

Hidden Markov Models with Multiple Observation Processes

James Yuanjie Zhao

Submitted in total fulfilment of the requirements
of the degree of Master of Philosophy

Submitted 18 August 2010

Revised 18 January 2011

Department of Mathematics and Statistics
Department of Electrical and Electronic Engineering
University of Melbourne

Abstract

We consider a hidden Markov model with multiple observation processes, one of which is chosen at each point in time by a policy—a deterministic function of the information state—and attempt to determine which policy minimises the limiting expected entropy of the information state. Focusing on a special case, we prove analytically that the information state always converges in distribution, and derive a formula for the limiting entropy which can be used for calculations with high precision. Using this formula, we find computationally that the optimal policy is always a threshold policy, allowing it to be easily found. We also find that the greedy policy is almost optimal.

Declaration

This is to certify that:

1. The thesis comprises only my original work towards the MPhil;
2. Due acknowledgement has been made to all other material used; and
3. The thesis is less than 50,000 words in length.



Acknowledgements

My deepest and sincerest appreciation goes to my supervisors, Bill Moran and Peter Taylor, for their countless hours of guidance, both in relation to this thesis and in more general matters.

Contents

1	Introduction	7
1.1	Motivation	9
1.2	Past Work	12
2	Analytic Results	15
2.1	Definitions	15
2.2	Convergence	25
2.3	Special Case	37
3	Computational Results	46
3.1	Limiting Entropy	46
3.2	Optimal Threshold Policy	51
3.3	General Policies	69
4	Conclusion and Future Work	76
	Bibliography	77

List of Figures

2.1	A hidden Markov model	17
3.1	Space of all threshold policies	57
3.2	Example of a Region I policy	58
3.3	Example of a Region II policy	59
3.4	Example of a Region III policy	60
3.5	Example of a Region IV policy	61
3.6	Example of a Region V policy	62
3.7	Example of a Region VI policy	63
3.8	Location of the optimal threshold	64
3.9	Proof of Proposition 3.8	66
3.10	Locally optimal policies	71
3.11	Optimality of the greedy policy	74

1 Introduction

A hidden Markov model is an underlying Markov chain together with an imperfect observation on this chain. In the case of multiple observations, the classical model assumes that they can be observed simultaneously, and considers them as a single vector of observations. However, the case where not all the observations can be used at each point in time often arises in practical problems, and in this situation, one is faced with the challenge of choosing which observation to use.

We consider the case where the choice is made as a deterministic function of the previous information state, which is a sufficient statistic for the sequence of past observations. This function is called the *policy*, which we rank according to the information entropy of the information state that arises due to that policy.

Our main results are:

- The information state converges in distribution for almost every underlying Markov chain, as long as each observation process gives a perfect information observation with positive probability;
- In a special case (see Section 2.3 for a precise definition), we can write down the limiting entropy explicitly as a rational function of subgeometric infinite series, which allows the calculation of limiting entropy to very good precision;
- Computational results suggest that the optimal policy is a threshold policy, hence finding the optimal threshold policy is sufficient for finding the optimal policy in general;

- Finding a locally optimal threshold policy is also sufficient, while finding a locally optimal general policy is sufficient with average probability 0.98; and
- The greedy policy is optimal 96% of the time, and close to optimal the remaining times, giving a very simple yet reasonably effective suboptimal alternative.

1.1 Motivation

The theory of hidden Markov models was first introduced in a series of papers from 1966 by Leonard Baum and others under the more descriptive name of *Probabilistic Functions of Markov Chains* [1]. An application of this theory was soon found in speech recognition, spurring development, and the three main problems—probability calculation, state estimation and parameter estimation—had essentially been solved by the time of Lawrence Rabiner’s influential 1989 tutorial paper [13].

The standard hidden Markov model consists of an underlying state which is described by a Markov chain, and an imperfect observation process which is a probabilistic function of this underlying state. In most practical examples, this single observation is equivalent to having multiple observations, since we can simply consider them as a single vector of simultaneous observations. However, this requires that these multiple observation can be made and processed simultaneously, which is often not the case.

Sometimes, physical constraints may prevent the simultaneous use of all of the available observations. This is most evident with a sensor which can operate in multiple modes. For example, a radar antenna must choose a waveform to transmit; each possible waveform results in a different distribution of observations, and only one waveform can be chosen for each pulse. Another example might be in studying animal populations, where a researcher must select locations for a limited pool of detection devices such as traps and cameras.

Even when simultaneous observations are physically possible, other con-

straints may restrict their availability. For example, in an application where processors are much more expensive than sensors, a sensor network might reasonably consist of a large number of sensors and insufficient processing power to analyse the data from every sensor, in which case the processor must choose a subset of sensors from which to receive data. Similarly, a system where multiple sensors share a limited communication channel must decide how to allocate bandwidth, in a situation where each bit of bandwidth can be considered a virtual sensor, not all of which can be simultaneously used.

Another example is the problem of searching for a target which moves according to a Markov chain, where observation processes represent possible sites to be searched. Indeed, MacPhee and Jordan's [11] special case of this problem exactly corresponds to the special case we consider in Section 2.3, although with a very different cost function. Johnston and Krishnamurthy [7] show that this search problem can be used to model file transfer over a fading channel, giving yet another application for an extended hidden Markov model with multiple observation processes.

Note that in the problem of choosing from multiple observation processes, it suffices to consider the case where only one observation is chosen, by considering an observation to be an allowable subset of sensors. The three main hidden Markov model problems of probability calculation, state estimation and parameter estimation remain essentially the same, as the standard algorithms can easily be adapted by replacing the parameters of the single observation process by those of whichever observation process is chosen at each point in time.

Thus, the main interesting problem in the hidden Markov model with multiple observation processes is that of determining the optimal choice of observation process, which cannot be adapted from the standard theory of hidden Markov models since it is a problem that does not exist in that framework. It is this problem which will be the focus of our work.

We will use information entropy of the information state as our measure of optimality. While Evans and Krishnamurthy [6] use a distance between the information state and the underlying state, it is not necessary to consider this underlying state explicitly, since the information state is by definition an unbiased estimator of the distribution of the underlying state. We choose entropy over other measures such as variance since it is a measure of uncertainty which requires no additional structure on the underlying set.

The choice of an infinite time horizon is made in order to simplify the problem, as is our decision to neglect sensor usage costs. These variables can be considered in future work.

1.2 Past Work

The theory of hidden Markov models is already well-developed [13]. On the other hand, very little research has been done into the extended model with multiple observation processes. The mainly algorithmic solutions in the theory of hidden Markov models with a single observation process cannot be extended to our problem, since the choice of observation process does not exist in the unextended model.

Similarly, there is a significant amount of work on the sensor scheduling literature, but mostly considering autoregressive Gaussian processes such as in [15]. The case of hidden Markov sensors was considered by Jamie Evans and Vikram Krishnamurthy in 2001 [6], using policies where an observation process is picked as a deterministic function of the previous observation, and with a finite time horizon. They transformed the problem of choosing an observation into a control problem in terms of the information state, thereby entering the framework of stochastic control. They were able to write down the optimal policy as an intractable dynamic programming problem, and suggested the use of approximations to find the solution.

Krishnamurthy [8] followed up this work by showing that this dynamic programming problem could be solved using the theory of Partially Observed Markov Decision Processes when the cost function is of the form

$$C(z) = \sum_i z(i) \|\delta(i) - z\|,$$

where z is the information state, $\delta(i) \in \mathcal{P}(S)$ is the Dirac measure and $\|\cdot\|$ is a piecewise constant norm. It was then shown that such piecewise linear cost functions could be used to approximate quadratic cost functions, in the

sense that a sufficiently fine piecewise linear approximation must have the same optimal policy. In particular, this includes the Euclidean norm on the information state space, which corresponds to the expected mean-square distance between the information state and the distribution of the underlying chain. However, no bounds were found on how fine an approximation is needed.

The problem solved by Evans and Krishnamurthy is a similar but different problem to ours. We consider policies based on the information state, which we expect to perform better than policies based on only the previous observation, as the information state is a sufficient statistic for the sample path of observations (see Proposition 2.8, also [17]). We also consider an infinite time horizon, and specify information entropy of the information state as our cost function. Furthermore, while Evans and Krishnamurthy consider the primary tradeoff as that between the precision of the sensors and the cost of using them, we do not consider usage costs and only aim to minimise the uncertainty associated with the measurements.

Further work by Krishnamurthy and Djonin [9] extended the set of allowable cost functions to a Lipschitz approximation to the entropy function, and proved that threshold policies are optimal under certain very restrictive assumptions. Their breakthrough uses lattice theory methods [16] to show that the cost function must be monotonic in a certain way with respect to the information state, and thus the optimal choice of observation process must be characterised by a threshold. However, this work still does not solve our problem, as their cost function, a time-discounted infinite sum of expected costs, differs significantly from our limiting expected entropy, and furthermore their assumptions are difficult to verify in practice.

