($a = 0$) then following an optimal policy;
- with the additional complexity that *optimal policy* here means optimal for the λ-price problem where λ equals the Whittle index in state x.

For examples of such expressions, see for instance: Niño-Mora (2002, equation (4.14), Theorem 4.7 and Section 6), Niño-Mora (2006, equation 19), Niño-Mora (2007, equation 6), Gittins et al. (2011, Theorem 6.4) and Larrañaga et al. (2016, Proposition 2). In particular, Niño-Mora (2002) appears to have been the first ever paper to use such marginal metrics to study questions of indexability, it also introduced resource metrics that are more general than those considered by Whittle (1988) and it established indexability of a general birth-death model. The intuition behind this expression is as follows. By definition, the Whittle index for state x corresponds to a price λ that makes both action 0 and action 1 optimal when in state x in the λ-price problem. Now, if both of these actions are optimal, then they are *equally good*. Let the reward-to-go by following policy π for a project $\mathcal{P} = \langle \mathcal{X}, c, r, \mathcal{Q}, \beta \rangle$ be denoted by

$$F(x, \pi) := \mathbb{E}_x^\pi \left[\sum_{t=0}^{\infty} \beta^t r(X_t, A_t) \right], \tag{7}$$

let the resource-to-go be denoted by

$$G(x,\pi) := \mathbb{E}_x^\pi\left[\sum_{t=0}^{\infty}\beta^t c(X_t, A_t)\right],$$

let $\langle a, \pi\rangle$ be the policy that first takes action a then follows policy π at subsequent times and let π^*_λ denote an optimal policy for the λ-price problem for project $\mathcal{P}$. Then the condition that the price λ makes actions 0 and 1 *equally good* in state x reads

$$F(x,\langle 0,\pi^*_\lambda\rangle) - \lambda G(x,\langle 0,\pi^*_\lambda\rangle) = F(x,\langle 1,\pi^*_\lambda\rangle) - \lambda G(x,\langle 1,\pi^*_\lambda\rangle).$$

This rearranges to give

$$\lambda = \frac{F(x,\langle 1,\pi^*_\lambda\rangle) - F(x,\langle 0,\pi^*_\lambda\rangle)}{G(x,\langle 1,\pi^*_\lambda\rangle) - G(x,\langle 0,\pi^*_\lambda\rangle)} \tag{8}$$

which is the ratio of a marginal reward to a marginal resource, as claimed. Unfortunately, this expression is usually only an *implicit* expression for λ, since the right-hand side involves π^*_λ. However, for projects of the multi-target tracking problem, this turns out to be an explicit expression, as we shall see below.

Literature on Restless Bandit Approach to Multi-Target Tracking. Whittle (1988) mentioned the problem of m aircraft tracking $n > m$ submarines as an example of a restless bandit problem. That problem seems rather an interesting challenge as it seems to couple multi-target tracking with a pursuit-evasion game. Nevertheless, the idea of taking a restless-bandit approach to multi-target tracking was discussed many times before anyone even experimented with such an approach, at least in the public literature. For instance, La Scala and Moran (2006) pointed out the potential interest of a restless multi-armed bandit approach to the multi-target tracking problem, but they did not pursue the Whittle index approach, rather focussing on trying to find conditions under which a one-step greedy policy is optimal when tracking a pair of targets with a single sensor. Also, Washburn (2008) reviewed applications of multi-armed bandit approaches to partially-observed sensor-management problems, expressing doubts as to whether Whittle's indexability conditions typically hold for such problems.

In contrast to such doubts, Le Ny et al. (2011) claimed to have found conditions under which the continuous-time version of the problem is indexable, at least for a scalar state, in the average-cost case, and with posterior variance as the cost function. A closer reading of that paper reveals that those authors present an incomplete argument leading to hypothetical values for optimal thresholds under the assumption that threshold policies are optimal. In particular, a key step of that argument is those authors' Theorem 1, which is Proposition 8 of Whittle (1988). This result gives an expression for the Whittle index, which the authors then invert to find a cubic equation for an optimal threshold. Now, as remarked by Whittle (1988, p. 295), *"that argument is formal, and conditions are certainly required for the calculations to make sense"*. However, Le Ny et al. (2011) do not try to find a suitable set of such conditions. Further, those authors *assume* the form of the optimal policy hoping that it can be verified a posteriori by substituting their expressions for an optimal threshold into the dynamic programming equation. However, they make no attempt to perform such a substitution and in fact such a verification may prove challenging since the thresholds are given by a cubic equation, and because the dynamic programming equation involves the trajectory of the posterior variance which is the solution of a Ricatti differential equation. In summary, the arguments presented in Le Ny et al. (2011) are incomplete and do not convincingly demonstrate the optimality of threshold policies and the indexability for continuous-time systems.

Meanwhile, discrete-time versions of the problem have proved to be far more challenging. Niño-Mora and Villar (2009) made the first empirical evaluation of Whittle's index policy applied to the multi-target tracking problem with a scalar state in discrete time. That paper claims that the problem is indexable and suggests an approach to studying its indexability based on partial conservation laws, which is the approach that we take in Section 3 of this paper. However, it provides no proof or argument to substantiate this claim. The dissertation of Villar (2012) began a theoretical investigation of the problem, without establishing indexability of the model of concern. More recently, Dance and Silander (2015) proved that the index function is a monotone function of the variance state, but they did so under the assumption that threshold policies are optimal.

A number of closely-related problems have also been explored. For instance, Le Ny et al. (2008) and Liu and Zhao (2010) both establish indexability results about the tracking of n targets with *binary* state-spaces given $m < n$ sensors. Meanwhile, Niño-Mora (2016) empirically explored a Whittle's index approach to a generalisation of the multi-target tracking problem in which measurements are randomly jammed.

Whittle Index for the Multi-Target Tracking Problem. The above discussion prompts the third question addressed in this paper: *is each project of the multi-target tracking problem indexable, and if so, what is a computationally-convenient expression for the Whittle index of a project?* To state the first part of this question explicitly, consider a single project of the multi-target tracking problem, corresponding to one of n time series to be tracked with $m < n$ sensors. For each $\lambda \in \mathbb{R}$ and initial state $x \in \mathcal{I}$, the corresponding λ-price problem then to minimise the performance functional

$$\mathbb{E}_x^\pi \left[\sum_{t=0}^{\infty} \beta^t \left(\lambda A_t + C(X_t) \right) \right] \tag{9}$$

with respect to the policy π for taking actions A_t where the state-sequence is given in terms of the variance updates of (4) as

$$X_0 = x, \qquad\qquad X_{t+1} = \phi_{A_t}(X_t),$$

for $t = 0, 1, \ldots$. The question is then whether there exists a simple-to-compute function $\lambda^* : \mathcal{I} \to \mathbb{R}$, for which action $A_t = 1$ is optimal if and only if $\lambda \leq \lambda^*(X_t)$, while action $A_t = 0$ is optimal if and only if $\lambda \geq \lambda^*(X_t)$, for all $t \in \mathbb{Z}_+$.

As suggested by Theorem 1 (Excerpt), to state a convenient expression for the Whittle index, some definitions concerning threshold policies should be useful. So, for any threshold $s \in \overline{\mathbb{R}}$, let the *s-threshold policy* be the policy that takes the action 1 if the state exceeds the threshold s and takes action 0 otherwise. Also, let $X_t(x, a; s)$ denote the state at time $t = 0, 1, \ldots$ if the system starts in state $x \in \mathcal{I}$ at time $t = 0$, then action $a \in \{0, 1\}$ is taken and the s-threshold policy is followed thereafter, so that $A_t(x,a;s) := \mathbf{1}_{X_t(x,a;s)>s}$ for $t > 0$. The answer to our third question is then given by the following theorem, whose proof is given in Section 3.

Theorem 1 *Suppose A1 holds. Then the family of λ-price problems, given by equation (9), is indexable and for each $x \in \mathcal{I}$ the Whittle index is*

$$\lambda^*(x) := \frac{\sum_{t=0}^{\infty} \beta^t (C(X_t(x,0;x)) - C(X_t(x,1;x)))}{\sum_{t=0}^{\infty} \beta^t (A_t(x,1;x) - A_t(x,0;x))}.$$

Furthermore,

1. *If $\lambda^*(s) = \lambda$ for some $s \in \mathcal{I}$ then both of the following threshold policies are optimal:*

$$A_t = \mathbf{1}_{X_t>s}, \qquad\qquad A_t = \mathbf{1}_{X_t \geq s};$$

2. *If $\lambda^*(s) > \lambda$ for all $s \in \mathcal{I}$ then the always-active policy is the unique optimal policy;*
3. *If $\lambda^*(s) < \lambda$ for all $s \in \mathcal{I}$ then the always-passive policy is the unique optimal policy.*

This paper thus generalises the work of Dance and Silander (2015) by demonstrating that threshold policies are in fact optimal for the single project problem, which was Assumption A1 of (Dance and Silander, 2015). It also generalises by considering the case of multipliers $\bar{A} < 1$ rather than only considering $\bar{A} = 1$, where $\bar{A}$ is as in equation (1), and by considering cost functions $C(x) \neq x$ other than the (scaled) posterior variance.

1.5. Intuitive Guide to the Paper

As with other work on Markov decision processes, we work with the cost-to-go $Q(x,a)$ when starting in initial state x and taking initial action a, but then following an optimal policy. A common way to prove that threshold policies are optimal when the state x is real-valued, is to show that the difference $Q(x,1) - Q(x,0)$ is a non-increasing function of x. Such approaches have been studied by Serfozo (1976), Altman and Stidham Jr. (1995) and Altman et al. (2000). Unfortunately, as shown in Figure 1, such an approach fails for the process considered in this paper, even when the cost equals the variance.

Instead, this paper proves the optimality of threshold policies using a new verification theorem by Niño-Mora (Niño-Mora, 2015, 2019). This theorem applies to Markov decision processes that satisfy the so-called *partial conservation law indexability* (PCLI) conditions (Section 3). The central concept underlying the verification theorem is the *marginal productivity index* which turns out to be equal to the ratio $\lambda^*(\cdot)$ given in Theorem 1.

One of the PCLI conditions requires that the marginal productivity index is a non-decreasing and continuous function of the state x and that it is bounded from below. This is the most challenging of the conditions to verify. As a quick check, we plot $\lambda^*(x)$ in Figure 2. Although $\lambda^*(x)$ is increasing, the numerator and denominator have a fractal structure, so it is surprising that the index is continuous. Furthermore, if we subtract a cubic fit to $\lambda^*(x)$, the residual has a complicated sequence of cusps. Therefore the paper then focusses on characterising the sequence of actions $A_t(x,a;x)$ resulting from applying

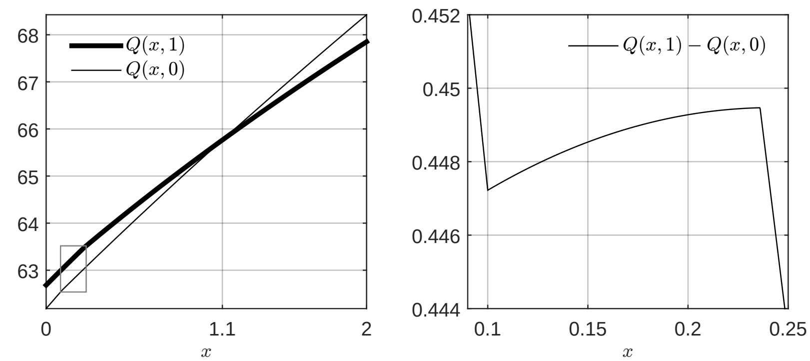

Figure 1: Counterexample to monotonicity of the difference in Q-functions. The functions $Q(x,0)$ and $Q(x,1)$ cross only a single time at $x = 1.1$ (*left plot*). However, the difference $Q(x,1) - Q(x,0)$ is increasing for some x (*right plot*, for x in the left plot's grey box). The model has $\beta = 0.95$, $C(x) = x$, $\phi_0(x) = x + 1$, $\phi_1(x) = 1/(\theta_1 + 1/(x+1))$ with $\theta_1 = 0.1$ and $\nu = 0.7647$.

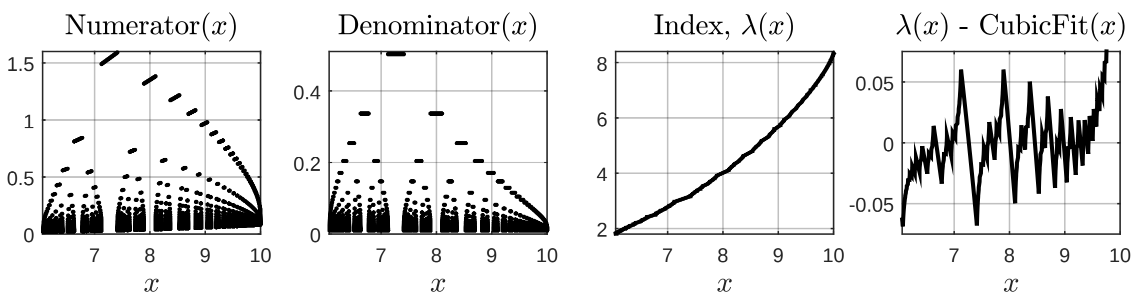

Figure 2: The numerator (*left*) and denominator (*mid-left*) of the index (*mid-right*), and the error in a cubic fit to the index (*right*). The model has cost $C(x) = x$, discount factor $\beta = 0.99$, map-with-a-gap $\phi_0(x) = r^2x+1$ and $\phi_1(x) = 1/(\theta_1+1/(r^2x+1))$ with $r^2 = 0.9$ and $\theta_1 = 0.01$.

an x-threshold policy, that give rise to this fractal pattern. We describe these sequences in terms of special binary strings that we call *$\mathcal{M}$-words* (Corollary 13 of Section 2). Such words include the set of all *Christoffel words* (Berstel et al., 2008) as well as a subset of the family of *Sturmian words* (Lothaire, 2002).

To understand what an $\mathcal{M}$-word is, consider the Cartesian coordinates $(x, y(x))$ of the points on a straight line $y(x) = \alpha x$ through the origin with slope $\alpha \in [0, 1]$. To draw this line on a screen with pixel coordinates that are integer pairs, for each integer $\tilde{x}$ one might fill the pixel $(\tilde{x}, \tilde{y}(\tilde{x}))$ where $\tilde{y}(\tilde{x}) = \lfloor \alpha\tilde{x} \rfloor$. Now consider the differences $\Delta_{\tilde{x}} := \tilde{y}(\tilde{x}) - \tilde{y}(\tilde{x} - 1)$ for positive integers $\tilde{x}$. As α is in the range $[0, 1]$ it follows that each such difference $\Delta_{\tilde{x}}$ is in the set $\{0, 1\}$. Also, the sequence $\Delta := (\Delta_1, \Delta_2, \Delta_3, \dots)$ is periodic when α is rational and aperiodic when α is irrational. Correspondingly, the *$\mathcal{M}$-word of rate* α is the finite sequence $(\Delta_1, \Delta_2, \dots, \Delta_n)$ if Δ is periodic with (shortest) period n, and otherwise it is the infinite sequence Δ itself.

The characterisation of the action sequences $(A_t(x, a; x))_{t=0}^{\infty}$ viewed as a function of $x \in \mathcal{I}$ (noting that x represents both the initial state and the threshold) shows that the interval $\mathcal{I}$ can be partitioned into intervals corresponding to finite $\mathcal{M}$-words and points corresponding to infinite aperiodic $\mathcal{M}$-words. In the most important case for our analysis, x lies in such an interval and the denominator of the index involves periodic sequences of the form

$$\left.\begin{aligned} (A_t(x,1;x))_{t=0}^{\infty} &= (1,0,p_1,\dots,p_n,1,0,p_1,\dots,p_n,\dots) \\ (A_t(x,0;x))_{t=0}^{\infty} &= (0,1,p_1,\dots,p_n,0,1,p_1,\dots,p_n,\dots) \end{aligned}\right\} \tag{10}$$

for some binary sequence $(p_1, \dots, p_n)$. Thus the denominator equals

$$\sum_{t=0}^{\infty} \beta^t (A_t(x,1;x) - A_t(x,0;x)) = 1 - \beta + \beta^{n+2} - \beta^{n+3} + \cdots = \frac{1-\beta}{1-\beta^{n+2}}.$$

Hence, the behaviour of the denominator in Figure 2 is given entirely by changes in the period $(n + 2)$ of this binary sequence and does not depend on the values of $p_1, \dots, p_n$.

As this denominator equals a positive constant on such intervals, to show that the index is non-decreasing on such an interval, we show that the derivative of the numerator is non-negative. One key to proving this is the fact that the mappings $\phi_0(x)$ and $\phi_1(x)$ are Möbius transformations of the form

$$\mu_B(x) := \frac{B_{11}x + B_{12}}{B_{21}x + B_{22}}$$

for some $B \in \mathbb{R}^{2\times 2}$. Now the composition of Möbius transformations is homomorphic to matrix multiplication, so that

$$\mu_B(\mu_D(x)) = \mu_{BD}(x)$$

for any $B, D \in \mathbb{R}^{2\times 2}$. Further, if $\det(B) = 1$ then the derivative of the corresponding Möbius transformation is

$$\frac{d}{dx}\mu_B(x) = \frac{1}{(B_{21}x + B_{22})^2}$$

which is a convex function for $x \in \mathbb{R}_+$ and $B_{21}, B_{22} \in \mathbb{R}_{++}$. So the derivative of the numerator of the index, in the case $C(x) = x$, is of the form

$$\sum_{t=0}^{\infty} \frac{\beta^t}{(B_{21}^{(0,t)} x + B_{22}^{(0,t)})^2} - \sum_{t=0}^{\infty} \frac{\beta^t}{(B_{21}^{(1,t)} x + B_{22}^{(1,t)})^2}$$

where $(B^{(a,t)})_{t=0}^{\infty}$, for $a \in \{0,1\}$, are sequences of 2×2 matrices that are determined by the action sequences $(A_t(x, a; x))_{t=0}^{\infty}$, but which otherwise do not depend on the value of x. This derivative is the difference of two sums, and each of these sums involves a sequence of convex functions $(z \mapsto \beta^t/z^2)_{t=0}^{\infty}$ applied to a sequence of linear functions of x. Inequalities involving such sums can be addressed by the theory of *majorisation* (Marshall et al., 2010), provided the sequences $(B_{21}^{(a,t)} x + B_{22}^{(a,t)})_{t=0}^{\infty}$ of linear functions of x satisfy certain *majorisation conditions* involving partial sums of those sequences. In the case where x lies in an interval such that the action sequences are given by (10), it turns out that those majorisation conditions are satisfied because the sequence $p_1, \ldots, p_n$ is a *palindrome*. That is, the sequence reads the same forwards as backwards, so that $p_k = p_{n-k}$ for $k = 1, \ldots, n-1$: see (11) for a proof of this well-known palindromic property of Christoffel words.

1.6. Structure of the Paper

First we relate the sequence of actions under threshold policies to $\mathcal{M}$-*words* (Section 2) before presenting proofs of our main results, Theorem 1 and Corollary 2 (Section 3). The proofs are based on Niño-Mora's theorem about the optimality of threshold policies (Section 3.1), which uses three *partial conservation law indexability (PCLI)* conditions. We use the properties of $\mathcal{M}$-words to demonstrate that each PCLI condition holds. These conditions concern the positivity of a *marginal resource* metric (Section 3.2), the continuity and non-decreasing nature of a *marginal productivity index* (Section 3.3), and a condition that characterises that index as a Radon-Nikodym derivative (Section 3.4). These results are then coupled into a proof of Theorem 1 (Section 3.5) and a proof of Corollary 2 is presented immediately thereafter (Section 3.6).

Having completed the proofs, we then turn to closed-form expressions for the index and numerical methods for evaluating it when such closed forms are not available (Section 4). We demonstrate the accuracy of such numerical methods and show how the index varies as its parameters change. Also, we compare the performance of Whittle's index policy with other well-known heuristics. Finally, we discuss interesting avenues for further work (Section 5). The appendices contain detailed proofs about the relation of itineraries to $\mathcal{M}$-words (Appendix A), of a key majorisation inequality (Appendix B) and about the linear systems orbits to which this majorisation result is applied (Appendix C).

2. Itineraries and Words

The transitions from state-to-state under threshold policies are given by a discontinuous mapping known as a *map-with-a-gap*, and the corresponding action sequence is known as the *itinerary* of that map. A detailed understanding of the properties of such itineraries is central to our proof of the optimality of threshold policies for the problems posed in the Introduction. So the main purpose of this section is to characterise these itineraries in terms of special binary strings that we call $\mathcal{M}$-*words* and to discuss relevant properties of such words.

This section is structured as follows. We begin by introducing definitions and notations for itineraries (Section 2.1). As itineraries can be viewed as infinite strings, we then remind the reader of notations commonly used in the study of combinatorics on strings (Section 2.2). This notation enables us to define $\mathcal{M}$-words and an important subset of such words called the *Christoffel words* (Section 2.3). It also enables us to discuss properties of such words that are important for our proof of the optimality of threshold policies, notably a palindromic property (11), a description of the lexicographic ordering of the cyclic rotations of such words which are called *conjugates* (Lemma 8), and a description of the prefixes of $\mathcal{M}$-words (Lemma 10). Finally, we introduce three results (Theorem 12, Corollary 13 and Theorem 14) which describe itineraries in terms of $\mathcal{M}$-words, for a variety of relations between the threshold and initial state, and for versions of the threshold policy that are either active or passive at the threshold (Section 2.4).

2.1. Maps-with-Gaps

Many phenomena involve the iterated application of discontinuous maps or *maps-with-gaps* in the terminology of Hogan et al. (2007). Such phenomena are important in control problems (Haddad et al., 2014), in physics, electronics and mechanics (Bernardo et al., 2008; Makarenkov and Lamb, 2012), economics (Tramontana et al., 2010), biology and medicine (Aihara and Suzuki, 2010). Such maps either arise directly from a discrete-time model or they may arise as the Poincaré maps of continuous-time systems. The *orbit* and *itinerary* of the following definition, which is illustrated in Figure 3, are the standard way to describe the behaviour of such maps in the study of dynamical systems (Devaney, 2008).

Definition 5 *Consider an interval $\mathcal{I}$ of $\mathbb{R}$, functions $\phi_0 : \mathcal{I} \to \mathcal{I}$ and $\phi_1 : \mathcal{I} \to \mathcal{I}$, an initial state $x \in \mathcal{I}$ and a threshold $s \in \overline{\mathbb{R}}$. The s-**threshold orbit** from x for ϕ_0 and ϕ_1, denoted by $\operatorname{orbit}(x|s, \phi_0, \phi_1)$, is the sequence $(x_k)_{k=1}^{\infty}$ with*

$$x_1 = x \quad \textit{and} \quad x_{k+1} = \begin{cases} \phi_0(x_k) & \textit{if } x_k \le s \\ \phi_1(x_k) & \textit{if } x_k > s. \end{cases}$$

*In terms of that sequence, the s-**threshold itinerary** from x for ϕ_0 and ϕ_1 is the infinite binary string $\sigma(x|s, \phi_0, \phi_1)$ with letters $\sigma(x|s, \phi_0, \phi_1)_k := \mathbf{1}_{x_k > s}$ for $k \in \mathbb{Z}_{++}$.*

Looking ahead, when we use the results of this section to prove Theorem 1, the mappings in the above definition will represent the transitions of a (relative) variance state. Thus ϕ_0, ϕ_1 and the interval $\mathcal{I}$ will satisfy Assumption A1 after a change of coordinates (as discussed straight after Theorem 12). However, the analysis of this section and Appendix A applies to *any* real interval $\mathcal{I}$ if not otherwise stated, and in most cases the analysis will make Assumption A2.

In most cases, the maps ϕ_0 and ϕ_1 are clear from the context, so we talk of the s-threshold orbit from x, denoted by $\operatorname{orbit}(x|s)$, and the s-threshold itinerary from x, denoted by $\sigma(x|s)$. Also, we use of the s^--*threshold orbit* from x, which is the left limit $\operatorname{orbit}(x|s^-) =$

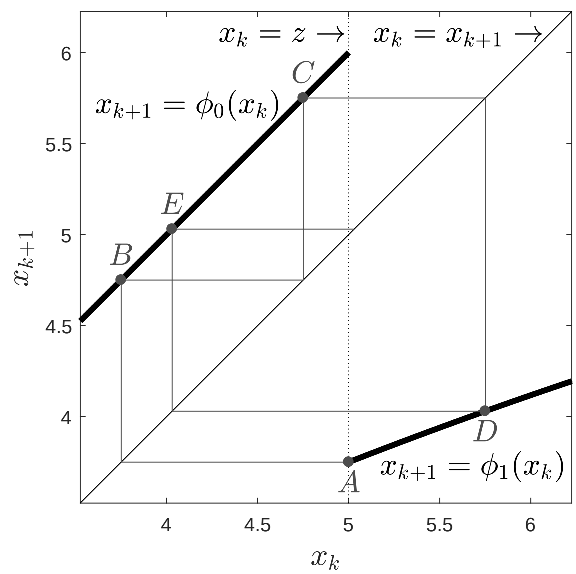

Figure 3: Map-with-a-gap. The s^--threshold orbit from $x = 5$ traces the path $ABCDE\ldots$ corresponding to the itinerary $10010\ldots$. The map has $\phi_0(x) = x + 1$, $\phi_1(x) = 1/(\theta_1 + 1/(x + 1))$ with $\theta_1 = 0.1$ and threshold $s = 5$.

$\lim_{s'\uparrow s} \text{orbit}(x|s') = (x_k)_{k=1}^{\infty}$ with

$$x_1 = x \quad \text{and} \qquad x_{k+1} = \begin{cases} \phi_0(x_k) & \text{if } x_k < s \\ \phi_1(x_k) & \text{if } x_k \geq s \end{cases}$$

and the corresponding *s^--threshold itinerary* from x, with letters $\sigma(x|s^-)_k := \mathbf{1}_{x_k \geq s}$ for $k \in \mathbb{Z}_{++}$. These left limits exist provided the maps ϕ_0, ϕ_1 are continuous, which is true when Assumption A2 holds (Lemma 32).

It is helpful to view such itineraries as *words* as we now explain.

2.2. Standard Definitions about Words

We remind the reader of standard definitions used in the study of the combinatorics of strings: see for instance Lothaire (2002) and Berstel et al. (2008).

In this paper, a *word* w is a string on the alphabet $\{0, 1\}$ and the empty word is denoted by ϵ. The *length* of a word w is the number of letters in the string, which is finite or countably infinite, and is denoted by $|w|$. The k^{th} letter of word w is w_k for $k \in \mathbb{Z}_{++}$ with $k \leq |w|$. Letters i-through-j of word w are denoted by $w_{i:j} := w_i w_{i+1} \ldots w_j$ for $i, j \in \mathbb{Z}_{++}$ with $i \leq j \leq |w|$. For $j < i$, we treat $w_{i:j}$ as the empty word. The *reverse* word of a finite word w is denoted by $w^R := w_{|w|} \ldots w_2 w_1$. A finite word satisfying $w^R = w$ is called a *palindrome*.

The *concatenation* of a finite word u and a word v is denoted by uv. For $n \in \mathbb{Z}_+$, the n-fold concatenation of a finite word w is denoted by w^n, with the convention that $w^0 = \epsilon$, and the word resulting from infinitely concatenating the word w is denoted by w^∞. For an infinite word w and $n \in \mathbb{Z}_{++}$ we define $w^n = w^\infty = w$.

A finite word f is a *factor* of a word w if $w = ufv$ for some finite word u and some word v. The number of times that word f appears in w, overlapping appearances included, is denoted by $|w|_f$. A finite word p is a *prefix* of word w if $w = ps$ for some word s and a word s is a *suffix* of word w if $w = ps$ for some finite word p.

We say that a word u is *lexicographically less than* a word v, written $u \prec v$, if either u is a finite word and $v = ua$ for some non-empty word a, or if $u = a0b$ and $v = a1c$ for some finite word a and some words b and c. We use $\succ, \preceq$ and $\succeq$ for the other lexicographic ordering relations.

We say an infinite word w is the *limit* of a sequence of words $(w^{(n)})_{n=1}^\infty$ and write $w = \lim_{n\to\infty} w^{(n)}$ if for each $i \in \mathbb{Z}_{++}$ there is an $n \in \mathbb{Z}_{++}$ such that $w_i = w_i^{(m)}$ for all $m \in \mathbb{Z}_{++}$ with $m \geq n$.

The *rate* of any finite non-empty word w is the ratio $\text{rate}(w) := |w|_1/|w|$ whereas for an infinite word w, when the limit exists, we define $\text{rate}(w) := \lim_{n\to\infty} |w_{1:n}|_1/n$. While some authors refer to such ratios as the "slope" of a word, we use the term "rate" as the "slope" of a word w is sometimes defined as the ratio $|w|_1/|w|_0$ and this seems justified from a geometrical point of view in terms of *digital straight lines* (Berstel et al., 2008).

Examples. For $w = 010111$ we have $|w| = 6$, $w_3 = 0$, $w_{2:4} = 101$, $|w|_{01} = |w|_{11} = 2$ and $w^2 = ww = 010111010111$. Also for $a = 01, b = 11$ we have $w = aab$ and $a \prec w \prec b$.

2.3. $\mathcal{M}$-Words

We characterise itineraries of maps-with-gaps in terms of the following type of words.

Definition 6 *The **$\mathcal{M}$-word** of rate $\alpha \in [0,1]$ is the shortest word w such that*

$$(w^\infty)_n = \lfloor \alpha n \rfloor - \lfloor \alpha(n-1) \rfloor \quad \text{for } n \in \mathbb{Z}_{++}.$$

*If α is rational then w is called a **Christoffel word**. If α is irrational then w is called a **Sturmian $\mathcal{M}$-word**.*

Example. The shortest $\mathcal{M}$-words are the words $0, 1$ and 01 with rates $0, 1$ and $\frac{1}{2}$.

The notion of a Christoffel word is a standard one (Berstel et al., 2008). Also, *Sturmian words* have been studied intensively, for instance Lothaire (2002, Chapter 2) says an infinite word w is Sturmian if it has $n+1$ distinct factors of length n for any integer $n \geq 0$. Lothaire (2002) then goes on to explore many other definitions of Sturmian words that turn out to be equivalent. Our intention in defining $\mathcal{M}$-words is to group together the Christoffel words with a specific subset of Sturmian words, so that together they characterise the possible itineraries of a large class of maps-with-gaps, as described in Theorem 16.

We call such words $\mathcal{M}$-words as our definition is closely related to the set of *mechanical words*. For a given *slope* $\alpha \in [0,1]$ and *intercept* $\rho \in \mathbb{R}$, Morse and Hedlund (1940) defined the *upper* and *lower mechanical words* to be the infinite sequences, for $n \in \mathbb{Z}_+$,

$$u_n = \lceil \alpha(n+1) + \rho \rceil - \lceil \alpha n + \rho \rceil$$

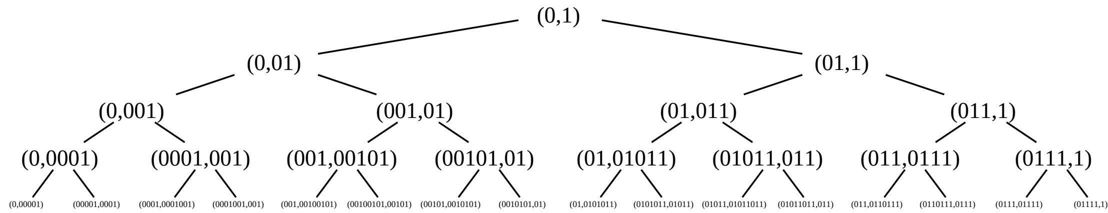

Figure 4: Part of the Christoffel tree.

$$l_n = \lfloor \alpha(n+1) + \rho \rfloor - \lfloor \alpha n + \rho \rfloor$$

Lothaire (2002) and Berstel et al. (2008) offer rich introductions to the mathematics of mechanical words, while Bousch and Mairesse (2002) and Altman et al. (2000) explore other optimisation problems that give rise to such words. Our $\mathcal{M}$-words are prefixes of lower-mechanical-words-of-zero-intercept up to a change of indexing from $n \in \mathbb{Z}_+$ to $n \in \mathbb{Z}_{++}$.

It is not hard to see that the Christoffel word of rate a/b, where a, b are relatively-prime integers, has length b. In contrast, Sturmian $\mathcal{M}$-words are infinite and aperiodic.

In general the $\mathcal{M}$-word w of rate α does have $\text{rate}(w) = \alpha$. Indeed if w is the $\mathcal{M}$-word of rate a/b for some $a, b \in \mathbb{Z}_{++}$, then

$$\text{rate}(w) = |w|_1/|w| = \lfloor (a/b)|w| \rfloor / |w| = a/b,$$

whereas, if w is an $\mathcal{M}$-word of irrational rate α, then

$$\text{rate}(w) = \lim_{n\to\infty} |w_{1:n}|_1/n = \lim_{n\to\infty} \lfloor \alpha n \rfloor / n = \alpha.$$

Furthermore, as remarked by Christoffel (1875), all Christoffel words other than the words 0 and 1 are of the form $0p1$ where the word p is a *palindrome*. Indeed for relatively-prime positive integers $m < n$, the letters of the Christoffel word w of rate m/n satisfy

$$w_{n-k} = \left\lfloor \frac{m}{n}(n-k) \right\rfloor - \left\lfloor \frac{m}{n}(n-k-1) \right\rfloor = \left\lfloor -\frac{m}{n}k \right\rfloor - \left\lfloor -\frac{m}{n}(k+1) \right\rfloor = w_{k+1} \qquad (11)$$

for $k = 1, 2, \ldots, n-2$.

The Christoffel words can be defined in other ways. In this paper the most important alternative-but-equivalent definition is in terms of the *Christoffel tree* (Figure 4), which is an infinite complete binary tree (Berstel et al., 2008) in which each node is labelled with a pair (u, v) of words, called a *Christoffel pair*. The root of the tree is labelled with the pair $(0, 1)$ and the left and right children of node (u, v) are the nodes (u, uv) and (uv, v) respectively. In fact the Christoffel words are the words $0, 1$ and the set of concatenations uv for all (u, v) in the Christoffel tree.

Another definition of Christoffel words is in terms of modular arithmetic, as in the following lemma, where we use a bar to denote the remainder modulo the length $n = |w|$ of a Christoffel word w, so that $\overline{x} := x \bmod n$ for $x \in \mathbb{Z}$, and for any positive integer n, the set of integers modulo n is denoted by $\mathbb{Z}_n := \{0, 1, \ldots, n-1\}$. The following lemma does not explicitly appear in Berstel et al. (2008) but is easily related to multiple discussions of modular arithmetic in that book.