Another similar problem was also considered by Mohammad Rezaeian [14], who redefined the information state as the posterior distribution of the underlying chain given the sample path of observations up to the previous, as opposed to current, time instant, which allowed for a simplification in the recursive formula for the information state. Rezaeian also transformed the problem into a Markov Decision Process, but did not proceed further in his description.

The model for the special case we consider in Section 2.3 is an instance of the problem of searching for a moving target, which was partially solved by MacPhee and Jordan [11] with a very different cost function – the expected cumulative sum of prescribed costs until the first certain observation. They proved that threshold policies are optimal for certain regions of parameter space by analysing the associated fractional linear transformations. Unfortunately, similar approaches have proved fruitless for our problem due to the highly non-algebraic nature of the entropy function.

Our problem as it appears here was first studied in unpublished work by Bill Moran and Sofia Suvorova, who conjectured that the optimal policy is always a threshold policy. More extensive work was done in [18], where it was shown that the information state converges in distribution in the same special case that we consider in Section 2.3. It was also conjectured that threshold policies are optimal in this special case, although the argument provided was difficult to work into a full proof. However, [18] contains a mistake in the recurrence formula for the information state distribution, a corrected version of which appears as Lemma 2.13. The main ideas of the convergence proof still work, and are presented in corrected and improved form in Section 2.2.

2 Analytic Results

2.1 Definitions

We begin by precisely defining the model we will use. In particular, we will make all our definitions within this section, in order to expediate referencing. For the reader's convenience, Table 2.1 at the end of this section lists the symbols we will use for our model.

For a sequence X_0, X_1, \dots and any non-negative integer $t \in \mathbb{Z}^+$, we will use the notation $X_{(t)}$ to represent the vector (X_0, \dots, X_t) .

Definition 2.1. A **Markov Chain** [12] is a stochastic process $(X_t)_{t \in \mathbb{Z}^+}$, such that for all times $t \geq s$, all states x and all measurable sets A ,

$$\mathbb{P}(X_t \in A | X_s = x, \mathcal{F}_s) = \mathbb{P}(X_t \in A | X_s = x),$$

where \mathcal{F}_t denotes the canonical filtration. We will consistently use the symbol X_t to refer to an underlying Markov chain, and $\pi_t = \mathbb{P}(X_t)$ to denote its distribution.

In the case of a time-homogeneous, finite state and discrete time Markov chain, this simplifies to a sequence of random variables $(X_t)_{t \in \mathbb{Z}^+}$ taking values in a common finite state space $S = \{1, \dots, n\}$, such that for all times $t \in \mathbb{Z}^+$, X_{t+1} is conditionally independent of $X_{(t-1)}$ given X_t , and the distributions of X_{t+1} given X_t does not depend on t .

In this case, there exists $n \times n$ matrix T , called the **Transition Matrix**,

such that for all $i, j \in S$ and $t \in \mathbb{Z}^+$,

$$T_{ij} = \mathbb{P}(X_{t+1} = j \mid X_t = i) = \mathbb{P}(X_{t+1} = j \mid X_t = i, X_{(t-1)}). \quad (2.1)$$

Since we mainly consider Markov chains which are time-homogeneous and finite state, we will henceforth refer to them as Markov chains without the additional qualifiers.

Definition 2.2. An **Observation Process** on the Markov chain $(X_t)_{t \in \mathbb{Z}^+}$ is a sequence of random variables $(Y_t)_{t \in \mathbb{Z}^+}$ given by $Y_t = c(X_t, W_t)$, where c is a deterministic function and $(W_t)_{t \in \mathbb{Z}^+}$ is a sequence of independent and identically distributed random variables which is also independent of the Markov chain $(X_t)_{t \in \mathbb{Z}^+}$ [5].

As before, we will only consider observation processes which take values in a finite set $V = \{1, \dots, m\}$. Similarly to before, there exists an $m \times n$ matrix M , which we call the **Observation Matrix**, such that for all $i, j \in S, k \in V$ and $t \in \mathbb{Z}^+$,

$$\begin{aligned} M_{jk} &= \mathbb{P}(Y_t = k \mid X_t = j) = \mathbb{P}(Y_t = k \mid X_t = j, X_{(t-1)}, Y_{(t-1)}); \\ T_{ij} &= \mathbb{P}(X_{t+1} = j \mid X_t = i) = \mathbb{P}(X_{t+1} = j \mid X_t = i, X_{(t-1)}, Y_{(t)}). \end{aligned} \quad (2.2)$$

Heuristically, these two conditions can be seen as requiring that observations depend only on the current state, and do not affect future states. A diagrammatic interpretation is provided in Figure 2.1.

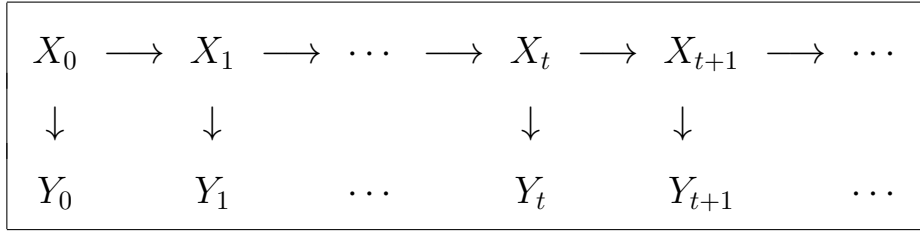


Figure 2.1: An observation process (Y_t) on a Markov chain (X_t) .
 At each node X_t , everything after X_t is conditionally independent of everything before X_t , given X_t .

Traditionally, a hidden Markov model is defined as the pair of a Markov chain and an observation process on that Markov chain. Since we will consider hidden Markov models with multiple observation processes, this definition does not suffice. We adjust it as follows.

Definition 2.3. A **Hidden Markov Model** is the triple of a Markov chain $(X_t)_{t \in \mathbb{Z}^+}$, a finite collection of observation processes $\{(Y_t^{(i)})_{t \in \mathbb{Z}^+}\}_{i \in O}$ on $(X_t)_{t \in \mathbb{Z}^+}$, and an additional sequence of random variables $(I_t)_{t \in \mathbb{Z}^+}$, called the **Observation Index**, mapping into the index set O .

Note that this amends the standard definition of a hidden Markov model. For convenience, we will no longer explicitly specify our hidden Markov models to have multiple observation processes.

It makes sense to think of $(X_t)_{t \in \mathbb{Z}^+}$ as the state of a system under observation, $\{(Y_t^{(i)})_{t \in \mathbb{Z}^+}\}_{i \in O}$ as a collection of potential observations that can be made on this system, and $(I_t)_{t \in \mathbb{Z}^+}$ as a choice of observation for each point in time.

Since our model permits only one observation to be made at each point in

time, and we will wish to determine which one to use based on past observations, it makes sense to define $(I_t)_{t \in \mathbb{Z}^+}$ as a sequence of random variables on the same probability space as the hidden Markov model.

We will discard the potential observations which are not used, leaving us with a single sequence of random variables representing the observations which are actually made.

Definition 2.4. The **Actual Observation** of a hidden Markov model is the sequence of random variables $(Y_t^{(I_t)})_{t \in \mathbb{Z}}$.

We will write Y_t to mean $Y_t^{(I_t)}$, noting that this is consistent with our notation for a hidden Markov model with a single observation process Y_t . On the other hand, for a hidden Markov model with multiple observation processes, the actual observation $(Y_t)_{t \in \mathbb{Z}^+}$ is not itself an observation process in general.

Since our goal is to analyse a situation in which only one observation can be made at each point in time, we will consider our hidden Markov model as consisting only of the underlying state $(X_t)_{t \in \mathbb{Z}^+}$ and the actual observation $(Y_t)_{t \in \mathbb{Z}^+}$. Where convenient, we will use the abbreviated terms *state* and *observation* at time t to mean X_t and Y_t respectively.

For any practical application of this model to a physical system, the underlying state cannot be determined, otherwise there would be no need to take non-deterministic observations. Therefore, we need a way of estimating the underlying state from the observations.

Definition 2.5. The **Information State Realisation** of a hidden Markov model at time t is the posterior distribution of X_t given the actual observa-

tions and observation indices up to time t .

To make this definition more precise, we introduce some additional notation.

First, recall that $(X_t)_{t \in \mathbb{Z}^+}$ has state space $S = \{1, \dots, n\}$, and define the set of probability measures on S ,

$$\mathcal{P}(S) \cong \left\{ (p_1, \dots, p_n) \in \mathbb{R}^n : p_i \geq 0 \forall i \in S, \sum_{i \in S} p_i = 1 \right\}. \quad (2.3)$$

Second, for a random variable X with state space S , and an event E , define the posterior distribution of X given E ,

$$\mathbb{P}(X | E) = \left(\mathbb{P}(X = 1 | E), \dots, \mathbb{P}(X = n | E) \right) \in \mathcal{P}(S). \quad (2.4)$$

Although we make this definition in general, we purposely choose the letters X and S , coinciding with the letters used to represent the underlying Markov chain and the state space, as this is the context in which we will use this definition. Then, the information state realisation is a function

$$\begin{aligned} z_t &: V^{t+1} \times O^{t+1} \longrightarrow \mathcal{P}(S), \\ z_t(y_{(t)}; i_{(t)}) &= \mathbb{P}(X_t | Y_{(t)} = y_{(t)}, I_{(t)} = i_{(t)}). \end{aligned} \quad (2.5)$$

This extends very naturally to a random variable.