Lemma 7 *Suppose w is a Christoffel word of length n. Let $m := |w|_1$ and $p := |w|_0$. Then*

$$w_{i+1} = \mathbf{1}_{\overline{mi} \geq p} \qquad (i \in \mathbb{Z}_n).$$

Proof As $n\lfloor mi/n \rfloor = mi - \overline{mi}$, the definition of Christoffel words gives

$$\begin{aligned} w_{i+1} &= -\lfloor mi/n \rfloor + \lfloor m(i+1)/n \rfloor \\ &= (-mi + \overline{mi} + m(i+1) - \overline{m(i+1)})/n \\ &= (-mi + \overline{mi} + m(i+1) - (\overline{mi} + m - n\mathbf{1}_{\overline{mi} \geq n-m}))/n, \end{aligned}$$

which simplifies to $\mathbf{1}_{\overline{mi} \geq p}$, as claimed. ■

Finally, we give two results about $\mathcal{M}$-words that play a key role elsewhere in the paper. The first result is about *conjugacy* and lexicographic order. In particular, we say two finite words a and b are conjugate if $a = uv$ and $b = vu$ for some words u and v. For instance, the words $a = 00011$ and $b = 01100$ are conjugate. Words that are conjugate can be viewed as cyclic shifts of each other, like the X86 assembler instruction ROL. The notion of conjugate words is also standard in the study of the combinatorics of strings (Lothaire, 2002; Berstel et al., 2008). The following lemma is rather similar to Berstel et al. (2008, Exercise 6.3, p. 49), where an outline of its proof is suggested, although the result that we state is more specific about the lexicographic ordering of the conjugates.

Lemma 8 *Suppose w is a Christoffel word of length n and that l satisfies $\overline{lm} = 1$ where $m = |w|_1$. Then the conjugates $u(i) := w_{(\bar{i}+1):n} w_{1:\bar{i}}$ satisfy*

$$w = u(0) \prec u(l) \prec u(2l) \prec \cdots \prec u((n-1)l) = w^R.$$

Furthermore, if the words c and d satisfy $w = 0dc1$, then $c01d$ and $c10d$ are lexicographically-consecutive conjugates of w.

Proof Let $x_i := \overline{mi-1}, y_i := \overline{mi}$ and $p := n - m$. Then $x_0 = n - 1$ and $x_{n-1} = p - 1$. As $\gcd(m, n) = 1$, the sequence $x_0, \ldots, x_{n-1}$ is a permutation of $\mathbb{Z}_n$. So, $x_i \notin \{p-1, n-1\}$ for $i \in \{1, \ldots, n-2\}$. As $y_i = \overline{x_i + 1}$ these results give

$$\begin{aligned} &\mathbf{1}_{x_i \geq p} > \mathbf{1}_{y_i \geq p} && \text{for } i = 0 \\ &\mathbf{1}_{x_i \geq p} = \mathbf{1}_{y_i \geq p} && \text{for } i = 1, \ldots, n-2 \\ &\mathbf{1}_{x_i \geq p} < \mathbf{1}_{y_i \geq p} && \text{for } i = n-1. \end{aligned}$$

But Lemma 7 gives $u(0)_{j+1} = \mathbf{1}_{y_j \geq p}$ and $u((n-1)l)_{j+1} = \mathbf{1}_{x_j \geq p}$ for $j \in \mathbb{Z}_n$. Thus $u(0) = 0a1$ and $u((n-1)l) = 1a0$ for some word a. But $u(0) = w$ and w is a Christoffel word, so a is a palindrome. Therefore $u((n-1)l) = w^R$.

Now for $i = 0, \ldots, n-2$, the conjugates $u(il)$ and $u((i+1)l)$ are related to $u((n-1)l)$ and $u(0)$ respectively by the same non-zero cyclic rotation. Thus $u(il) = c01d$ and $u((i+1)l) = c10d$ for some words c and d with $dc = a$. Therefore $u(il) \prec u((i+1)l)$. ■

To illustrate this lemma, consider the Christoffel word $w = 00101$. As $n = |w| = 5$ and $m = |w|_1 = 2$, we set $l = 3$ so that $\overline{lm} = 1$. The sequence $(u(kl))_{k=0}^{n-1}$ of conjugates of w as defined in the lemma is then

$$\begin{aligned} u(0) &= 00101 \\ u(3) &= 01001 \\ u(6) &= 01010 \\ u(9) &= 10010 \\ u(12) &= 10100 \end{aligned}$$

which is indeed the same as the conjugates arranged in increasing lexicographic order. This lemma plays a key role in showing that the index function in Theorem 1 is non-decreasing in the case that the cost function $C(\cdot)$ is convex, by enabling the application of a rearrangement inequality at Steps (29) and (30) of the corresponding proof.

The second result shows how the prefixes of $\mathcal{M}$-words vary as a function of their rates. It requires one more definition, which is well known: see for instance Graham et al. (1994).

Definition 9 *For each positive integer n, the **Farey sequence** F_n is the sequence of rational numbers on $[0, 1]$ whose denominator is at most n, sorted in increasing order.*

Clearly, it is implicit in this definition that the numerator and denominator of each such rational number are relatively prime. For example, the Farey sequence F_5 is

$$0, \frac{1}{5}, \frac{1}{4}, \frac{1}{3}, \frac{2}{5}, \frac{1}{2}, \frac{3}{5}, \frac{2}{3}, \frac{3}{4}, \frac{4}{5}, 1.$$

Lemma 10 *Suppose $n \in \mathbb{Z}_{++}$ and $q \in [0, 1]$. Let $q_1 < q_2 < \cdots < q_m$ be the Farey sequence F_n. Let $p(s)$ be the first n letters of the word w^∞ where w is the $\mathcal{M}$-word of rate $s \in [0, 1]$. Then $p(q) = p(q_i)$ if and only if either $q = q_i = 1$ or $q \in [q_i, q_{i+1})$ for some $1 \le i < m$.*

Proof Let $b(q) := (\lfloor q \rfloor, \lfloor 2q \rfloor, \ldots, \lfloor nq \rfloor)$ and consider the intervals $\mathcal{Q}_i := [q_i, q_{i+1})$ for $i < m$ and $\mathcal{Q}_m := \{1\}$. As the line $y = qx$ hits an integer point $(x, y) \in \mathbb{Z}^2$ with $1 \le x \le n$ and $0 \le y \le x$ if and only if q is an element of F_n, it follows that $b(q) = b(q_i)$ if and only if $q \in \mathcal{Q}_i$. Let $g(x_1, x_2, \ldots, x_n) := (x_1, x_2 - x_1, \ldots, x_n - x_{n-1})$. By definition of $\mathcal{M}$-words, $p(q) = g(b(q))$. As g is invertible it follows that $p(q) = p(q_i)$ if and only if $q \in \mathcal{Q}_i$. ■

2.4. Characterising Itineraries with $\mathcal{M}$-Words

Our aim here is to characterise the itineraries of maps-with-gaps. We first set up some notation and then demonstrate a simple result about the lexicographical ordering of itineraries. Then we state an assumption under which Theorem 12 guarantees that itineraries correspond to specific $\mathcal{M}$-words.

Let $\mathcal{I}$ be an interval of $\mathbb{R}$ and consider two mappings $\phi_0 : \mathcal{I} \to \mathcal{I}$ and $\phi_1 : \mathcal{I} \to \mathcal{I}$. For any finite word w, the *composition* $\phi_w : \mathcal{I} \to \mathcal{I}$ is the mapping

$$\phi_w(x) := \phi_{w_{|w|}} \circ \cdots \circ \phi_{w_2} \circ \phi_{w_1}(x) \quad \text{and} \quad \phi_\epsilon(x) := x.$$

A simple application of compositions gives the following result about lexicographic ordering of the itineraries $\sigma(\cdot|s)$ of a map-with-a-gap given by mappings $\phi_0 : \mathcal{I} \to \mathcal{I}$ and $\phi_1 : \mathcal{I} \to \mathcal{I}$ and threshold s. We remind the reader that this paper uses *increasing* and *decreasing* in the strict sense.

Lemma 11 *Suppose ϕ_0, ϕ_1 are increasing mappings and that $x, y \in \mathcal{I}$ with either $\sigma(x|s^-) \prec \sigma(y|s^-)$ or $\sigma(x|s) \prec \sigma(y|s)$. Then $x < y$.*

Proof If $\sigma(x|s^-) \prec \sigma(y|s^-)$ then $\sigma(x|s^-) = a0b$ and $\sigma(y|s^-) = a1c$ for some finite word a and some infinite words b, c, by the definition of lexicographic order. So, the definition of $\sigma(\cdot|s^-)$ gives $\phi_a(x) < s \leq \phi_a(y)$. But $\phi_a(\cdot)$ increasing as it is a finite composition of increasing functions. It follows that $x < y$. The proof for $\sigma(\cdot|s)$ is similar. ■

However, without additional assumptions about the mappings ϕ_0, ϕ_1, it is not possible to precisely characterise the itineraries of the associated maps-with-gaps. In order to state such an assumption, we shall say that a map $f : \mathcal{I} \to \mathcal{I}$, where $\mathcal{I}$ is an interval of $\mathbb{R}$ is *contractive* if for all $x, y \in \mathcal{I}$ with $x \neq y$, we have

$$|f(y) - f(x)| < |y - x|.$$

Also, a *fixed point* of such a map f is any $x \in \mathcal{I}$ for which the equation $x = f(x)$ is satisfied.

Assumption A2
Functions $\phi_0 : \mathcal{I} \to \mathcal{I}$ and $\phi_1 : \mathcal{I} \to \mathcal{I}$, where $\mathcal{I}$ is an interval of $\mathbb{R}$, are increasing, contractive and have unique fixed points y_0 and y_1 on $\mathcal{I}$ which satisfy $y_1 < y_0$.

Not all the maps-with-gaps considered in this paper satisfy Assumption A2 directly. For instance, the function $\phi_0(x) = x + 1$, which corresponds to Kalman-filter update of equation (3) in the case of an uninformative observation (for $A = 1, \Sigma_X = 1$ and $\Sigma_Y(0) \to \infty$), is not contractive. This can be addressed by a change of coordinates, and is discussed in the remark following Theorem 12.

Before using Assumption A2 to characterise the itineraries of maps-with-gaps, we must first clarify the notion of fixed points. A *fixed point* for a finite word w, is a solution to the equation $x = \phi_w(x)$. If Assumption A2 holds, then Lemma 32 in Appendix A.1 shows that there is a unique such fixed point on $\mathcal{I}$ for any finite non-empty word w and we shall denote it by y_w.

In general, it is not clear what a "fixed point" corresponding to an *infinite* word w might mean. One approach might be to consider a sequence $(w^{(n)})_{n=1}^{\infty}$ of words with $w = \lim_{n\to\infty} w^{(n)}$ and to define "y_w" as $\lim_{n\to\infty} y_{w^{(n)}}$ if that limit exists. However, for any word b, the sequence with elements $w^{(n)}b$ also converges to w and it is not hard to find examples where

$$\lim_{n\to\infty} y_{w^{(n)}} \neq \lim_{n\to\infty} y_{w^{(n)}b}.$$

Therefore we shall only define fixed points for a particular class of infinite words, as follows. Let $0s$ be the Sturmian $\mathcal{M}$-word of rate α. Consider the sequence of Christoffel words $0w^{(n)}1$ that lie on the following path through the Christoffel tree. We start from the root,

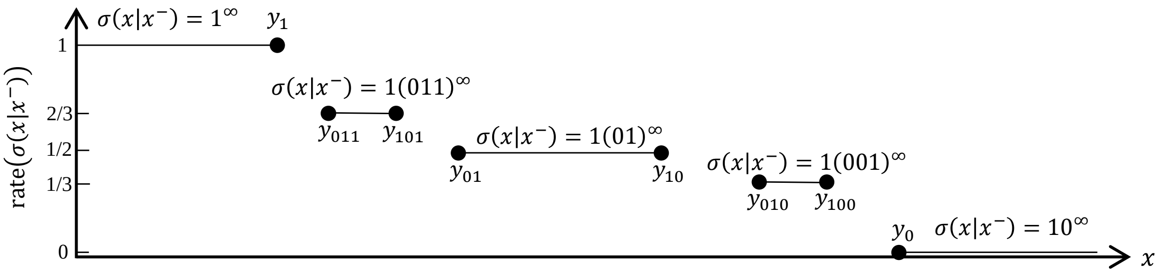

Figure 5: The itinerary $\sigma(x|x^-)$ of Theorem 12. The filled circles in this plot emphasise that the corresponding intervals are closed.

so that $w^{(1)} = \epsilon$. Then for $n \in \mathbb{Z}_{++}$, we set $0w^{(n+1)}1$ equal to the left child of $0w^{(n)}1$ if the rate of $0w^{(n)}1$ exceeds α and equal to the right child otherwise. We call

$$y_s := \lim_{n\to\infty} y_{01w^{(n)}} = \lim_{n\to\infty} y_{10w^{(n)}}$$

the *fixed point* of the Sturmian $\mathcal{M}$-word $0s$. The fact that these limits exist and are equal is proved as Lemma 55 in Appendix A.3.

We are now ready to fully characterise the itineraries of maps-with-gaps, or equivalently to describe the action sequences resulting from applying threshold policies, under Assumption A2. We do so in three cases. Theorem 12 considers the itinerary $\sigma(x|x^-)$ of an active-at-threshold map with initial state equal to the threshold. Corollary 13 considers the pair of itineraries for a passive-at-threshold map with initial states $\phi_0(x)$ or $\phi_1(x)$ for threshold x. Finally, Theorem 14 characterises the itinerary $\sigma(x|s)$ in the general case where the initial state x is unrelated to the threshold s.

Let us begin with an intuitive description of our characterisation of the itineraries $\sigma(x|x^-)$ of active-at-threshold maps with initial state x equal to the threshold. This description is illustrated by Figure 5. Firstly, all such itineraries begin with the letter or action 1, since the initial state equals the threshold. For low values of x, at or below the fixed point y_1 of the map $\phi_1(\cdot)$, taking action 1 leaves the state at or above the threshold, so the itinerary repeats 1 for ever. Similarly, for large values of x, at or above the fixed point y_0 of the map $\phi_0(\cdot)$, after taking action 1, the orbit never goes as high as its initial state again, so the itinerary repeats 0 for ever. Meanwhile, at intermediate values of x, there is an interval $[y_{01}, y_{10}]$ for which the itinerary repeats the Christoffel word 01 for ever. Given that $y_{10} > y_{01}$, such a ping-pong behaviour is very intuitive for x in the set $\{y_{01}, y_{10}\}$, since $\phi_0(y_{01}) = y_{10}$ and $\phi_1(y_{10}) = y_{01}$ by definition of the fixed points.

For other values of x, the itinerary is more complex. In fact, for any Christoffel word, there is an interval of positive length for which the itinerary simply repeats that word over and over again. See for instance the interval $[y_{011}, y_{101}]$ in Figure 5 on which the itinerary repeats the Christoffel word 011, or the interval $[y_{010}, y_{100}]$ on which it repeats the word 001. Now, between these intervals of positive length, there are individual *points* at which the itinerary equals any given Sturmian $\mathcal{M}$-word. Such points are not shown in Figure 5 but lie somewhere in the gaps between the intervals that are shown. Furthermore, as x

increases through the interval $\mathcal{I}$, the itinerary goes through all the $\mathcal{M}$-words in order of decreasing rate, and it never takes a value that is not given by an $\mathcal{M}$-word.

In the following theorem, let $\{0,1\}^\infty$ denote the set of all infinite binary strings and let $\{0,1\}^+$ denote the set of all finite non-empty binary strings. Also, the *image* of a function is the subset of the function's range consisting of those values that the function takes for some point in the function's domain. Although in general an itinerary $\sigma(x|x^-)$ is in $\{0,1\}^\infty$, the following theorem shows that, under Assumption A2, the itinerary comes from a very specific subset of such binary strings that is generated by the set of $\mathcal{M}$-words.

Theorem 12 *Suppose A2 holds,* $0p1$ *is a Christoffel word and* $0s$ *is a Sturmian* $\mathcal{M}$*-word. Then the fixed points* y_{01p}, y_{10p}, y_s *exist in* $\mathcal{I}$*. Also, the itinerary* $\sigma(x|x^-)$ *is a lexicographically non-increasing function of* $x \in \mathcal{I}$ *and is of the form* $\sigma(x|x^-) = 1\ell(x)^\infty$ *for some mapping* $\ell : \mathcal{I} \to \{0,1\}^\infty \cup \{0,1\}^+$ *whose image is the set of* $\mathcal{M}$*-words. Specifically,*

$$\sigma(x|x^-) = \begin{cases} 1^\infty & \textit{if and only if } x \le y_1 \\ (10p)^\infty & \textit{if and only if } x \in [y_{01p}, y_{10p}] \\ 10s & \textit{if and only if } x = y_s \\ 10^\infty & \textit{if and only if } x \ge y_0. \end{cases}$$

See Appendix A.4 for a proof. This result is previously known for *linear* maps-with-gaps (Rajpathak et al., 2012), although those authors do not draw any relation to mechanical words. Dance and Silander (2015) previously extended those authors' proof to the nonlinear case under Assumption A2. The proof presented in Appendix A of this paper can be seen as a simplification of that extension. On the other hand, it is known that itineraries of a broader class of nonlinear maps-with-gaps that do not necessarily satisfy Assumption A2 also correspond to mechanical words (Kozyakin, 2003). However such generality comes at a cost, as it is not clear in that work which range of thresholds gives rise to which words.

Remark. Not all the maps-with-gaps considered in this paper satisfy Assumption A2. However, this does not always prevent the application of Theorem 12. Notably for $\mathcal{I} := [0,\infty)$ and $\theta \in (0,\infty)$, the pair

$$\phi_0(x) = x + 1, \qquad \phi_1(x) = 1/(\theta + 1/(x+1))$$

involves the non-contractive map ϕ_0. Nevertheless, after the change of coordinates

$$g : x \mapsto x/(x+1),$$

the transformed functions

$$\tilde{\phi}_0(x) := g(\phi_0(g^{(-1)}(x))) = 1/(2-x), \qquad \tilde{\phi}_1(x) := g(\phi_1(g^{(-1)}(x))) = 1/(2+\theta-x)$$

and the interval $\tilde{\mathcal{I}} := [0,1]$ do satisfy Assumption A2. Indeed

$$\frac{d\tilde{\phi}_1(x)}{dx} = 1/(2+\theta-x)^2 \in (0,1]$$

for $x \in \tilde{\mathcal{I}}$ and $\theta \in [0, \infty)$, and this derivative only equals 1 for the endpoint $x = 1$. Thus $\tilde{\phi}_1$ is increasing and contractive on $\tilde{\mathcal{I}}$. Noting that $\tilde{\phi}_0(x) = \lim_{\theta \to 0} \tilde{\phi}_1(x)$, the same holds for $\tilde{\phi}_0$. Also $\tilde{\phi}_1$ has a fixed point at $y(\theta) = (2 + \theta - \sqrt{\theta^2 + 4\theta})/2$ which lies in $\tilde{\mathcal{I}}$ for $\theta \in [0, \infty)$, and $\tilde{\phi}_0$ has a fixed point at $y(0) = 1 > y(\theta)$. As g is an increasing function, all conclusions of Theorem 12 still hold for the original functions ϕ_0, ϕ_1.

Pairs of Itineraries. As we are particularly interested in expressions like the Whittle index in Theorem 1, it is important to have a result characterising itineraries starting from $\phi_0(x)$ and $\phi_1(x)$, as in the following Corollary of Theorem 12, whose proof is given as Appendix A.5.

Corollary 13 *Suppose A2 holds,* $0p1$ *is a Christoffel word and* $0s$ *is a Sturmian* $\mathcal{M}$*-word. Then the pair of itineraries* $(\sigma(\phi_0(x)|x), \sigma(\phi_1(x)|x))$ *is given by*

$$\begin{cases} (1^\infty, 1^\infty) & \text{if } x < y_1 \\ (1^\infty, 01^\infty) & \text{if } x = y_1 \\ ((1p0)^\infty, (0p1)^\infty) & \text{if } x \in [y_{01p}, y_{10p}) \\ ((1p0)^\infty, 0p(01p)^\infty) & \text{if } x = y_{10p} \\ (1s, 0s) & \text{if } x = y_s \\ (0^\infty, 0^\infty) & \text{if } x \geq y_0. \end{cases}$$

Looking ahead, applying this corollary to the problem addressed in Theorem 1 shows that the interval $\mathcal{I}$ can be partitioned into intervals corresponding to Christoffel words on which the marginal resource $g(x, x)$ of Section 3 is a constant function of x and points corresponding to Sturmian $\mathcal{M}$-words. In particular, this partition of the interval $\mathcal{I}$ is central to our analysis of the the marginal productivity index $m^*(x)$ in Section 3.3.

Discontinuities of the Itinerary. We conclude this section by considering the number of discontinuities of the mapping from the threshold to the length-n prefix of the itinerary, for any given positive integer n. We show that this number of discontinuities can be upper-bounded by a polynomial function of n. Now for any finite word w and state $x \in \mathcal{I}$, the map $s \mapsto \mathbf{1}_{s < \phi_w(x)}$ is right continuous. Also, if Assumption A2 holds, then the map ϕ_w is continuous (Lemma 32), and it then follows from Definition 5 that the itinerary satisfies $\sigma(x|s)_{1:n} = \sigma(x|s^+)_{1:n}$ for any $n \in \mathbb{Z}_+$, and that the limit $\sigma(x|s^-)_{1:n}$ exists. Thus, in the following theorem, a *discontinuity* of the mapping $s \mapsto \sigma(x|s)_{1:n}$ is a point $d \in \mathcal{I}$ at which its left and right limits disagree, so that $\sigma(x|d^-)_{1:n} \neq \sigma(x|d^+)_{1:n}$. Also, in the following theorem, we say that a finite word w is a *factor of a lower mechanical word* if there exists a slope $\alpha \in [0, 1]$ and an intercept $\rho \in \mathbb{R}$ such that $w_n = \lfloor \alpha(n+1) + \rho \rfloor - \lfloor \alpha n + \rho \rfloor$ for $n = 1, 2, \ldots, |w|$.

Theorem 14 *Suppose* ϕ_0, ϕ_1 *satisfy A2, that* $n \in \mathbb{Z}_{++}$ *and* $x \in \mathcal{I}$*. Then* $\sigma(x|s)$ *is a lexicographically non-increasing function of* $s \in \mathcal{I}$*. Also, for any fixed* $x, s \in \mathcal{I}$*, we have*

$$\sigma(x|s)_{1:n} = l^m w$$

for some $l \in \{0, 1\}$*, some* $m \in \{0, 1, \ldots, n\}$*, and some factor* w *of a lower mechanical word. Furthermore, for any* $x \in \mathcal{I}$*, the mapping* $s \mapsto \sigma(x|s)_{1:n}$ *for* $s \in \mathcal{I}$ *has at most a polynomial number* $p(n)$ *of discontinuities.*

This result is an immediate consequence of work by Kozyakin (2003) and Mignosi (1991). Its proof is given as Appendix A.6. Looking ahead, this result is important for changing the order of certain summations in the proof of Proposition 30, which shows that the Radon-Nikodym condition (PCLI3) of Section 3 holds.

3. Proof of Main Result

First we present a theorem due to Niño-Mora (Niño-Mora, 2015, 2019) which guarantees the indexability and optimality of threshold policies for projects satisfying three so-called partial conservation law indexability (PCLI) conditions (Section 3.1). The following three subsections show that these three conditions in turn are satisfied for projects satisfying Assumption A1, by making intensive use of our characterisation of the itineraries of maps-with-a-gap from Section 2. Then we couple the results together to prove our main result, Theorem 1, concerning threshold policies for observing time-series and the associated Whittle indices (Section 3.5). Finally, we prove Corollary 2 concerning the structure of optimal policies for the linear-quadratic Gaussian problem with costly observations (Section 3.6).

3.1. Verification Theorem

We present a theorem which guarantees the optimality of threshold policies for two-action Markov decision problems called *projects* under certain hypotheses. Under the same hypotheses, the theorem also guarantees the indexability of such projects. This theorem is due to Niño-Mora (2015) and extends previous work on countable state spaces (Niño-Mora, 2001, 2002, 2006) to problems where the state space is an interval of the real line. Niño-Mora calls the theorem's hypotheses the *partial conservation law indexability (PCLI)* conditions. This terminology was chosen to contrast with the *strong* conservation law conditions of Shanthikumar and Yao (1992) and the *generalised* conservation law conditions of Bertsimas and Niño-Mora (1996), which have their roots in *conservation laws* for queueing systems under which waiting times are invariant to the queueing discipline (Kleinrock, 1965). Although, at the time of writing, Niño-Mora (2015) is not published in a peer-reviewed publication, we have developed our own proofs of a version of the verification theorem and associated results for the special case addressed in this paper in Dance and Silander (2017).

Projects and λ-Price Problems. A *project* $\mathcal{P} = \langle \mathcal{X}, c, r, \mathcal{Q}, \beta \rangle$ describes a stochastic process with state space $\mathcal{X}$, resource function $c : \mathcal{X} \times \{0, 1\} \to \mathbb{R}$, reward function $r : \mathcal{X} \times \{0, 1\} \to \mathbb{R}$, transition law $\mathcal{Q}$ and discount factor $\beta \in [0, 1)$, with the same interpretation as on page 10. Let Π denote the class of all history-dependent randomized policies. Let $\mathbb{E}_x^\pi$ denote the expectation over sequences $(X_t, A_t)_{t=0}^\infty$ with initial state $X_0 = x$, where actions are taken according to policy π and transitions are according to transition law $\mathcal{Q}$. Then the *λ-price problem* for project $\mathcal{P}$ and price $\lambda \in \mathbb{R}$ is to is to find a policy $\pi \in \Pi$ that minimises

$$V_\lambda(x, \pi) := \mathbb{E}_x^\pi \left[\sum_{t=0}^\infty \beta^t \left(r(X_t, A_t) - \lambda c(X_t, A_t) \right) \right]$$

for all initial states $x \in \mathcal{X}$. Such a policy is said to be *optimal* for the λ-price problem. We say action $a \in \{0,1\}$ is *optimal in state* x if there is a policy $\pi^* \in \Pi$ such that $V(x, \langle a, \pi^* \rangle) = \sup_{\pi \in \Pi} V(x, \pi)$, where $\langle a, \pi \rangle$ denotes the policy which first takes action a then follows policy π. Niño-Mora (2015) makes the following assumption about projects.

Assumption A3

(i) The state space is $\mathcal{X} = \{x \in \mathbb{R} : l \le x \le u\}$ *for some endpoints* $-\infty \le l < u \le \infty$.

(ii) For each action $a \in \{0,1\}$*, the reward function* $r(\cdot, a)$ *and resource function* $c(\cdot, a)$ *are continuous on* $\mathcal{X}$*, and* $c(\cdot, a)$ *satisfies* $0 \le c(x,0) < c(x,1)$ *for all* $x \in \mathcal{X}$.

(iii) There exists a measurable weight function $w : \mathcal{X} \to [1, \infty)$ *and constants* $M > 0$ *and* $\gamma \in [\beta, 1)$*, such that for every state* $x \in \mathcal{X}$ *and action* $a \in \{0,1\}$*,*

(a) $\max\{|C(x,a)|, |c(x,a)|\} \le Mw(x)$*; and*

(b) $\beta \int_{\mathcal{X}} w(y) \mathcal{Q}(dy|x,a) \le \gamma w(x)$.

In part (i) of this assumption, the state space $\mathcal{X}$ is of one of the forms $(-\infty, \infty)$, $(-\infty, u']$, $[l', \infty)$ or $[l', u']$ for some $l', u' \in \mathbb{R}$. Although this does not include the state space $\mathcal{I} = (0, \infty)$ in Assumption A1, we show that the result still holds for $\mathcal{X} = (0, \infty)$ in the proof of Theorem 1. The assumption allows c and r to be unbounded functions of the state, however part (iii) ensures that the weighted supremum approach to discounted problems (Wessels, 1977; Hernández-Lerma and Lasserre, 1999) can be applied. This assumption guarantees that discounted sums involving c or r are well-defined for all policies. Furthermore, although Bellman's equations may have multiple unbounded solutions, this assumption identifies the solution corresponding to the optimal value function $V^*(\cdot)$ as the unique solution for which the ratio $|V^*(x)/w(x)|$ is bounded on $\mathcal{X}$.

Optimality of Threshold Policies. For a threshold $s \in \bar{\mathbb{R}}$ and assuming a state-space consisting of real numbers, let π_s denote the *s-threshold policy* that takes actions $A_t = \mathbf{1}_{X_t > s}$ and let π_{s^-} denote the s^-*-threshold policy* that takes actions $A_t = \mathbf{1}_{X_t \ge s}$. As on page 11, given a policy π for a project $\mathcal{P}$ in initial state x, we define the reward-to-go or *reward metric* $F(\cdot,\cdot)$ and the resource-to-go or *resource metric* $G(\cdot,\cdot)$ by

$$F(x,\pi) := \mathbb{E}_x^\pi \left[\sum_{t=0}^{\infty} \beta^t r(X_t, A_t) \right], \qquad G(x,\pi) := \mathbb{E}_x^\pi \left[\sum_{t=0}^{\infty} \beta^t c(X_t, A_t) \right].$$

Motivated by expression (8) for the Whittle index, we define the *marginal reward* $f : \mathcal{X} \times \bar{\mathbb{R}} \to \mathbb{R}$ and the *marginal resource* $g : \mathcal{X} \times \bar{\mathbb{R}} \to \mathbb{R}$ by

$$f(x,s) := F(x, \langle 1, \pi_s \rangle) - F(x, \langle 0, \pi_s \rangle), \qquad g(x,s) := G(x, \langle 1, \pi_s \rangle) - G(x, \langle 0, \pi_s \rangle).$$

Finally, we define the *marginal productivity index* $m^*(\cdot)$ to be the ratio of marginal reward to marginal resource when the threshold equals the state, so that

$$m^*(x) := \frac{f(x,x)}{g(x,x)}.$$

Note that the marginal productivity index has the same form as the right-hand side of (8) under the assumption that the x-threshold policy is optimal in the λ-price problem. Niño-Mora (2015) makes the following assumption about the behaviour of the above metrics.

Assumption A4 (Partial Conservation Law Indexability Conditions)

PCLI1. *The marginal resource $g(x,s)$ is positive for every state $x \in \mathcal{X}$ and threshold $s \in \bar{\mathbb{R}}$.*

PCLI2. *The marginal productivity index $m^*(\cdot)$ is non-decreasing, continuous and bounded from below on $\mathcal{X}$.*

PCLI3. *The marginal productivity index $m^*(\cdot)$ and metrics $F(x,\cdot)$ and $G(x,\cdot)$ satisfy*

$$F(x,\pi_{s_2}) - F(x,\pi_{s_1}) = \int_{(s_1,s_2]} m^*(s) G(x,\pi_{ds})$$

for each pair of thresholds $-\infty < s_1 < s_2 < \infty$ and for each state $x \in \mathcal{X}$.

In this assumption, PCLI1 ensures that the marginal productivity index is well defined in the sense that its denominator does not vanish. To clarify the notation in PCLI3, for fixed $x \in \mathcal{X}$, let $\tilde{F}(s) := F(x,\pi_s)$ and $\tilde{G}(s) := G(x,\pi_s)$ for $s \in \bar{\mathbb{R}}$. Then PCLI3 requires that $\tilde{F}(s_2) - \tilde{F}(s_1) = \int_{(s_1,s_2]} m^*(s)\tilde{G}(ds)$ where the right-hand side is a Lebesgue-Stieltjes integral (Carter and van Brunt, 2000). Thus, PCLI3 requires that marginal productivity index is a Radon-Nikodym derivative of the reward metric for an s-threshold policy with respect to the resource metric for an s-threshold policy, so that $m^* = \frac{d\tilde{F}}{d\tilde{G}}$. In analyses of discrete-state projects (Niño-Mora, 2001, 2002, 2006), PCLI3 is not required as it is implied by PCLI1 and PCLI2. Niño-Mora (2015, Appendix B) gives sufficient conditions for PCLI3 to hold, but it is not clear that those are satisfied for the problem addressed in this paper.

We now state the verification theorem, which is central to our proof of Theorem 1.

Theorem 15 (Niño-Mora, 2015, Theorem 2.1, p. 13) *Suppose project $\mathcal{P}$ satisfies A3 and A4. Then project $\mathcal{P}$ is indexable with Whittle index equal to its marginal productivity index $m^*(\cdot)$. Furthermore, for every $\lambda \in \mathbb{R}$, there exists an $s \in \bar{\mathbb{R}}$ such that both the s-threshold and s^--threshold policies are optimal for the λ-price problem for project $\mathcal{P}$.*

Remark. For the project in Theorem 1 we have $c(x,a) = a$ and $r(x,a) = -C(x)$. Let us give some expressions for the metrics F, G, f and g for that project that we shall use in the rest of this section. Firstly, let $(X_t(x,a;s))_{t=0}^{\infty}$ be the sequence of states starting from $X_0(x,a;s) = x$ then taking action $a \in \{0,1\}$ to give $X_1(x,a;s) = \phi_a(x)$ then acting according to the s-threshold policy, so that the action sequence is $A_0(x,a;s) = a$ and $A_t(x,a;s) = \mathbf{1}_{X_t(x,a;s)>s}$ for $t > 0$. Let $X_t(x,a;s^-)$ and $A_t(x,a;s^-)$ be the state and actions from initial state x and initial action a when the s^--threshold policy is followed thereafter, so that $A_t(x,a;s) = \mathbf{1}_{X_t(x,a;s)\geq s}$ for $t > 0$. Also, let $X_t(x;s) := X_t(x,\mathbf{1}_{x>s};s)$, $A_t(x;s) = A_t(x,\mathbf{1}_{x>s};s)$, $A_t(x,a;s^-) = A_t(x,\mathbf{1}_{x\geq s};s^-)$ for $t \in \mathbb{Z}_+$. Then

$$f(x,s) = \sum_{t=0}^{\infty} \beta^t \left(C(X_t(x,0;s)) - C(X_t(x,1;s))\right), \; g(x,s) = \sum_{t=0}^{\infty} \beta^t \left(A_t(x,1;s) - A_t(x,0;s)\right).$$

Given the notation for itineraries (Definition 5) and for compositions of maps ϕ_0, ϕ_1 we can write more-explicit expressions for the reward and cost metric for that project under a threshold policy, for instance

$$F(x, \langle a, \pi_s \rangle) = \sum_{n=1}^{\infty} \beta^{n-1} C(\phi_{(a\sigma(\phi_0(x)|s))_{1:n}}(x)).$$

Furthermore, in the light of Corollary 13, we can find expressions for the marginal reward and marginal resource in terms of $\mathcal{M}$-words. For instance if $0p1$ is a Christoffel word of length N and $x \in [y_{01p}, y_{10p})$ then

$$g(x,x) = \sum_{n=1}^{\infty} \beta^{n-1} \left[((10p)^\infty)_n - ((01p)^\infty)_n\right] = \sum_{k=0}^{\infty} \beta^{kN}(1-\beta) = \frac{1-\beta}{1-\beta^N} \tag{12}$$

since $((10p)^\infty)_n$ and $((01p)^\infty)_n$ only differ when n is 1 or 2 modulo N, whereas for $x = y_{10p}$ we have

$$g(x,x) = \sum_{n=1}^{\infty} \beta^{n-1} \left[(10p(01p)^\infty)_n - ((01p)^\infty)_n\right] = 1-\beta.$$

In spite of this discontinuity in $g(x,x)$, Lemma 19 shows that the marginal productivity index does not have a discontinuity at points such as y_{10p}. Similarly, if $0s$ is a Sturmian $\mathcal{M}$-word and $x = y_s$, then the same argument gives $g(x,x) = 1-\beta$.