Definition 2.6. The **Information State Random Variable** is

$$Z_t = z_t(Y_{(t)}; I_{(t)}) = \mathbb{P}(X_t | Y_{(t)}, I_{(t)}).$$

Its distribution is the **Information State Distribution** $\mu_t = \mathbb{P}(Z_t)$, taking values in $\mathcal{P}(\mathcal{P}(S))$, the space of Radon probability measures on $\mathcal{P}(S)$, which is a subset of the real Banach space of signed Radon measures on $\mathcal{P}(S)$.

Thus, the information state realisation is exactly a realisation of the information state random variable. It is useful because it represents the maximal information we can deduce about the underlying state from the observation index and the actual observation, as shown in Proposition 2.8.

For the purpose of succinctness, we will refer to any of z_t , Z_t and μ_t as simply the **Information State** when the context is clear.

Definition 2.7. A random variable Z is a **sufficient statistic** for a parameter X given data Y if for any values y and z of Y and Z respectively, the probability $\mathbb{P}(Y = y \mid Z = z, X = x)$ is independent of x [3]. As before, we make the definition in general, but purposely choose the symbols X , Y and Z to coincide with symbols already defined.

In our case, X , which is a random variable, is used in the context of a parameter. Our problem takes place in a Bayesian framework, where the information state represents our belief about the underlying state, and is updated at each observation.

Proposition 2.8. *The information random variable Z_t is a sufficient statistic for the underlying state X_t , given the actual observations $Y_{(t)}$ and the observation indices $I_{(t)}$.*

Proof. By Definition 2.7, we need to prove that for all $y \in V^{t+1}$ and $i \in O^{t+1}$,

$$\mathbb{P}(Y, I \mid Z, x) = \mathbb{P}(Y_{(t)} = y, I_{(t)} = i \mid Z_t = z_t(y; i), X_t = x) \quad (2.6)$$

is independent of x .

First, note that the event $\{Z_t = z_t(y; i)\}$ is the disjoint union of events

$\{Y_{(t)} = y', I_{(t)} = i'\}$ over all $(y', i') \in V^{t+1} \times O^{t+1}$ such that $z_t(y'; i') = z_t(y; i)$.

Next, if $z_t(y; i) = z_t(y'; i')$, then by definition of z_t , for all $x \in S$,

$$\mathbb{P}(X_t = x \mid Y_{(t)} = y, I_{(t)} = i) = \mathbb{P}(X_t = x \mid Y_{(t)} = y', I_{(t)} = i'). \quad (2.7)$$

Then, by definition of conditional probability,

$$\frac{\mathbb{P}(X_t = x, Y_{(t)} = y', I_{(t)} = i')}{\mathbb{P}(X_t = x, Y_{(t)} = y, I_{(t)} = i)} = \frac{\mathbb{P}(Y_{(t)} = y', I_{(t)} = i')}{\mathbb{P}(Y_{(t)} = y, I_{(t)} = i)}. \quad (2.8)$$

Hence,

$$\begin{aligned} \mathbb{P}(Y, I \mid Z, x) &= \mathbb{P}(Y_{(t)} = y, I_{(t)} = i \mid Z_t = z_t(y; i), X_t = x) \\ &= \frac{\mathbb{P}(Y_{(t)} = y, I_{(t)} = i, Z_t = z_t(y; i), X_t = x)}{\mathbb{P}(Z_t = z_t(y; i), X_t = x)} \\ &= \frac{\mathbb{P}(Y_{(t)} = y, I_{(t)} = i, X_t = x)}{\sum_{y', i'} \mathbb{P}(Y_{(t)} = y', I_{(t)} = i', X_t = x)} \\ &= \left(\sum_{y', i'} \frac{\mathbb{P}(Y_{(t)} = y', I_{(t)} = i', X_t = x)}{\mathbb{P}(Y_{(t)} = y, I_{(t)} = i, X_t = x)} \right)^{-1} \\ &= \left(\sum_{y', i'} \frac{\mathbb{P}(Y_{(t)} = y', I_{(t)} = i')}{\mathbb{P}(Y_{(t)} = y, I_{(t)} = i)} \right)^{-1}. \end{aligned} \quad (2.9)$$

Each sum above is taken over all $(y', i') \in V^{t+1} \times O^{t+1}$ such that $z_t(y'; i') = z_t(y; i)$. This expression is clearly independent of x , which completes the proof that Z_t is a sufficient statistic for X_t . \square

Since the information state represents all information that can be deduced from the past, it makes sense to use it to determine which observation process to use in future.

Definition 2.9. A **policy** on a hidden Markov model is a deterministic function $g : \mathcal{P}(S) \rightarrow O$, such that for all $t \in \mathbb{Z}^+$, $I_{t+1} = g(Z_t)$. We will use

the symbol $A_i = g^{-1}\{i\}$ to denote the preimage of the observation method i under the policy, that is, the subset of $\mathcal{P}(\mathcal{P}(S))$ on which observation method i is prescribed by the policy. We will always consider the policy g as fixed.

Since Z_t is a function of $Y_{(t)}$ and $I_{(t)}$, this means that $I_{(t+1)}$ is a function of $Y_{(t)}$ and $I_{(t)}$. Then by induction, we see that $I_{(t+1)}$ is a function of $Y_{(t)}$ and I_0 . Therefore, if we prescribe some fixed I_0 , then I_t is a function of $Y_{(t)}$.

For fixed I_0 , we can write

$$Z_t = \mathbb{P}(X_t \mid I_{(t)}, Y_{(t)}) = \mathbb{P}(X_t \mid Y_{(t)}). \quad (2.10)$$

Hence, the information random variable is a deterministic function of only $Y_{(t)}$. In particular, the information state z_t can be written with only one argument, that is, $Z_t = z_t(Y_{(t)})$.

Since our aim is to determine the underlying state with the least possible uncertainty, we need to introduce a quantifier of uncertainty. There are many possible choices, especially if the state space has additional structure. For example, variance would be a good candidate in an application where the state space embeds naturally into a real vector space.

However, in the general case, there is no particular reason to suppose our state space has any structure; our only assumption is that it is finite, in which case information entropy is the most sensible choice, being a natural, axiomatically-defined quantifier of uncertainty for a distribution on a countable set without any additional structure [4].

Definition 2.10. The **Information Entropy** of a discrete probability

measure $(p_1, \dots, p_n) \in \mathcal{P}(S)$ is given by

$$H((p_1, \dots, p_n)) = - \sum_j p_j \log p_j.$$

We will use the natural logarithm, and define $0 \log 0 = 0$ in accordance with the fact that $p \log p \rightarrow 0$ as $p \rightarrow 0$.

Since Z_t takes values in $\mathcal{P}(S)$, $H(Z_t)$ is well-defined, and by definition measures the uncertainty in X_t given $Y_{(t)}$, and therefore by Proposition 2.8, measures the uncertainty in our best estimate of X_t . Thus, the problem of minimising uncertainty becomes quantified as one of minimising $H(Z_t)$.

We are particularly interested in the limiting behaviour, and thus, the main questions we will ask are:

- Under what conditions, and in particular what policies, does $H(Z_t)$ converge as $t \rightarrow \infty$?
- Among the policies under which $H(Z_t)$ converges, which policy gives the minimal limiting value of $H(Z_t)$?
- Are there interesting cases where $H(Z_t)$ does not converge, and if so, can we generalise the above results?

Symbol	Value	Meaning
$\mathcal{P}(S)$	set	probability measures on S
$\mathcal{P}(\mathcal{P}(S))$	set	probability measures on $\mathcal{P}(S)$
A_i	set	region of observation process i
H	function	information entropy
I_t	random variable	observation index
O	finite set	set of observation processes
S	finite set	state space of Markov chain
V	finite set	observation space
W_t	random variable	observation randomness
X_t	random variable	Markov chain
$Y_t^{(i)}$	random variable	observation process
Y_t	random variable	actual observation $Y_t^{(I_t)}$
Z_t	random variable	information state random variable
g	function	policy
m	integer	number of observation values
n	integer	number of states
t	integer	position in time
z_t	distribution	information state realisation
π_t	distribution	Markov chain distribution
μ_t	distribution	information state distribution

Table 2.1: List of symbols, ordered alphabetically.

2.2 Convergence

In this section, we will prove that under certain conditions, the information state converges in distribution. This fact is already known for classical hidden Markov models, and is quite robust: LeGland and Mevel [10] prove geometric ergodicity of the information state even when calculated from incorrectly specified parameters, while Cappé, Moulines and Rydén [2] prove Harris recurrence of the information state for certain uncountable state underlying chains. We will present a mostly elementary proof of convergence in the case of multiple observation processes.