3.2. Positivity of Marginal Resource Metric (PCLI1)

We prove that condition PCLI1 holds. The argument is based on the following notion of *swapping*, which is partly inspired by results about the Burrows-Wheeler transform of Christoffel words (Berstel et al., 2008, Chapter 6).

Definition 16 *A finite word a **swaps to** a finite word b if either $a = b$ or there exist words $p_1, q_1, p_2, q_2, \ldots, p_n, q_n$ for some $n \in \mathbb{Z}_{++}$ with*

$$a = p_1 10 q_1, \qquad p_1 01 q_1 = p_2 10 q_2, \qquad \ldots, \qquad p_n 01 q_n = b.$$

*We call a transformation $p_k 10 q_k \to p_k 01 q_k$ an **exchange**.*

Example. The word 1100 swaps to the word 0101 via the exchanges

$$1100 \to 1010 \to 0110 \to 0101$$

for which $p_1 = 1$, $q_1 = 0$, $p_2 = \epsilon$, $q_2 = 10$, $p_3 = 01$ and $q_3 = \epsilon$.

The idea of our proof is as follows. First we find conditions on the number of 1's in prefixes of two words that make it possible to swap one word for another (Lemma 17). Proposition 18 shows that those conditions are satisfied by the dynamical system (statement 1), and they imply positivity of the marginal resource metric (statement 2).

Lemma 17 *Suppose a, b are words of common length $|a| = |b| = n \in \mathbb{Z}_+$ with*

$$|a|_1 = |b|_1 \quad and \quad |a_{1:k}|_1 \geq |b_{1:k}|_1 \quad for\ k < n.$$

Then a swaps to b.

Proof Given any words u, v of length n, consider the distance

$$d(u,v) := \sum_{i=1}^{n} \left| |u_{1:i}|_1 - |v_{1:i}|_1 \right|.$$

If $a = b$ then a swaps to b after $d(a,b) = 0$ exchanges. Otherwise $a \neq b$ and we shall show that there exists a word a' such that

$$a \text{ and } a' \text{ differ by a single exchange,} \tag{13}$$

$$d(a', b) = d(a,b) - 1, \tag{14}$$

$$\text{and } a' \text{ and } b \text{ satisfy the hypotheses of this Lemma.} \tag{15}$$

Repeating this argument shows that a swaps to b after $d(a,b)$ exchanges.

We now define an appropriate word a'. As $a \neq b$ and $|a_{1:i}|_1 \geq |b_{1:i}|_1$ for $i = 1, 2, \ldots, n-1$, there must exist a first index i such that $|a_{1:i}|_1 > |b_{1:i}|_1$. Also, as $|a|_1 = |b|_1$, there must exist a first index $j > i$ such that $a_j = 0$. As i, j are the first such indices, it follows that $a_k = 1$ for $i \leq k < j$. Thus

$$a = a_{1:(j-2)} 10 a_{(j+1):n}$$

with the convention that $a_{1:0} = a_{(n+1):n} = \epsilon$. Now consider the word

$$a' := a_{1:(j-2)} 01 a_{(j+1):n}.$$

It is immediate that (13) holds. Furthermore, as $a_k = 1$ for $i \leq k < j$, we have

$$\left|a_{1:(j-1)}\right|_1 - \left|b_{1:(j-1)}\right|_1 \geq |a_{1:i}|_1 - |b_{1:i}|_1 > 0$$

where the second inequality follows from the definition of i. Thus the definition of a' gives

$$\left|a'_{1:l}\right|_1 = |a_{1:l}|_1 - \mathbf{1}_{l=j-1} \geq |b_{1:l}|_1 \quad \text{for } l = 1, 2, \ldots, n, \tag{16}$$

so that

$$d(a', b) = \sum_{i=1}^{n} \left(\left|a'_{1:i}\right|_1 - |b_{1:i}|_1 \right) = d(a,b) - 1.$$

Therefore (14) holds.

Finally, combining (16) with the fact that $|a'|_1 = |a|_1 = |b|_1$, we conclude that (15) holds. This completes the proof. ■

The next proposition shows that PCLI1 holds if $\langle \mathcal{I}, C, \phi_0, \phi_1, \beta \rangle$ satisfy Assumption A1. To see that the hypotheses of this proposition follow from A1, note that A1 gives

$$\frac{d\phi_a(x)}{dx} = \frac{r^2}{(\theta_a(r^2x+1)+1)^2} > 0$$

for $a \in \{0,1\}$. Also, putting $u := r^2x+1$, the inequality

$$\phi_{10}(x) = \frac{u\theta_1 \quad + \quad (r^2u+1)}{(r^2u+1)\theta_0 + u\theta_1 + (u\theta_0\theta_1+1)} > \frac{u\theta_0 \quad + \quad (r^2u+1)}{(r^2u+1)\theta_1 + u\theta_0 + (u\theta_0\theta_1+1)} = \phi_{01}(x)$$

follows directly from the mediant inequality, as $\theta_1 > \theta_0$.

Proposition 18 *Suppose $\mathcal{I}$ is an interval of $\mathbb{R}$ and that $\phi_0 : \mathcal{I} \to \mathcal{I}, \phi_1 : \mathcal{I} \to \mathcal{I}$ satisfy*

(i) $\phi_0(\cdot), \phi_1(\cdot)$ are increasing functions

(ii) $\phi_{01}(z) < \phi_{10}(z)$ for all $z \in \mathcal{I}$.

Also suppose that $x \in \mathcal{I}$, $s \in \overline{\mathbb{R}}$ and consider the itineraries

$$a := 1\sigma(\phi_1(x)|s) \quad \textit{and} \quad b := 0\sigma(\phi_0(x)|s).$$

Then the following statements hold:

1. *For any $n \in \mathbb{Z}_{++}$, we have $|a_{1:n}|_1 \geq |b_{1:n}|_1$.*
2. *For any $\beta \in (0,1)$, the marginal resource metric $g(x,s)$ is positive.*

Proof We prove statement 1 by induction. In the base case $|a_1|_1 = 1 \geq 0 = |b_1|_1$. For the inductive step, suppose $|a_{1:k}|_1 \geq |b_{1:k}|_1$ for all $k \leq m$ for some $m \in \mathbb{Z}_{++}$. This induction hypothesis shows that either $|a_{1:m}|_1 > |b_{1:m}|_1$ or $|a_{1:m}|_1 = |b_{1:m}|_1$. In the first case, $\left|a_{1:(m+1)}\right|_1 \geq \left|b_{1:(m+1)}\right|_1$ as we are only adding one letter to $a_{1:m}$ and $b_{1:m}$. In the second case, the induction hypothesis shows that the words $a_{1:m}$ and $b_{1:m}$ satisfy the assumptions of Lemma 17, so there is a sequence of swaps that transforms $a_{1:m}$ into $b_{1:m}$. Consider any swap $p10q$ to $p01q$ in this sequence. Then hypothesis (ii) gives $\phi_{10}(\phi_p(x)) > \phi_{01}(\phi_p(x))$ and hypothesis (i) implies that $\phi_q(\cdot)$ is increasing, so

$$\phi_{p10q}(x) > \phi_{p01q}(x).$$

Repeating this argument over the sequence of swaps gives

$$\phi_{a_{1:m}}(x) > \phi_{b_{1:m}}(x).$$

Thus, it follows from the definition of itineraries that the last letters of $a_{1:(m+1)}, b_{1:(m+1)}$ have $a_{m+1}b_{m+1} \in \{00, 10, 11\}$. Hence $\left|a_{1:(m+1)}\right|_1 \geq \left|b_{1:(m+1)}\right|_1$. This proves statement 1.

To prove statement 2, note that the definition of $g(x,s)$ gives

$$g(x,s) = \sum_{k=1}^{\infty} \beta^{k-1}(a_k - b_k)$$

$$
\begin{aligned}
&= \sum_{k=1}^{\infty} \beta^{k-1}(|a_{1:k}|_1 - |a_{1:(k-1)}|_1 - |b_{1:k}|_1 + |b_{1:(k-1)}|_1) \\
&= (1-\beta) \sum_{k=1}^{\infty} \beta^{k-1}(|a_{1:k}|_1 - |b_{1:k}|_1) \\
&\geq 1 - \beta
\end{aligned}
$$

where the second line follows as $a_k = |a_{1:k}|_1 - |a_{1:(k-1)}|_1$ and similarly for b_k, the third line holds because $|a_{1:0}|_1 = |b_{1:0}|_1 = 0$, and the last line follows from statement 1 and the fact that $a_1 = 1, b_1 = 0$. As $\beta < 1$, this completes the proof. ■

3.3. Behaviour of the Marginal Productivity Index (PCLI2)

Condition PCLI2 requires that the marginal productivity index $m^*(x) = f(x,x)/g(x,x)$ is continuous, non-decreasing and bounded from below for $x \in \mathcal{I}$. We begin by showing that this index is *Lipschitz* continuous on any finite interval of $\mathcal{I}$ in Proposition 21. To show the index is non-decreasing, we couple Lipschitz continuity with a piecewise analysis of the index, driven by Theorem 12, which shows that the interval $\mathcal{I}$ can be divided up into intervals corresponding to Christoffel words and points corresponding to Sturmian $\mathcal{M}$-words. The central result of this analysis is Lemma 25, which shows that the marginal reward with threshold equal to initial state $x \mapsto f(x,x)$ is non-decreasing for x in the interval corresponding to any given Christoffel word of the form $0p1$. Since the marginal resource metric $g(x,x)$ equals a positive constant on such intervals, given by (12), this implies that $m^*(x)$ is non-decreasing on such intervals. Since $m^*(\cdot)$ is also Lipschitz continuous on any finite interval, Proposition 27 concludes that $m^*(\cdot)$ is non-decreasing on $\mathcal{I}$. We note that Proposition 27 would not follow if we were only to have shown that $m^*(\cdot)$ is continuous (rather than Lipschitz continuous). Indeed, there are functions that are continuous and have a positive derivative almost everywhere, but are not non-decreasing. One example is the function $x \mapsto x - \text{Cantor}(x)$ where $\text{Cantor}(\cdot)$ is the Cantor-Lebesgue function (Royden and Fitzpatrick, 2010, page 50). Finally, Proposition 28 shows that the index is bounded below, so PCLI2 is satisfied.

A related proof was given by Dance and Silander (2015). However, that proof only covers systems for which the multiplier r in Assumption A1 is $r = 1$, rather than multipliers $r \in (0,1]$ as addressed here. For $r = 0$, we have $\phi_0(x) = \phi_1(x) = 1$ for all $x \in \mathcal{I}$, so the sum in Lemma 25 vanishes, and the analysis presented here is unnecessary. Also, the proof of Dance and Silander (2015) only addressed the cost function $C(x) = x$, whereas here we generalise to any cost function satisfying Assumption A1, which includes any cost function of the form x^q/q for $q \in [-1, \infty)$. A counterexample in Figure 7, shows that the marginal reward is not necessarily increasing on intervals corresponding to Christoffel words for $C(x) = x^q/q$ with $q < -1$.

3.3.1. The Marginal Productivity Index is Continuous

We demonstrate Proposition 21 which shows that the marginal productivity index m^* for systems satisfying Assumption A1 is Lipschitz continuous on any finite interval. We begin by showing that m^* is continuous.

Lemma 19 *If A1 holds then the marginal productivity index $m^*(\cdot)$ is continuous on $\mathcal{I}$.*

Proof Consider the itineraries $h_a(x) := a\sigma(\phi_a(x)|x)$ for $a \in \{0,1\}$ and $x \in \mathcal{I}$. In terms of these itineraries the marginal productivity index is

$$m^*(x) = \frac{\sum_{n=1}^{\infty} \beta^{n-1} \left(C(\phi_{[h_0(x)]_{1:n}}(x)) - C(\phi_{[h_1(x)]_{1:n}}(x)) \right)}{\sum_{n=1}^{\infty} \beta^{n-1} \left([h_1(x)]_n - [h_0(x)]_n \right)}.$$

Since A1 holds, this is a continuous function of x at any point where $h_a(x^-) = h_a(x) = h_a(x^+)$. By following paths in the Christoffel tree to find the limiting word just lexicographically higher or lower than a given Christoffel word, and using Corollary 13, we obtain the following limits:

- $h_0(x^-) = 01p(10p)^\infty$ if $0p1$ is a Christoffel word and $x = y_{01p}$; $h_0(x^-) = 010^\infty$ if $x = y_0$; and $h_0(x^-) = h_0(x)$ otherwise.
- $h_1(x^-) = 1^\infty$ if $x = y_1$; $h_1(x^-) = (10p)^\infty$ if $0p1$ is a Christoffel word and $x = y_{10p}$; and $h_1(x^-) = h_1(x)$ otherwise.
- $h_a(x^+) = h_a(x)$ for all $x \in \mathcal{I}$ and all $a \in \{0,1\}$.

In view of these limits, it remains to show that $m^*(x^-) = m^*(x)$ when: (i) $x = y_{01p}$ for some Christoffel word $0p1$; (ii) $x = y_{10p}$ for some Christoffel word $0p1$; (iii) $x = y_0$; and (iv) $x = y_1$. We only consider (i) as the proofs for (ii)-(iv) are similar.

Say $x = y_{01p}$ for a Christoffel word $0p1$. The itineraries above and Corollary 13 give

$$h_0(y_{01p}^-) = 01p(10p)^\infty, \quad h_0(y_{01p}) = (01p)^\infty, \quad h_1(y_{01p}^-) = (10p)^\infty, \quad h_1(y_{01p}) = (10p)^\infty.$$

So the marginal resource satisfies

$$\lim_{x \uparrow y_{01p}} g(x,x) = \sum_{n=1}^{\infty} \beta^{n-1} \left([(10p)^\infty]_n - [01p(10p)^\infty]_n \right) = 1 - \beta, \tag{17}$$

$$g(y_{01p}, y_{01p}) = \sum_{n=1}^{\infty} \beta^{n-1} \left([(10p)^\infty]_n - [(01p)^\infty]_n \right) = \frac{1-\beta}{1-\beta^{|0p1|}}. \tag{18}$$

Consider the sum $S(w,x) := \sum_{n=1}^{|w|} \beta^{n-1} C(\phi_{w_{1:n}}(x))$ for any (possibly infinite) word w. Then the marginal reward $f(\cdot,\cdot)$ satisfies

$$\lim_{x \uparrow y_{01p}} f(x,x) = S(01p(10p)^\infty, y_{01p}) - S((10p)^\infty, y_{01p}) \tag{19}$$

$$= S(01p, y_{01p}) + \beta^{|01p|} S((10p)^\infty, \phi_{01p}(y_{01p})) - S((10p)^\infty, y_{01p})$$

$$= S(01p, y_{01p}) - (1 - \beta^{|0p1|}) S((10p)^\infty, y_{01p}) \tag{20}$$

$$= (1 - \beta^{|0p1|}) S((01p)^\infty, y_{01p}) - (1 - \beta^{|0p1|}) S((10p)^\infty, y_{01p}) \tag{21}$$
$$= (1 - \beta^{|0p1|}) f(y_{01p}, y_{01p}). \tag{22}$$

In more detail, (19) follows from the fact that for any word w, we have $S(w, x^-) = S(w, x)$ for any $x \in \mathcal{I}$ as A1 holds; (20) holds as $\phi_{01p}(y_{01p}) = y_{01p}$; (21) follows from the expression for the sum of a geometric progression.

Combining (17), (18) and (22), shows the marginal productivity index $m^*(\cdot)$ satisfies

$$m^*(y_{01p}^-) = \lim_{x \uparrow y_{01p}} \frac{f(x,x)}{g(x,x)} = \frac{(1 - \beta^{|0p1|}) f(y_{01p}, y_{01p})}{1 - \beta} = \frac{f(y_{01p}, y_{01p})}{g(y_{01p}, y_{01p})} = m^*(y_{01p}).$$

This completes the proof. ∎

In order to show that the marginal productivity index $m(\cdot)$ is in fact Lipschitz continuous, Proposition 21 must consider how much $m(\cdot)$ changes over any finite interval $[a, b]$. The analysis proceeds by cutting up any such interval into pieces on each of which the length-k prefixes of the itineraries $[\sigma(\phi_a(u)|u)]_{1:k}$ for $a \in \{0, 1\}$ are constant. The following lemma bounds how much $m(\cdot)$ can change over any one such interval.

Lemma 20 *Suppose A1 holds, that $k \in \mathbb{Z}_{++}$ and $0 \le x < y$ with $\sigma(\phi_a(x)|x)_{1:k} = \sigma(\phi_a(y)|y)_{1:k}$ for $a \in \{0, 1\}$. Let $K := \sup_{z \in [\phi_1(x), \phi_0(y)]} C'(z)$ where $C'(\cdot)$ is the derivative of $C(\cdot)$. Then*

$$|m^*(x) - m^*(y)| \le \frac{K(3\beta^k (y+1) + 2(y - x))}{(1-\beta)^2}.$$

Proof Let $1p := [\sigma(\phi_0(x)|x)]_{1:k}$ and let s, s' be the shortest words such that $0\sigma_0(\phi_0(x)|x) = (01ps)^\infty$ and $0\sigma_0(\phi_0(y)|y) = (10ps')^\infty$. Also let

$$a_n := C(\phi_{((01ps)^\infty)_{1:n}}(x)) - C(\phi_{((10ps)^\infty)_{1:k}}(x)), \qquad e := (1 - \beta^{|01ps|})/(1 - \beta),$$
$$b_n := C(\phi_{((01ps')^\infty)_{1:n}}(y)) - C(\phi_{((10ps')^\infty)_{1:k}}(y)), \qquad f := (1 - \beta^{|01ps'|})/(1 - \beta).$$

Then the simple bounds

$$\sup_{m \ge 1} |a_m| \le K(y+1), \qquad \sup_{m \ge 1} |a_m - b_m| \le 2K(y+1), \qquad \sup_{m \le k} |a_m - b_m| \le 2K(y - x)$$

follow from the facts that: (i) the lowest and highest points on the z-threshold orbit are $\phi_1(z)$ and $\phi_0(z)$ (by Lemma 39); (ii) $y \ge x$; (iii) function $C(\cdot)$ is non-decreasing and differentiable; (iv) function $\phi_w(\cdot)$ is non-expansive for any word w; (v) for any $z \in \mathbb{R}_+$, $\phi_0(z) \le z + 1$ and $\phi_1(z) > 0$.

Also, as $|0p| = k$ it follows that $|e - f| \le \beta^k/(1 - \beta)$. Therefore

$$\beta |m^*(x) - m^*(y)|$$
$$= \left| (e - f) \sum_{n=1}^{\infty} \beta^n a_n + f \sum_{n=1}^{k} \beta^n (a_n - b_n) + f \sum_{n=k+1}^{\infty} \beta^n (a_n - b_n) \right|$$

$$\leq |e - f| \sum_{n=1}^{\infty} \beta^n \sup_{m \geq 1} |a_m| + f \sum_{n=1}^{k} \beta^n \sup_{m \leq k} |a_m - b_m| + f \sum_{n=k+1}^{\infty} \beta^n \sup_{m > k} |a_m - b_m|$$
$$\leq \frac{\beta^k}{1-\beta} \frac{\beta}{1-\beta} K(y+1) + \frac{1}{1-\beta} \frac{\beta}{1-\beta} 2K(y-x) + \frac{1}{1-\beta} \frac{\beta^{k+1}}{1-\beta} 2K(y+1)$$

which rearranges to the inequality claimed. ■

Proposition 21 *Suppose A1 holds. Then the marginal productivity index m^* is Lipschitz continuous on any finite interval $[x, y]$ of $\mathcal{I}$.*

Proof Let $k \in \mathbb{Z}_{++}$. Now, the length-k prefixes of the itineraries $[\sigma(\phi_a(u)|u)]_{1:k}$ for $a \in \{0, 1\}$ are piecewise constant functions of $u \in \mathcal{I}$. Say there are n_k pieces as u increases from x to y and let $[a, b] \subseteq [x, y]$. Then Lemmas 19 and 20 together show that

$$|m^*(b) - m^*(a)| \leq \frac{3K(y+1)}{(1-\beta)^2} n_k \beta^k + \frac{2K}{(1-\beta)^2}(b-a)$$

where

$$K := \sup_{z \in [\phi_1(x), \phi_0(y)]} C'(z) < \infty.$$

However, the total number of pieces of f_k is at most a polynomial function $p(k)$ of k, by Theorem 14. As k was an arbitrary positive integer, it follows that

$$|m^*(b) - m^*(a)| \leq \frac{3K(y+1)}{(1-\beta)^2} \underbrace{\lim_{k \to \infty} \left(p(k) \beta^k \right)}_{=0} + \frac{2K}{(1-\beta)^2}(y-x).$$

As $[a, b] \subseteq [x, y]$ was arbitrary, we conclude that m^* is Lipschitz continuous on $[x, y]$. ■

3.3.2. The Marginal Productivity Index is Non-Decreasing

We use the following well-known result about *majorisation* (Marshall et al., 2010). A proof is given in Appendix B.

Lemma 22 *Suppose that:*

1. *The sequences $a_{1:n}$ and $b_{1:n}$ are non-decreasing sequences on $\mathbb{R}_{++}$*
2. *The inequality $\sum_{i=1}^{k} a_i \leq \sum_{i=1}^{k} b_i$ holds for $k = 1, 2, \ldots, n$*
3. *For $i = 1, 2, \ldots, n$, the function $f_i : \mathbb{R}_{++} \to \mathbb{R}$ is non-increasing and convex*
4. *For $i = 2, 3, \ldots, n$, the difference $f_{i-1}(x) - f_i(x)$ is non-increasing for $x \in \mathbb{R}_{++}$.*

Then

$$\sum_{i=1}^{n} f_i(a_i) \geq \sum_{i=1}^{n} f_i(b_i).$$

We apply this majorisation result to the sequences appearing in Lemma 24 below. To state that Lemma, we first define some matrices that are motivated by the form of the Kalman Filter variance updates.

Definition 23 *Let I be the 2-by-2 identity matrix. For $r \in (0,1]$ and $0 \leq a \leq b$, let*

$$F := \begin{pmatrix} r & 1/r \\ ar & (a+1)/r \end{pmatrix}, \qquad G := \begin{pmatrix} r & 1/r \\ br & (b+1)/r \end{pmatrix}.$$

Let $M(\epsilon) = I, M(0) = F, M(1) = G$ and for any finite non-empty word w let

$$M(w) = M(w_{|w|}) \cdots M(w_2) M(w_1).$$

Thus, if $a = \theta_0, b = \theta_1$ and ϕ_0, ϕ_1 are as in Assumption A1, we have

$$\phi_0(x) = \frac{F_{11}x + F_{12}}{F_{21}x + F_{22}}, \qquad \phi_1(x) = \frac{G_{11}x + G_{12}}{G_{21}x + G_{22}}.$$

As remarked in Section 3, the central portion of Christoffel words are palindromes. The following result holds for any palindromes, not just palindromes generated by Christoffel words.

Lemma 24 *Suppose p is a palindrome, $r \in (0,1]$, $n \in \mathbb{Z}_+$ and x satisfies*

$$\phi_p(0) \leq x \leq \phi_p\left(\frac{1}{1-r^2}\right).$$

Let $m := |01p|$ and for $k = 1, 2, \ldots, m$ let

$$\begin{pmatrix} a_k(x) \\ c_k(x) \end{pmatrix} := M\Big((01p)^n (01p)_{1:k}\Big) \begin{pmatrix} x \\ 1 \end{pmatrix}, \qquad \begin{pmatrix} b_k(x) \\ d_k(x) \end{pmatrix} := M\Big((10p)^n (10p)_{1:k}\Big) \begin{pmatrix} x \\ 1 \end{pmatrix}.$$

Then

1. *The sequences $a_{1:m}(x), b_{1:m}(x), c_{1:m}(x)$ and $d_{1:m}(x)$ are non-decreasing and positive*
2. *The inequality $a_k(x) \leq b_k(x)$ holds for $k = 1, 2, \ldots, m$*
3. *The inequality $\sum_{i=1}^{k} c_i(x) \leq \sum_{i=1}^{k} d_i(x)$ holds for $k = 1, 2, \ldots, m$*
4. *The inequalities $c_1(x) \leq d_1(x)$ and $c_k(x) \geq d_k(x)$ hold for $k = 2, 3, \ldots, m$.*
5. *The fixed points y_{01p} and y_{10p} satisfy*

$$\phi_p(0) \leq y_{01p} < y_{10p} \leq \phi_p\left(\frac{1}{1-r^2}\right)$$

so Claims 1-4 hold if $x \in [y_{01p}, y_{10p}]$.

The proof of the above lemma is given in Appendix C.

Now we are ready to demonstrate that marginal reward is non-decreasing for a wide range of cost functions on the intervals $[y_{01p}, y_{10p}]$ corresponding to a Christoffel word $0p1$ appearing in Theorem 12.

Lemma 25 *Suppose that A1 holds, that $0p1$ is a Christoffel word and that $N \in \mathbb{Z}_+$. Then*

$$\sum_{k=1}^{n} \beta^k (C(\phi_{(01p)^N(01p)_{1:k}}(x)) - C(\phi_{(10p)^N(10p)_{1:k}}(x)))$$

is a non-decreasing function of x for $y_{01p} \le x \le y_{10p}$, where $n = |0p1|$.

We prove the result for functions satisfying parts C1 and C2 of Assumption A1 separately, as sums of non-decreasing functions are non-decreasing, the result holds for any function satisfying Assumption A1.

Proof [Proof when $C(\cdot)$ satisfies part C1 of Assumption A1 (iv).] Let $a_k(x), b_k(x), c_k(x), d_k(x)$ be as defined in Lemma 24. Then the proposition is proved by the following inequalities (as justified immediately below):

$$\begin{aligned}
&\frac{d}{dx}\sum_{k=1}^{n} \beta^k (C(\phi_{(01p)^N(01p)_{1:k}}(x)) - C(\phi_{(10p)^N(10p)_{1:k}}(x))) \\
&\quad = \sum_{k=1}^{n} \left(\frac{\beta^k}{c_k(x)^2} C'\left(\frac{a_k(x)}{c_k(x)}\right) - \frac{\beta^k}{d_k(x)^2} C'\left(\frac{b_k(x)}{d_k(x)}\right)\right) && (23) \\
&\quad \ge \sum_{k=1}^{n} \left(\frac{\beta^k}{c_k(x)^2} C'\left(\frac{b_k(x)}{c_k(x)}\right) - \frac{\beta^k}{d_k(x)^2} C'\left(\frac{b_k(x)}{d_k(x)}\right)\right) && (24) \\
&\quad = \sum_{k=1}^{n} (f_k(c_k(x)) - f_k(d_k(x))) \quad \text{where } f_k(u) := \beta^k C'(b_k(x)/u)/u^2 \\
&\quad \ge 0. && (25)
\end{aligned}$$

Step (23) follows from the chain rule as the homomorphism of matrix multiplication and composition of Möbius transformations gives

$$\phi_{(01p)^n(01p)_{1:k}}(x) = \frac{\left[M((01p)^N(01p)_{1:k})\begin{pmatrix} x \\ 1 \end{pmatrix}\right]_1}{\left[M((01p)^N(01p)_{1:k})\begin{pmatrix} x \\ 1 \end{pmatrix}\right]_2} = \frac{a_k(x)}{c_k(x)}$$

while the fact that matrices F and G have unit determinant implies that the matrix $M((01p)^N(01p)_{1:k})$ also has unit determinant so that

$$\frac{d}{dx}\phi_{(01p)^N(01p)_{1:k}}(x) = \frac{1}{c_k(x)^2}.$$

Step (24) follows as $\beta^k, c_k(x) > 0$, as $C(\cdot)$ is concave and as $a_k(x) \le b_k(x)$ by Lemma 24.

Step (25) follows from Lemma 22. In particular, Lemma 24 shows that the sequences $c_{1:n}(x)$ and $d_{1:n}(x)$ satisfy the hypotheses of Lemma 22. Also, the fact that $C(\cdot)$ satisfies part C1 shows that the functions $f_i(\cdot)$ for $i = 1, \ldots, n$ satisfy hypotheses 1 and 2 of Lemma 22. Indeed, $f_i(\cdot)$ is non-increasing and convex as $\frac{1}{u^2} C'\left(\frac{1}{u}\right)$ is non-increasing and convex and $\beta^i, b_i(x) > 0$. Also, as $\frac{1}{u^3} C''\left(\frac{b}{u}\right)$ is non-decreasing in u for $b > 0$ and $0 < b_{i-1}(x) \le b_i(x)$ for $i = 2, \ldots, n$, by Claim 1 of Lemma 24, the following integral is also non-decreasing in u:

$$\int_{b_{i-1}(x)}^{b_i(x)} \frac{1}{u^3} C''\left(\frac{b}{u}\right) db = \frac{1}{u^2} C'\left(\frac{b_i(x)}{u}\right) - \frac{1}{u^2} C'\left(\frac{b_{i-1}(x)}{u}\right) = \frac{1}{\beta^i}\left(f_i(u) - \beta f_{i-1}(u)\right)$$

So $f_{i-1}(u) - f_i(u)$ is the sum of the non-increasing functions $\beta f_{i-1}(u) - f_i(u)$ and $(1 - \beta) f_{i-1}(u)$.

This completes the proof. ■

Proof [Proof when $C(\cdot)$ satisfies part C2 of Assumption A1 (iv).] For $k \in \mathbb{Z}_n$ let

$$a'_k := \beta^{k+1} \frac{d}{dx} \phi_{(01p)^N (01p)_{1:(k+1)}}(x), \qquad b'_k := \beta^{k+1} \frac{d}{dx} \phi_{(10p)^N (10p)_{1:(k+1)}}(x),$$
$$a_k := C'(y_{(1p0)_{(k+1):n} (1p0)_{1:k}}), \qquad b_k := C'(y_{(0p1)_{(k+1):n} (0p1)_{1:k}}).$$

Let $r_{[0]} \ge \cdots \ge r_{[n-1]}$ denote real numbers $r_0, \ldots, r_{n-1}$ in non-increasing numerical order. Let $\overline{x}$ denote x modulo n and let l satisfy $\overline{l|0p1|_1} = 1$. Then the proposition is proved by the following inequalities (as justified immediately below):

$$\frac{d}{dx} \sum_{k=1}^{n} \beta^k \Big(C\left(\phi_{(01p)^N (01p)_{1:k}}(x)\right) - C\left(\phi_{(10p)^N (10p)_{1:k}}(x)\right)\Big)$$

$$= \sum_{k=1}^{n} \Big(C'\left(\phi_{(01p)^N (01p)_{1:k}}(x)\right) a'_{k-1} - C'\left(\phi_{(10p)^N (10p)_{1:k}}(x)\right) b'_{k-1}\Big) \tag{26}$$

$$\ge \sum_{k=1}^{n} \Big(C'\left(y_{(1p0)_{k:n} (1p0)_{1:(k-1)}}\right) a'_{k-1} - C'\left(y_{(0p1)_{k:n} (0p1)_{1:(k-1)}}\right) b'_{k-1}\Big) \tag{27}$$

$$= \sum_{k=0}^{n-1} \left(a_k a'_k - b_k b'_k\right) \tag{28}$$

$$= \sum_{k=0}^{n-1} a_{[k]} \left(a'_{\overline{l(n-k)}} - b'_{\overline{l(n-k-1)}}\right) \tag{29}$$

$$\ge \sum_{k=0}^{n-1} a_{[k]} \left(a'_{\overline{l(n-k)}} - b'_{\overline{l(n-k)}}\right) \tag{30}$$

$$\ge \sum_{k=0}^{n-1} a_{[0]} \left(a'_{\overline{l(n-k)}} - b'_{\overline{l(n-k)}}\right) \tag{31}$$

$$\ge 0 \tag{32}$$

Step (26) follows from the chain rule and definition of a'_k, b'_k.

Step (27) follows as $C(\cdot)$ is convex, as $a'_k, b'_k \geq 0$, and as, for $k = 1, 2, \ldots, n$, we have

$$\begin{aligned}\phi_{(01p)^N(01p)_{1:k}}(x) &= \phi_{(01p)_{1:k}}(\phi_{(01p)^N}(x)) \\ &\geq \phi_{(01p)_{1:k}}(y_{01p}) \\ &= y_{(01p)_{(k+1):n}(01p)_{1:k}} \\ &= y_{(1p0)_{k:n}(1p0)_{1:(k-1)}}\end{aligned}$$

where the inequality holds as $x \geq y_{01p}$ so that $\phi_{(01p)^N}(x) \geq y_{01p}$, and as $\phi_{(01p)_{1:k}}(\cdot)$ is increasing. The same argument using $x \leq y_{10p}$ gives an upper bound on $C'(\phi_{(10p)^N(10p)_{1:k}}(x))$.

Step (28) follows by shifting the summation indices and from the definition of a_k, b_k.

Step (29) follows from Lemmas 8 and 11 and the convexity of $C(\cdot)$. Let $w_{[0]} \succeq \cdots \succeq w_{[n-1]}$ denote words $w(0), \ldots, w(n-1)$ in non-increasing lexicographic order and let $c_k := (1p0)_{(k+1):n}(1p0)_{1:k}, d_k := (0p1)_{(k+1):n}(0p1)_{1:k}$ for $k \in \mathbb{Z}_n$. Then Lemma 8 shows that $c_{[i]} = d_{[i]} = c_{\overline{l(n-i)}} = d_{\overline{l(n-i-1)}}$. Thus Lemma 11 gives $y_{c_{[i]}} = y_{d_{[i]}} = y_{c_{\overline{l(n-i)}}} = y_{d_{\overline{l(n-i-1)}}}$. Therefore the convexity of $C(\cdot)$ gives

$$a_{[i]} = b_{[i]} = a_{\overline{l(n-i)}} = b_{\overline{l(n-i-1)}}.$$

Step (30) follows as $C(\cdot)$ is non-decreasing, so that $a_{[i]} \geq 0$, and as $\phi_0(\cdot), \phi_1(\cdot)$ are non-decreasing and non-expansive, so that b'_i is a product of derivatives where each derivative is in $[0, 1]$. Thus $a_{[0]}, \ldots, a_{[n-1]}$ and $b'_0, \ldots, b'_{n-1}$ are non-negative non-increasing sequences. Therefore the rearrangement inequality $a_{[i]}b'_j + a_{[i+1]}b'_0 \leq a_{[i]}b'_0 + a_{[i+1]}b'_j$ holds for all $i \in \mathbb{Z}_{n-1}$ and all $j \in \mathbb{Z}_n$. But $b'_{\overline{l(n-(n-1)-1)}} = b'_0$ so repeated application of the rearrangement inequality gives

$$\sum_{k=0}^{n-1} a_{[k]} b'_{\overline{l(n-k-1)}} \leq \sum_{k=0}^{n-1} a_{[k]} b'_{\overline{l(n-k)}}.$$

Step (31) follows as $a_{[0]} = a_{\overline{ln}} = a_0$, as a_0 is non-negative, and from Claim 4 of Lemma 24, as for the $c_i(x), d_i(x)$ defined in that lemma, we have

$$a'_{i-1} - b'_{i-1} = \frac{\beta^i}{c_i(x)^2} - \frac{\beta^i}{d_i(x)^2} \begin{cases} \geq 0 & \text{for } i = 1 \\ \leq 0 & \text{for } i = 2, 3, \ldots, n. \end{cases}$$

Step (32) follows from this lemma using the function $\tilde{C}(x) = x$ which satisfies part C1 and which was already proved above.

This completes the proof. ■

It is much simpler to show that the marginal reward is non-decreasing when the itinerary is 0^∞ (that is, $x \geq y_0$) or when the itinerary is 01^∞ (that is, $x \leq y_1$). Now the sums in the following lemma correspond to the numerator of the marginal productivity index m^* in these cases, up to a multiple of β. Also, the denominator of m^* in these cases is

$$\sum_{k=1}^{\infty} \beta^{k-1}\left((10^\infty)_k - (0^\infty)_k\right) = 1 - \beta = \sum_{k=1}^{\infty} \beta^{k-1}\left((1^\infty)_k - (01^\infty)_k\right)$$

which is positive. So, it follows from this lemma that m^* is non-decreasing on such intervals.

Lemma 26 *Suppose A1 holds. Then*

$$\sum_{k=1}^{\infty} \beta^k \left(C(\phi_{0^k}(x)) - C(\phi_{10^{k-1}}(x))\right) \qquad \textit{and} \qquad \sum_{k=1}^{\infty} \beta^k \left(C(\phi_{01^{k-1}}(x)) - C(\phi_{1^k}(x))\right)$$

are non-decreasing functions of $x \in \mathcal{I}$.

Proof Consider the first sum. As in Lemma 24, for $k \in \mathbb{Z}_{++}$, we define

$$\begin{pmatrix} a_k(x) & b_k(x) \\ c_k(x) & d_k(x) \end{pmatrix} := \left(M(0^k) \begin{pmatrix} x \\ 1 \end{pmatrix} \quad M(10^{k-1}) \begin{pmatrix} x \\ 1 \end{pmatrix} \right).$$

As $G \geq F$, the entries of F are non-negative and $x \geq 0$, it follows that

$$\begin{pmatrix} b_k(x) - a_k(x) \\ d_k(x) - c_k(x) \end{pmatrix} = F^{k-1}(G - F) \begin{pmatrix} x \\ 1 \end{pmatrix} \geq 0. \tag{33}$$

If C satisfies part C1 of Assumption A1, then for any $k \in \mathbb{Z}_{++}$,

$$\begin{aligned} \frac{d}{dx} \left(C(\phi_{0^k}(x)) - C(\phi_{10^{k-1}}(x))\right) &= \frac{1}{c_k(x)^2} C' \left(\frac{a_k(x)}{c_k(x)}\right) - \frac{1}{d_k(x)^2} C' \left(\frac{b_k(x)}{d_k(x)}\right) \\ &\geq \frac{1}{c_k(x)^2} C' \left(\frac{b_k(x)}{c_k(x)}\right) - \frac{1}{d_k(x)^2} C' \left(\frac{b_k(x)}{d_k(x)}\right) \geq 0. \end{aligned}$$

The first inequality holds as C' is concave and $a_k(x) \leq b_k(x)$ by (33). The second inequality holds as $C'(1/x)/x^2$ is non-increasing and $c_k(x) \leq d_k(x)$ by (33).

If C satisfies part C2 of Assumption A1, then for any $k \in \mathbb{Z}_{++}$,

$$\frac{d}{dx} \left(C(\phi_{0^k}(x)) - C(\phi_{10^{k-1}}(x))\right) = \frac{1}{c_k(x)^2} C'(\phi_{0^k}(x)) - \frac{1}{d_k(x)^2} C'(\phi_{10^{k-1}}(x)) \geq 0.$$

The inequality is justified as follows. As $a \leq b$ in the definition of F, G, we have $\phi_0(x) \geq \phi_1(x)$. As $\phi_{0^{k-1}}(\cdot)$ is an increasing function, it follows that $\phi_{0^k}(x) \geq \phi_{10^{k-1}}(x)$. As C is convex, it follows that $C'(\phi_{0^k}(x)) \geq C'(\phi_{10^{k-1}}(x))$. Furthermore, $c_k(x) \leq d_k(x)$, by (33).

Thus, if C satisfies Assumption A1, the sum

$$\sum_{k=1}^{\infty} \beta^k \left(C(\phi_{0^k}(x)) - C(\phi_{10^{k-1}}(x))\right)$$

is the sum of non-decreasing functions. Therefore this sum is non-decreasing.

The proof for the second sum is similar. This completes the proof. ■

Proposition 27 *Suppose Assumption A1 holds. Then the marginal productivity index m^* is non-decreasing.*

Proof Consider any $a, b \in \mathcal{I}$ with $a < b$. By Proposition 21, the index m^* is Lipschitz continuous on any finite interval of $\mathcal{I}$, such as the interval $[a, b]$. Now, if a function is Lipschitz continuous on a closed interval then it is absolutely continuous on that interval (Royden and Fitzpatrick, 2010, Proposition 7, p. 120). Also, an absolutely continuous function m^* has a derivative $m^{*\prime}$ almost everywhere, which is Lebesgue integrable, and which satisfies (Pugh, 2002, Theorem 38, p. 399)

$$m^*(b) = m^*(a) + \int_a^b m^{*\prime}(t)\ dt. \tag{34}$$

Now, for any Christoffel word $0p1$ and for t in the interval $[y_{01p}, y_{10p}]$, the derivative $m^{*\prime}(t)$ is non-negative, since $m^*(t) = f(t,t)/g(t,t)$, in which $f'(t) \geq 0$ by Lemma 25, and in which $g(t,t)$ is a positive constant depending on $0p1$. Similarly, Lemma 26 shows that the marginal productivity index is non-decreasing on $\mathcal{I} \backslash (y_1, y_0)$. Therefore, $m^*(a) \leq m^*(b)$. This completes the proof. ■

Remark. Having proved that the marginal productivity index is non-decreasing, one might wonder if it is actually always strictly increasing. It is not. In particular, Assumption A1 allows the cost function $C(\cdot)$ to be a constant, in which case the marginal productivity index is a constant. In other cases, the marginal productivity index is not a constant, yet it is constant over a finite interval of the state space $\mathcal{I}$. For instance, this is the case for sufficiently large precision parameter θ_1 that $y_1 < 1$, and for the cost function

$$C(x) = \begin{cases} 0 & x < 1 \\ (x-1)^2 & x \geq 1, \end{cases}$$

which satisfies A1.

3.3.3. The Marginal Productivity Index is Bounded Below

Proposition 28 *Suppose Assumption A1 holds. Then the marginal productivity index m^* is bounded from below.*

Proof Let $[l, u] := \mathcal{I}$. By Proposition 27, m^* is non-decreasing, so

$$\inf_{x \in [l,u]} m^*(x) = f(l,l)/g(l,l).$$

Now $l \leq y_1$ so $\sigma(l|l) = 01^\infty$ and $\sigma(l|l^-) = 1^\infty$. Thus $g(l,l) = 1 - \beta$ and

$$f(l,l) = \sum_{t=0}^{\infty} \beta^t \left(C(\phi_{(01^\infty)_{1:t}}(l)) - C(\phi_{1^\infty_{1:t}}(l)) \right) \geq \frac{C(l) - C(y_1)}{1-\beta}$$

where the inequality follows as $\phi_{(01^\infty)_{1:t}}(l) \geq l$ and $\phi_{(1^\infty)_{1:t}}(l) \leq y_1$ for all t, and as C is increasing. As $C(l) > -\infty$, $C(y_1) < \infty$ and $\beta < 1$, it follows that $m^*(x) > -\infty$ for $x \in \mathcal{I}$. ■

3.4. Radon-Nikodym Condition (PCLI3)

We demonstrate that PCLI3 holds. The argument is based on the fact that the reward metric $F(x, \pi_s)$ and resource metric $G(x, \pi_s)$ for threshold policies are *jump functions* when viewed as a function of the threshold s, because the number of discontinuities of the itinerary map $s \mapsto \sigma(x|s)_{1:n}$ is bounded by a polynomial function of n by Theorem 14. First we recall what a jump function is.

Definition 29 (Kolmogorov and Fomin, 1975, p. 341) *A function $h : \mathbb{R} \to \mathbb{R}$ is a* ***jump function*** *if $h(x) = h_0 + \sum_{n=1}^{\infty} h_n \mathbf{1}_{x_n < x} + \sum_{n=1}^{\infty} h'_n \mathbf{1}_{x'_n \leq x}$ for some real-valued sequences $(h_i)_{i=0}^{\infty}$, $(h'_i)_{i=1}^{\infty}$, $(x_i)_{i=1}^{\infty}$ and $(x'_i)_{i=1}^{\infty}$ satisfying $\sum_{i=1}^{\infty} |h_i| < \infty$ and $\sum_{i=1}^{\infty} |h'_i| < \infty$.*

Proposition 30 *Suppose $\langle \mathcal{I}, C, \phi_0, \phi_1, \beta \rangle$ satisfy Assumption A1. Then PCLI3 holds.*

Proof Throughout the proof, let us consider a fixed initial state $x \in \mathcal{I}$. Recall the notation for the state $X_t(x; s)$ and action $A_t(x; s)$ from initial state x under the s-threshold policy discussed in the remark on page 30. Also recall also that a function $h : \mathbb{R} \to \mathbb{R}$ is *càdlàg* if both limits $h(s^-)$ and $h(s^+)$ exist and $h(s) = h(s^+)$ for all $s \in \mathbb{R}$.

First we use induction to show that $A_t(x; s)$ is a càdlàg function of s for any $t \in \mathbb{Z}_+$. In the base case, we have $X_0(x; s) = x$ and $A_0(x; s) = \mathbf{1}_{x>s}$, which are càdlàg in s. Now say $t \geq 0$ and $X_t(x; s)$ and $A_t(x; s)$ are càdlàg in s. Then, $\phi_a(X_t(x; s))$ is càdlàg in s for $a \in \{0, 1\}$ as it is the composition of a continuous function and a càdlàg function. Thus

$$X_{t+1}(x; s) = A_t(x; s)\phi_1(X_t(x; s)) + (1 - A_t(x; s))\phi_0(X_t(x; s))$$

is càdlàg in s, since sums and products of càdlàg functions are càdlàg. Hence $s - X_{t+1}(x; s)$ is càdlàg in s. Noting that $z \mapsto \mathbf{1}_{z<0}$ is càdlàg, it follows that the composition

$$A_{t+1}(x; s) = \mathbf{1}_{s - X_{t+1}(x;s) < 0}$$

is càdlàg in s, completing the induction.

Also, note that $\lim_{s \to -\infty} A_t(x; s) = 1$, since under Assumption A1, the state is never less than zero, so for any $s < 0$, the s-threshold policy always takes action 1.

For $t \in \mathbb{Z}_+$, let $\mathcal{D}_t(x)$ be a sequence consisting of all the values of the threshold s at which $A_t(x; s)$ is discontinuous. As $A_t(x; s)$ is a càdlàg function of s, taking values in $\{0, 1\}$, with $\lim_{s \to -\infty} A_t(x; s) = 1$, it follows from the definition of $\mathcal{D}_t(x)$ that

$$A_t(x; s) = 1 + \sum_{d \in \mathcal{D}_t(x)} (A_t(x; d) - A_t(x; d^-))\mathbf{1}_{d \leq s},$$

noting that this sum is finite as $\operatorname{card}(\mathcal{D}_t(x)) \leq p(t)$ for some polynomial $p(t)$ by Theorem 14. Now let $d_1, d_2, \ldots$ be the sequence generated by first concatenating $\mathcal{D}_0(x)$ then $\mathcal{D}_1(x)$ then $\mathcal{D}_2(x)$ and so forth, then deduplicating that sequence by including only the first occurrence of each distinct element. Since the sequence $d_1, d_2, \ldots, d_k$ contains all of $\mathcal{D}_t(x)$ for some finite k, and because $A_t(x; d_i) - A_t(x; d_i^-) = 0$ for any d_i in the sequence which is not an element of $\mathcal{D}_t(x)$, it follows that

$$A_t(x; s) = 1 + \sum_{i=1}^{\infty} (A_t(x; d_i) - A_t(x; d_i^-))\mathbf{1}_{d_i \leq s}.$$

Thus the resource metric is

$$G(x,\pi_s) = (1-\beta)^{-1} + \sum_{t=0}^{\infty}\sum_{i=1}^{\infty} \beta^t (A_t(x;d_i) - A_t(x;d_i^-))\mathbf{1}_{d_i \le s}.$$

As $\text{card}(\mathcal{D}_t(x)) \le p(t)$ for some polynomial $p(t)$ by Theorem 14 and $A_t(x;s) \in \{0,1\}$, we have

$$\sum_{t=0}^{\infty}\sum_{i=1}^{\infty} \left|\beta^t (A_t(x;d_i) - A_t(x;d_i^-))\mathbf{1}_{d_i \le s}\right| \le \sum_{t=0}^{\infty} \beta^t p(t) < \infty.$$

So Fubini's theorem gives

$$G(x,\pi_s) = (1-\beta)^{-1} + \sum_{i=1}^{\infty} \left(\sum_{t=0}^{\infty} \beta^t (A_t(x;d_i) - A_t(x;d_i^-)) \right) \mathbf{1}_{d_i \le s}.$$

Thus the mapping $s \mapsto G(x,\pi_s)$ is a càdlàg jump function.

Let $\tau := \inf\{t > 0 : X_t(s;s^-) = s\}$ with $\tau = \infty$ if $X_t(s;s^-) \ne s$ for all $t > 0$. Then

$$\begin{aligned} G(s,\pi_{s^-}) &= \sum_{t=0}^{\infty} \beta^t A_t(s;s^-) \\ &= \sum_{t=0}^{\infty} \beta^t A_t(s,1;s) + \sum_{t=\tau}^{\infty} \beta^t (A_t(s;s^-) - A_t(s,0;s)) \\ &= G(s,\langle 1,\pi_s\rangle) + \beta^\tau (G(s,\pi_{s^-}) - G(s,\langle 0,\pi_s\rangle)) \\ &= \frac{G(s,\langle 1,\pi_s\rangle) - \beta^\tau G(s,\langle 0,\pi_s\rangle)}{1-\beta^\tau}. \end{aligned}$$

Let $\tau_1 := \inf\{t \ge 0 : X_t(x;s) = s\}$ with $\tau_1 = \infty$ if $X_t(x;s) \ne s$ for all $t \ge 0$. Then the expression for $G(x,\pi_{s^-})$ just derived gives

$$\begin{aligned} G(x,\pi_s) - G(x,\pi_{s^-}) &= \beta^{\tau_1} (G(s,\pi_s) - G(s,\pi_{s^-})) \\ &= \beta^{\tau_1} \left(G(s,\langle 0,\pi_s\rangle) - \frac{G(s,\langle 1,\pi_s\rangle) - \beta^\tau G(s,\langle 0,\pi_s\rangle)}{1-\beta^\tau} \right) \\ &= -\frac{\beta^{\tau_1}}{1-\beta^\tau} g(s,s) \end{aligned}$$

using definition of the marginal resource metric $g(\cdot,\cdot)$.

A similar argument shows that the mapping $s \mapsto F(x,\pi_s)$ for the reward metric is also a càdlàg jump function which can be described in terms of the same sequence of discontinuities $d_1, d_2, \ldots$ as $G(x,\pi_s)$. Furthermore, it is related to the marginal reward $f(s,s)$ through

$$F(x,\pi_s) - F(x,\pi_{s^-}) = -\frac{\beta^{\tau_1}}{1-\beta^\tau} f(s,s).$$

Recalling that $m^*(s) = f(s,s)/g(s,s)$ gives

$$F(x,\pi_s) - F(x,\pi_{s^-}) = m^*(s)(G(x,\pi_s) - G(x,\pi_{s^-})).$$

As $F(x,\pi_s), G(x,\pi_s)$ are càdlàg jump functions of s, the Lebesgue-Stieltjes integral of this expression over any interval $(s_1, s_2]$ of $\mathbb{R}$ is

$$F(x,\pi_{s_2}) - F(x,\pi_{s_1}) = \int_{(s_1,s_2]} m^*(s)\ dG(x,\pi_{ds}).$$

Therefore PCLI3 holds. ■

In the above proof, the expressions for the jumps $G(x,\pi_s) - G(x,\pi_{s^-})$ and $F(x,\pi_s) - F(x,\pi_{s^-})$ fit with a general pattern for such jumps given in Lemma 7.5 of Niño-Mora (2015). Also, Section 7.1 of that paper shows that the metrics $F(x,\pi_s)$ and $G(x,\pi_s)$ are càdlàg functions of the threshold s. However, while Appendix B of that paper also presents sufficient conditions for PCLI3 to hold, we were unable to apply those conditions to the model at hand.

3.5. Proof of Theorem 1

Proof First we show that the system considered in Theorem 1 can be seen as a project in the sense of Theorem 15, except that the former system may have a state space of the form $(0,\infty)$ which is not explicitly addressed by the verification theorem (Step 1). So, we then demonstrate that this choice of interval is in fact implicitly addressed by the verification theorem (Step 2). Finally, we recall that the PCLI conditions and A3 hold (Step 3) and explain the characterisation of the optimal thresholds (Step 4).

Step 1. The system given by $\mathcal{I}, C(\cdot), \phi_0(\cdot), \phi_1(\cdot)$ and factor β that is addressed in Theorem 1 is a project $\mathcal{P} = \langle \mathcal{X}, r, c, \mathscr{Q}, \beta\rangle$ with the same discount factor β, with $\mathcal{X} = \mathcal{I}$, $r(x,a) = -C(x)$, $c(x,a) = a$ and a transition law corresponding to deterministic transitions $\mathscr{Q}(\mathcal{B}|x,a) = \mathbf{1}_{\phi_a(x)\in\mathcal{B}}$ for each Borel subset $\mathcal{B}$ of $\mathcal{X}$, for all $x \in \mathcal{X}$ and $a \in \{0,1\}$.

Step 2. Assumption A1 allows the interval $\mathcal{I} = [0,\infty)$, which explicitly matches the assumptions of Theorem 15. The interval $\mathcal{I} = (0,\infty)$ is also allowed by A1. If Assumption A3 and the PCLI conditions are satisfied for a project $\mathcal{P} = \langle \mathcal{X}, r, c, \mathscr{Q}, \beta\rangle$ with state space $\mathcal{X} = (0,\infty)$ then they are also satisfied by the equivalent project $\tilde{\mathcal{P}} = \langle \tilde{\mathcal{X}}, \tilde{r}, \tilde{c}, \tilde{\mathscr{Q}}, \beta\rangle$ after the change of coordinates $\tilde{x} := \log(x)$. Now project $\tilde{\mathcal{P}}$ has state space $\tilde{\mathcal{X}} := (-\infty,\infty)$, so the conclusions of Theorem 15 apply to project $\tilde{\mathcal{P}}$ and therefore to project $\mathcal{P}$.

(To see this, note that continuity of $\tilde{r}(\tilde{x},a) = r(\exp(\tilde{x}),a)$, $\tilde{c}(\tilde{x},a) = c(\exp(\tilde{x}),a)$ follows from continuity of $\exp(\cdot), r(\cdot,a)$ and $c(\cdot,a)$. Also, for any Borel set $\tilde{\mathcal{B}}$ we have $\tilde{\mathscr{Q}}(\tilde{\mathcal{B}}|\tilde{x},a) = \mathscr{Q}(\{\exp(\tilde{y}) : \tilde{y} \in \tilde{\mathcal{B}}\}|\exp(\tilde{x}),a)$ and this is a valid transition law since the mapping $x \mapsto \log(x)$ is one-to-one. Further, if the weight function $w(\cdot)$ satisfies A3 for $\mathcal{P}$ then $\tilde{w}(\tilde{x}) := w(\exp(\tilde{x}))$ satisfies A3 for $\tilde{\mathcal{P}}$. Finally, for a threshold policy $\pi_s(x) = 1_{x>s}$ consider the threshold policy $\tilde{\pi}_{\tilde{s}}(\tilde{x}) := 1_{\tilde{x}>\tilde{s}}$ where $\tilde{s} := \log(s)$. Then the metrics $\tilde{F}(\cdot,\cdot), \tilde{G}(\cdot,\cdot), \tilde{f}(\cdot,\cdot), \tilde{g}(\cdot,\cdot)$ for $\tilde{\mathcal{P}}$ and $F(\cdot,\cdot), G(\cdot,\cdot), f(\cdot,\cdot), g(\cdot,\cdot)$ for $\mathcal{P}$ satisfy

$$\tilde{F}(\tilde{x},\langle a,\tilde{\pi}_{\tilde{s}}\rangle) = F(\exp(\tilde{x}),\langle a,\pi_s\rangle), \qquad \tilde{G}(\tilde{x},\langle a,\tilde{\pi}_{\tilde{s}}\rangle) = G(\exp(\tilde{x}),\langle a,\pi_s\rangle),$$
$$\tilde{f}(\tilde{x},\tilde{s}) = f(\exp(\tilde{x}),\exp(\tilde{s})), \qquad \tilde{g}(\tilde{x},\tilde{s}) = g(\exp(\tilde{x}),\exp(\tilde{s})).$$

As $\exp(\cdot)$ is increasing and continuous, and as the PCLI conditions hold for $\mathcal{P}$, it follows that the PCLI conditions hold for $\tilde{\mathcal{P}}$.)

Step 3. Clearly the assumptions about C and w in A1 imply that A3 holds. Furthermore, the PCLI conditions hold by Propositions 18, 21, 27, 28 and 30. Thus it follows from Theorem 15 that the family of λ-price problems in equation (9) is indexable, that the Whittle index has the form claimed and that a threshold policy is optimal.

Step 4. Say $\lambda^*(s) = \lambda$. If $x < s$ then $\lambda^*(x) \le \lambda$, as $\lambda^*(\cdot)$ is non-decreasing by PCLI2. Yet if $\lambda^*(x) \le \lambda$ then action 0 is optimal, by the definition of indexability (Definition 3). A similar argument shows that action 1 is optimal for $x > s$. Also, indexability shows that both actions 0 and 1 are optimal at s. Therefore policies $\mathbf{1}_{X_t > s}$ and $\mathbf{1}_{X_t \ge s}$ are both optimal.

If $\lambda^*(s) > \lambda$ for all $s \in \mathcal{I}$ then action 0 is never optimal, by definition of indexability. Therefore the always-active policy is the unique optimal policy. The argument when $\lambda^*(s) < \lambda$ for all $s \in \mathcal{I}$ is similar.

This completes the proof. ■

3.6. Proof of Corollary 2

Proof As $\mathbb{E}[Z_t^2 | H_t] = z_t^2 + v_t$, it is not hard to see that the problem corresponds to the dynamic program

$$V(z_t, v_t) = \min_{u_t \in \mathbb{R}, a_t \in \{0,1\}} \left\{ D z_t^2 + D v_t + F u_t^2 + \lambda a_t + \beta \mathbb{E}[V(z_{t+1}, v_{t+1}) | z_t, v_t, a_t, u_t] \right\} \quad (35)$$

where the expectation is over the following Markovian transitions of the information state:

$$\begin{aligned} z_{t+1} | z_t, v_t, a_t, u_t &\sim \mathcal{N}(\bar{A} z_t + \bar{B} u_t, \bar{A}^2 v_t + \Sigma_Z - \Phi_{a_t}(v_t)) \\ v_{t+1} | z_t, v_t, a_t, u_t &= \Phi_{a_t}(v_t). \end{aligned}$$

For trial solutions of the form $V(z, v) = Rz^2 + Rv + g(v)$, where $R \in \mathbb{R}$ and $g : \mathbb{R}_+ \to \mathbb{R}$ are to be determined, the expectation in (35) is

$$\begin{aligned} &\mathbb{E}[V(z_{t+1}, v_{t+1}) | z_t, v_t, a_t, u_t] \\ &\qquad = R((\bar{A} z_{t+1} + \bar{B} u_t)^2 + \bar{A}^2 v_t + \Sigma_Z - \Phi_{a_t}(v_t) + \Phi_{a_t}(v_t)) + g(\Phi_{a_t}(v_t)). \end{aligned}$$

Thus (35) is solved if

$$\begin{aligned} R z_t^2 + R v_t + g(v_t) = \min_{u_t \in \mathbb{R}} &\left\{ D z_t^2 + F u_t^2 + \beta R (\bar{A} z_{t+1} + \bar{B} u_t)^2 \right\} \\ &+ \min_{a_t \in \{0,1\}} \left\{ \lambda a_t + \beta R \Sigma_Z + (D + \beta R \bar{A}^2) v_t + \beta g(\Phi_{a_t}(v_t)) \right\}. \quad (36) \end{aligned}$$

Now the minimum with respect to u_t is achieved if the coefficient $(F + \beta \bar{B}^2 R)$ of u_t^2 is positive, in which case the minimiser is

$$u_t = -\frac{\beta \bar{A} \bar{B} R}{F + \beta \bar{B}^2 R} z_t.$$

So (36) is solved if R satisfies

$$R = D + F\left(\frac{\beta\bar{A}\bar{B}R}{F+\beta\bar{B}^2R}\right)^2 + \beta R\left(\bar{A} - \bar{B}\frac{\beta\bar{A}\bar{B}R}{F+\beta\bar{B}^2R}\right)^2$$

and if $g(\cdot)$ satisfies the dynamic program

$$g(v) = \min_{a\in\{0,1\}}\left\{\lambda a + \beta R\Sigma_Z + \alpha v + \beta g(\Phi_a(v))\right\} \tag{37}$$

where $\alpha := D - (1-\beta\bar{A}^2)R$.

After simple algebra, the condition on R is equivalent to the quadratic equation

$$-\beta\bar{B}^2R^2 + (\beta\bar{B}^2D + \beta\bar{A}^2F - F)R + DF = 0.$$

Using Descartes' rule of signs and considering the cases $F = 0$ and $F > 0$ separately, we see that this equation has a unique positive root for $\beta\bar{B}^2 > 0$ and $D > 0$.

To apply Theorem 1 to the dynamic program for $g(\cdot)$ we must ensure that Assumption A1 holds, which is not hard to see provided that $\alpha \geq 0$. Noting that $m := 1-\beta\bar{A}^2 > 0$ by the hypotheses about $\bar{A}, \beta$, we see that $\alpha \geq 0$ if $R \leq D/m$. But, substituting $R = y+(D/m)$ in the equation for R gives

$$\begin{aligned} 0 &= [-\beta\bar{B}^2R^2 + (\beta\bar{B}^2D + \beta\bar{A}^2F - F)R + DF]_{R=y+(D/m)} \\ &= -\ \beta\bar{B}^2y^2\ -\ ((\beta^2\bar{A}^2+\beta)\bar{B}^2(D/m) + mF)y\ -\ \beta^2\bar{A}^2\bar{B}^2(D/m)^2 \end{aligned}$$

in which the coefficients of y^0, y^1, y^2 are all negative by the hypotheses about $\bar{A}, \bar{B}, D, F$ and β. So this quadratic equation for y has no positive roots and it follows that $\alpha \geq 0$. Therefore Theorem 1 shows that there is an optimal policy for (37) that sets $a_t = 1$ if and only if $v_t \geq s$ for some $s \in \overline{\mathbb{R}}$. This completes the proof. ∎

4. Numerical Experiments

We discuss algorithms for computing the Whittle index given in Theorem 1, we present closed-form expressions for that index and compare the performance of the Whittle index policy with two previously-proposed heuristics.

4.1. Approximating the Index for Discount Factor $\beta \leq 0.999$

Truncating the sums in the Whittle index $\lambda^*(x)$ after a suitably large number of terms T suggests the approximation

$$\hat{\lambda}(x) = \frac{\sum_{t=0}^{T}\beta^t\big(C(X_t(x,0;x)) - C(X_t(x,1;x))\big)}{\sum_{t=0}^{T}\beta^t\big(A_t(x,1;x) - A_t(x,0;x)\big)}. \tag{38}$$

Assuming accurate calculation of the terms in the numerator and denominator, as well as continuity of the cost function C, this approximation requires $O(T)$ basic arithmetical and

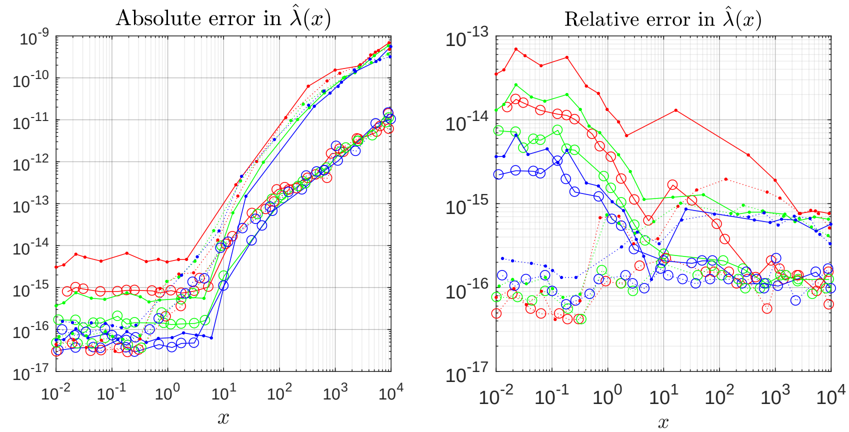

Figure 6: Errors in approximating the index. The cost $C(x)$ corresponds to variance ($C(x) = x$, red), to entropy ($C(x) = \log(x)$, green) or to negative precision ($C(x) = -1/x$, blue). The discount factor β is either 0.9 (open circles) or 0.999 (filled circles). The map-with-a-gap has $\phi_0(x) = x + 1$ and $\phi_1(x) = 1/(\theta_1 + 1/(x + 1))$ with θ_1 equal to 0.001 (solid lines) or 1 (dotted lines), where $\phi_a(\cdot)$ and θ_a for $a \in \{0, 1\}$ are as in Assumption A1.

comparison operations, and setting T to $\Omega(\log(\epsilon)/\log(\beta))$ guarantees absolute errors in the numerator and denominator of $O(\epsilon)$. Of course, the constants hidden by the $O(\cdot)$ or $\Omega(\cdot)$ depend on the detailed properties of C and choice of x.

However, this approximation faces a potential complication. Indeed, some of the iterates $X_t(x, a; x)$ may be so close to the threshold x that an arbitrarily small tolerance is required to correctly decide whether $X_t(x, a; x) \geq x$. This might be problematic as errors in such decisions can result in large changes to the numerator and denominator of this approximation.

Dance and Silander (2015) overcame this problem by constraining the sequence of decisions to correspond to $\mathcal{M}$-words. This resulted in a polynomial-time algorithm for approximating the index $\lambda^*(x)$ using variable-precision arithmetic. Specifically, suppose that basic arithmetic operations to tolerance 2^{-m} on positive numbers less than 2^n require at most $M(n + m)$ operations, and that $C(x) = x$. Then, the approximate index $\hat{\lambda}(x)$ output by their algorithm satisfies $\left|\hat{\lambda}(x) - \lambda^*(x)\right| < \epsilon$ when $x < 2^n$ and $\epsilon > 2^{-m}$ and is computed in $O((n + m)^3 M(n + m))$ operations. Nevertheless, that algorithm requires the tabulation of the fixed points of all Christoffel words of at most a given length.

Here we suggest that such tabulation is an unnecessary expense, and conjecture that standard floating-point approximations of the decision sequences $A_t(x, a; x)$ correspond to $\mathcal{M}$-words (at least for the vast majority of floating point values of x). Perhaps such a conjecture might be proven by extending results by Kozyakin (2003). His results concern mappings ϕ_0, ϕ_1 which are strictly increasing but potentially discontinuous. In the floating-point case one would only require those mappings to be non-decreasing and piecewise constant, but might perhaps impose additional conditions.

Rather than attempting to prove such a conjecture here, we simply evaluate the accuracy of approximation (38). Figure 6 shows the accuracy based on comparing double-precision and quadruple-precision implementations, with $T = \lceil \log 10^{-17} / \log \beta \rceil$ and $T = \lceil \log 10^{-34} / \log \beta \rceil$ respectively. As the difference between these approximations is a highly variable function of x, we only show points that are local maxima of the error. Specifically, we used a logarithmically-spaced grid with $10^{-2} = x_1, x_2, \ldots, x_{1000} = 10^4$ and plot the error $e(x_i)$ only for points with $e(x_i) = \max_{i-20 \le j \le i+20} e(x_j)$. The plot shows no line for x less than the first such point or greater than the last such point.

The worst absolute and relative errors are below 10^{-6} and 10^{-11} respectively. In any practical application, such errors would be swamped by imprecision in the time-series models. The absolute error remains small as x increases to the fixed point y_1 of the mapping ϕ_1, and then it increases due to round-off in computing iterates of the map-with-a-gap for large x. Overall, the worst results are for large discount factors and for variance as the cost function.

Finally, it is possible to substantially accelerate the convergence of the numerator, for instance with Aitken acceleration (Brezinski, 2000), particularly if one has high accuracy requirements. For instance, if the itinerary $\sigma(x|x)$ is periodic with small enough period n, one may accumulate n terms of the sum at a time and apply acceleration methods to such partial sums. Having experimented with such approaches, we find that further work is required in selecting appropriate termination conditions if one is interested in accuracy guarantees for a wide range of problem instances. The difficulty we encountered is that *two* types of linear convergence are going on simultaneously, namely the convergence due to ϕ_0 (when $\theta_0 > 0$ or $r < 1$) and ϕ_1, and the convergence due to β. In such situations, what looks like a healthy stopping time to existing termination criteria can actually be a misleading and unhealthy prematurity.