To determine the limiting behaviour of the information state, we begin by finding an explicit form for its one-step time evolution.

Definition 2.11. For each observation process i and each observed state y , the **r -function** is the function $r_{i,y} : \mathcal{P}(S) \rightarrow \mathcal{P}(S)$ given by

$$r_{i,y}(z) = \frac{\sum_{x,j \in S} M_{x,y}^{(i)} T_{j,x} z_j \delta(x)}{\sum_{x,j \in S} M_{x,y}^{(i)} T_{j,x} z_j},$$

where $\delta : S \rightarrow \mathcal{P}(S)$ is the Dirac measure on S and z_j is the j th component of $z \in \mathcal{P}(S) \subset \mathbb{R}^n$.

Lemma 2.12. *In a hidden Markov model with multiple observation processes and a fixed policy g , the information state satisfies the recurrence relation*

$$z_{t+1}(y_{(t+1)}) = r_{g(z_t(y_{(t)})), y_{t+1}}(z_t(y_{(t)})). \quad (2.11)$$

Proof. Let $i_{t+1} = g(z_t(y_{(t)}))$ and $k_t = \mathbb{P}(Y_{(t)}^{(i_{t+1})} = y_{(t)})$. By the Markov

property as in Definition 2.2, and the simplification (2.10),

$$\begin{aligned}
z_{t+1}(y_{(t+1)})_x &= \mathbb{P}(X_{t+1} = x \mid Y_{(t+1)}^{(i_{(t+1)})} = y_{(t+1)}) \\
&= \frac{1}{\mathbb{P}(Y_{(t+1)}^{(i_{(t+1)})} = y_{(t+1)})} \sum_j \mathbb{P}(X_{t+1} = x, Y_{(t+1)}^{(i_{(t+1)})} = y_{(t+1)}, X_t = j) \\
&= \frac{1}{k_{t+1}} \sum_j \mathbb{P}(Y_{t+1}^{(i_{t+1})} = y_{t+1} \mid X_{t+1} = x) \\
&\quad \times \mathbb{P}(X_{t+1} = x \mid X_t = j) \mathbb{P}(X_t = j, Y_{(t)}^{(i_{(t)})} = y_{(t)}) \\
&= \frac{1}{k_{t+1}} \sum_j M_{x, y_{t+1}}^{(i_{t+1})} T_{j, x} k_t z_t(y_{(t)})_j \\
&= \frac{k_t}{k_{t+1}} \sum_j M_{x, y_{t+1}}^{(i_{t+1})} T_{j, x} z_t(y_{(t)})_j \\
&= \frac{\sum_j M_{x, y_{t+1}}^{(i_{t+1})} T_{j, x} z_t(y_{(t)})_j}{\sum_x \sum_j M_{x, y_{t+1}}^{(i_{t+1})} T_{j, x} z_t(y_{(t)})_j}, \tag{2.12}
\end{aligned}$$

since k_t/k_{t+1} does not depend on x and $\sum_x z_{t+1}(y_{(t+1)})_x = 1$. \square

Note that for each information state z and each observation process i , there are at most m possible information states at the next step, which are given explicitly by $r_{i, y}(z)$ for each observation $y \in V$.

Lemma 2.13. *The information distribution satisfies the recurrence relation*

$$\mu_{t+1} = \sum_{i \in O} \sum_{y \in V} \int_{A_i} (z \cdot T \cdot M^{(i)})_y \delta(r_{i, y}(z)) d\mu_t(z),$$

where the sum is taken over all observation processes i and all observation states y , $\delta : \mathcal{P}(S) \rightarrow \mathcal{P}(\mathcal{P}(S))$ is the Dirac measure on $\mathcal{P}(S)$, and \cdot is the matrix product considering $z \in \mathcal{P}(S) \subset \mathbb{R}^n$ as a row vector.

Proof. Since $Z_t = \mathcal{P}(X_t \mid Y_{(t)})$ is a deterministic function of $Y_{(t)}$, given that $Y_{(t+1)} = y_{(t+1)}$,

$$Z_{t+1} = z_{t+1}(y_{(t+1)}) = r_{g(z_t(y_{(t)}), y_{t+1})}(z_t(y_{(t)})). \tag{2.13}$$

This depends only on $z_t(y_{(t)})$ and y_{t+1} , so given that $Z_t = z$ and $Y_{t+1} = y$,

$$Z_{t+1} = r_{g(z),y}(z). \quad (2.14)$$

Integration over $(Z_t, Y_{t+1}) \in \mathcal{P}(S) \times V$ gives

$$\begin{aligned} \mu_{t+1} &= \int_{\mathcal{P}(S)} \sum_{y \in V} \delta(r_{g(z),y}(z)) \mathbb{P}(Y_{t+1} = y | Z_t = z) d\mu_t(z) \\ &= \sum_{i \in O} \sum_{y \in V} \int_{A_i} \delta(r_{g(z),y}(z)) \mathbb{P}(Y_{t+1} = y | Z_t = z) d\mu_t(z). \end{aligned} \quad (2.15)$$

By Definition 2.5, Z_t is the posterior distribution of X_t given the observations up to time t , so $\mathbb{P}(X_t = x | Z_t = z) = z_x$, the x th coordinate of the vector $z \in \mathcal{P}(S) \subset \mathbb{R}^n$. Since Z_t is a function of $Y_{(t)}$, which is a function of $X_{(t)}$ and the observation randomness $W_{(t)}$, by the Markov property as in Definition 2.2,

$$\begin{aligned} \mathbb{P}(Y_{t+1} = y | Z_t = z) &= \sum_{x \in S} (Y_{t+1} = y | X_t = x) \mathbb{P}(X_t = x | Z_t = z) \\ &= \sum_{x \in S} (T \cdot M^{(i)})_{x,y} z_x = (z \cdot T \cdot M^{(i)})_y. \end{aligned} \quad (2.16)$$

Substituting (2.16) into (2.15) completes the proof. \square

Note that Lemma 2.13 shows that the information distribution is given by a linear dynamical system on $\mathcal{P}(\mathcal{P}(S))$, and therefore the information state is a Markov chain with state space $\mathcal{P}(S)$. We will use tools in Markov chain theory to analyse the convergence of the information state, for which it will be convenient to give a name to this recurrence.

Definition 2.14. The **transition function** of the information distribution is the deterministic function $F : \mathcal{P}(\mathcal{P}(S)) \rightarrow \mathcal{P}(\mathcal{P}(S))$ given by $F(\mu_t) = \mu_{t+1}$, extended linearly to all of $\mathcal{P}(\mathcal{P}(S))$ by the recurrence in Lemma 2.13. The coefficients $\alpha_{i,y}(z) = (z \cdot T \cdot M^{(i)})_y$ are called the **α -functions**.

We now give a criterion under which the information state is always positive recurrent.

Definition 2.15. A discrete state Markov Chain X_t is called **Ergodic** if it is irreducible, aperiodic and positive recurrent. Such a chain has a unique invariant measure π , which is a limiting distribution in the sense that X_t converges to π in total variation norm [12].

Definition 2.16. A discrete state Markov Chain X_t is called **Positive** if every transition probability is strictly positive, that is, for all $i, j \in S$, $\mathbb{P}(X_{t+1} = i | X_t = j) > 0$. This is a stronger condition than ergodicity.

Definition 2.17. We shall call a hidden Markov model **Anchored** if the underlying Markov chain X_t is ergodic, and for each observation process i , there is a state x_i and an observation y_i such that $M_{x_i, y_i}^{(i)} > 0$ and $M_{x, y_i}^{(i)} = 0$ for all $x \neq x_i$. The pair (x_i, y_i) is called an **Anchor Pair**.

Heuristically, the latter condition allows for perfect information $\delta(x_i)$ whenever the observation y_i is made using observation process i . This anchors the information chain in the sense that this state can be reached with positive probability from any other state, thus resulting in a recurrent atom in the uncountable state chain Z_t . On the other hand, since each information state can make a transition to only finitely many other information states, starting the chain at $\delta(x_i)$ results in a discrete state Markov chain, for which it is much easier to prove positive recurrence.