4.2. Approximating the Index for Discount Factor $\beta \to 1$

For discount factors $\beta > 0.999$ the number of terms T required for accuracy in approximation (38) becomes prohibitively large. In such cases, it makes sense to Taylor expand the numerator of λ as a function of β. For brevity, we only do so here for the case $\beta \to 1$.

Suppose the actions $A_t(x, a; x)$ in the expression in Theorem 1 for the index $\lambda^*(x)$ correspond to a Christoffel word of length n. By (12) the denominator of the index is

$$\lim_{\beta \to 1} \frac{1-\beta}{1-\beta^n} = \frac{1}{n}.$$

Let $X_t^{a,\infty} := \lim_{k\to\infty} X_{kn+t}(x, a; x)$. Now, as $C(\cdot)$ is continuous by Assumption A1, we have

$$\lim_{k\to\infty} C(X_{kn+t}(x, a; x)) = C(\lim_{k\to\infty} C(X_{kn+t}(x, a; x)) = C(X_t^{a,\infty}). \tag{39}$$

Thus, putting $\Delta_t := C(X_t^{1,\infty}) - C(X_t^{0,\infty})$, the numerator of index has the limit

$$\lim_{\beta\to 1} \sum_{t=0}^{\infty} \beta^t \left(C(X_t(x, 0; x)) - C(X_t(x, 1; x)) \right)$$

$$
\begin{aligned}
&= \underbrace{\sum_{t=0}^{\infty} \left(C(X_t(x,0;x)) - C(X_t(x,1;x)) + \Delta_t)\right)}_{=:T_1} - \lim_{\beta \to 1} \sum_{k=0}^{\infty} \beta^{kn} \sum_{t=0}^{n-1} \beta^t \Delta_t \\
&= T_1 - \lim_{\beta \to 1} \frac{\sum_{t=0}^{n-1} \beta^t \Delta_t}{1-\beta^n} \\
&= T_1 + \frac{1}{n} \sum_{t=0}^{n-1} t \Delta_t
\end{aligned}
$$

having applied L'Hôpital's rule. Now (39) shows that the sequences $C(X_t(x,a;x)) - C(X_t^{a,\infty})$ appearing in sum T_1 converge to zero. This suggests the approximation

$$
X_{kn+t}^{a,\infty} \approx X_{Tn-n+t}(x,a;x) \qquad (a \in \{0,1\}, k \in \mathbb{Z}_+, t = 0,1,\ldots,n-1)
$$

for a suitably large positive integer T. This also suggests approximating the first sum by truncating it after Tn terms. While the itinerary for x might have a very large and possibly infinite period n, we did not encounter such situations when tabulating the index. If this were an issue, it is possible to find a good rational approximation to the rate of the x-threshold itinerary from initial state x, for instance as in Dance and Silander (2015). To summarise, given an appropriate choice of n and T, the discussion above suggests approximating the index by

$$
\hat{\lambda}(x) := n \sum_{k=0}^{T-1} \sum_{t=0}^{n-1} \left(\tilde{C}_{kn+t}^0 - \tilde{C}_{Tn-n+t}^0 - \tilde{C}_{kn+t}^1 + \tilde{C}_{Tn-n+t}^1\right) + \sum_{t=0}^{n-1} t \left(\tilde{C}_{Tn-n+t}^1 - \tilde{C}_{Tn-n+t}^0\right)
$$

where $\tilde{C}_k^a := C(X_t(x,a;x))$ for $k \in \mathbb{Z}_+$ and $a \in \{0,1\}$.

4.3. Closed Form Expressions and Graphs

We analyse the behaviour of the index as the cost function $C(\cdot)$ and parameters $\beta, r, \theta_0, \theta_1$ vary.

Given noise free observations for action $a = 1$ and totally uninformative observations for action $a = 0$, it is easy to find a closed form for the index.

Proposition 31 *Suppose the cost function is $C(x) = x$ and the precision $\theta_0 = 0$. Then*

$$
\lim_{\theta_1 \to \infty} \lambda^*(x) = \frac{1-\beta^{n+1}}{1-\beta} \left(rx + 1 - \frac{\beta}{1-\beta^{n+1}} \left(\frac{1-(\beta r)^n}{1-\beta r} - \beta^n \frac{1-r^n}{1-r} \right) \right),
$$

for all $\beta \in [0,1)$, $r \in [0,1)$ and $x \in [0, 1/(1-r))$, where $n := \left\lceil \frac{\log(1-(1-r)x)}{\log r} \right\rceil$. Thus

$$
\lim_{\beta \to 1} \lim_{r \to 1} \lim_{\theta_1 \to \infty} \lambda^*(x) = \lceil x+1 \rceil \left(x + 1 - \frac{\lceil x \rceil}{2} \right) = \int_0^{x+1} \lceil u \rceil \, du, \qquad \text{for all } x \in \mathbb{R}_+.
$$

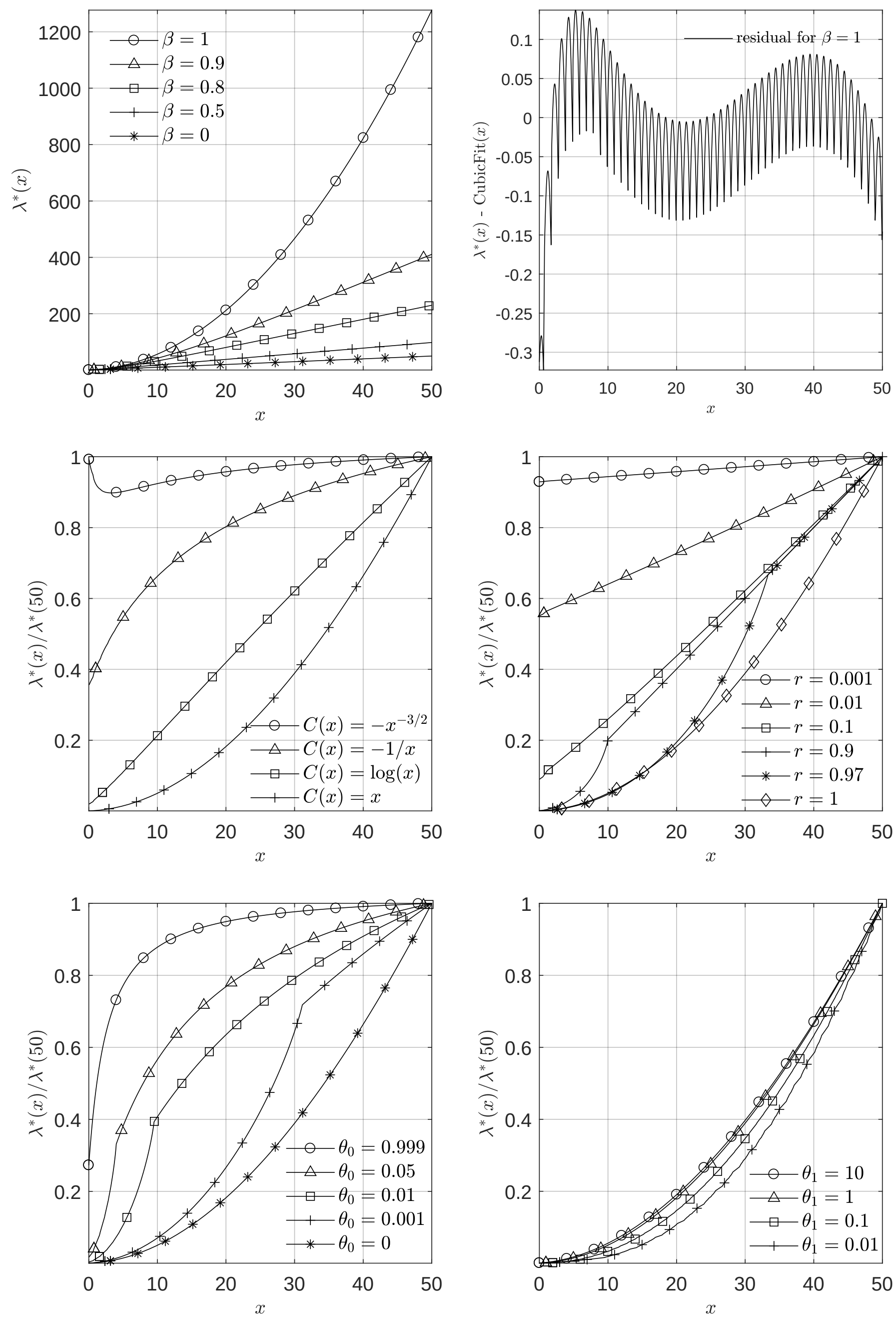

Figure 7: Behaviour of the index. The index as a function of the discount factor β (*top-left*) and the residual after fitting a cubic to this curve in the case $\beta = 1$ (*top-right*). The index, normalised by $\lambda^*(50)$, as the cost function $C(x)$, the multiplier r and the observation precisions θ_0 and θ_1 are varied (*other plots*). In all plots, all parameters (or function) other than that varied are set to $\beta = 0.99$, $C(x) = x$, $\phi_0(x) = 1/(\theta_0+1/(rx+1))$ and $\phi_1(x) = 1/(\theta_1+1/(rx+1))$, with $\theta_0 = 0$, $\theta_1 = 1$ and $r = 1$.

Proof As $\lim_{\theta_1 \to \infty} \phi_1(x) = 0$, the orbits involve the sequence $0, 1, r+1, r^2+r+1, \ldots, r^{n-1} + \cdots + r + 1$, where n is the least integer for which $r^{n-1} + \cdots + r + 1 \geq x$. The result then follows from the definition of the index, using well-known summation formulae. ■

Other closed forms exist, for instance in the limit $\beta \to 0$, or whenever the cost is a polynomial function of x or whenever the process tends to the continuous-time process analysed in Le Ny et al. (2011). However, the usefulness of the above proposition to provide practical approximations in circumstances where observations are not perfectly precise, so that $\theta_1 < \infty$, is limited. For instance, if $C(x) = x, \beta \to 1, r = 1, \theta_0 = 0$ and $\theta_1 = 4$, then we have

$$\lambda^*(0.2) = 1.01\cdots, \qquad \int_0^{1.2} \lceil u \rceil \, du = 1.4$$

which corresponds to a relative error of over 38%.

Therefore, Figure 7 graphs the index using the algorithms of the previous subsection. Looking at these graphs, one notices that the index is increasing in all-but-one of the cases shown: indeed $C(x) = -x^{-3/2}$ is *not* covered by Assumption A1. Also, the index has cusps at the fixed point $x = \phi_0(x)$ which are clearly visible as θ_0 and r vary. Finally, the index becomes increasingly serrated as $\beta \to 1$ and $\theta_1 \to 0$. One would anticipate such serrations on the basis of the above proposition as

$$\int_0^{x+1} \lceil u \rceil \, du \; - \; \frac{1}{2}(x+1)(x+2) = \frac{1}{2}(\lceil x \rceil - x)(x - \lfloor x \rfloor)$$

and they are visible in the plot of the residual after subtracting a cubic fit, for $\theta_1 = 1$ and $\beta \to 1$. However, in general, the serrations have a complex non-periodic pattern and give a slightly ragged appearance to the plot with $\theta_1 = 0.01$.

4.4. Performance Relative to Heuristic Policies

Two heuristic policies have been commonly used for the problem of multi-sensor scheduling. The *myopic* policy observes the m time series with the highest current cost $C_i(x_{i,t})$ and has been used in radar systems (Moran et al., 2008). Meanwhile, the *round robin* policy chooses an ordering of the n projects and observes the next m projects, according to this ordering, at each time.

Figure 8 compares the costs incurred by these heuristics in a simple scenario in which uncertainty about one of the projects is more expensive than for the other projects. In detail, there are $n = 10$ projects, the number of observations per round is $m = 1$, the cost at time t is $40x_{1,t} + \sum_{i=2}^{n} x_{i,t}$, and the initial variance is $x_{i,0} = 4$ for all projects. The discount factor is $\beta = 0.99$ and the "total cost" for a given method is computed as $\sum_{t=0}^{199} \beta^t (40x_{1,t} + \sum_{i=2}^{n} x_{i,t})$. The projects have the same map-with-a-gap given by $\phi_0(x) = x + 1$ and $\phi_1(x) = 1/(0.1 + 1/(x+1))$. The colours in this plot represent the variance state $x_{i,t}$.

Clearly, the myopic policy is overeager to observe project $i = 1$ and does so at the expense of projects $i = 2, \ldots, 10$. In contrast, the round robin policy makes no special effort to observe project $i = 1$ and incurs substantial expense for that project. Meanwhile,

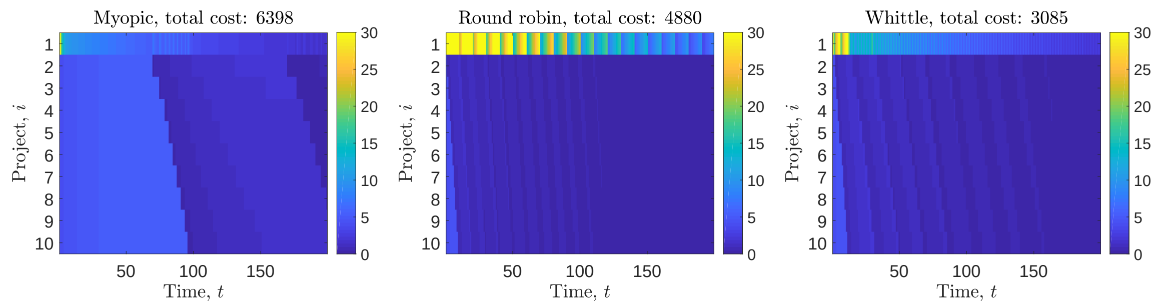

Figure 8: Comparison of heuristic policies. Colour represents the variance state $x_{i,t}$.

the Whittle index policy's behaviour lies between these extremes and it is by far the least costly policy. Indeed the ratio of the total cost of the round robin policy to the total cost of the Whittle index policy is just over 1.5, whereas the ratio of the total cost for the myopic policy to the total cost for the Whittle index policy is just over 2.

5. Further Work

This paper presented conditions under which threshold policies are optimal for observing a single time series with costly observations. It also explored the implications of this result by showing that it leads to optimal policies for the linear-quadratic Gaussian problem with costly observations and that it demonstrates the indexability of related restless bandit problems, which were both long-standing open questions.

It would be natural to extend this work to situations where more than two observation actions are available, perhaps using known generalisations of Sturmian words (Glen and Justin, 2009). Additionally, it would be interesting to complete the work by Le Ny et al. (2011) analysing optimal policies and indexability for scalar-valued time series in continuous time. In particular, there are a number of gaps in their argument to close, as discussed on page 12, and there is as yet no analysis of the discounted case nor of the class of cost function for which threshold policies are optimal. There are also truly-stochastic versions of the one-dimensional problem considered here. These include situations where the costs depend on the posterior mean rather than just the posterior variance (Kuhn et al., 2014), situations where the quality of the observation is a random variable, as well as situations involving non-Gaussian time series. It is also important to understand the structure of optimal policies for making costly observations with discrete-time Kalman filters in *multiple* dimensions. One attack on this problem would begin by extending the verification theorem of Niño-Mora (2015) to multi-dimensional state spaces for a variety of extensions of the notion of a "threshold policy". Finally, we cannot claim the asymptotic optimality of Whittle's index policy for the problem studied here as the results of Verloop (2016) only apply to *countable* state spaces. Furthermore, little is known about the performance of policies for restless bandits in *non-asymptotic* situations involving finite numbers of projects and finite time horizons.

Acknowledgments

We thank José Niño-Mora for pointing out his work on the verification theorem, Sofia Michel for her careful reading of an initial draft of the manuscript, and the reviewers for their tenacity and numerous helpful suggestions. Much of the work described in this paper was done while the authors were with Xerox Research Centre Europe.

Appendix A. Itineraries as Words

We begin with some elementary properties of compositions of functions satisfying Assumption A2 and their fixed points (Section A.1). This enables us to present a proof of Theorem 12 which is based on the Christoffel tree. First we consider the case where itineraries correspond to Christoffel words (Section A.2), then the case where itineraries correspond to Sturmian $\mathcal{M}$-words (Section A.3) and finally we couple these results together to prove the theorem (Subsection A.4). We also present proofs of Corollary 13 concerning x-threshold itineraries from initial state $\phi_0(x)$ and $\phi_1(x)$ (Section A.5) and of Theorem 14 about itineraries from general initial points not equal to the threshold (Section A.6) .

A.1. Compositions and Fixed Points

The following fundamental result about compositions is well known.

Lemma 32 *Suppose A2 holds and that w is a finite non-empty word. Then ϕ_w is increasing, contractive, continuous and has a unique fixed point y_w on $\mathcal{I}$.*

Proof First we show that $\phi_w(x)$ is increasing, by induction on the length of word w. In the base case, as w is non-empty, we suppose that $|w| = 1$. The claim then follows immediately from A2. For the inductive step, assume $\phi_u(x)$ is increasing, where $w = au$ for some letter $a \in \{0, 1\}$ and some finite non-empty word u. Then for any $x, y \in \mathcal{I}$ such that $x < y$,

$$\begin{aligned} \phi_w(y) &= \phi_u(\phi_a(y)) \\ &> \phi_u(\phi_a(x)) && \text{as } \phi_a(y) > \phi_a(x) \text{ and } \phi_u \text{ is increasing} \\ &= \phi_w(x). \end{aligned}$$

Therefore ϕ_w is increasing.

Now we show that $\phi_w(x)$ is contractive, again by induction on $|w|$. If $|w| = 1$ then this follows immediately from A2. Else, say $\phi_u(x)$ is contractive where $w = ua$ and $a \in \{0, 1\}$. Then for any $x, y \in \mathcal{I}$ such that $x < y$,

$$\begin{aligned} \phi_w(y) - \phi_w(x) &= \phi_a(\phi_u(y)) - \phi_a(\phi_u(x)) \\ &< \phi_u(y) - \phi_u(x) && \text{as } \phi_u(y) > \phi_u(x) \text{ and } \phi_a \text{ is contractive} \\ &< y - x && \text{as } \phi_u \text{ is contractive.} \end{aligned}$$

Therefore ϕ_w is contractive.

As ϕ_w is contractive, for any $\epsilon > 0$ and $c \in \mathcal{I}$, it follows that

$$|\phi_w(x) - \phi_w(c)| < |x - c| < \epsilon \quad \text{for any } x \in \mathcal{I} \text{ with } |x - c| < \epsilon.$$

Therefore ϕ_w is continuous.

Now we show that the fixed point y_w exists, using the intermediate value theorem applied to the function $g(x) := x - \phi_w(x)$. First we show that $g(y_0) \geq 0$. Indeed, as ϕ_1 is contractive, the definition of y_1 gives

$$y_0 - y_1 > \phi_1(y_0) - \phi_1(y_1) = \phi_1(y_0) - y_1, \quad \text{so that} \quad \phi_1(y_0) < y_0.$$

So, it follows from the definition of y_0 that the upper bound $\psi(x) := \max\{\phi_0(x), \phi_1(x)\}$ satisfies $\psi(y_0) = \phi_0(y_0) = y_0$. As $\phi_u(x)$ is increasing for any finite word u, it follows that

$$\phi_w(y_0) = \phi_{w_{2:|w|}} \circ \phi_{w_1}(y_0) \leq \phi_{w_{2:|w|}} \circ \psi(y_0) = \phi_{w_{2:|w|}}(y_0) \leq \cdots \leq y_0,$$

so that

$$g(y_0) = y_0 - \phi_w(y_0) \geq y_0 - y_0 = 0.$$

A similar argument, using the lower bound $\min\{\phi_0(x), \phi_1(x)\}$, gives $g(y_1) \leq 0$. In summary, $g(y_1) \leq 0 \leq g(y_0)$, where $y_1 < y_0$ by A2, and $g(x)$ is continuous as $\phi_w(x)$ is continuous. So the intermediate value theorem shows that $g(y) = 0$ for some $y \in [y_1, y_0]$. Therefore a fixed point y_w exists on $\mathcal{I}$.

Now we show that the fixed point y_w is unique. Suppose both y and z are fixed points of ϕ_w with $y > z$. This leads to the following contradiction: as ϕ_w is contractive we have

$$\frac{\phi_w(y) - \phi_w(z)}{y - z} < 1,$$

yet as $\phi_w(y) = y$ and $\phi_w(z) = z$ we have

$$\frac{\phi_w(y) - \phi_w(z)}{y - z} = 1.$$

Therefore the fixed point is unique. This completes the proof. ■

We make widespread use of the following simple lemma. Given a word w, this lemma gives necessary and sufficient conditions for $\phi_w(x)$ to be greater than or less than x.

Lemma 33 *Suppose A2 holds, that $x \in \mathcal{I}$ and w is a finite non-empty word. Then*

$$x < \phi_w(x) \;\Leftrightarrow\; \phi_w(x) < y_w \;\Leftrightarrow\; x < y_w \quad \textit{and} \quad x > \phi_w(x) \;\Leftrightarrow\; \phi_w(x) > y_w \;\Leftrightarrow\; x > y_w.$$

Proof We use Lemma 32 and the definition of y_w throughout without further mention.

As ϕ_w is increasing, if $x < y_w$ then $\phi_w(x) < \phi_w(y_w) = y_w$. Similarly, if $x > y_w$ then $\phi_w(x) > y_w$. Thus if $\phi_w(x) \leq y_w$ then $x \leq y_w$, by the contrapositive. But if $\phi_w(x) \neq y_w$ then $x \neq y_w$, as ϕ_w is increasing and therefore injective. So if $\phi_w(x) < y_w$ then $x < y_w$. Therefore

$$x < y_w \quad \Leftrightarrow \quad \phi_w(x) < y_w.$$

As ϕ_w is contractive, if $x < y_w$ then $\phi_w(y_w) - \phi_w(x) < y_w - x$. As $\phi_w(y_w) = y_w$, this rearranges to give $x < \phi_w(x)$. Similarly, if $x > y_w$ then $x > \phi_w(x)$. Thus if $x \leq \phi_w(x)$ then $x \leq y_w$, by the contrapositive. But if $x \neq \phi_w(x)$ then x is not a fixed point, so $x \neq y_w$. So if $x < \phi_w(x)$ then $x < y_w$. Therefore

$$x < y_w \quad \Leftrightarrow \quad x < \phi_w(x).$$

A similar argument shows that $x > y_w \Leftrightarrow \phi_w(x) > y_w$ and $x > y_w \Leftrightarrow x > \phi_w(x)$. ■

Lemma 34 *Suppose A2 holds and w is a finite word with $|w|_0|w|_1 > 0$. Then*

$$y_1 < y_w < y_0.$$

Proof As $|w|_0 > 0$ we have $w =: s01^q$ for some finite word s and some $q \in \mathbb{Z}_+$. As $y_0 > y_1$ by A2, Lemma 33 gives $\phi_0(y_1) > y_1$. Thus the definition of y_1 and the fact that ϕ_s is increasing give

$$\phi_w(y_1) = \phi_{s01^q}(y_1) = \phi_{s0}(y_1) = \phi_s(\phi_0(y_1)) > \phi_s(y_1) \geq y_1$$

where the last step follows by repeating the same argument. But if $\phi_w(y_1) > y_1$ then Lemma 33 shows that $y_w > y_1$.

A similar argument leads to the conclusion that $y_w < y_0$. ■

Lemma 35 *Suppose A2 holds, $x \in \mathcal{I}$ and w is a finite non-empty word. Then*

$$\lim_{n\to\infty} \phi_{w^n}(x) = y_w.$$

Proof The sequence with elements $x_n := \phi_{w^n}(x)$ for $n \in \mathbb{Z}_{++}$ is monotone and bounded, by Lemma 33. So the monotone convergence theorem for sequences of real numbers shows that $l := \lim_{n\to\infty} x_n$ exists. But as ϕ_w is continuous, by Lemma 32, the limit l satisfies

$$\phi_w\big(l\big) = \phi_w\big(\lim_{n\to\infty} x_n\big) = \lim_{n\to\infty} \phi_w\big(x_n\big) = \lim_{n\to\infty} x_{n+1} = l.$$

So l is a fixed point of ϕ_w. By Lemma 32, y_w is the unique such fixed point. ■

As a contractive function is not necessarily a contraction mapping, some additional work is required to prove the following result which is essential in the Sturmian case addressed in Subsection A.3.

Lemma 36 *Suppose A2 holds, w is an infinite word and $y_0 \geq a > b \geq y_1$. Then*

$$\lim_{n\to\infty} (\phi_{w_{1:n}}(a) - \phi_{w_{1:n}}(b)) = 0.$$

Proof Let $a_n := \phi_{w_{1:n}}(a)$ and $b_n := \phi_{w_{1:n}}(b)$ for $n \in \mathbb{Z}_{++}$. By Assumption A2, ϕ_{w_n} is a contractive function, so $(a_n - b_n)_{n=1}^{\infty}$ is a decreasing sequence and as $a, b \in [y_1, y_0]$ Lemma 33 shows that this is also a bounded sequence. Therefore the monotone convergence theorem for real-valued sequences shows that

$$\lim_{n\to\infty} (a_n - b_n) \text{ exists.}$$

Now we argue that the limit is zero, by contradiction. Assume that $a_n - b_n \geq \epsilon$ for all $n \in \mathbb{Z}_{++}$, where ϵ is a positive real number. Let the domain $\mathcal{D} \subset \mathbb{R}^2$ be

$$\mathcal{D} := \{(h, l) \ : \ h, l \in [y_1, y_0], \ h \geq l + \epsilon\},$$

let the functions $f_c : \mathcal{D} \to \mathbb{R}$ for letters $c \in \{0, 1\}$ be

$$f_c(h, l) := \frac{\phi_c(h) - \phi_c(l)}{h - l}$$

where $(h, l) \in \mathcal{D}$ and define the number $q \in \mathbb{R}$ by

$$q := \sup_{(h,l)\in\mathcal{D}} \left\{ \max_{c\in\{0,1\}} f_c(h, l) \right\}.$$

Now the functions f_0, f_1 are continuous on their domain $\mathcal{D}$ by Lemma 32. Also, the domain $\mathcal{D}$ is a non-empty, bounded and closed set. So the extreme value theorem for functions of several variables shows that the maximum equals the supremum. Thus

$$q = \max_{(h,l)\in\mathcal{D}} \left\{ \max_{a\in\{0,1\}} f_a(h, l) \right\}.$$

As ϕ_c is contractive for $c \in \{0, 1\}$ it follows that

$$q < 1$$

and as ϕ_c is increasing we have $q > 0$. So the definition of q and hypothesis that $a > b$ give

$$a_n - b_n \leq q^n(a - b) < \epsilon \quad \text{for } n > \log((a - b)/\epsilon)/ \log(1/q).$$

Thus there is an $n \in \mathbb{Z}_{++}$ with $a_n - b_n < \epsilon$. This contradicts the assumption that $a_n - b_n \geq \epsilon$ for all $n \in \mathbb{Z}_{++}$. Since $\epsilon > 0$ was arbitrary, we conclude that

$$\lim_{n\to\infty} (a_n - b_n) = 0.$$

This completes the proof. ■

A.2. x-Threshold Words as Christoffel Words

We begin with some definitions that go beyond the main text.

In Definition 5 and the text immediately following it, we introduced the s^--threshold itinerary $\sigma(x|s^-, \phi_0, \phi_1)$ from initial state x for a pair of maps ϕ_0 and ϕ_1. Let us focus on the case where the threshold equals the initial state as in Theorem 12. Now, the first letter of this itinerary is always 1 and as that theorem suggests, the remainder of the itinerary is frequently periodic, motivating the following definition.

Definition 37 *Consider an interval $\mathcal{I} \subseteq \mathbb{R}$, maps $\phi_0 : \mathcal{I} \to \mathcal{I}$ and $\phi_1 : \mathcal{I} \to \mathcal{I}$, and a state $x \in \mathcal{I}$. The x-**threshold word** for ϕ_0 and ϕ_1 is the shortest word $\ell(x, \phi_0, \phi_1)$ such that $\sigma(x|x^-, \phi_0, \phi_1) = 1\ell(x, \phi_0, \phi_1)^\infty$.*

If the maps are evident from the context, we talk of the x-threshold word $\ell(x)$ and if we are also focusing on a particular x, we talk of the x-threshold word ℓ.

The set $\{0, 1\}^*$ consists of all finite words on the alphabet $\{0, 1\}$, including the empty string ϵ. The *morphism* $\mathscr{M} : \{0, 1\}^* \to \{0, 1\}^*$ generated by a mapping $s : \{0, 1\} \to \{0, 1\}^*$, substitutes $s(w_k)$ for each letter w_k of a word w, so that

$$\mathscr{M}(\epsilon) = \epsilon \text{ and } \mathscr{M}(w) = s(w_1)s(w_2)\cdots s(w_{|w|}).$$

We work with the morphisms $\mathscr{L} : \{0, 1\}^* \to \{0, 1\}^*$ and $\mathscr{R} : \{0, 1\}^* \to \{0, 1\}^*$ given by

$$\mathscr{L} : \begin{cases} 0 \mapsto 0 \\ 1 \mapsto 01 \end{cases} \qquad \mathscr{R} : \begin{cases} 0 \mapsto 01 \\ 1 \mapsto 1 \end{cases}.$$

Let $\circ$ denote composition of morphisms, for example $\mathscr{R} \circ \mathscr{L}(1) = \mathscr{R}(01) = 011$.

Remark 38 *These morphisms generate the Christoffel tree through pre-composition. Say (u, v) is a Christoffel pair and consider the morphism*

$$\mathscr{M} : \begin{cases} 0 \mapsto u \\ 1 \mapsto v \end{cases}.$$

Then pre-composition maps the node (u, v) of the Christoffel tree to its children:

$$\begin{aligned} (\mathscr{M} \circ \mathscr{L}(0), \mathscr{M} \circ \mathscr{L}(1)) &= (\mathscr{M}(0),\ \mathscr{M}(01)) = (u, uv) \\ (\mathscr{M} \circ \mathscr{R}(0),\ \mathscr{M} \circ \mathscr{R}(1)) &= (\mathscr{M}(01),\ \mathscr{M}(1)) = (uv, v). \end{aligned} \tag{40}$$

Now let us give a simple upper and lower bound on the elements of x^--threshold orbits.

Lemma 39 *Suppose A2 holds and $(x_k)_{k=1}^\infty = \text{orbit}(\phi_1(x)|x^-)$. Then*

$$x \in [y_1, y_0] \;\Rightarrow\; \phi_1(x) \le x_k < \phi_0(x) \quad \textit{for } k \in \mathbb{Z}_{++}.$$

Proof Say $z \in [y_1, y_0)$. Then Lemma 33 gives

$$y_1 \le \phi_1(z) \le z < \phi_0(z) < y_0.$$

An induction using this fact, immediately shows that for $k \in \mathbb{Z}_{++}$,

$$y_1 \le x_k < y_0.$$

Now we prove the claim by induction with hypothesis

$$H_k : \quad \phi_1(x) \le x_k < \phi_0(x)$$

for $k \in \mathbb{Z}_{++}$. The base case H_1 is true as $x_1 = \phi_1(x)$ by definition of $\text{orbit}(\phi_1(x)|x^-)$. For the inductive step, say H_k is true for some $k \in \mathbb{Z}_{++}$. Then there are two cases to consider: $x_k \in [\phi_1(x), x)$ and $x_k \in [x, \phi_0(x))$. If $x_k \in [\phi_1(x), x)$ then $x_{k+1} = \phi_0(x_k)$ and

$$\begin{aligned} \phi_1(x) &\le x_k \\ &< \phi_0(x_k) && \text{by Lemma 33 as } x_k < y_0 \\ &< \phi_0(x) && \text{as } x_k < x \text{ and } \phi_0(\cdot) \text{ is increasing,} \end{aligned}$$

so H_{k+1} is true. If $x_k \in [x, \phi_0(x))$ then $x_{k+1} = \phi_1(x_k)$ and

$$\begin{aligned} \phi_1(x) &\le \phi_1(x_k) && \text{as } x \le x_k \text{ and } \phi_1(\cdot) \text{ is increasing} \\ &\le x_k && \text{by Lemma 33 as } y_1 \le x_k \\ &< \phi_0(x), \end{aligned}$$

so H_{k+1} is true. This completes the proof. ∎

Now we show that x-threshold words are Christoffel words in three important special cases and then we find the general conditions on x for which x-threshold words are Christoffel words (Lemma 44).

Lemma 40 *Suppose ℓ is the x-threshold word for ϕ_0, ϕ_1 satisfying A2. Then*

$$\begin{aligned} \ell = 1 &\Leftrightarrow x \le y_1 \\ \ell = 01 &\Leftrightarrow x \in [y_{01}, y_{10}] \\ \ell = 0 &\Leftrightarrow x \ge y_0. \end{aligned}$$

Proof If $\ell = 1$ then the definition of the x-threshold word shows that $\phi_1(x) \ge x$. Therefore Lemma 33 shows that $x \le y_1$.

If $x \le y_1$ then Lemma 33 shows that $x \le \phi_{1^n}(x) \le y_1$ for all $n \in \mathbb{Z}_{++}$. Therefore $\ell = 1$.

If $\ell = 01$ then the next letter of ℓ^∞ after $(01)^n$ is 0 for any $n \in \mathbb{Z}_+$. So $x > \phi_{(01)^n}(\phi_1(x))$. Therefore Lemma 35 gives

$$x \ge \lim_{n \to \infty} \phi_{(01)^n}(\phi_1(x)) = y_{01}.$$

If $\ell = 01$ then $\ell_2 = 1$. So $x \leq \phi_{10}(x)$. Therefore Lemma 33 gives

$$x \leq y_{10}.$$

If $x \geq y_{01}$ then $x > y_1$ by Lemma 34. So $\phi_1(x) < x$ by Lemma 33. As ϕ_{01} is increasing by Lemma 32, it follows that for any $n \in \mathbb{Z}_+$,

$$\phi_{(01)^n}(\phi_1(x)) < \phi_{(01)^n}(x) \leq x$$

where the second inequality follows from Lemma 33 as $x \geq y_{01}$. Therefore, if ℓ^∞ begins with $(01)^n$ then the next letter is 0.

If $x \leq y_{10}$ then Lemma 33 shows that $\phi_{(10)^n}(x) \geq x$ for all $n \in \mathbb{Z}_+$. Therefore, if ℓ^∞ begins with $(01)^n 0$ then the next letter is 1.

The proof for the case $\ell = 0$ is symmetric to that for $\ell = 1$. This completes the proof. ∎

Lemma 41 *Suppose ϕ_0, ϕ_1 satisfy A2, $x \in \mathcal{I}$ and ℓ is the x-threshold word. Then*

$$\begin{cases} |\ell|_{11} > 0 & \Rightarrow \quad x < y_{01} \\ |\ell|_{00} > 0 & \Rightarrow \quad x > y_{10}. \end{cases}$$

Proof If $|\ell|_{11} > 0$ then $x < y_0$ by Lemma 40. So, either $x \leq y_1$, in which case Lemma 34 shows that $x < y_{01}$, or $x \in (y_1, y_0)$. In the latter case, let $(x_k)_{k=1}^{\infty} = \text{orbit}(\phi_1(x)|x^-)$. As $|\ell|_{11} > 0$ there exists a $k \in \mathbb{Z}_{++}$ with $\phi_1(x_k) \geq x$ by definition of the x-threshold word. Now $x_k < \phi_0(x)$ by Lemma 39, so that $\phi_1(\phi_0(x)) > \phi_1(x_k) \geq x$ as ϕ_1 is increasing. But $\phi_{01}(x) > x$ implies that $x < y_{01}$ by Lemma 33.

The proof for $|\ell|_{00} > 0$ is symmetric. This completes the proof. ∎

Lemma 42 *Suppose ϕ_0, ϕ_1 satisfy A2. Then for $x \in [y_{10}, y_0]$, there is a unique $z \in \mathcal{I}$ with*

$$\phi_0(z) = x.$$

Furthermore

$$\ell(x, \phi_0, \phi_1) = \begin{cases} \mathscr{L}(\ell(\phi_0^{(-1)}(x), \phi_0, \phi_{01})) & \text{if } x \in [y_{10}, y_0] \\ \mathscr{R}(\ell(x, \phi_{01}, \phi_1)) & \text{if } x \in [y_1, y_{01}]. \end{cases}$$

Proof In the first claim, existence of z follows from the intermediate value theorem, as ϕ_0 is continuous by Lemma 32, as $y_{01} \in [y_0, y_1] \subseteq \mathcal{I}$ by Lemma 34, as $\phi_0(y_{01}) = y_{10} \leq x$ by definition of y_{01}, and as $\phi_0(y_0) = y_0 \geq x$ by definition of y_0. Uniqueness follows as ϕ_0 is increasing.

Now say $x \in [y_{10}, y_0]$ consider the claim involving the morphism $\mathscr{L}$. Let

$$V := \ell(\phi_0^{(-1)}(x), \phi_0, \phi_{01})^\infty \quad \text{and} \quad W := \ell(x, \phi_0, \phi_1)^\infty.$$

We show by induction that the hypothesis

$$H_i \ : \ \mathscr{L}(V_{1:i}) \text{ is a prefix of } W$$

is true for all $i \in \mathbb{Z}_+$, noting that this proves the claim.

In the base case, $\mathscr{L}(\epsilon) = \epsilon$ is a prefix of W, so H_0 is true. For the inductive step, say H_{i-1} is true for some $i \in \mathbb{Z}_{++}$. Let

$$(x_n)_{n=1}^{\infty} := \text{orbit}(\phi_1(x)|x^-, \phi_0, \phi_1),$$
$$(\tilde{x}_n)_{n=1}^{\infty} := \text{orbit}(\psi_1(\phi_0^{(-1)}(x))|\phi_0^{(-1)}(x)^-, \psi_0, \psi_1)$$

where $\psi_0 := \phi_0, \psi_1 := \phi_{01}$ and let $k := \left|\mathscr{L}(V_{1:(i-1)})\right| + 1$. Then

$$\tilde{x}_1 = \psi_1(\phi_0^{(-1)}(x)) = \phi_1(x) = x_1$$

and, letting ψ_w denote the composition of ψ_0, ψ_1 corresponding to a word w,

$$\begin{aligned}
\tilde{x}_i &= \psi_{V_{1:(i-1)}}(\tilde{x}_1) && \text{by definition of } V \\
&= \phi_{\mathscr{L}(V_{1:(i-1)})}(\tilde{x}_1) && \text{by definition of } \psi_0, \psi_1 \\
&= \phi_{W_{1:(k-1)}}(\tilde{x}_1) && \text{as } H_{i-1} \text{ is true} \\
&= \phi_{W_{1:(k-1)}}(x_1) && \text{as } \tilde{x}_1 = x_1 \\
&= x_k && \text{by definition of } W.
\end{aligned}$$

As $x \in [y_{10}, y_0]$, we have $\tilde{x} := \phi_0^{(-1)}(x) \in [y_{01}, y_0]$. Letting $\tilde{y}_0, \tilde{y}_1$ be the fixed points of ψ_0, ψ_1, this reads $\tilde{x} \in [\tilde{y}_1, \tilde{y}_0]$. But ψ_0, ψ_1 satisfy A2, as Lemma 32 shows that these functions are increasing and contractive, and Lemma 34 shows that $\tilde{y}_1 < \tilde{y}_0$. Thus Lemma 39 shows that

$$\tilde{x}_i < \psi_0(\tilde{x}) = \phi_0(\phi_0^{(-1)}(x)) = x.$$

But we already showed that $x_k = \tilde{x}_i$ so this gives $x_k < x$. Therefore $W_k = 0$, by definition of the x-threshold word. If $V_i = 0$ then we can conclude that H_i is true. Otherwise $V_i = 1$ so that $\tilde{x}_i \geq \phi_0^{(-1)}(x)$. But we already showed that $W_k = 0$ and $x_k = \tilde{x}_i$, so $x_{k+1} = \phi_0(x_k) = \phi_0(\tilde{x}_i) \geq x$. Therefore $W_{k+1} = 1$ and we conclude that H_i is true.

The proof for the claim involving $\mathscr{R}$ is similar. This completes the proof. ■

Lemma 43 *Suppose ϕ_0, ϕ_1 satisfy A2, $x \in \mathcal{I}$ and $0v1$ is a finite word. Then*

$$\begin{cases} \ell(\phi_0(x), \phi_0, \phi_1) = \mathscr{L}(0v1) & \Leftrightarrow \quad \ell(x, \phi_0, \phi_{01}) = 0v1 \\ \ell(x, \phi_0, \phi_1) = \mathscr{R}(0v1) & \Leftrightarrow \quad \ell(x, \phi_{01}, \phi_1) = 0v1. \end{cases}$$

Proof Consider the claim involving the morphism $\mathscr{L}$.

If $\ell(\phi_0(x), \phi_0, \phi_1) = \mathscr{L}(0v1)$ then $|\ell(\phi_0(x), \phi_0, \phi_1)|_{00} > 0$, as $|0v1|_{01} > 0$ for any finite word v and $\mathscr{L}(01) = 001$. So Lemma 41 shows that $\phi_0(x) > y_{10}$. Thus $x > y_{01}$. Also,

$|\ell(\phi_0(x), \phi_0, \phi_1)|_1 = |\mathscr{L}(0v1)|_1 > 0$, so Lemma 40 shows that $\phi_0(x) < y_0$. Thus $x < y_0$. Hence Lemma 42 gives

$$\mathscr{L}(0v1) = \mathscr{L}(\ell(x, \phi_0, \phi_{01})).$$

As $\mathscr{L}$ is injective, it follows that

$$0v1 = \ell(x, \phi_0, \phi_{01}).$$

If $\ell(x, \phi_0, \phi_{01}) = 0v1$ then $|\ell(x, \phi_0, \phi_{01})|_0 > 0$ and $|\ell(x, \phi_0, \phi_{01})|_1 > 0$. So Lemma 40 shows that $x \in (y_{01}, y_0)$. Hence Lemma 42 gives

$$\ell(\phi_0(x), \phi_0, \phi_1) = \mathscr{L}(0v1).$$

The argument for the claim involving $\mathscr{R}$ is symmetric. This completes the proof. ■

Lemma 44 *Suppose ϕ_0, ϕ_1 satisfy A2 and $0w1$ is a Christoffel word. Then*

$$x \in [y_{01w}, y_{10w}] \;\Leftrightarrow\; \ell(x, \phi_0, \phi_1) = 0w1.$$

Proof We use induction on the depth of $0w1$ in the Christoffel tree, with hypothesis

$$H_n \;:\; \left\{ \begin{array}{c} \text{If } 0w1 \text{ is at depth } n \text{ of the tree and } \phi_0, \phi_1 \text{ satisfy A2, then} \\ x \in [y_{01w}, y_{10w}] \;\Leftrightarrow\; \ell(x, \phi_0, \phi_1) = 0w1. \end{array} \right.$$

Lemma 40 shows that the base case (H_1) with $0w1 = 01$ is true. For the inductive step, let $0w1$ be a word at depth $n+1$ of the tree and assume H_n is true. Then either $0w1 = \mathscr{L}(0v1)$ or $0w1 = \mathscr{R}(0v1)$ for some word $0v1$ which is at depth n of the tree. If $0w1 = \mathscr{L}(0v1)$ then Lemma 43 gives

$$\ell(x, \phi_0, \phi_1) = 0w1 \quad \Leftrightarrow \quad \ell(\phi_0^{(-1)}(x), \phi_0, \phi_{01}) = 0v1.$$

Now ϕ_{01} is increasing and contractive by Lemma 32 and $y_{01} < y_0$ by Lemma 34. Thus ϕ_0, ϕ_{01} satisfy A2. So the assumption that H_n is true shows that

$$\ell(\phi_0^{(-1)}(x), \phi_0, \phi_{01}) = 0v1 \quad \Leftrightarrow \quad \phi_0^{(-1)}(x) \in [y_{\mathscr{L}(01v)}, y_{\mathscr{L}(10v)}].$$

But as $0w1 = \mathscr{L}(0v1) = 0\mathscr{L}(v)01$, we have

$$\begin{aligned} \phi_0(y_{\mathscr{L}(01v)}) &= \phi_0(y_{001\mathscr{L}(v)}) = y_{01\mathscr{L}(v)0} = y_{01w} \\ \phi_0(y_{\mathscr{L}(10v)}) &= \phi_0(y_{010\mathscr{L}(v)}) = y_{10\mathscr{L}(v)0} = y_{10w}. \end{aligned}$$

As ϕ_0 is increasing and continuous, it follows that

$$\phi_0^{(-1)}(x) \in [y_{\mathscr{L}(01v)}, y_{\mathscr{L}(10v)}] \quad \Leftrightarrow \quad x \in [y_{01w}, y_{10w}].$$

Therefore H_{n+1} is true.

If $0w1 = \mathscr{R}(0v1)$ then the proof is similar. This completes the proof. ■

A.3. x-Threshold Words as Sturmian $\mathcal{M}$-Words

It turns out that Lemma 44 characterises x-threshold words for nearly all $x \in (y_1, y_0)$. However, its proof cannot be extended to all values of x as it is based on induction on the depth $n \in \mathbb{Z}_{++}$ in the Christoffel tree, and we need to take the limit as $n \to \infty$ to address the remaining values of x. Those remaining values correspond to Sturmian $\mathcal{M}$-words, as we show in this subsection.

First we show how Sturmian $\mathcal{M}$-words can be written as the limit of a sequence of words.

Lemma 45 *A word w is a Sturmian $\mathcal{M}$-word if and only if it is of the form*

$$w = \lim_{n\to\infty} \mathscr{M}_1 \circ \mathscr{M}_2 \circ \cdots \circ \mathscr{M}_n(01)$$

for some sequence of morphisms $(\mathscr{M}_n)_{n=1}^{\infty}$ with $\mathscr{M}_n \in \{\mathscr{L}, \mathscr{R}\}$ which is not eventually constant (i.e. there is no $m \in \mathbb{Z}_{++}$ such that $\mathscr{M}_n = \mathscr{M}_m$ for all $n \in \mathbb{Z}_{++}$ with $n > m$).

The above result is well known, for instance, noting the isomorphism of the Christoffel tree and the Stern-Brocot tree, see Chapter 4 of Graham et al. (1994).

Remark 46 *The requirement that the sequence of morphisms generating Sturmian $\mathcal{M}$-words is not eventually constant rules out words such as*

$$\lim_{n\to\infty} \mathscr{L}^n(01) = 0^{\infty}.$$

Indeed, this corresponds to the Christoffel word 0. It also rules out words such as

$$\lim_{n\to\infty} \mathscr{R}^n(01) = 01^{\infty}.$$

Indeed, this is not an $\mathcal{M}$-word because there is no $\alpha \in [0,1]$ with

$$0 = \lfloor \alpha n \rfloor - \lfloor \alpha(n-1) \rfloor \text{ for } n = 1 \text{ and } 1 = \lfloor \alpha n \rfloor - \lfloor \alpha(n-1) \rfloor \text{ for } n = 2, 3, \ldots.$$

Now we show that every x-threshold word corresponds to an $\mathcal{M}$-word.

Lemma 47 *Suppose ϕ_0, ϕ_1 satisfy A2 and $x \in \mathcal{I}$. Then $\ell(x, \phi_0, \phi_1)$ is an $\mathcal{M}$-word.*

Proof Lemma 40 shows that $\ell(x, \phi_0, \phi_1)$ is an $\mathcal{M}$-word for $x \in \mathcal{I} \backslash (y_1, y_0)$. So, for the rest of this proof we assume that $x \in (y_1, y_0)$. We shall now define procedure which generates a sequence of morphisms $\mathscr{M}_k$, thresholds $x_k \in \mathcal{I}$, mappings $\phi_{k,a} : \mathcal{I} \to \mathcal{I}$ for $a \in \{0, 1\}$ and we denote the compositions of those mappings by

$$\phi_{k,w} = \phi_{k,w_{|w|}} \circ \cdots \phi_{k,w_2} \circ \phi_{k,w_1}$$

for any finite word w. The procedure is as follows:

1. **let** $k \leftarrow 1$
2. **let** $(x_k, \phi_{k,0}, \phi_{k,1}) \leftarrow (x, \phi_0, \phi_1)$
3. **let** $y_{k,01}, y_{k,10}$ be the fixed points of $\phi_{k,01}, \phi_{k,10}$

while $x_k \notin [y_{k,01}, y_{k,10}]$
  if $x_k < y_{k,01}$
    set $(x_{k+1}, \phi_{k+1,0}, \phi_{k+1,1}, \mathscr{M}_k) \leftarrow (x_k, \phi_{k,01}, \phi_{k,1}, \mathscr{R})$
  else
    set $(x_{k+1}, \phi_{k+1,0}, \phi_{k+1,1}, \mathscr{M}_k) \leftarrow (\phi_{k,0}^{(-1)}(x), \phi_{k,0}, \phi_{k,01}, \mathscr{L})$
  end
  let $k \leftarrow k + 1$
  let $y_{k,01}, y_{k,10}$ be the fixed points of $\phi_{k,01}, \phi_{k,10}$
end

The sequence of morphisms $\mathscr{M}_1, \mathscr{M}_2, \ldots$ generated by this procedure is either empty, of finite length $n \in \mathbb{Z}_{++}$ or of infinite length. We consider each of these cases in turn.

If the sequence is empty, then $x \in [y_{1,01}, y_{1,10}] = [y_{01}, y_{10}]$. So Lemma 40 shows that $\ell(x, \phi_0, \phi_1) = 01$, which is an $\mathcal{M}$-word.

If the sequence has finite length, let n be that length. As the morphisms $\mathscr{L}$ and $\mathscr{R}$ generate the Christoffel tree by pre-composition, as remarked at (40), the word

$$w := \mathscr{M}_1 \circ \cdots \circ \mathscr{M}_n(01)$$

is a Christoffel word. Also, when the procedure sets $\mathscr{M}_k = \mathscr{R}$ for some $k \in \mathbb{Z}_{++}$, we have $x_k < y_{k,01} < y_{k,10}$ so Lemma 42 shows that

$$\ell(x_k, \phi_{k,0}, \phi_{k,1}) = \mathscr{R}(\ell(x_{k+1}, \phi_{k+1,0}, \phi_{k+1,1})).$$

Similarly, when the procedure sets $\mathscr{M}_k = \mathscr{L}$ we have

$$\ell(x_k, \phi_{k,0}, \phi_{k,1}) = \mathscr{L}(\ell(x_{k+1}, \phi_{k+1,0}, \phi_{k+1,1})).$$

Therefore the x-threshold word is $\ell(x, \phi_0, \phi_1) = w$ which is an $\mathcal{M}$-word.

Finally, if the sequence does not terminate, then Lemma 45 shows that the word

$$w := \lim_{n \to \infty} \mathscr{M}_1 \circ \cdots \circ \mathscr{M}_n(01)$$

is a Sturmian $\mathcal{M}$-word *provided* there is no $m \in \mathbb{Z}_{++}$ such that $\mathscr{M}_n = \mathscr{M}_m$ for all $n \in \mathbb{Z}_{++}$ with $n \geq m$. Also, the argument given for finite sequences above shows that $w = \ell(x, \phi_0, \phi_1)$. If there were such an m and $\mathscr{M}_m = \mathscr{R}$, then

$$\begin{aligned} w &= \mathscr{M}_1 \circ \cdots \circ \mathscr{M}_{m-1} \circ \lim_{k \to \infty} \mathscr{R}^k(01) \\ &= \mathscr{M}_1 \circ \cdots \circ \mathscr{M}_{m-1}(01^\infty). \end{aligned}$$

Thus the word at stage m is $\ell(x_m, \phi_{m,0}, \phi_{m,1}) = 01^\infty$. But this is impossible, as the fact that the first letter is 0 requires $\phi_{m,1}(x_m) < x_m$, so that $x_m > y_{m,1}$, whereas the fact that remaining letters are 1 requires $\phi_{m,101^n}(x_m) \geq x_m$ for all $n \in \mathbb{Z}_+$, so that $y_{m,1} \geq x_m$, which is a contradiction. If there were such an m and $\mathscr{M}_m = \mathscr{L}$, then

$$w = \mathscr{M}_1 \circ \cdots \circ \mathscr{M}_{m-1} \circ \lim_{k \to \infty} \mathscr{L}^k(01)$$

$$= \mathscr{M}_1 \circ \cdots \circ \mathscr{M}_{m-1}(0^\infty)$$

Therefore at stage m the word $\ell(x_m, \phi_{m,0}, \phi_{m,1}) = 0$, but this is impossible as $x_m < y_{m,0}$. In conclusion, $\ell(x, \phi_0, \phi_1)$ is an $\mathcal{M}$-word.

This completes the proof. ■

Lemma 48 *Suppose* $0 \leq \alpha < \beta \leq 1$ *and let* a, b *be the* $\mathcal{M}$*-words of rate* α, β *respectively. Then* $a^\infty \prec b^\infty$*.*

Proof Consider the first $n \in \mathbb{Z}_{++}$ with $(a^\infty)_n \neq (b^\infty)_n$. Then

$$\lfloor \alpha(n-1) \rfloor = \left|(a^\infty)_{1:(n-1)}\right|_1 = \left|(b^\infty)_{1:(n-1)}\right|_1 = \lfloor \beta(n-1) \rfloor,$$

by definition of $\mathcal{M}$-words, so that

$$(a^\infty)_n = \lfloor \alpha n \rfloor - \lfloor \alpha(n-1) \rfloor = \lfloor \alpha n \rfloor - \lfloor \beta(n-1) \rfloor < \lfloor \beta n \rfloor - \lfloor \beta(n-1) \rfloor = (b^\infty)_n$$

where the inequality holds as $(a^\infty)_n \neq (b^\infty)_n$, as $\alpha < \beta$ and as the floor function is non-decreasing. Therefore $a^\infty \prec b^\infty$. ■

Consider a tree and a node x of the tree. Recall that the *subtree rooted at* x is the tree of all descendants of x that has node x as a root. If the tree is a binary tree, the *left subtree of* x is the subtree rooted at the left child of x, and the *right subtree of* x is the subtree rooted at the right child of x. Thus the subtree rooted at x contains node x, but the left and right subtrees of x do not contain node x.

Lemma 49 *Suppose* (u, v) *is a Christoffel pair. Considering* uv *as a node of the Christoffel tree, let* l *and* r *be Christoffel words in the left and right subtrees of* uv*. Then*

$$\text{rate}(u) < \text{rate}(l) < \text{rate}(uv) < \text{rate}(r) < \text{rate}(v).$$

Proof First we use induction to show that for any Christoffel pair (u, v), we have

$$\text{rate}(u) < \text{rate}(v). \tag{41}$$

In the base case $(u, v) = (0, 1)$ and (41) is true. For the inductive step, say (u, v) is the left child of a Christoffel pair (a, b) with $\text{rate}(a) < \text{rate}(b)$. Then $(u, v) = (a, ab)$ so that

$$\text{rate}(u) = \frac{|a|_1}{|a|} < \frac{|a|_1 + |b|_1}{|a| + |b|} = \text{rate}(v)$$

by the mediant inequality. The proof for right children is similar. Therefore (41) is true.

As l is in the left subtree of uv, it follows from the construction of the Christoffel tree that l consists of m copies of u and n copies of uv concatenated in some order, for some $m, n \in \mathbb{Z}_{++}$. Thus the mediant inequality and (41) give

$$\text{rate}(l) = \text{rate}(u^m (uv)^n) = \frac{m|u|_1 + n|uv|_1}{m|u| + n|uv|} \in \left(\frac{|u|_1}{|u|}, \frac{|uv|_1}{|uv|}\right)$$

Therefore $\text{rate}(u) < \text{rate}(l) < \text{rate}(uv)$, as claimed.

The proof for a node r in the right subtree of uv is similar. This completes the proof. ■

Lemma 50 *Suppose* $(0a1, 0b1)$ *is a Christoffel pair. Then* $a10b = b01a$.

Proof As $a, b, a10b$ are palindromes, we have $a10b = (a10b)^R = b^R01a^R = b01a$. ■

Lemma 51 *Suppose the word* $0c1$ *is in the subtree of the Christoffel tree rooted at* $0p1$. *Then* p *is both a prefix and suffix of* c.

Proof There are four cases to consider:

1. The word $0c1$ has no parent, in which case $c = p = \epsilon$ and the claim holds.
2. The word c is of one of the forms 0^m or 0^m10^m for some $m \in \mathbb{Z}_{++}$. In that case either $p = c$ or $p = 0^l$ for some $l \in \mathbb{Z}_+$ with $l < m$ and the claim holds.
3. The word c is of one of the forms 1^m or 1^m01^m for some $m \in \mathbb{Z}_{++}$. This is similar to the previous case.
4. The Christoffel pair of the parent of $0c1$ is of the form $(0a1, 0b1)$.

In the last case, we use induction on the length of the path $n \in \mathbb{Z}_+$ through the Christoffel tree from $0p1$ to $0c1$. In the base case, $n = 0$, we have $c = p$ and the claim is true. For the inductive step, say $0c1$ is a child of the node with Christoffel pair $(0a1, 0b1)$, that $0a10b1$ is n steps along the path from $0p1$, and that p is both a prefix and suffix of $a10b$, so that $a10b = pq = rp$ for some words q, r. If $0c1$ is a left-child, then $c = a10a10b = a10rp$ and Lemma 50 gives

$$c = a10a10b = a10b01a = pq01a$$

so p is a prefix and suffix of c. Similarly, if $0c1$ is a right child, then $c = a10b10b$ and

$$pq10b = a10b10b = b01a10b = b01rp.$$

This completes the proof. ■

Lemma 52 *Suppose* ϕ_0, ϕ_1 *satisfy A2. Then*

$$\ell(x, \phi_0, \phi_1)^\infty \text{ is a lexicographically non-increasing function of } x \in \mathcal{I}.$$

Proof If $a, b \in \mathcal{I}$ and $\ell(a, \phi_0, \phi_1)^\infty \succ \ell(b, \phi_0, \phi_1)^\infty$, then there is a finite word u such that

$$\ell(a, \phi_0, \phi_1)^\infty = u1v, \qquad \ell(b, \phi_0, \phi_1)^\infty = u0w,$$

for some words v, w. So the definition of threshold orbits and Lemma 33 give

$$y_{1u} \geq \phi_{1u}(a) \geq a, \qquad\qquad y_{1u} < \phi_{1u}(b) < b.$$

Therefore $a < b$. This completes the proof. ■

In the main text, we defined the fixed point y_s of a Sturmian $\mathcal{M}$-word as the limit of a sequence of fixed points $y_{01w^{(n)}}$ or $y_{10w^{(n)}}$ where the words $(0w^{(n)}1)_{n=1}^{\infty}$ correspond to a particular path in the Christoffel tree. We now define a subsequence associated with each of these sequences of words and fixed points. Consider a Sturmian $\mathcal{M}$-word

$$0s = \lim_{n\to\infty} \mathscr{L}^{a_1} \circ \mathscr{R}^{b_1} \circ \cdots \circ \mathscr{L}^{a_n} \circ \mathscr{R}^{b_n}(01)$$

where $a_1 \in \mathbb{Z}_+$ and $a_{n+1}, b_n \in \mathbb{Z}_{++}$ for $n \in \mathbb{Z}_{++}$. We define the sequences $(u^{(n)})_{n=1}^{\infty}$ and $(l^{(n)})_{n=1}^{\infty}$ of the central portions of the Christoffel words as

$$0u^{(n)}1 := \mathscr{L}^{a_1} \circ \mathscr{R}^{b_1} \circ \cdots \circ \mathscr{L}^{a_n}(01), \quad 0l^{(n)}1 := \mathscr{L}^{a_1} \circ \mathscr{R}^{b_1} \circ \cdots \circ \mathscr{L}^{a_n} \circ \mathscr{R}^{b_n}(01). \tag{42}$$

These words form a subsequence of $(0w^{(n)}1)_{n=1}^{\infty}$, so if $\lim_{n\to\infty} y_{01w^{(n)}}$ exists, then so does $\lim_{n\to\infty} y_{01l^{(n)}}$ and these limits are equal.

Lemma 53 *Suppose A2 holds, $n \in \mathbb{Z}_{++}$ and that $0u^{(n)}1$ and $0l^{(n)}1$ are as in (42). Then*

$$y_{01l^{(n)}} < y_{01l^{(n+1)}} < y_{10u^{(n+1)}} < y_{10u^{(n)}}.$$

Proof By definition, $0l^{(n+1)}1$ is in the left subtree of $0l^{(n)}1$, so Lemma 49 gives

$$\text{rate}(0l^{(n+1)}1) < \text{rate}(0l^{(n)}1).$$

Hence Lemma 44 and Lemma 48 give

$$\ell(y_{01l^{(n+1)}})^{\infty} = (0l^{(n+1)}1)^{\infty} \prec (0l^{(n)}1)^{\infty} = \ell(y_{10l^{(n)}})^{\infty}.$$

Therefore Lemma 52 shows that

$$y_{01l^{(n)}} < y_{01l^{(n+1)}}.$$

But $0l^{(n+1)}1$ is in the right subtree of $0u^{(n+1)}1$. So the same argument gives

$$y_{01l^{(n+1)}} < y_{10u^{(n+1)}}.$$

Similarly, $0u^{(n+1)}1$ is in the right subtree of $0u^{(n)}1$, so

$$y_{10u^{(n+1)}} < y_{10u^{(n)}}.$$

This completes the proof. ■

Note that the argument in the above proof also shows that any word $0w^{(m)}1$ lying strictly between $0u^{(n)}1$ and $0l^{(n)}1$ on the path through the Christoffel tree from $0u^{(n)}1$ to $0l^{(n)}1$ has

$$y_{10l^{(n)}} < y_{01w^{(m)}} < y_{10w^{(m)}} < y_{01u^{(n)}}.$$

Similarly, any word $0w^{(m)}1$ lying strictly between $0l^{(n)}1$ and $0u^{(n+1)}1$ on the path from $0l^{(n)}1$ to $0u^{(n+1)}1$ has

$$y_{10l^{(n)}} < y_{01w^{(m)}} < y_{10w^{(m)}} < y_{01u^{(n+1)}}.$$

Thus if the subsequences $(y_{01l^{(n)}})_{n=1}^{\infty}$ and $(y_{10u^{(n)}})_{n=1}^{\infty}$ have limits, then so do the full sequences $(y_{01w^{(m)}})_{n=1}^{\infty}$ and $(y_{10w^{(m)}})_{n=1}^{\infty}$.

Lemma 54 *Suppose $(0w^{(n)}1)_{n=1}^{\infty}$ is the sequence of words traversed along an infinite path down the Christoffel tree. Let $(n_i)_{i=1}^{\infty}$ be an increasing sequence on $\mathbb{Z}_{++}$. Then there exists an infinite word s and an increasing sequence $(k_i)_{i=1}^{\infty}$ on $\mathbb{Z}_{++}$ such that $s_{1:k_i}$ is a suffix of both $w^{(n_i)}$ and $w^{(n_{i+1})}$, for all $i \in \mathbb{Z}_{++}$.*

Proof We show that $s_{1:k_i} := w^{(n_i)}$ for $i \in \mathbb{Z}_{++}$, is well-defined and satisfies this claim. As $w^{(n_i)}$ is a prefix of all of its descendants by Lemma 51, it follows that $s_{1:k_i}$ is a prefix of $s_{1:k_{i+1}}$. Also, as $w^{(n_i)}$ is a suffix of all of its descendants, it follows that $s_{1:k_i}$ is a suffix of both $w^{(n_i)}$ and $w^{(n_{i+1})}$. ∎

Lemma 55 *Suppose ϕ_0, ϕ_1 satisfy A2 and $0s$ is a Sturmian $\mathcal{M}$-word. Consider the sequence of Christoffel words $(0w^{(n)}1 : n \in \mathbb{Z}_{++})$ traversed on the infinite path through the Christoffel tree towards $0s$ (as defined in the main text just before Theorem 12). Then the fixed points*

$$y_{01s} := \lim_{n\to\infty} y_{01w^{(n)}} \qquad \textit{and} \qquad y_{10s} := \lim_{n\to\infty} y_{10w^{(n)}}$$

exist and are equal.

Proof Let $a_n, b_n, 0u^{(n)}1, 0l^{(n)}1$ for $n \in \mathbb{Z}_{++}$ be as in (42).

Existence of the fixed points follows from the monotone convergence theorem for real-valued sequences. Indeed, Lemma 53 shows that $(y_{01l^{(n)}})_{n=1}^{\infty}$ is an increasing sequence, that $(y_{10u^{(n)}})_{n=1}^{\infty}$ is a decreasing sequence, and that these sequences are bounded.

By Lemma 54 we have $u^{(n)} = c^{(n)}w_{1:k_n}$ and $l^{(n)} = d^{(n)}w_{1:k_n}$ for all $n \in \mathbb{Z}_{++}$ for some sequences of finite words $(c^{(n)})_{n=1}^{\infty}$ and $(d^{(n)})_{n=1}^{\infty}$, for some infinite word w and for some increasing sequence $(k_n)_{n=1}^{\infty}$ on $\mathbb{Z}_{++}$.

As $\phi_a(x)$ is an increasing function of $x \in \mathcal{I}$ for any finite word a, by Lemma 32,

$$\begin{aligned} y_{10u^{(n)}} &= \phi_{10u^{(n)}}(y_{10u^{(n)}}) \\ &< \phi_{10u^{(n)}}(y_0) \\ &= \phi_{10c^{(n)}w_{1:k_n}}(y_0) \\ &= \phi_{w_{1:k_n}}(\phi_{10c^{(n)}}(y_0)) \end{aligned}$$

$$< \phi_{w_{1:k_n}}(y_0).$$

Similarly, we have

$$y_{01l^{(n)}} > \phi_{w_{1:k_n}}(y_1).$$

Thus

$$\lim_{n\to\infty}(y_{10u^{(n)}} - y_{01l^{(n)}}) \le \lim_{n\to\infty}(\phi_{w_{1:k_n}}(y_0) - \phi_{w_{1:k_n}}(y_1)) = 0$$

where the last step is Lemma 36. But $y_{10u^{(n)}} > y_{01l^{(n)}}$ for $n \in \mathbb{Z}_{++}$ by Lemma 53. Therefore

$$y_{10s} = \lim_{n\to\infty} y_{10u^{(n)}} = \lim_{n\to\infty} y_{01l^{(n)}} = y_{01s}.$$

This completes the proof. ■

In view of the above Lemma, from now on we shall write $y_s = y_{10s} = y_{01s}$.

Lemma 56 *Suppose ϕ_0, ϕ_1 satisfy A2 and $0s$ is a Sturmian $\mathcal{M}$-word. Then*

$$\ell(x, \phi_0, \phi_1) = 0s \;\Leftrightarrow\; x = y_s.$$

Proof Let $(u^{(n)})_{n=1}^{\infty}$ and $(l^{(n)})_{n=1}^{\infty}$ be the sequences (42) appearing in the definition of the fixed point of $0s$. For fixed ϕ_0, ϕ_1 let us write $\ell(x)$ in place of $\ell(x, \phi_0, \phi_1)$.

Say $x = y_s$. Recall that $x = y_{10s} = y_{01s}$ by Lemma 55. As $y_{10u^{(n)}} > y_{10s}$ for all $n \in \mathbb{Z}_{++}$ by Lemma 53, and $\ell(z)^\infty$ is a lexicographically non-increasing function of z, by Lemma 52, it follows that

$$\ell(x)^\infty \succ \ell(y_{10u^{(n)}})^\infty = (0u^{(n)}1)^\infty$$

where the equality follows from Lemma 44. A similar argument gives $\ell(x)^\infty \prec (0l^{(n)}1)^\infty$ for $n \in \mathbb{Z}_{++}$. Therefore

$$0s = \lim_{n\to\infty}(0u^{(n)}1)^\infty \preceq \ell(x)^\infty \preceq \lim_{n\to\infty}(0l^{(n)}1)^\infty = 0s.$$

Now say $\ell(x) = 0s$ for some $x \in \mathcal{I}$. Then $y_{01l^{(n)}} < x < y_{10u^{(n)}}$ as $\ell(z)^\infty$ is a lexicographically non-increasing function of z and $(0u^{(n)}1)^\infty \prec 0s \prec (0l^{(n)}1)^\infty$ for all $n \in \mathbb{Z}_{++}$. Therefore

$$y_s = \lim_{n\to\infty} y_{01l^{(n)}} \le x \le \lim_{n\to\infty} y_{10u^{(n)}} = y_s.$$

This completes the proof. ■

A.4. Proof of Theorem 12

Proof The existence of fixed points y_{01p}, y_{10p} follows from Lemma 32 and the existence of y_s follows from Lemma 55.

The fact that $\sigma(z|z^-) = 1\ell(z, \phi_0, \phi_1)^\infty$ is a lexicographically non-decreasing function of $z \in \mathcal{I}$ follows from Lemma 52.

Lemma 40 addresses the value of $\sigma(z|z^-)$ for $z \leq y_1$ and $z \geq y_0$, Lemma 44 addresses the case $z \in [y_{01p}, y_{10p}]$ and Lemma 56 address the case $z = y_s$. Thus the image of $\ell(z, \phi_0, \phi_1)$ as z ranges over $\mathcal{I}$ contains all $\mathcal{M}$-words. Using Lemma 47, it follows that this image is exactly the set of $\mathcal{M}$-words.