Lemma 2.18. *In an anchored hidden Markov model, for any anchor pair (x_i, y_i) , $r_{i, y_i}(z) = \delta(x_i)$ for all $z \in \mathcal{P}(S)$.*

Proof. When $x \neq x_i$, $M_{x,y_i}^{(i)} = 0$ by Definition 2.17, so every term in the numerator of Definition 2.11 is zero except the coefficient of $\delta(x_i)$. Since we know the coefficients have sum 1, it follows that $r_{i,y_i} = \delta(x_i)$. \square

Lemma 2.19. *In a positive, anchored hidden Markov model, the α -functions α_{i,y_i} , for each $i \in O$, are uniformly bounded below by some $\epsilon > 0$, that is, $\alpha_{i,y_i}(z) \geq \epsilon$ for all i and z .*

Proof. We can write $\alpha_{i,y_i}(z) = \sum_x z_x T_{x,x_i} M_{x_i,y_i}^{(i)}$ by Definitions 2.14 and 2.17, which is bounded below by $\min_x T_{x,x_i} M_{x_i,y_i}^{(i)}$ since $\sum_x z_x = 1$. Since each $M_{x_i,y_i}^{(i)} > 0$, if all the entries of T are positive, then $\alpha_{i,y_i}(z)$ is bounded below uniformly in z for fixed i , which then implies a uniform bound in z and i since there are only finitely many i . \square

Definition 2.20. For each state $x \in S$, the **Orbit** R_x of $\delta(x) \in \mathcal{P}(S)$ under the r -functions is

$$\begin{aligned} R_x = & \{ \delta(x) \} \cup \{ r_{i,y}(\delta(x)) : \alpha_{i,y}(\delta(x)) > 0 \} \\ & \cup \{ r_{i',y'} \circ r_{i,y}(\delta(x)) : \alpha_{i',y'}(r_{i,y}(\delta(x))) \alpha_{i,y}(\delta(x)) > 0 \} \\ & \cup \dots \end{aligned}$$

By requiring the α -functions to be positive, we exclude points in the orbit which are reached with zero probability. Let $R = \bigcup_x R_x$.

Proposition 2.21. *In a positive, anchored hidden Markov model, there exists a constant $0 < \lambda < 1$ such that for all measures $Z \in \mathcal{P}(\mathcal{P}(S))$, the mass of the measure $F^t(Z)$ outside R is bounded by λ^t , that is, $F^t(Z)(R^c) \leq \lambda^t$.*

Proof. We can rewrite Definition 2.14 as

$$F(Z) = \int_{\mathcal{P}(S)} Q(z) dZ(z), \quad (2.17)$$

where

$$Q(z) = \sum_i \sum_y \mathbb{1}_{A_i}(z) \alpha_{i,y}(z) \delta(r_{i,y}(z)). \quad (2.18)$$

In this notation, the integral is the Lebesgue integral of the function Q with respect to the measure Z . Since Q takes values in the $\mathcal{P}(\mathcal{P}(S))$ and Z is a probability, the integral also takes values in $\mathcal{M}(\mathcal{P}(S))$, thus F maps the information state space $\mathcal{P}(\mathcal{P}(S))$ to itself.

Since $Q(z)$ is a measure supported on the set of points reachable from z via an r -function, and R is a union of orbits of r -functions and therefore closed under r -functions, it follows that all mass in R is mapped back into R under the evolution function, that is

$$\left(\int_R Q(z) dZ(z) \right) (R) = Z(R) = 1 - Z(R^c). \quad (2.19)$$

On the other hand, by Lemma 2.18, $r_{i,y_i}(z) = \delta(x_i) \in R$ for all z , hence

$$\begin{aligned} \left(\int_{R^c} Q(z) dZ(z) \right) (R) &\geq \inf_z (Q(z))(R) \\ &\geq \inf_z \left(\mathbb{1}_{A_{g(z)}}(z) \alpha_{g(z),y_{g(z)}}(z) \delta(r_{g(z),y_{g(z)}}(z)) \right) (R) \\ &= \inf_z \alpha_{g(z),y_{g(z)}}(z) \geq \inf_{i,z} \alpha_{i,y_i}(z). \end{aligned} \quad (2.20)$$

Putting these together gives

$$\begin{aligned} F(Z)(R) &= \left(\int_R Q(z) dZ(z) + \int_{R^c} Q(z) dZ(z) \right) (R) \\ &\geq 1 - \left(1 - \inf_{i,z} \alpha_{i,y_i}(z) \right) Z(R^c). \end{aligned} \quad (2.21)$$

Setting $\lambda = 1 - \inf_{i,z} \alpha_{i,y_i}(z)$ gives $F(Z)(R^c) \leq \lambda Z(R^c)$, hence $F^t(Z)(R^c) \leq \lambda^t$ by induction. By Lemma 2.19, $\lambda < 1$, while $\lambda > 0$ since we can always choose a larger value. \square

Up to this point, we have considered the evolution function as a deterministic function $F : \mathcal{P}(\mathcal{P}(S)) \rightarrow \mathcal{P}(\mathcal{P}(S))$. However, we can also consider it as a probabilistic function $F : \mathcal{P}(S) \rightarrow \mathcal{P}(S)$. By Definition 2.14, F maps points in R to R , hence the restriction $F|_R : R \rightarrow R$ gives a probabilistic function, and therefore a Markov chain, with countable state space R .

By Proposition 2.21, the limiting behaviour of the information chain takes place almost entirely in R in some sense, so we would expect that convergence of the restricted information chain $F|_R$ is sufficient for convergence of the full information chain F . This is proved below.

Proposition 2.22. *In a positive, anchored hidden Markov model, under any policy, the chain $F|_R$ has at least one state of the form $\delta(x_i)$ which is positive recurrent, that is, whose expected return time is finite.*

Proof. Construct a Markov chain P on the set $\{A_i\} \cup \{R_x\}$, with transition probabilities $P(A_i, R_{x_i}) > 0$ for all i , $P(R_x, A_i) > 0$ whenever $R_x \cap A_i$ is nonempty, and all other transition probabilities zero. We note that this is possible since we allow each state a positive probability transition to some other state.

Since P is a finite state Markov chain, it must have a recurrent state. Each state R_x can reach some state A_i , so some state A_i is recurrent; call it A_1 .

Consider a state $z_0 = r^{k_2}(\delta(x_2)) \in R_{x_2} \subseteq R$ of the chain F which is reachable from $\delta(x_1)$, where r^{k_2} is a composition of k_2 r -functions with corresponding α -functions nonzero. Since the A_i partition $\mathcal{P}(S)$, one of them must contain z_0 ; call it A_3 . We will assume $A_3 \neq A_1$; the proof follows the same argument

and is simpler in the case when $z_0 \in A_1$.

By definition of the α -functions,

$$\mathbb{P}(Z_1 = \delta(x_3) | Z_0 = z_0) = \alpha_{3,y_3}(z_0) > 0. \quad (2.22)$$

This means that $\delta(x_3)$ is reachable from $\delta(x_1)$ in the chain F , hence in the chain P , A_3 is reachable from A_1 , and by recurrence of A_1 , A_1 must also be reachable from A_3 via some sequence of positive probability transitions

$$A_3 \rightarrow R_{x_3} \rightarrow A_4 \rightarrow R_{x_4} \rightarrow \cdots \rightarrow A_1. \quad (2.23)$$

By definition of $P(R_{x_3}, A_4) > 0$, $R_{x_3} \cap A_4$ is nonempty, and thus contains some point $r^{k_3}(\delta(x_3))$, where r^{k_3} is a composition of k_3 r -functions with corresponding α nonzero.

By Definition 2.20, each transition $r^\ell(\delta(x_3))$ to $r^{\ell+1}(\delta(x_3))$ in the information chain occurs with positive probability, so

$$\mathbb{P}(Z_{k_3+1} = r^{k_3}(\delta(x_3)) | Z_1 = \delta(x_3)) = \beta_3 > 0. \quad (2.24)$$

Since $r^{k_3}(\delta(x_3)) \in A_4$, by anchoredness and positivity,

$$\mathbb{P}(Z_{k_3+2} = \delta(x_4) | Z_{k_3+1} = r^{k_3}(\delta(x_3))) = \gamma_3 > 0. \quad (2.25)$$

The Markov property then gives

$$\mathbb{P}(Z_{k_3+2} = \delta(x_4) | Z_0 = z_0) = \alpha_{3,y_3}(z_0)\beta_3\gamma_3 > 0. \quad (2.26)$$

Continuing via the sequence (2.23), we obtain

$$\mathbb{P}(Z_{k_1+\dots+k_j+j+1} = \delta(x_1) | Z_0 = z_0) = \alpha_{3,y_3}(z_0)\beta_3\gamma_3 \cdots \beta_j\gamma_j > 0. \quad (2.27)$$

Thus, for every state $z \in R$ reachable from $\delta(x_1)$, we have found constants $s_z \in \mathbb{N}$ and $c_z > 0$ such that

$$\mathbb{P}(Z_{s_z} = \delta(x_1) | Z_0 = z) = c_z.$$

By Lemma 2.19, $\alpha_{3,y_3}(z)$ is uniformly bounded below, while $\beta_3\gamma_3\cdots\beta_j\gamma_j$ depends only on the directed path (2.23) and not on z , and thus is also uniformly bounded below since there are only finitely many A_i , and hence it suffices to choose finitely many such paths.

Similarly, s_z also depends only on the directed path (2.23), and thus is uniformly bounded above. In particular, it is possible to pick s_z and c_z such that $s = \sup_z s_z < \infty$ and $c = \inf_z c_z > 0$.

Let τ be the first entry time into the state $\delta(x_1)$. By the above bound, we have $\mathbb{P}(\tau > s | Z_0 = z) \leq 1 - c$ for any initial state z reachable from $\delta(x_1)$. Letting τ' and Z' be independent copies of τ and Z , the time-homogeneous Markov property gives

$$\begin{aligned}
\frac{\mathbb{P}(\tau > (k+1)s | Z_0 = z)}{\mathbb{P}(\tau > ks | Z_0 = z)} &= \mathbb{P}(\tau > (k+1)s | \tau > ks, Z_0 = z) \\
&= \sum_{z'} \mathbb{P}(\tau > (k+1)s | \tau > ks, Z_{ks} = z', Z_0 = z) \\
&\quad \times \mathbb{P}(Z_{ks} = z' | \tau > ks, Z_0 = z) \\
&\leq \sup_{z'} \mathbb{P}(\tau > (k+1)s | \tau > ks, Z_{ks} = z') \\
&= \sup_{z'} \mathbb{P}(\tau' > s | Z'_0 = z') \\
&\leq 1 - c.
\end{aligned} \tag{2.28}$$

By induction, $\mathbb{P}(\tau > ks | Z_0 = z) \leq (1 - c)^k$ for all k . Dropping the condition on the initial distribution Z_0 for convenience, we have

$$\begin{aligned}
\mathbb{E}[\tau] &= \sum_{k \in \mathbb{Z}^+} \mathbb{P}(\tau > k) = \sum_{k \in \mathbb{Z}^+} \sum_{0 \leq t < s} \mathbb{P}(\tau > ks + t) \\
&\leq \sum_{k \in \mathbb{Z}^+} \sum_{0 \leq t < s} \mathbb{P}(\tau > ks) \leq \sum_{k \in \mathbb{Z}^+} s(1 - c)^k = \frac{s}{c} < \infty.
\end{aligned} \tag{2.29}$$

In particular, $\mathbb{E}[\tau | Z = \delta(x_1)] < \infty$, so $\delta(x_1)$ is a positive recurrent state. \square

Lemma 2.23. *The transition function F , considered as an operator on the real Banach space $\mathcal{M}(\mathcal{P}(S))$ of signed Radon measures on $\mathcal{P}(S)$ with the total variation norm, is linear with operator norm $\|F\| = 1$.*

Proof. Linearity follows immediately from the fact that F is defined as a finite sum of integrals.

For each $\mu \in \mathcal{M}(\mathcal{P}(S))$, let $\mu = \mu^+ - \mu^-$ be the Hahn decomposition, so that $\|\mu\| = \|\mu^+\| + \|\mu^-\| = \mu^+(\mathcal{P}(S)) + \mu^-(\mathcal{P}(S))$ by definition of the total variation norm.

If $\mu^+(\mathcal{P}(S)) = 0$, then $F(\mu^+)(\mathcal{P}(S)) = 0$ by linearity of F . Otherwise, let $c = \mu^+(\mathcal{P}(S))$, so that $\frac{1}{c}\mu^+ \in \mathcal{P}(\mathcal{P}(S))$. Since F maps probability measures to probability measures, $\|F(\mu^+)\| = c\|F(\frac{1}{c}\mu^+)\| = c = \|\mu^+\|$, and similarly for μ^- , hence by linearity of F and the triangle inequality,

$$\|F(\mu)\| \leq \|F(\mu^+)\| + \|F(\mu^-\|) = \|\mu^+\| + \|\mu^-\| = \|\mu\|. \quad (2.30)$$

This shows that $\|F\| \leq 1$. Picking μ to be any probability measure gives $\|F(\mu)\| = 1 = \|\mu\|$, hence $\|F\| = 1$. \square

Theorem 2.24. *In a positive, anchored hidden Markov model, irreducibility and aperiodicity of the restricted information chain $F|_R$ is a sufficient condition for convergence in distribution of the information state Z_t to some discrete invariant measure $\mu_\infty \in \mathcal{P}(R) \subset \mathcal{P}(\mathcal{P}(S))$.*

Proof. In this case, the restricted chain $F|_R$ is irreducible, aperiodic and positive recurrent—that is, ergodic—and hence has a unique invariant probability measure μ_∞ , which can be considered as an element of $\mathcal{P}(\mathcal{P}(S))$ supported on $R \subset \mathcal{P}(S)$.

We will show that the information distribution μ_t converges in total variation norm to μ_∞ . Fix $\mu_0 \in \mathcal{P}(\mathcal{P}(S))$ and $\epsilon > 0$, and pick s such that $\lambda^s < \frac{\epsilon}{3}$, with λ as defined in Proposition 2.21.

Let $\mu_s|_R$ be the restriction of the probability measure μ_s to R , which is a positive measure with total mass $m = \mu_s(R) \geq 1 - \lambda^s > 0$, hence we can divide to obtain the probability measure $\mu' = \frac{1}{m}\mu_s|_R$ supported on R .

Since $F|_R$ is ergodic and μ' is a probability measure, $F^t(\mu') = (F|_R)^t(\mu')$ converges to μ_∞ in total variation norm, so pick K such that for all $k \geq K$, $\|F^k(\mu') - \mu_\infty\| < \frac{\epsilon}{3}$.

For any $t \geq K + s$, we have the triangle inequality bound

$$\|\mu_t - \mu_\infty\| \leq \|\mu_t - F^{t-s}(m\mu')\| + \|F^{t-s}(m\mu') - m\mu_\infty\| + \|m\mu_\infty - \mu_\infty\|. \quad (2.31)$$

By Lemma 2.23, F is linear with operator norm 1, so by Proposition 2.21,

$$\begin{aligned} \|\mu_t - F^{t-s}(m\mu')\| &\leq \|F\|^{t-s} \|\mu_s - \mu_s|_R\| \\ &= \|\mu_s|_{R^c}\| = \mu_s(R^c) \leq \lambda^s < \frac{\epsilon}{3}. \end{aligned} \quad (2.32)$$

Since $t - s \geq K$, again using linearity and the fact that $m \leq 1$,

$$\|F^{t-s}(m\mu') - m\mu_\infty\| = m \|F^{t-s}(\mu') - \mu_\infty\| < \frac{\epsilon}{3}. \quad (2.33)$$

Finally, again using Proposition 2.21

$$\|m\mu_\infty - \mu_\infty\| = (1 - m) \|\mu_\infty\| = 1 - \mu_s(R) \leq \lambda^s < \frac{\epsilon}{3}. \quad (2.34)$$

We see that for all $t \geq K + s$, $\|\mu_t - \mu_\infty\| < \epsilon$, so the information state μ_t converges to μ_∞ in total variation norm and therefore in distribution. \square

Conjecture 2.25. *The conditions of Theorem 2.24 can be weakened to the case when the underlying chain is only ergodic.*

Idea of Proof. Given an ergodic finite transition matrix T , some power T^k will be positive, and thus we should still have convergence by taking time in k -step increments, since the information state will not fluctuate too much within those k steps. The difficulty lies in the fact that the information state taken in k -step increments is not the same as the information state of the k -step chain. □

It is our belief that the information state converges in all but a small number of pathological examples; however, we are only able to prove it in the above cases. If the information state does not converge, then it does not make sense to consider a limiting expected entropy. However, it is possible that a Cesàro sum of the expected entropy converges, and the limsup and liminf will certainly exist. Alternatively, we could simply work with a finite time horizon.

2.3 Special Case

Continuing further within the general case has proven to be quite difficult, so we will restrict the remainder of our results to a special case, where there are two states and two observation processes with two possible observations each, and each observation process observes a different underlying state perfectly. Formally, $|S| = |V| = |O| = 2$, and the transition and observation matrices are

$$T = \begin{bmatrix} a & 1-a \\ 1-b & b \end{bmatrix} \quad M^{(0)} = \begin{bmatrix} 1 & 0 \\ p & 1-p \end{bmatrix} \quad M^{(1)} = \begin{bmatrix} 1-q & q \\ 0 & 1 \end{bmatrix}. \quad (2.35)$$

In order to exclude trivial cases, we will require that the parameters a, b, p and q are contained in the open interval $(0, 1)$, and $a + b \neq 1$.

Note that this special case exactly corresponds to the problem of searching for a moving target studied by MacPhee and Jordan [11], although our cost function, limiting expected entropy, is very different from theirs, expected cumulative sum of prescribed costs until the first zero-entropy observation.

We will begin by proving that given this restriction, the information state always converges in distribution, except for one case which is pathological in a sense that will be explained later. This proof is both a correction and an improvement of the proof given in [18].

In the special case, the space $\mathcal{P}(S)$ is a 1-simplex embedded in \mathbb{R}^2 , which we can identify with the interval $[0, 1]$, via equating each point $z \in [0, 1]$ with the point $z\delta(0) + (1-z)\delta(1) \in \mathcal{P}(S)$, so that z represents the mass at 0 in the distribution.