This completes the proof. ∎

A.5. Proof of Corollary 13

Proof First we give a simple result about the relation between $\mathcal{M}$-words of rates r and $1-r$. Then we demonstrate the result for itineraries of the form $\sigma(\phi_0(x)|x)$ and finally the result for itineraries of the form $\sigma(\phi_1(x)|x)$.

Relation between $\mathcal{M}$-words. Let bitflip$(\cdot)$ denote the bitwise NOT operation, that maps any finite or infinite word w to the word of the same length whose n^{th} letter is $1 - w_n$. We now show that

if $0s$ is a Sturmian $\mathcal{M}$-word then $0\,\text{bitflip}(s)$ is also a Sturmian $\mathcal{M}$-word.

Let $r := \text{rate}(0s)$. By definition, the word w of rate $1 - r$ has

$$w_n = \lfloor (1-r)n \rfloor - \lfloor (1-r)(n-1) \rfloor = 1 - \lceil rn \rceil + \lceil r(n-1) \rceil$$

for any positive integer n. As $0s$ is Sturmian, its rate r is irrational, so rm is not an integer for any positive integer m and it follows that $\lceil rm \rceil = \lfloor rm \rfloor$. Thus

$$w_n = 1 - (\lfloor rn \rfloor - \lfloor r(n-1) \rfloor) = \text{bitflip}(0s)_n$$

as claimed. A similar argument shows that

if $0p1$ is a Christoffel word then $0\,\text{bitflip}(p)1$ is also a Christoffel word.

Result for $\sigma(\phi_0(x)|x)$. Let $\tilde{\phi}_0(x) := -\phi_1(-x)$, $\tilde{\phi}_1(x) := -\phi_0(-x)$ and $\tilde{\mathcal{I}} := \{-x : x \in \mathcal{I}\}$. For a word w, let $\tilde{y}_w$ be the fixed point $y_w = \tilde{\phi}_w(y_w)$ of the composition $\tilde{\phi}_w$ of $\tilde{\phi}_0, \tilde{\phi}_1$ corresponding to w. Since A2 holds for $\mathcal{I}, \phi_0$ and ϕ_1, it is easy to see that it also holds for $\tilde{\mathcal{I}}, \tilde{\phi}_0$ and $\tilde{\phi}_1$. Let $\tilde{\sigma}(x|s^-)$ denote the s^--threshold itinerary from $x \in \tilde{\mathcal{I}}$ for the maps $\tilde{\phi}_0, \tilde{\phi}_1$. From these definitions, for any $x \in \mathcal{I}$ we have

$$\sigma(\phi_0(x)|x) = \text{bitflip}(\tilde{\sigma}(\tilde{\phi}_1(-x)|(-x)^-)).$$

But if $x \leq y_1$ then $-x \geq \tilde{y}_0$, so Theorem 12 gives $\tilde{\sigma}(-x|(-x)^-) = 10^\infty$ and thus $\sigma(\phi_0(x)|x) = 1^\infty$. The argument that $\sigma(\phi_0(x)|x) = 0^\infty$ for $x \geq y_0$ is similar. Also if $x \in [y_{01p}, y_{10p}]$ then $-x \in [\tilde{y}_{01\,\text{bitflip}(p)}, \tilde{y}_{10\,\text{bitflip}(p)}]$, and the relation between $\mathcal{M}$-words above shows that

$0\,\text{bitflip}(p)1$ is a Christoffel word, so Theorem 12 gives $\tilde{\sigma}(-x|(-x)^-) = (10\,\text{bitflip}(p))^\infty$ and thus $\sigma(\phi_0(x)|x) = (1p0)^\infty$. Finally if $x = y_s$ then $-x = \tilde{y}_{\text{bitflip}(s)}$, and the relation between $\mathcal{M}$-words above shows that $0\,\text{bitflip}(s)$ is a Sturmian $\mathcal{M}$-word, so Theorem 12 gives $\tilde{\sigma}(-x|(-x)^-) = 10\,\text{bitflip}(s)$ and thus $\sigma(\phi_0(x)|x) = 1s$.

Result for $\sigma(\phi_1(x)|x)$. Let orbit$(x|s^-)$ denote the sequence of states in the s^--threshold orbit from initial state x under maps ϕ_0 and ϕ_1. The core of our argument is that if orbit$(x|s^-)$ does not contain s then $\sigma(x|s) = \sigma(x|s^-)$. Indeed if s is not in orbit$(x|s^-) =: (x_1, x_2, \dots)$, the decisions $\mathbf{1}_{x_k > s}$ and $\mathbf{1}_{x_k \geq s}$ are the same for each $k = 1, 2, \dots$.

Say $x \in [y_{01p}, y_{10p})$. First we show that assuming that x is is in orbit$(\phi_1(x)|x^-)$ leads to a contradiction. If indeed x were in that orbit, then either $x = \phi_{(10p)^N}(x)$ for some $N \in \mathbb{Z}_{++}$ or $x = \phi_{(10p)^N(10p)_{1:m}}(x)$ for some $N \in \mathbb{Z}_+$ and $1 \leq m < |10p|$. The first case gives $x = y_{10p}$ which contradicts $x < y_{10p}$. The second case gives $\sigma(x|x^-) = (10p)^N(10p)_{1:m}\sigma(x|x^-) = ((10p)^N(10p)_{1:m})^\infty$ and as m is not a multiple of $|10p|$, this contradicts Theorem 12 which shows that $\sigma(x|x^-) = (10p)^\infty$. It follows that

$$\sigma(\phi_1(x)|x) = \sigma(\phi_1(x)|x^-).$$

Since $\sigma(x|x^-) = (10p)^\infty$ by Theorem 12, we conclude that $\sigma(\phi_1(x)|x) = (0p1)^\infty$.

If $x = y_{10p}$ then $\sigma(x|x^-) = (10p)^\infty$ by Theorem 12 and the first time orbit$(x|x^-)$ returns to x is after the letters $10p$. Thus $\sigma(\phi_1(x)|x)$ and $\sigma(\phi_1(x)|x^-)$ agree up to that point, so $\sigma(\phi_1(x)|x) = 0p0\sigma(\phi_0(x)|x) = 0p(0p1)^\infty$ using the result about $\sigma(\phi_0(x)|x)$ proved earlier.

Say $x = y_s$. Theorem 12 gives $\sigma(x|x^-) = 10s$. Say $\phi_{(10s)_{1:k}}(x) = x$ for some $k \in \mathbb{Z}_{++}$. Then $\sigma(x|x^-) = (10s)_{1:k}\sigma(x|x^-) = ((10s)_{1:k})^\infty$ which contradicts the fact that Sturmian $\mathcal{M}$-words are not periodic. Therefore orbit$(x|x^-)$ does not return to x and it follows that $\sigma(\phi_1(x)|x) = \sigma(\phi_1(x)|x^-) = 0s$.

This completes the proof. ■

A.6. Proof of Theorem 14

Based on the work of Kozyakin (2003), we begin by showing that maps-with-gaps formed from functions satisfying Assumption A2 are *locally-growing relaxation functions* (Lemmas 58 and 59) and that the itineraries of such functions are 1-*balanced words* (Lemmas 61 and 62). As 1-balanced words correspond to factors of lower mechanical words (Lemma 63), this enables us to describe the itineraries of maps-with-gaps, for arbitrary initial states and thresholds (Lemma 64). We couple this description with an easy result about lexicographic ordering (Lemma 65) and with a result about the number of factors of mechanical words (Lemma 66) to bound the number of discontinuities of the itinerary of a map-with-a-gap as a function of its threshold (Lemma 67). Finally, we prove Theorem 14.

Definition 57 *Let $\mathcal{I}$ be an interval of $\mathbb{R}$. A function $f : \mathcal{I} \to \mathcal{I}$ is a* ***locally-growing relaxation function*** *with threshold $z \in \mathcal{I}$ if*

1. $f(z) < z < f(z^-) < \infty$
2. $f(f(z^-)) \leq f(f(z))$

3. *$f(x)$ is increasing for $x \in [f(z), z)$*
4. *$f(x)$ is increasing for $x \in [z, f(z^-))$.*

In Kozyakin (2003), this terminology was used for a smaller class of functions f. In particular, the domain and range were restricted to the interval $[0, 1)$ and there was a requirement that f is continuous on each of the intervals $[f(z), z)$ and $[f(z^-), f(z))$.

The following two Lemmas show that maps-with-gaps whose parts satisfy Assumption A2 lead to locally-growing relaxation functions.

Lemma 58 *Suppose ϕ_0, ϕ_1 satisfy A2 and $x \in [y_1, y_0]$. Then $\phi_{01}(x) < \phi_{10}(x)$.*

Proof Suppose that $x \in [y_1, y_0)$. Then using Lemma 33 gives

$$\begin{aligned} \phi_1(\phi_0(x)) - \phi_1(x) &< \phi_0(x) - x && \text{as } \phi_1 \text{ is contractive and } x < y_0 \\ \phi_0(x) - \phi_0(\phi_1(x)) &\leq x - \phi_1(x) && \text{as } \phi_0 \text{ is contractive and } x \geq y_1 \\ \phi_{01}(x) &< \phi_{10}(x) && \text{by adding these inequalities.} \end{aligned}$$

A symmetric argument holds if $x \in (y_1, y_0]$. This completes the proof. ■

Lemma 59 *Suppose ϕ_0, ϕ_1 satisfy A2, that $z \in (y_1, y_0)$ and let $f(x) := \phi_{\mathbf{1}_{x \geq z}}(x)$ for $x \in \mathcal{I}$. Then f is a locally-growing relaxation function with threshold z.*

Proof By definition $f(z^-) = \phi_0(z^-) = \phi_0(z)$, as ϕ_0 is continuous by Lemma 32. Also $f(z) = \phi_1(z)$, $f(f(z^-)) = \phi_{01}(z)$ and $f(f(z)) = \phi_{10}(z)$. So the conditions defining a locally-growing relaxation function read as follows:

1. $\phi_1(z) < z < \phi_0(z) < \infty$, which is true for $z \in (y_1, y_0)$ by Lemma 33
2. $\phi_{01}(z) < \phi_{10}(z)$, which is the result of Lemma 58
3. $\phi_0(x)$ is increasing for $x \in [\phi_1(z), z)$, which holds by A2
4. $\phi_1(x)$ is increasing for $x \in [z, \phi_0(z))$, which holds by A2.

This completes the proof. ■

The following definition is due to Morse and Hedlund (1940).

Definition 60 *An word w is 1-**balanced** if*

$$\left| \, |u|_1 - |v|_1 \right| \leq 1$$

for all factors u, v of w with $|u| = |v|$.

The next two Lemmas use an argument from Kozyakin (2003) to show that the itineraries of a locally-growing relaxation function are 1-balanced. We denote the fractional part of a real number x by $\{x\} := x - \lfloor x \rfloor$.

Lemma 61 *Suppose f is a locally-growing relaxation function with threshold z. Let*

$$F(x) := \frac{f(t\{x\} - t\mathbf{1}_{\{x\}\in[\alpha,1)} + z) - z}{t} + \lfloor x \rfloor + 1 \qquad \text{for } x \in \mathbb{R}$$

where $t := f(z^-) - f(z)$ and $\alpha := (f(z^-) - z)/t$. Then

F1. $F(0) \in [0, 1)$

F2. $F(x + 1) = F(x) + 1$ *for all* $x \in \mathbb{R}$

F3. F *is increasing*

F4. The itineraries of f and F agree in the sense that

$$f^{(n-1)}(x) \geq z \;\Leftrightarrow\; \left\{F^{(n-1)}\left(\frac{x-z}{t}\right)\right\} \in [0, F(0))$$

for all $n \in \mathbb{Z}_{++}$ and all $x \in [f(z), f(z^-))$.

Proof First, note that the function F is well-defined as $f(z) \in \mathbb{R}$ and $f(z) < z < f(z^-) < \infty$ so that $t \in \mathbb{R}_{++}$ and

$$\alpha = \frac{f(z^-) - z}{f(z^-) - f(z)} \in (0, 1). \tag{43}$$

It is easy to see that F1 and F2 hold. Indeed

$$F(0) = \frac{f(z) - z}{t} + 1 = \frac{f(z) - z + f(z^-) - f(z)}{f(z^-) - f(z)} = \alpha \in (0, 1)$$

and as the ratio in the definition of $F(x)$ only depends on $\{x\}$, we have

$$F(x + 1) - F(x) = \lfloor x + 1 \rfloor - \lfloor x \rfloor = 1 \qquad \text{for } x \in \mathbb{R}. \tag{44}$$

Now we show that F3 holds.

- If $x \in [0, \alpha)$ then $tx + z \in [z, f(z))$. But $f(\cdot)$ is increasing on $[z, f(z))$ and $t \in \mathbb{R}_{++}$. It follows that $F(x) = (f(tx + z) - z)/t$ is increasing for $x \in [0, \alpha)$.
- As $f(f(z^-)) \leq f(f(z))$ and f is increasing on $[z, f(z^-))$, it follows that

$$F(\alpha^-) = \frac{\lim_{x \uparrow f(x^-)} f(u) - z}{t} + 1 \leq \frac{f(f(z^-)) - z}{t} + 1 \leq \frac{f(f(z)) - z}{t} + 1 = F(\alpha).$$

- If $x \in [\alpha, 1)$ then $tx - t + z \in [f(z^-), z)$. But $f(\cdot)$ is increasing on $[f(z^-), z)$ and $t \in \mathbb{R}_{++}$. It follows that $F(x) = (f(tx - t + z) - z)/t$ is increasing for $x \in [\alpha, 1)$.
- The definition of F gives

$$F(1^-) = \frac{f(z^-) - z}{t} + 1 = \alpha + 1 = F(1).$$

In summary, $F(x)$ is increasing for $x \in [0,1]$. But as $F(x+1) = F(x)+1$ for $x \in \mathbb{R}$ we conclude that F is increasing for $x \in \mathbb{R}$. Therefore F3 holds.

Now, we note that for any $m \in \mathbb{Z}_+$ and any $x \in [\alpha-1, \alpha)$ we have

$$\frac{f^{(m)}(tx+z)-z}{t} \in [\alpha-1, \alpha). \tag{45}$$

Indeed, for $m=0$ we have $(f^{(m)}(tx+z)-z)/t = x \in [\alpha-1,\alpha)$. Also if $y \in [f(z), f(z^-))$ then the assumptions about f show that $f(y) \in [f(z), f(z^-))$, while if $x \in [\alpha-1,\alpha)$ then the definitions of t and α show that $tx+z \in [f(z), f(z^-))$. Thus $(f^{(m)}(tx+z)-z)/t \in [\alpha-1,\alpha)$.

Now, we show by induction that for any for any $x \in [\alpha-1,\alpha)$ and $n \in \mathbb{Z}_+$ we have

$$F^{(n)}(\{x\}) - \frac{f^{(n)}(tx+z)-z}{t} \in \mathbb{Z}. \tag{46}$$

In the base case $n=0$ we have $F^{(0)}(\{x\}) - (f^{(0)}(tx+z)-z)/t = \{x\} - ((tx+z)-z)/t \in \mathbb{Z}$. For the inductive step, let $p := F^{(n)}(\{x\})$ and $q := (f^{(n)}(tx+z)-z)/t$ and suppose that $p-q \in \mathbb{Z}$. Then we have

$$F^{(n+1)}(\{x\}) = F(p) = F(q) + p - q$$

as F satisfies F2 and by the assumption that $p-q \in \mathbb{Z}$. Furthermore, $q \in [\alpha-1,\alpha)$, by (45) so that for some $k \in \mathbb{Z}$,

$$\begin{aligned}
F^{(n+1)}(\{x\}) &= p - q + \begin{cases} \dfrac{f(tq+z)-z}{t} + 1 & \text{if } q \in [0,\alpha) \\ \dfrac{f(t(q+1)-t+z)-z}{t} & \text{if } q \in [\alpha-1, 0) \end{cases} \\
&= \frac{f(tq+z)-z}{t} + k \\
&= \frac{f(t\frac{f^{(n)}(tx+z)-z}{t} - z)}{t} + k \\
&= \frac{f^{(n+1)}(tx+z)-z}{t} + k.
\end{aligned}$$

Therefore (46) is true.

From (45) and (46) it follows that for any $x \in [\alpha-1,\alpha)$ and $n \in \mathbb{Z}_+$ we have

$$\begin{aligned}
\{F^{(n-1)}(\{x\})\} \in [0,\alpha) &\Leftrightarrow \left\{\frac{f^{(n-1)}(tx+z)-z}{t}\right\} \in [0,\alpha) \\
&\Leftrightarrow \frac{f^{(n-1)}(tx+z)-z}{t} \in [0,\alpha) \\
&\Leftrightarrow f^{(n-1)}(tx+z) \geq z.
\end{aligned} \tag{47}$$

Therefore F4 holds. This completes the proof. ∎

Lemma 62 *Suppose $F : \mathbb{R} \to \mathbb{R}$ satisfies Claims F1-F3 of Lemma 61. Let $\alpha := F(0)$ and for $x \in \mathbb{R}$ let $s(x)$ denote the infinite word with letters*

$$s_n(x) := \begin{cases} 1 & \text{if } \{F^{(n-1)}(x)\} \in [0, \alpha) \\ 0 & \text{if } \{F^{(n-1)}(x)\} \in [\alpha, 1) \end{cases}$$

for $n \in \mathbb{Z}_{++}$. Then $s(x)$ is 1-balanced.

Proof For any $n \in \mathbb{Z}_{++}$ and any $x \in [0, 1)$ we have

$$\begin{aligned}
&\lfloor F^{(n)}(x) - \alpha \rfloor - \lfloor F^{(n-1)}(x) - \alpha \rfloor \\
&\quad = \lfloor F(r + z) - \alpha \rfloor - \lfloor r + z - \alpha \rfloor && \text{for } r := \{F^{(n-1)}(x)\}, z := \lfloor F^{(n-1)}(x) \rfloor \\
&\quad = \lfloor F(r) + z - \alpha \rfloor - \lfloor r + z - \alpha \rfloor && \text{as } F(u) + 1 = F(u) \text{ for } u \in \mathbb{R} \\
&\quad = \lfloor F(r) - \alpha \rfloor - \lfloor r - \alpha \rfloor && \text{as } \lfloor r + z \rfloor = \lfloor r \rfloor + z \text{ since } z \in \mathbb{Z} \\
&\quad = -\lfloor r - \alpha \rfloor && \text{as } \alpha = F(0) \le F(r) < F(1) = \alpha + 1 \\
& && \text{since } F(\cdot) \text{ is increasing and } r \in [0, 1) \\
&\quad = \lceil \alpha - r \rceil \\
&\quad = \begin{cases} 1 & \text{if } r \in [0, \alpha) \\ 0 & \text{if } r \in [\alpha, 1) \end{cases} \\
&\quad = s_n(x).
\end{aligned}$$

Therefore, for any $m \in \mathbb{Z}_+$,

$$\sum_{k=1}^{m} s_k(x) = \sum_{k=1}^{m} \Big(\lfloor F^{(k)}(x) - \alpha \rfloor - \lfloor F^{(k-1)}(x) - \alpha \rfloor \Big) = \lfloor F^{(m)}(x) - \alpha \rfloor - \lfloor x - \alpha \rfloor. \tag{48}$$

Now we show by induction that for any $m \in \mathbb{Z}_+$,

$$F^{(m)}(z) - F^{(m)}(y) \in [0, 1) \qquad \text{for any } y, z \in \mathbb{R} \text{ with } z - y \in [0, 1). \tag{49}$$

The base case with $m = 0$ reads $z - y \in [0, 1)$ which is true. For the inductive step, let $z' := F^{(m-1)}(z)$ and $y' := F^{(m-1)}(y)$ for $m - 1 \in \mathbb{Z}_+$ and assume that $z' - y' \in [0, 1)$. Then $F(y') \le F(z')$ as $y' \le z'$ and F is increasing. Also $F(z') < F(y' + 1) = F(y') + 1$ as $z' - y' < 1$, as F is increasing and as $F(u + 1) = F(u) + 1$ for any $u \in \mathbb{R}$. Therefore $F^{(m)}(z) - F^{(m)}(y) = F(z') - F(y') \in [0, 1)$.

For any choice of $m \in \mathbb{Z}_+$ and $x, y \in [0, 1)$, we wish to prove that $|\Delta| \le 1$ for

$$\begin{aligned}
\Delta &:= |s_1(x)s_2(x) \ldots s_m(x)|_1 - |s_1(y)s_2(y) \ldots s_m(y)|_1 \\
&= \sum_{k=1}^{m} (s_k(x) - s_k(y)) = \underbrace{\Big(\lfloor F^{(m)}(x) - \alpha \rfloor - \lfloor F^{(m)}(y) - \alpha \rfloor \Big)}_{=:A} - \underbrace{(\lfloor x - \alpha \rfloor - \lfloor y - \alpha \rfloor)}_{=:B}
\end{aligned}$$

using equation (48). But if $x \ge y$ then $x - y \in [0, 1)$ so (49) shows that A and B are both of the form $\lfloor a \rfloor - \lfloor b \rfloor$ for some $a - b \in [0, 1)$ and it follows that both A and B are in $\{0, 1\}$.

Whereas if $x \le y$ then $y - x \in [0, 1)$ so both A and B are in $\{-1, 0\}$. We conclude that $|\Delta| = |A - B| \le |1 - 0| = 1$. ■

The following is Theorem 3.1 of Dulucq and Gouyou-Beauchamps (1990) and provides the missing link between 1-balanced words and mechanical words.

Lemma 63 *Suppose w is a word of length $n \in \mathbb{Z}_{++}$ with $|w|_0 |w|_1 > 0$. Then w is 1-balanced if and only if $w_k = \left\lfloor \frac{pk+r}{q} \right\rfloor - \left\lfloor \frac{p(k-1)+r}{q} \right\rfloor$ for $k = 1, 2, \ldots, n$, where $p, q, r \in \mathbb{Z}$ satisfy $0 < p < q \le n$, $0 \le r < q \le n$ and $\gcd(p, q) = 1$.*

Now we are ready to describe the itineraries of maps-with-gaps, for arbitrary initial states and thresholds.

Lemma 64 *Suppose ϕ_0, ϕ_1 satisfy A2 and that $n \in \mathbb{Z}_{++}$, $x \in \mathcal{I}$, $s \in \overline{\mathbb{R}}$ and $a \in \{0, 1\}$. Then*

$$\sigma(x|s^-)_{1:n} = l^m w$$

for some $l \in \{0, 1\}$, some $m \in \{0, 1, \ldots, n\}$, and some factor w of a lower mechanical word.

Proof First, note that if $a, b \in \mathcal{I}$, then Lemma 35 shows that

$$b > y_1 \Rightarrow \phi_{1^m}(a) < b \text{ for some } m \in \mathbb{Z}_{++} \tag{50}$$

$$b \le y_0 \Rightarrow \phi_{0^m}(a) \ge b \text{ for some } m \in \mathbb{Z}_{++}. \tag{51}$$

We consider seven cases.

1. Say $s \le y_1$ and $x \ge s$. Then $\sigma(x|s^-) = 1\sigma(\phi_1(x)|s^-)$. But if $x > y_1$ then Lemma 33 gives $\phi_1(x) > y_1 \ge s$, whereas if $x \le y_1$ then $\phi_1(x) \ge x \ge s$. In both cases, $\sigma(\phi_1(x)|s^-) = 1\sigma(\phi_{11}(x)|s^-)$. Repeating this argument gives $\sigma(x|s^-) = 1^\infty$, so the claim is true.

2. Say $s \le y_1$ and $x < s$. Then $\sigma(x|s^-)$ begins with 0, and as $s < y_0$ it follows from (51) that there is a least $m \in \mathbb{Z}_{++}$ with $\phi_{0^m}(x) \ge s$. Thus $\sigma(x|s^-) = 0^m \sigma(\phi_{0^m}(x)|s^-) = 0^m 1^\infty$ by Case 1, so the claim is true.

3. Say $s \in (y_0, y_1)$ and $x \in [\phi_1(s), \phi_0(x))$. As ϕ_0, ϕ_1 satisfy A2, Lemmas 59, 61 and 62 together show that $\sigma(x|s^-)$ is a 1-balanced word. Thus Lemma 63 shows that $\sigma(x|s^-)_{1:n} = w$ for some factor w of a lower mechanical word, so the claim is true.

4. Say $s \in (y_0, y_1)$ and $x < \phi_1(s)$. Then Lemma 33 shows that $\phi_1(s) < s$. Thus $x < s$, so $\sigma(x|s^-)$ begins with 0, and $\phi_1(s) < y_0$, so (51) shows that there is a least $m \in \mathbb{Z}_{++}$ with $\phi_{0^m}(x) \ge \phi_1(s)$. Thus $\sigma(x|s^-) = 0^m \sigma(\phi_{0^m}(x)|s^-)$ where $\sigma(\phi_{0^m}(x)|s^-)_{1:n} = w$ for some factor w of a lower mechanical word, by Case 3, so the claim is true.

5. Say $s \in (y_0, y_1)$ and $x \ge \phi_0(s)$. Then arguing as in Case 4, but using (50), shows that $\sigma(x|s^-)_{1:n} = 1^m w$ for some $m \in \mathbb{Z}_{++}$ and some factor w of a lower mechanical word.

6. Say $s \geq y_0$ and $x < s$. Then arguing as in Case 1 shows that $\sigma(x|s^-) = 0^\infty$.
7. Say $s \geq y_0$ and $x \geq s$. Then arguing as in Case 2 shows that $\sigma(x|s^-) = 1^m 0^\infty$ for some $m \in \mathbb{Z}_{++}$.

This completes the proof. ■

The following is an analogue of Lemma 52 in which only the threshold varies.

Lemma 65 *Suppose ϕ_0, ϕ_1 satisfy A2 and $x \in \mathcal{I}$. Then $\sigma(x|s^-)$ is a lexicographically non-increasing function of $s \in \overline{\mathbb{R}}$.*

Proof If $s, t \in \overline{\mathbb{R}}$ and $\sigma(x|s^-) \succ \sigma(x|t^-)$, then for some finite word u and words v, w,

$$\sigma(x|s^-) = u1v, \qquad\qquad \sigma(x|t^-) = u0w.$$

So the definition of the itinerary gives $t > \phi_u(x) \geq s$. This completes the proof. ■

The following is Theorem 17 and part of Corollary 18 of Mignosi (1991).

Lemma 66 *Let $\mathcal{A}_n$ be the set of factors of length $n \in \mathbb{Z}_{++}$ of lower mechanical words. Then*

$$\operatorname{card}(\mathcal{A}_n) = 1 + \sum_{i=1}^{n} (n - i + 1) \operatorname{EulerTotient}(i) = \frac{2n^3}{\ell^2} + O(n^2 \log n).$$

Now we are ready to bound the number of discontinuities of the itinerary of a map-with-a-gap as a function of its threshold.

Lemma 67 *Suppose ϕ_0, ϕ_1 satisfy A2, that $t \in \mathbb{Z}_+$, $x \in \mathcal{I}$ and $a \in \{0, 1\}$. Then the mapping $s \mapsto A_{1:t}(x, a; s)$ for $s \in \overline{\mathbb{R}}$ has at most a polynomial number $p(t)$ of discontinuities.*

Proof Let $\mathcal{A}_n$ be the set of all factors of length n of lower mechanical words. Let $\mathcal{F}_t$ be the set of all words of the form $l^m w$ for some $l \in \{0, 1\}$, some $m \in \{0, 1, \ldots, n\}$ and some factor w of a lower mechanical word. By Lemma 64, the word $A_{1:t}(x, a; s)$ is in $\mathcal{F}_t$. Also, Lemma 65, shows that the mapping $s \mapsto A_{1:t}(x, a; s)$ is lexicographically non-increasing. Thus, the number of discontinuities of this mapping is at most

$$\operatorname{card}(\mathcal{F}_t) = \operatorname{card}\left(\{l^m w : l \in \{0, 1\}, m \in \mathbb{Z}_+, m \leq t, w \in \mathcal{A}_{t-m}\}\right) = 2 \sum_{m=0}^{t} \operatorname{card}(\mathcal{A}_{t-m}).$$

Finally, Lemma 66 shows the right-hand side is $O(t^4)$. This completes the proof. ■

Proof [Proof of Theorem 14.] The theorem simply couples together Lemmas 64, 65 and 67. ■

Appendix B. Proof of Lemma 22

Demonstrating Lemma 22 by combining results in Marshall et al. (2010) requires as much text as a direct proof.

Proof Let $\mathcal{X}$ be the set of sequences X with components $X_k = \sum_{i=1}^{k} x_i$ for $k = 1, 2, \ldots, n$ where $x_1, x_2, \ldots, x_n$ is a non-decreasing sequence on $\mathbb{R}_{++}$. Let $g : \mathcal{X} \to \mathbb{R}$ be the function

$$g(X) := f_1(X_1) + \sum_{i=2}^{n} f_i(X_i - X_{i-1}).$$

Let $f_i'(\cdot)$ denote the (sub)gradient of $f_i(\cdot)$. For $i = 1, 2, \ldots, n-1$, the (sub)gradients of $g(\cdot)$ are

$$\frac{\partial g(X)}{\partial X_i} = f_i'(x_i) - f_{i+1}'(x_{i+1}) \leq f_{i+1}'(x_i) - f_{i+1}'(x_{i+1}) \leq 0$$

where the first inequality holds as $f_i'(x) \leq f_{i+1}'(x)$ for $x \in \mathbb{R}_{++}$ (by Hypothesis 4) and the second holds as $x_i \leq x_{i+1}$ and $f_{i+1}(\cdot)$ is convex (by Hypothesis 3). Also,

$$\frac{\partial g(X)}{\partial X_n} = f_n'(x_n) \leq 0$$

as $f_n(\cdot)$ is non-increasing (by Hypothesis 3). Therefore $g(\cdot)$ is non-increasing in all of its arguments. As the sequences A, B with components $A_k := \sum_{i=1}^{k} a_i, B_k := \sum_{i=1}^{k} b_k$ are in $\mathcal{X}$ (by Hypothesis 1) and $A_k \leq B_k$ for $k = 1, 2, \ldots, n$ (by Hypothesis 2), it follows that $g(A) \geq g(B)$. So the definition of $g(\cdot)$ gives

$$\sum_{i=1}^{n} f(a_i) = g(A) \geq g(B) = \sum_{i=1}^{n} f(b_i)$$

as claimed. ■

Appendix C. Proof of Lemma 24

We start by recalling the definition of the matrix $M(w)$ corresponding to a given finite word w, which corresponds to the composition of Kalman-Filter variance updates, and introducing some related matrices $K, S(w)$ and X. We then prove Claims 1 to 5 of Lemma 24 in turn.

Definition 68 *Let I be the 2-by-2 identity matrix. For $r \in (0, 1]$ and $0 \leq a \leq b$, let*

$$F := \begin{pmatrix} r & 1/r \\ ar & (a+1)/r \end{pmatrix}, \qquad G := \begin{pmatrix} r & 1/r \\ br & (b+1)/r \end{pmatrix}, \qquad K := \begin{pmatrix} r & 1/r \\ r - r^3 & -r \end{pmatrix}.$$

Let $M(\epsilon) = I, M(0) = F, M(1) = G$ and for any finite word w let

$$M(w) = M(w_{|w|}) \cdots M(w_2) M(w_1), \qquad S(w) = \sum_{i=0}^{|w|} M(w_{1:i}).$$

Let

$$X := \begin{pmatrix} -r/(1-r^2) & 0 \\ 0 & 1/r \end{pmatrix}.$$

Remark 69 *We use the the following facts repeatedly without mention. Clearly* $\det(F) = \det(G) = 1$*, so that* $\det(M(w)) = 1$ *for any word* w*. Also,* $KF = F^{-1}K, KG = G^{-1}K$ *and* $K^2 = I$*. Thus for* $A \in \{KF, KG, K\}$ *we have* $A^2 = I$*, so* A *is an involutory matrix. Thus* $KM(w)^{-1}K = M(w^R)$*, where* w^R *denotes the reverse* $w_n \dots w_2 w_1$ *of a word* $w = w_1 w_2 \dots w_n$.

Notation. For a vector $v \in \mathbb{R}^m$ where $m \in \mathbb{Z}_{++}$, we write $v > 0$ if $v_i > 0$ for $i = 1, 2, \dots, m$ and $v \geq 0$ if $v_i \geq 0$ for $i = 1, 2, \dots, m$. Similarly, for a matrix $P \in \mathbb{R}^{m \times n}$ where $m, n \in \mathbb{Z}_{++}$, we write $P > 0$ ($P \geq 0$) if $P_{ij} > 0$ ($P_{ij} \geq 0$) for $i = 1, 2, \dots, m$ and $j = 1, 2, \dots, n$.

C.1. Claim 1

We require one simple Lemma.

Lemma 70 *Suppose* $a \in \mathbb{R}_+$ *and* $r \in (0, 1]$*. Then the fixed point* y_0 *satisfies*

$$y_0 \leq \frac{1}{1-r^2}.$$

Proof As the positive root of

$$y_0 = \frac{r^2 y_0 + 1}{a r^2 y_0 + 1 + a}$$

is a decreasing function of a for $a \in \mathbb{R}_+$, setting $a = 0$ gives an upper bound. This upper bound u satisfies $u = r^2 u + 1$, so that $u = 1/(1-r^2)$. ■

Proof [Proof of Claim 1 of Lemma 24.] We prove the claim for x satisfying

$$\phi_p(0) \leq x \leq \frac{1}{r - r^2},$$

noting that

$$\phi_p\left(\frac{1}{1-r^2}\right) \leq \frac{1}{1-r^2} \leq \frac{1}{r-r^2}$$

where the first inequality follows from Lemma 33 (in Appendix A) as Lemma 70 gives $1/(1-r^2) \geq y_0 \geq y_p$, and the second inequality holds as $r \in (0, 1]$.