By substituting these parameters into into Definition 2.14, the transition function in the special case is

$$\begin{aligned}
F(\mu) = & \int_{A_0} \alpha_0(z)\delta(r_0(z))d\mu(z) + \int_{A_1} \alpha_1(z)\delta(r_1(z))d\mu(z) & (2.36) \\
& + \int_{A_0} (1 - \alpha_0(z))d\mu(z)\delta(0) + \int_{A_1} (1 - \alpha_1(z))d\mu(z)\delta(1),
\end{aligned}$$

where:

- $\alpha_0(z) = (1 - p)(a + b - 1)z + 1 - b + pb$;
- $\alpha_1(z) = (1 - q)(1 - a - b) + b + q - bq$;
- $r_0(z) = \frac{(a + b - 1)z + 1 - b}{\alpha_0(z)}$; and
- $r_1(z) = \frac{q(a + b - 1)z + q - qb}{\alpha_1(z)}$.

Note that the second line of (2.36) consists of two point masses at 0 and 1, which is a feature of the anchoredness condition. In the special case, it allows us to represent the location of masses by only two r -functions.

We will continue to use the symbols α_0 , α_1 , r_0 and r_1 in their meaning above for the remainder of our discourse. Note that we have simplified the notation, in that r_0 represents the r -function $r_{0,0}$, while $r_{0,1}$ does not appear since $r_{0,1}(z) = 0$ identically, and similarly for the other symbols.

Since the special case satisfies the conditions of positivity and anchoredness, Proposition 2.22 applies, so irreducibility and aperiodicity of the information chain are sufficient for ergodicity. We now show that this occurs in all but two exceptional cases:

- Case 1: Each orbit is contained entirely in the policy region which maps to that orbit, that is, $R_0 \subseteq A_0$ and $R_1 \subseteq A_1$. Note that by

(2.36), the α -functions are always strictly positive, so we must have $R_0 = \{0, r(0), r^2(0), \dots\}$ and $R_1 = \{1, r(1), r^2(1), \dots\}$.

- Case 2: The orbits alternate periodically between policy regions, that is, $\{1, r(0), r^2(1), r^3(0), \dots\} \in A_0$ and $\{0, r(1), r^2(0), r^3(1), \dots\} \in A_1$, where r is the combined r -function $r(z) = r_0(z)\mathbb{1}_{A_0}(z) + r_1(z)\mathbb{1}_{A_1}(z)$.

Let Case 0 denote the general case when neither Case 1 nor Case 2 occurs.

Lemma 2.26. *The chain $F|_R$ has only one irreducible recurrent class, except in Case 1, where it splits into two irreducible recurrent classes, both of which are positive recurrent.*

Proof. By Proposition 2.22, without loss of generality, assume that the state 0 is positive recurrent.

If 1 is reachable from 0, that is, there is some t such that $(F|_R)^t(0, 1) > 0$, then 0 is also reachable from 1 since 0 is recurrent, hence 0 and 1 are in the same irreducible recurrent class. By (2.36), either 0 or 1 is reachable from every state, so there cannot be any other recurrent classes.

Otherwise, if 0 is reachable from 1 but 1 is not reachable from 0, then 1 is transient, and furthermore, all of R_1 is transient since any $r^k(1) \in R_1$ is reachable only via the transient state 1, while all of R_0 is reachable from the recurrent state 0 and hence forms the only irreducible recurrent class.

Finally, if 0 and 1 are both unreachable from each other, then it must be the case that $R_0 \subseteq A_0$ and $R_1 \subseteq A_1$, in which case the chain splits into two irreducible classes R_0 and R_1 , both of which are positive recurrent by the argument in Proposition 2.22. \square

Lemma 2.27. *The recurrent classes are aperiodic, except in Case 2, where the chain consists of a single irreducible recurrent class with period 2.*

Proof. If any recurrent class is periodic, then at least one of 0 or 1 is recurrent and periodic; without loss of generality suppose it is 0. Since 0 is periodic, it cannot reach itself in 1 step, so must be contained in A_1 , thus 1 is reachable from 0 and hence is in the same irreducible recurrent class. Note that by the same argument, $1 \in A_0$, thus 0 reaches itself in 2 steps and hence the period must be 2.

This means 0 cannot reach itself in an odd number of steps, so $r^k(0) \in A_1$ when k is even and $r^k(0) \in A_0$ when k is odd, and similarly for the orbit of 1, which is the only possibility of periodicity. \square

Thus, there is one exception in which the information chain is periodic, and another in which it is reducible. As will be evident, both exceptions can be remedied. We begin by giving a class of policies under which reducibility cannot occur. This class of policies is simple enough to be easily analysed, and as will be conjectured later, always includes the optimal policy in any hidden Markov model within the special case.

Definition 2.28. A policy g is called a **Threshold Policy** if its preimages $A_0 = g^{-1}\{0\}$ and $A_1 = g^{-1}\{1\}$ are both intervals.

A threshold policy is indeed given by a threshold, since there must be some boundary point between A_0 and A_1 , such that one observation process is used on one side and the other is used on the other side. Outside the threshold case, it is unclear what the equivalent definition would be, since the concept

of an interval does not generalise easily to higher dimensions.

Lemma 2.29. *The linear fractional transformations r_0 and r_1 satisfy the inequality $r_0(z) > r_1(z)$ for all $z \in [0, 1]$.*

Proof. We can write $r_0(z) = \frac{(a+b+1)z+1-b}{p((1-a)z+b(1-z))+ (a+b-1)z+1-b}$. Note that the coefficient $(1-a)z + b(1-z)$ of p is strictly positive and in the denominator, while $r_1(z)$ is exactly the same with q^{-1} instead of p . Since $p, q \in (0, 1)$, $p < 1 < q^{-1}$, hence $r_0(z) > r_1(z)$. \square

Lemma 2.30. *The linear fractional transformations r_0 and r_1 are both strictly increasing when $a+b > 1$ and both strictly decreasing when $a+b < 1$.*

Proof. The derivative of r_0 is $r'_0(z) = \frac{p(a+b-1)}{((1-p)(a+b-1)z+1-b+pb)^2}$, which is positive everywhere if $a+b > 1$ and negative everywhere if $a+b < 1$. The same holds for r_1 , since it is identical with q^{-1} instead of p . \square

Lemma 2.31. *The linear fractional transformations r_0 and r_1 have unique fixed points η_0 and η_1 , which are global attractors of their respective dynamical systems.*

Proof. Split the interval $[0, 1]$ into open subintervals at its interior fixed points, noting that by inspection, the two boundary points are not fixed. Since linear fractional transformations have at most two fixed points, there must be between one and three such subintervals, all non-degenerate.

By continuity, in any such subinterval I , either $r_0(z) > z$ for all $z \in I$, or $r_0(z) < z$ for all $z \in I$. Since $r_0(0) > 0$, $r_0(z) > z$ for all points z in the

leftmost subinterval, and since $r_0(1) < 1$, $r_0(z) < z$ for all points z in the rightmost subinterval. This means that there are at least two subintervals.

If there were three subintervals, I_1 , I_2 and I_3 in that order, then either $r_0(z) > z$ for all $z \in I_2$, or $r_0(z) < z$ for all $z \in I_2$. In the first case, the fixed point between I_1 and I_2 is attracting on the left and repelling on the right, and in the second case, the fixed point between I_2 and I_3 is repelling on the left and attracting on the right. However, such fixed points can only occur for parabolic linear fractional transformations with only one fixed point. This is a contradiction, hence there cannot be three subintervals.

Thus, there are exactly two subintervals, and exactly one fixed point, which must then be an attractor across all of $[0, 1]$. \square

Lemma 2.32. *The fixed points satisfy $\eta_1 < \eta_0$.*

Proof. By Lemma 2.29, $r_1(\eta_0) < r_0(\eta_0) = \eta_0$.

First consider the case $a + b > 1$. Applying Lemma 2.30 k times gives $r_1^{k+1}(\eta_0) < r_1^k(\eta_0)$, so the orbit of η_0 under r_1 is monotonically decreasing, but it also converges to η_1 by Lemma 2.31, hence $\eta_1 < \eta_0$.

In the remaining case $a + b < 1$, suppose $\eta_1 \geq \eta_0$. Then by Lemma 2.30, $\eta_1 = r_1(\eta_1) \leq r_1(\eta_0) < \eta_0$, which is a contradiction, hence $\eta_1 < \eta_0$. \square

Proposition 2.33. *The first exception to Theorem 2.24, Case 1, cannot occur under a threshold policy.*

Proof. Suppose $R_0 \subseteq A_0$, $R_1 \subseteq A_1$, and the policy is threshold. Since $0 \in R_0$

and $1 \in R_1$, it follows that every point of R_0 is less than every point of R_1 . Since R_0 and R_1 are the orbits of 0 and 1 under r_0 and r_1 respectively, by Lemma 2.31, they have limit points η_0 and η_1 respectively, hence $\eta_0 \leq \eta_1$. This contradicts Lemma 2.32, hence this cannot occur. \square

The remaining exception is when the information chain is periodic with period 2, in which case the expected entropy oscillates between two limit points. The limiting expected entropy can still be defined in a sensible way, by taking the average, minimum or maximum of the two limit points, depending on which is most appropriate for the situation. Thus, for threshold policies, it is possible to define optimality without exception.