For any word w and for $k = 1, 2, \dots, |w|$, let

$$\begin{pmatrix} u_k \\ v_k \end{pmatrix} := M(w_{1:k}) \begin{pmatrix} x \\ 1 \end{pmatrix}.$$

Clearly u_k, v_k are positive as $x \geq \phi_p(0) \geq 0$ and

$$M(w_{1:k}) \begin{pmatrix} x \\ 1 \end{pmatrix} \geq \begin{pmatrix} r & \frac{1}{r} \\ 0 & \frac{1}{r} \end{pmatrix}^k \begin{pmatrix} 0 \\ 1 \end{pmatrix} > 0.$$

Now, for any $0 \leq z \leq 1/(r - r^2)$ and $H = M(w_k)$, noting that $a, b, r \geq 0$ and $r \leq 1$ gives

$$\frac{H_{11}z + H_{12}}{H_{21}z + H_{22}} \leq r^2 z + 1 \leq \frac{r^2}{r - r^2} + 1 \leq \frac{1}{r - r^2}.$$

For $k = 1, 2, \ldots, |w|$, induction using this inequality proves that

$$\frac{u_k}{v_k} \leq \frac{1}{r - r^2}.$$

(For $k = 1$, put $z = x \leq 1/(r - r^2)$. For $k > 1$, assume that $z = u_{k-1}/v_{k-1} \leq 1/(r - r^2)$.) Thus

$$\begin{pmatrix} u_{k+1} - u_k \\ v_{k+1} - v_k \end{pmatrix} = (M(w_{k+1}) - I) \begin{pmatrix} \frac{u_k}{v_k} \\ 1 \end{pmatrix} v_k \geq \begin{pmatrix} r - 1 & \frac{1}{r} \\ 0 & \frac{1}{r} - 1 \end{pmatrix} \begin{pmatrix} \frac{1}{r - r^2} \\ 1 \end{pmatrix} v_k \geq 0.$$

But both $a_k(x), b_k(x)$ are of the form u_k and $c_k(x), d_k(x)$ are of the form v_k for appropriate w. Thus $a_{1:m}(x), b_{1:m}(x), c_{1:m}(x)$ and $d_{1:m}(x)$ are non-decreasing and positive. ■

C.2. Claim 2

To prove Claim 2 we need two simple Lemmas.

Lemma 71 *If w is a word, then $M(w) - M(w^R) = tr(KM(w))K$.*

Proof For any $C \in \mathbb{R}^{2\times 2}$, direct calculation gives $C - K\text{adj}(C)K = \text{tr}(KC)K$. But $\det(M(w)) = 1$, so $M(w) - M(w^R) = M(w) - KM(w)^{-1}K = M(w) - K\text{adj}(M(w))K = \text{tr}(KM(w))K$. ■

Lemma 72 *Suppose p is a palindrome, $r \in (0, 1]$ and $n \in \mathbb{Z}_+$. Then for some $x \geq 0$,*

$$M((10p)^n 10) - M((01p)^n 01) = xK.$$

Proof For any matrices $P, Q \in \mathbb{R}^{2\times 2}$ with $\det(P) = 1$ and $Q \geq 0$, direct calculation gives

$$\begin{aligned} &[QGFP]_{22}[QFGP]_{21} - [QFGP]_{22}[QGFP]_{21} \\ &\quad = (b - a)\det(P)((1 - r^2 + a + b + ab)Q_{22}^2 + (2 + a + b)Q_{22}Q_{21} + Q_{21}^2) \geq 0 \end{aligned}$$

as $1 \geq r^2, b \geq a \geq 0$. Therefore, for any words w, w',

$$\frac{M(w01w')_{22}}{M(w01w')_{21}} \geq \frac{M(w10w')_{22}}{M(w10w')_{21}}. \tag{52}$$

Let $A = M((10p)^n10), B = M((01p)^n01)$. Repeated application of (52) gives

$$\frac{B_{22}}{B_{21}} = \frac{M(01p01p\cdots01)_{22}}{M(01p01p\cdots01)_{21}} \geq \frac{M(10p01p\cdots01)_{22}}{M(10p01p\cdots01)_{21}} \geq \frac{M(10p10p\cdots10)_{22}}{M(10p10p\cdots10)_{21}} = \frac{A_{22}}{A_{21}}.$$

As $A, B \geq 0$ and Lemma 71 gives $A = B + xK$ for some $x \in \mathbb{R}$, it follows that

$$(B + xK)_{21}B_{22} \geq (B + xK)_{22}B_{21} \quad \Rightarrow \quad K_{21}B_{22}x \geq K_{22}B_{21}x.$$

Finally, the fact that $K_{22} \leq 0 < K_{21}$ and $B \geq 0$ gives $x \geq 0$. ■

Proof [Proof of Claim 2 of Lemma 24.] For some $t \geq 0$, Lemma 72 gives

$$\begin{aligned}(d/dx)(b_1(x) - a_1(x)) &= [M((10p)^n1) - M((01p)^n0)]_{11} \\ &= [(F^{-1} - G^{-1})M((01p)^n01) + tF^{-1}K]_{11} \\ &= [(0,0)M((01p)^n01)]_1 + t[KF]_{11} \\ &= t\,(rF_{11} + (1/r)F_{21}) \\ &\geq 0\end{aligned}$$

and

$$\begin{aligned}b_1(\phi_p(0)) - a_1(\phi_p(0)) &= [M(p(10p)^n1) - M(p(01p)^n0)]_{12} \\ &= [(F^{-1} - G^{-1})M(p(01p)^n01) + tF^{-1}KM(p)]_{12} \\ &= t[KFM(p)]_{12} \\ &= t(rM(p0)_{12} + (1/r)M(p0)_{22}) \\ &\geq 0.\end{aligned}$$

Also, if $k > 1$ and $w = p_{1:(k-2)}$, then

$$\begin{aligned}(d/dx)(b_k(x) - a_k(x)) &= [M((10p)^n10w) - M((01p)^n01w)]_{11} \\ &= t[M(w)K]_{11} \\ &= t(M(w)_{11}r + M(w)_{12}r(1 - r^2)) \\ &\geq 0\end{aligned}$$

and as p is a palindrome, $p = sw^R$ for some word s, so

$$\begin{aligned}b_k(\phi_p(0)) - a_k(\phi_p(0)) &= [M(p(10p)^n10w) - M(p(01p)^n01w)]_{12} \\ &= t[M(w)KM(p)]_{12} \\ &= t[KM(w^R)^{-1}M(w^R)M(s)]_{12} \\ &= t[KM(s)]_{12} \\ &= t(rM(s)_{12} + (1/r)M(s)_{22}) \\ &\geq 0.\end{aligned}$$

This completes the proof. ■

C.3. Claim 3

Claim 3 of Lemma 24 is more challenging than Claims 1 and 2. We begin with six Lemmas.

Lemma 73 *Suppose p is any palindrome and $r \in (0,1]$. Then for any $k \in \mathbb{Z}_+$,*

$$\Delta_k := [M(((10p)^\infty)_{1:k}) - M(((01p)^\infty)_{1:k})]_{21} \geq 0.$$

Proof If $k = 1$ then

$$\begin{aligned}\Delta_k &= [M(1) - M(0)]_{21} \\ &= [G - F]_{21} \\ &= (b-a)r \\ &\geq 0.\end{aligned}$$

If $k = (n+1)|01p| + 1$ for some $n \in \mathbb{Z}_+$ then Lemma 72 shows there is an $x \geq 0$ such that

$$\begin{aligned}\Delta_k &= [M((10p)^n 10p1) - M((01p)^n 01p0)]_{21} \\ &= [M(p1)(xK + M((01p)^n 01)) - M(p0)M((01p)^n 01)]_{21} \\ &= [xM(p1)K + (G-F)M((01p)^{n+1})]_{21} \\ &= x(M(p1)_{21} r + M(p1)_{22}(r - r^3)) \\ &\qquad + (b-a)(rM((01p)^{n+1})_{12} + (1/r)M((01p)^{n+1})_{22}) \\ &\geq 0.\end{aligned}$$

Otherwise, there is a prefix w of p and an $n \in \mathbb{Z}_+$ such that

$$\begin{aligned}\Delta_k &= [M((10p)^n 10w) - M((01p)^n 01w)]_{21} \\ &= [M(w)(M((10p)^n 10) - M((01p)^n 01))]_{21} \\ &= [xM(w)K]_{21} \qquad \text{for some } x \geq 0 \text{ by Lemma 72} \\ &= x(M(w)_{21} r + M(w)_{22}(r - r^3)) \\ &\geq 0.\end{aligned}$$

This completes the proof. ∎

Lemma 74 *Suppose $p = ws$ is a palindrome, $n \in \mathbb{Z}_+$ and $r \in (0,1]$. Then*

$$[M(p(10p)^n 10w) - M(p(01p)^n 01w)]_{22} \leq 0.$$

Proof First note that for any finite word u,

$$M(u)_{22} \geq M(u)_{21}. \tag{53}$$

Indeed, if $u = \epsilon$ then $M(\epsilon) = I$ so the inequality holds. Otherwise, for some $c \in \{a, b\}$

$$M(u)_{22} - M(u)_{21} = \left[M(u_{2:|u|})M(u_1)\begin{pmatrix}-1 \\ 1\end{pmatrix}\right]_2 = \left[M(u_{2:|u|})\frac{1}{r}\begin{pmatrix}1 - r^2 \\ 1 + (1-r^2)c\end{pmatrix}\right]_2 \geq 0$$

as the definition of $M(\cdot)$ assumes that $0 \leq a < b$ so that $c \geq 0$, and as $r \in (0,1]$ and $M(u_{2:|u|}) \geq 0$.

Also, by Lemma 72, for some $x \geq 0$

$$\begin{aligned}
[M(p(10p)^n10w) - M(p(01p)^n01w)]_{22} &= [M(s)^{-1}(M((10p)^{n+1}) - M((01p)^{n+1}))M(p)]_{22} \\
&= x[M(s)^{-1}KM(p)]_{22} \\
&= x[KM(ps^R)]_{22} \\
&= xr\left((1-r^2)M(ps^R)_{21} - M(ps^R)_{22}\right) \\
&\leq xr\left((1-r^2)M(ps^R)_{22} - M(ps^R)_{22}\right) \\
&\leq 0
\end{aligned}$$

where the penultimate line is (53). ■

Lemma 75 *Suppose p is a palindrome, $n \in \mathbb{Z}_+$ and $r \in (0,1]$. Then*

$$\left[\Big(M((10p)^n10)X - M((01p)^n01)X + M((10p)^n1) - M((01p)^n0)\Big)M(p)\right]_{22} = 0.$$

Proof Let $P = M(p)$. Solving $KP = P^{-1}K$ shows that there exist $f, h \in \mathbb{R}$ such that

$$P = \begin{pmatrix} \dfrac{1 - f^2r^2 + fhr^2 + f^2r^4}{fr^2+h} & f \\ \dfrac{-1 - fh + h^2 + fhr^2}{fr^2+h}r^2 & h \end{pmatrix}.$$

Directly substituting this expression into

$$Q_n := \left[\Big(FG(PFG)^nX - GF(PGF)^nX + G(PFG)^n - F(PGF)^n\Big)P\right]_{22}$$

shows that $Q_0 = Q_1 = 0$. (Showing that $Q_1 = 0$ directly is algebraically tedious and we had to check this with computer algebra. The authors would be interested in a short demonstration that $Q_1 = 0$ as this may give insight into related problems.)

Lemma 71, then the fact that $\mathrm{tr}(K) = 0$, then the cyclic property of the trace give

$$\mathrm{tr}(PFG) = \mathrm{tr}(GFP + \mathrm{tr}(KPFG)K) = \mathrm{tr}(GFP) = \mathrm{tr}(PGF).$$

So PFG and PGF are 2-by-2 matrices with unit determinant whose traces are equal. For some matrices of eigenvectors U, V and some eigenvalue $\lambda \geq 1$, such matrices may be written in the form

$$PFG = U\Lambda U^{-1}, \qquad PGF = V\Lambda V^{-1}, \qquad \Lambda := \begin{pmatrix} \lambda & 0 \\ 0 & 1/\lambda \end{pmatrix}.$$

Thus, for some $\alpha, \beta \in \mathbb{R}$,

$$Q_n = \alpha\lambda^n + \beta\lambda^{-n}$$

But $Q_0 = 0$ implies that $\beta = -\alpha$. Thus $Q_1 = \alpha(\lambda - 1/\lambda) = 0$ implies that either $\alpha = 0$ or $\lambda = 1$. In either case, $Q_n = \alpha(\lambda^n - \lambda^{-n}) = 0$ for all $n \geq 0$. ■

Lemma 76 *Suppose p is a palindrome. Then $S(p)KM(p) = KS(p)$.*

Proof We proceed by induction on the length of p. In the base case, $S(\epsilon)KM(\epsilon) = IKI = KS(\epsilon)$ and if a is a letter then $S(a)KM(a) = (I + M(a))KM(a) = K(I + M(a)^{-1})M(a) = KS(a)$. Otherwise, say $p = aqa$ where a is a letter and $S(q)KM(q) = KS(q)$. Then

$$\begin{aligned}
S(p)KM(p) &= (I + M(a) + S(q)M(a) + M(aqa))KM(aqa) \\
&= (I + M(a) + M(aqa))KM(aqa) + S(q)KM(a)^{-1}M(a)M(q)M(a) \\
&= K(I + M(a)^{-1} + M(aqa)^{-1})M(aqa) + S(q)KM(q)M(a) \\
&= K(M(aqa) + M(aq) + I) + KS(q)M(a) \\
&= KS(p).
\end{aligned}$$

■

Lemma 77 *Suppose w is a finite word and $r \in (0, 1]$. Then $[K(S(w) - XM(w))]_{22} \geq 0$.*

Proof We proceed by induction on the length of w. In the base case,

$$[K(S(\epsilon) - XM(\epsilon))]_{22} = [K(I - X)]_{22} = 1 - r \geq 0.$$

Otherwise, say $w = ul$ where $|l| = 1$, $|u| < \infty$ and $[K(S(u) - XM(u))]_{22} \geq 0$. Then

$$\begin{aligned}
[K(S(w) - XM(w))]_{22} &= [K(S(u) + M(w) - X(M(w) - M(u)) - XM(u))]_{22} \\
&= [K(S(u) - XM(u))]_{22} + [K(I - X(I - M(l)^{-1}))M(w)]_{22} \\
&\geq [K(I - X(I - M(l)^{-1}))M(w)]_{22} \\
&= (1 - r)(r^2 M(w)_{12} + M(w)_{22}) \\
&\geq 0
\end{aligned}$$

where in the penultimate line we substituted the definitions of K and X, noting that $M(l) \in \{F, G\}$, and where the final inequality follows as $r \leq 1$ and $M(w) \geq 0$. ■

Lemma 78 *Suppose p is a palindrome and $r \in (0, 1]$. Then for any $n \in \mathbb{Z}_+$*

$$[(S(10p) - I)M(p(10p)^n) - (S(01p) - I)M(p(01p)^n)]_{22} \geq 0.$$

Proof Let $P := M(p)$. Then

$$\begin{aligned}
&[(S(10p) - I)M(p(10p)^n) - (S(01p) - I)M(p(01p)^n)]_{22} \\
&= [S(p)(FG(PFG)^n - GF(PGF)^n)P + (G(PFG)^n - F(PGF)^n)P]_{22} \\
&= [S(p)(FG(PFG)^n - GF(PGF)^n)P - (FG(PFG)^n - GF(PGF)^n)XP]_{22} \\
&= [S(p)xKP - xKXP]_{22} \\
&= [xK(S(p) - X)P]_{22}
\end{aligned}$$

$$\geq 0$$

where the second equality uses Lemma 75, the third holds for some $x \geq 0$ by Lemma 72, the fourth follows from Lemma 76 and the final inequality is Lemma 77. ■

Proof [Proof of Claim 3 of Lemma 24.] Say $1 \leq k \leq m$. Using Lemmas 73, 74 and 78 successively gives

$$\begin{aligned}
\sum_{i=1}^{k}(d_i(x) - c_i(x)) &= \sum_{i=1}^{k}[M((10p)^n(10p)_{1:i}) - M((01p)^n(01p)_{1:i})]_{21}x \\
&\quad + \sum_{i=1}^{k}[M((10p)^n(10p)_{1:i}) - M((01p)^n(01p)_{1:i})]_{22} \\
&\geq \sum_{i=1}^{k}[M((10p)^n(10p)_{1:i}) - M((01p)^n(01p)_{1:i})]_{21}\frac{M(p)_{12}}{M(p)_{22}} \\
&\quad + \sum_{i=1}^{k}[M((10p)^n(10p)_{1:i}) - M((01p)^n(01p)_{1:i})]_{22} \\
&= \frac{1}{M(p)_{22}}\sum_{i=1}^{k}[M(p(10p)^n(10p)_{1:i}) - M(p(01p)^n(01p)_{1:i})]_{22} \\
&\geq \frac{1}{M(p)_{22}}\sum_{i=1}^{m}[M(p(10p)^n(10p)_{1:i}) - M(p(01p)^n(01p)_{1:i})]_{22} \\
&= \frac{1}{M(p)_{22}}[(S(10p) - I)M(p(10p)^n) - (S(01p) - I)M(p(01p)^n)]_{22} \\
&\geq 0.
\end{aligned}$$

This completes the proof. ■

C.4. Claim 4

The proof of Claim 4 requires only one preparatory Lemma.

Lemma 79 *Suppose w is any finite word. Then*

$$[M(w)^{-1}]_{21} \leq 0 \leq [M(w)^{-1}]_{22}.$$

Proof We use induction on the length of w. In the base case, $M(\epsilon) = I$, for which

$$[I^{-1}]_{21} = 0 \leq 1 = [I^{-1}]_{22}.$$

Otherwise, suppose $w = 0u$ (the case $w = 1u$ is similar), let $U := M(u)^{-1}$ and assume that

$$U_{21} \leq 0 \leq U_{22}.$$

Then the induction assumption and and fact that $a, r \geq 0$ give

$$[M(w)^{-1}]_{21} = [UF^{-1}]_{21} = \frac{1}{r}((1+a)U_{21} - ar^2 U_{22}) \leq 0$$
$$[M(w)^{-1}]_{22} = [UF^{-1}]_{22} = \frac{1}{r}(U_{22}r^2 - U_{21}) \geq 0.$$

This completes the proof. ■

Proof [Proof of Claim 4 of Lemma 24.] Claim 3 of Lemma 24 applied for $k = 1$ shows that

$$c_1(x) \leq d_1(x).$$

For $k = 2, 3, \ldots, m$, as Lemma 73 shows that $c_k(x) - d_k(x)$ is a decreasing function of x, it suffices to prove that

$$\left[M((01p)^n (01p)_{1:k}) \begin{pmatrix} \phi_p\left(\frac{1}{1-r^2}\right) \\ 1 \end{pmatrix} \right]_2 \geq \left[M((10p)^n (10p)_{1:k}) \begin{pmatrix} \phi_p\left(\frac{1}{1-r^2}\right) \\ 1 \end{pmatrix} \right]_2 .$$

The left-hand side minus the right-hand side, up to a positive factor, is

$$\begin{aligned}
&\left[(M((01p)^n (01p)_{1:k}) - M((10p)^n (10p)_{1:k}))\, M(p) \begin{pmatrix} 1 \\ 1 - r^2 \end{pmatrix} \right]_2 \\
&\quad = -z \left[M(p_{1:(k-2)}) K M(p) \begin{pmatrix} 1 \\ 1 - r^2 \end{pmatrix} \right]_2 \quad \text{for some } z \geq 0 \text{ by Lemma 72} \\
&\quad = -z \left[M(p_{1:(|p|-k+2)})^{-1} K \begin{pmatrix} 1 \\ 1 - r^2 \end{pmatrix} \right]_2 \\
&\quad = -z \left[M(p_{1:(|p|-k+2)})^{-1} \begin{pmatrix} \frac{1}{r} \\ 0 \end{pmatrix} \right]_2 \\
&\quad = -\frac{z}{r} [M(p_{1:(|p|-k+2)})^{-1}]_{21} \\
&\quad \geq 0
\end{aligned}$$

where the last line is Lemma 79. Therefore

$$c_k(x) \geq d_k(x) \quad \text{for } k = 2, 3, \ldots, m \text{ and } x \leq \phi_p\left(\frac{1}{1-r^2}\right).$$

This completes the proof. ■

C.5. Claim 5

Proof [Proof of Claim 5 of Lemma 24.] We show that

$$\phi_w(0) \leq y_{01w} < y_{10w} \leq \phi_w\left(\frac{1}{1-r^2}\right)$$

for any finite word w (not just for palindromes p).

The first inequality follows as

$$\phi_w(0) \le \phi_w(\phi_{01}(0)) = \phi_{01w}(0) \le y_{01w}$$

as $0 \le \phi_{01}(0)$, as ϕ_w is increasing, as $0 \le y_{01w}$ and by Lemma 33 (in Appendix A).

The second inequality holds as $\phi_{01}(x) < \phi_{10}(x)$ for $x \in \mathbb{R}_+$. Thus

$$y_{01w} = \phi_w(\phi_{01}(y_{01w})) < \phi_w(\phi_{10}(y_{01w})) = \phi_{10w}(y_{01w})$$

so applying Lemma 33 gives

$$y_{01w} < y_{10w}.$$

The third inequality holds as

$$y_{10w} = \phi_w(y_{w10}) < \phi_w(y_0) \le \phi_w\left(\frac{1}{1-r^2}\right)$$

by definition of y_{10w}, as $y_{w10} < y_0$, as ϕ_w is increasing and by Lemma 70.

This completes the proof. ■

References

Kazuyuki Aihara and Hideyuki Suzuki. Theory of hybrid dynamical systems and its applications to biological and medical systems. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 368(1930):4893–4914, 2010.

Eitan Altman and Shaler Stidham Jr. Optimality of monotonic policies for two-action Markovian decision processes, with applications to control of queues with delayed information. *Queueing Systems*, 21(3-4):267–291, 1995.

Eitan Altman, Bruno Gaujal, and Arie Hordijk. Multimodularity, convexity, and optimization properties. *Mathematics of Operations Research*, 25(2):324–347, 2000.

Michael Athans. On the determination of optimal costly measurement strategies for linear stochastic systems. *Automatica*, 8(4):397–412, 1972.

Ashwinkumar Badanidiyuru, Baharan Mirzasoleiman, Amin Karbasi, and Andreas Krause. Streaming submodular maximization: Massive data summarization on the fly. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 671–680, 2014.

Mario Bernardo, Chris Budd, Alan R. Champneys, and Piotr Kowalczyk. *Piecewise-Smooth Dynamical Systems: Theory and Applications.* Springer Science and Business Media, 2008.

Jean Berstel, Aaron Lauve, Christophe Reutenauer, and Franco Saliola. *Combinatorics on Words: Christoffel Words and Repetitions in Words.* CRM Monograph Series, 2008.

Dimitris Bertsimas and José Niño-Mora. Conservation laws, extended polymatroids and multiarmed bandit problems; a polyhedral approach to indexable systems. *Mathematics of Operations Research*, 21(2):257–306, 1996.

Thierry Bousch and Jean Mairesse. Asymptotic height optimization for topical IFS, Tetris heaps, and the finiteness conjecture. *Journal of the American Mathematical Society*, 15 (1):77–111, 2002.

Claude Brezinski. Convergence acceleration during the 20th century. *Journal of Computational and Applied Mathematics*, 122(1):1–21, 2000.

Michael Carter and Bruce van Brunt. *The Lebesgue-Stieltjes Integral.* Springer, 2000.

Yuxin Chen, Hiroaki Shioi, Cesar Montesinos, Lian Pin Koh, Serge Wich, and Andreas Krause. Active detection via adaptive submodularity. In *Proceedings of The 31st International Conference on Machine Learning*, pages 55–63, 2014.

Erwin B. Christoffel. Observatio arithmetica. *Annali di Matematica Pura ed Applicata*, 6: 148–152, 1875.

Christopher R. Dance and Tomi Silander. When are Kalman-filter restless bandits indexable? In *Advances in Neural Information Processing Systems*, pages 1711–1719, 2015.

Christopher R. Dance and Tomi Silander. Optimal policies for observing time series and related restless bandit problems. *arXiv preprint arXiv:1703.10010*, 2017.

Ramez M. Daoud, Hassanein H. Amer, Hany M. Elsayed, and Y. Sallez. Ethernet-based car control network. In *Electrical and Computer Engineering, 2006. CCECE'06. Canadian Conference on*, pages 1031–1034. IEEE, 2006.

Abhimanyu Das and David Kempe. Algorithms for subset selection in linear regression. In *Proceedings of the 40th ACM Symposium on the Theory of Computing*, pages 45–54. ACM, 2008.

Geoffrey Davis, Stéphane Mallat, and Marco Avellaneda. Adaptive greedy approximations. *Constructive Approximation*, 13(1):57–98, 1997.

Robert Devaney. *An Introduction to Chaotic Dynamical Systems.* Westview Press, 2008.

Serge Dulucq and Dominique Gouyou-Beauchamps. Sur les facteurs des suites de Sturm. *Theoretical Computer Science*, 71(3):381–400, 1990.

John Gittins, Kevin Glazebrook, and Richard Weber. *Multi-Armed Bandit Allocation Indices.* John Wiley and Sons, 2011.

Kevin Glazebrook, Hodge David, and Kirkbride Chris. General notions of indexability for queueing control and asset management. *Annals of Applied Probability*, 21(3):876–907, 2011.

Amy Glen and Jacques Justin. Episturmian words: a survey. *RAIRO-Theoretical Informatics and Applications*, 43(3):403–442, 2009.

Ronald Graham, Donald Knuth, and Oren Patashnik. *Concrete Mathematics: A Foundation for Computer Science*. Addison-Wesley, 1994.

Sudipto Guha, Kamesh Munagala, and Peng Shi. Approximation algorithms for restless bandit problems. *Journal of the ACM*, 58(1):3, 2010.

Vijay Gupta, Timothy H. Chung, Babak Hassibi, and Richard M. Murray. On a stochastic sensor selection algorithm with applications in sensor scheduling and sensor coverage. *Automatica*, 42(2):251–260, 2006.

Wassim M. Haddad, Vijay S. Chellaboina, and Sergey G. Nersesov. *Impulsive and Hybrid Dynamical Systems: Stability, Dissipativity, and Control*. Princeton University Press, 2014.

Onésimo Hernández-Lerma and Jean Bernard Lasserre. *Further Topics in Discrete-Time Markov Control Processes*. Springer, 1999.

Alfred O. Hero and Douglas Cochran. Sensor management: Past, present, and future. *IEEE Sensors Journal*, 11(12):3064–3075, 2011.

Sandra Hirche, Manuel Ferre, Jordi Barrio, Claudio Melchiorri, and Martin Buss. Bilateral control architectures for telerobotics. In *Advances in Telerobotics*, pages 163–176. Springer, 2007.

Stephen John Hogan, L. Higham, and T.C.L. Griffin. Dynamics of a piecewise linear map with a gap. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, volume 463, pages 49–65. The Royal Society, 2007.

Rudolph E. Kalman. A new approach to linear filtering and prediction problems. *Journal of basic Engineering*, 82(1):35–45, 1960.

Thomas H. Kerr. Modeling and evaluating an empirical INS difference monitoring procedure used to sequence SSBN navaid fixes. *Navigation*, 28(4):263–285, 1981.

Leonard Kleinrock. A conservation law for a wide class of queueing disciplines. *Naval Research Logistics Quarterly*, 12(2):181–192, 1965.

Andrey N. Kolmogorov and Sergei V. Fomin. *Introductory Real Analysis, translated by Richard A. Silverman*. Dover, New York, 1975.

Victor S. Kozyakin. Sturmian sequences generated by order-preserving circle maps. Available as http://cool.iitp.ru/personal/kozyakin/koz-11-5.pdf., 2003.

Andreas Krause, Ajit Singh, and Carlos Guestrin. Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9(Feb):235–284, 2008.

Andreas Krause, Ram Rajagopal, Anupam Gupta, and Carlos Guestrin. Simultaneous optimization of sensor placements and balanced schedules. *IEEE Transactions on Automatic Control*, 56(10):2390–2405, 2011.

Julia Kuhn, Michel Mandjes, and Yoni Nazarathy. Exploration vs exploitation with partially observable gaussian autoregressive arms. In *Proceedings of the 8th International Conference on Performance Evaluation Methodologies and Tools*, pages 209–216. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2014.

Harold J. Kushner. On the optimum timing of observations for linear control systems with unknown initial state. *IEEE Transactions on Automatic Control*, 9(2):144–150, 1964.

Barbara La Scala and Bill Moran. Optimal target tracking with restless bandits. *Digital Signal Processing*, 16(5):479–487, 2006.

Maialen Larrañaga, Urtzi Ayesta, and Ina Verloop. Dynamic control of birth-and-death restless bandits: application to resource-allocation problems. *IEEE/ACM Transactions on Networking*, 24(6):3812–3825, 2016.

Jérôme Le Ny, Munther Dahleh, and Eric Feron. Multi-UAV dynamic routing with partial observations using restless bandit allocation indices. In *American Control Conference, 2008*, pages 4220–4225. IEEE, 2008.

Jérôme Le Ny, Eric Feron, and Munther Dahleh. Scheduling continuous-time Kalman filters. *IEEE Transactions on Automatic Control*, 56(6):1381–1394, 2011.

Keqin Liu and Qing Zhao. Indexability of restless bandit problems and optimality of Whittle index for dynamic multichannel access. *IEEE Transactions on Information Theory*, 56 (11):5547–5567, 2010.

M. Lothaire. *Algebraic Combinatorics on Words*. Cambridge University Press, 2002.

Oleg Makarenkov and Jeroen S. W. Lamb. Dynamics and bifurcations of nonsmooth systems: A survey. *Physica D: Nonlinear Phenomena*, 241(22):1826–1844, 2012.

Albert Marshall, Ingram Olkin, and Barry Arnold. *Inequalities: Theory of Majorization and its Applications*. Springer Science and Business Media, 2010.

Lewis Meier, John Peschon, and Robert Dressler. Optimal control of measurement subsystems. *IEEE Transactions on Automatic Control*, 12(5):528–536, 1967.

Filippo Mignosi. On the number of factors of Sturmian words. *Theoretical Computer Science*, 82(1):71–84, 1991.

Adam Molin and Sandra Hirche. On LQG joint optimal scheduling and control under communication constraints. In *Proceedings of the 48th IEEE Conference on Decision and Control*, pages 5832–5838. IEEE, 2009.

William Moran, Sofia Suvorova, and Stephen Howard. Application of sensor scheduling concepts to radar. In *Foundations and Applications of Sensor Management*, pages 221–256. Springer, 2008.

Marston Morse and Gustav A. Hedlund. Symbolic dynamics II. Sturmian trajectories. *American Journal of Mathematics*, 62(1):1–42, 1940.

Anastasios I. Mourikis and Stergios I. Roumeliotis. Optimal sensor scheduling for resource-constrained localization of mobile robot formations. *IEEE Transactions on Robotics*, 22 (5):917–931, 2006.

José Niño-Mora. Restless bandits, partial conservation laws and indexability. *Advances in Applied Probability*, 33(1):76–98, 2001.

José Niño-Mora. Dynamic allocation indices for restless projects and queueing admission control: a polyhedral approach. *Mathematical Programming, Series A*, 93(3):361–413, 2002.

José Niño-Mora. Restless bandit marginal productivity indices, diminishing returns, and optimal control of make-to-order/make-to-stock M/G/1 queues. *Mathematics of Operations Research*, 31(1):50–84, 2006.

José Niño-Mora. A $(2/3)n^3$ fast-pivoting algorithm for the Gittins index and optimal stopping of a Markov chain. *INFORMS Journal on Computing*, 19(4):596–606, 2007.

José Niño-Mora. A dynamic page-refresh index policy for web crawlers. In *Proceedings of the 19th International Conference on Analytical and Stochastic Modeling Techniques and Applications, LNCS 8499*, pages 46–60. Springer, 2014.

José Niño-Mora. A verification theorem for indexability of discrete time real state discounted restless bandits. *arXiv preprint arXiv:1512.04403*, 2015.

José Niño-Mora. Whittle's index policy for multi-target tracking with jamming and nondetections. In *International Conference on Analytical and Stochastic Modeling Techniques and Applications*, pages 210–222. Springer, 2016.

José Niño-Mora. A verification theorem for threshold-indexability of real-state discounted restless bandits. *Mathematics of Operations Research*, 2019. Forthcoming.

José Niño-Mora and Sofía Villar. Multi-target tracking via restless bandit marginal productivity indices and Kalman filter in discrete time. In *Proceedings of the 48th IEEE Conference on Decision and Control*, pages 2905–2910, 2009.

Yaakov Oshman. Optimal sensor selection strategy for discrete-time state estimators. *IEEE Transactions on Aerospace and Electronic Systems*, 30(2):307–314, 1994.

Christos H. Papadimitriou and John N. Tsitsiklis. The complexity of optimal queuing network control. *Mathematics of Operations Research*, 24(2):293–305, 1999.

Charles Chapman Pugh. *Real Mathematical Analysis*. Springer, 2002.

Bhooshan Rajpathak, Harish Pillai, and Santanu Bandyopadhyay. Analysis of stable periodic orbits in the one dimensional linear piecewise-smooth discontinuous map. *Chaos*, 22 (3):033126, 2012.

Halsey L. Royden and Patrick M. Fitzpatrick. *Real Analysis*. Pearson, fourth edition, 2010.

Peter Joseph Seiler. *Coordinated control of unmanned aerial vehicles*. PhD thesis, University of California, Berkeley, 2001.

Richard F. Serfozo. Monotone optimal policies for Markov decision processes. In *Stochastic Systems: Modeling, Identification and Optimization, II*, pages 202–215. Springer, 1976.

J. George Shanthikumar and David D. Yao. Multiclass queueing systems: Polymatroidal structure and optimal scheduling control. *Operations Research*, 40(3-supplement-2):S293–S299, 1992.

Thorvald Thiele. *Sur la compensation de quelques erreurs quasi-systématiques par la méthode des moindres carrés*. CA Reitzel, 1880.

Fabio Tramontana, Frank Westerhoff, and Laura Gardini. On the complicated price dynamics of a simple one-dimensional discontinuous financial market model with heterogeneous interacting traders. *Journal of Economic Behavior and Organization*, 74(3):187–205, 2010.

Ina Verloop. Asymptotically optimal priority policies for indexable and nonindexable restless bandits. *Annals of Applied Probability*, 26(4):1947–1995, 2016.

Sofía Villar. *Restless bandit index policies for dynamic sensor scheduling optimization*. PhD thesis, Statistics Department, Universidad Carlos III de Madrid, 2012.

Robert B. Washburn. Application of multi-armed bandits to sensor management. In *Foundations and Applications of Sensor Management*, pages 153–175. Springer, 2008.

Richard Weber and Gideon Weiss. On an index policy for restless bandits. *Journal of Applied Probability*, pages 637–648, 1990.

Jaap Wessels. Markov programming by successive approximations with respect to weighted supremum norms. *Journal of Mathematical Analysis and Applications*, 58(2):326–335, 1977.

Peter Whittle. Restless bandits: Activity allocation in a changing world. *Journal of Applied Probability*, pages 287–298, 1988.

Wei Wu and Ari Arapostathis. Optimal sensor querying: general Markovian model and LQG models with controlled observations. *IEEE Transactions on Automatic Control*, 53 (6):1392–1405, 2008.

Lin Zhao, Wei Zhang, Jianghai Hu, Alessandro Abate, and Claire J. Tomlin. On the optimal solutions of the infinite-horizon linear sensor scheduling problem. *IEEE Transactions on Automatic Control*, 59(10):2825–2830, 2014.