We conclude this section by writing down a closed form general expression for the limiting expected entropy.

Proposition 2.34. *Under the conditions of Theorem 2.24, that is, in Case 0, the limiting expected entropy of a policy is given by*

$$H_\infty = \frac{C^{(0)}(H) + C^{(1)}(H)}{C^{(0)}(\mathbf{1}) + C^{(1)}(\mathbf{1})}, \quad (2.37)$$

where, for $i \in \{0, 1\}$:

- $H(z) = -z \log z - (1-z) \log(1-z)$ is the entropy function and $\mathbf{1}(z) = 1$ is the constant function with value 1;
- $r(z) = r_0(z) \mathbb{1}_{A_0}(z) + r_1(z) \mathbb{1}_{A_1}(z)$ and $\alpha(z) = \alpha_0(z) \mathbb{1}_{A_0}(z) + \alpha_1(z) \mathbb{1}_{A_1}(z)$ are the combined r -function and combined α -function respectively, with r_0, r_1, α_0 and α_1 defined as in (2.36);
- $z_k^{(i)} \in \mathcal{P}(S) = [0, 1]$ are defined by the recursion

$$z_0^{(i)} = i, \quad z_{k+1}^{(i)} = r(z_k^{(i)}); \quad (2.38)$$

- $c_k^{(i)} \in \mathbb{R}$ are defined by the recursion

$$c_0^{(i)} = 1, \quad c_{k+1}^{(i)} = \alpha(z_k^{(i)})c_k^{(i)}; \quad (2.39)$$

- $C^{(i)} : \mathcal{C}(\mathcal{P}(S)) \rightarrow \mathbb{R}$ is the linear functional

$$C^{(i)}(f) = \sum_{k \in \mathbb{Z}^+} c_k^{(1-i)} (1 - \alpha_i(z_k^{(1-i)})) \mathbb{1}_{\{z_k^{(1-i)} \in A_i\}} \sum_{k \in \mathbb{Z}^+} c_k^{(i)} f(z_k^{(i)}). \quad (2.40)$$

Proof. By Theorem 2.24, a unique invariant probability measure Z_∞ exists, so it suffices to show that

$$Z_\infty = \frac{C^{(0)}(\delta) + C^{(1)}(\delta)}{C^{(0)}(\mathbf{1}) + C^{(1)}(\mathbf{1})}, \quad (2.41)$$

where $\delta : \mathcal{P}(S) \rightarrow \mathcal{M}(\mathcal{P}(S))$ is the Dirac measure, and $C^{(i)}$ is extended in the obvious way to a linear functional $\mathcal{L}(\mathcal{P}(S), \mathcal{M}(\mathcal{P}(S))) \rightarrow \mathcal{M}(\mathcal{P}(S))$.

By Proposition 2.21, the invariant measure is supported in the combined orbit set R , and $z_k^{(i)}$ is the only point which can make a one-step transition under the restricted information chain $F|_R$ to $z_{k+1}^{(i)}$, with probability $\alpha(z_k^{(i)})$. Since the masses at $z_k^{(i)}$ and $z_{k+1}^{(i)}$ are invariant, we must have

$$\mathbb{P}(Z_\infty = z_{k+1}^{(i)}) = \alpha(z_k^{(i)})\mathbb{P}(Z_\infty = z_k^{(i)}). \quad (2.42)$$

It then follows that for some constants $B^{(0)}, B^{(1)} \in \mathbb{R}$,

$$Z_\infty = B^{(0)} \sum_{k \in \mathbb{Z}^+} c_k^{(0)} \delta(z_k^{(0)}) + B^{(1)} \sum_{k \in \mathbb{Z}^+} c_k^{(1)} \delta(z_k^{(1)}). \quad (2.43)$$

The mass at $z_k^{(i)}$ is $B^{(i)}c_k^{(i)}$, which makes a transition under $F|_R$ to 0 in one step with probability $(1 - \alpha_0(z_k^{(i)}))\mathbb{1}_{\{z_k^{(i)} \in A_0\}}$. Since Z_∞ is an invariant measure, mass is conserved at $z_0^{(0)} = 0$, hence

$$\sum_{i \in \{0,1\}} B^{(i)} \sum_{k \in \mathbb{Z}^+} c_k^{(i)} (1 - \alpha(z_k^{(i)})) \mathbb{1}_{\{z_k^{(i)} \in A_0\}} = B^{(0)}. \quad (2.44)$$

Since $c_k^{(i)} \rightarrow 0$ as $k \rightarrow \infty$, by telescoping series,

$$\sum_{k \in \mathbb{Z}^+} c_k^{(i)} (1 - \alpha(z_k^{(i)})) = \sum_{k \in \mathbb{Z}^+} (c_k^{(i)} - c_{k+1}^{(i)}) = 1. \quad (2.45)$$

Multiplying the right hand side of (2.44) by (2.45) with $i = 0$, then collecting coefficients of $B^{(0)}$ and $B^{(1)}$, yields

$$B^{(0)} \sum_{k \in \mathbb{Z}^+} c_k^{(0)} (1 - \alpha(z_k^{(0)})) \mathbb{1}_{\{z_k^{(0)} \in A_1\}} = B^{(1)} \sum_{k \in \mathbb{Z}^+} c_k^{(1)} (1 - \alpha(z_k^{(1)})) \mathbb{1}_{\{z_k^{(1)} \in A_0\}}. \quad (2.46)$$

Note that conservation of mass at $\delta(0) = 1$ is now automatic, since mass is conserved on all of R and at every other point in R . The second equation comes from requiring Z_∞ to have total mass 1, that is,

$$B^{(0)} \sum_{k \in \mathbb{Z}^+} c_k^{(0)} + B^{(1)} \sum_{k \in \mathbb{Z}^+} c_k^{(1)} = 1. \quad (2.47)$$

The solution to (2.46) and (2.47) is exactly the required result. We note that the denominator is zero only when $z_k^{(0)} \in A_0$ and $z_k^{(1)} \in A_1$ for all $k \in \mathbb{Z}^+$, which is exactly the excluded case $R_0 \subseteq A_0$ and $R_1 \subseteq A_1$. \square

This proof can be generalised easily to the case of more than two observation processes, as long as each one is anchored with only two states.

3 Computational Results

3.1 Limiting Entropy

We now present computational results, again in the special case, which will illustrate the nature of the optimal policy in the case of minimising limiting expected entropy. Since any such discussion must cover how best to calculate the limiting expected entropy given a hidden Markov model and a policy, this is a natural place to start.

The simplest approach is to use the formula in Proposition 2.34, noting that each of $C^{(0)}(H)$, $C^{(1)}(H)$, $C^{(0)}(\mathbf{1})$ and $C^{(1)}(\mathbf{1})$ is a product of two infinite series, each of which is bounded by a geometric series and hence well-approximated by truncation. Specifically, we write the limiting expected entropy as

$$H_\infty = \frac{C_1 H_0 + C_0 H_1}{C_1 I_0 + C_0 I_1},$$

where:

- $I_i = \sum_{k \in \mathbb{Z}^+} c_k^{(i)}$;
- $H_i = \sum_{k \in \mathbb{Z}^+} c_k^{(i)} H(z_k^{(i)})$;
- $C_i = \sum_{k \in \mathbb{Z}^+} c_k^{(i)} (1 - \alpha_{1-i}(z_k^{(i)})) \mathbb{1}_{\{z_k^{(i)} \in A_{1-i}\}}$.

We can simplify the calculation slightly by recursively updating $c_k^{(i)}$ and $z_k^{(i)}$, storing them as real-valued variables c_i and z_i .

Algorithm 3.1. Estimation of limiting expected entropy.

1. Define as functions the entropy $H(z) = -z \log z - (1 - z) \log(1 - z)$, the mixed r -function $r(z) = \mathbb{1}_{A_0}(z)r_0(z) + \mathbb{1}_{A_1}(z)r_1(z)$ and the mixed α -function $\alpha(z) = \mathbb{1}_{A_0}(z)\alpha_0(z) + \mathbb{1}_{A_1}(z)\alpha_1(z)$;
2. Pick a large number N ;
3. Initialise variables $\widehat{C}_0, \widehat{C}_1, \widehat{H}_0, \widehat{H}_1, \widehat{I}_0$ and \widehat{I}_1 to 0, c_0 and c_1 to 1, z_0 to 0 and z_1 to 1;
4. Repeat the following $N + 1$ times:
 - (a) Add c_0 to \widehat{I}_0 and c_1 to \widehat{I}_1 ;
 - (b) Add $c_0 H(z_0)$ to \widehat{H}_0 and $c_1 H(z_1)$ to \widehat{H}_1 ;
 - (c) If $z_0 \in A_1$, then add $c_0(1 - \alpha_1(z_0))$ to \widehat{C}_0 , and if $z_1 \in A_0$, then add $c_1(1 - \alpha_0(z_1))$ to \widehat{C}_1 ;
 - (d) Let $z_0 = r(z_0)$ and $z_1 = r(z_1)$;
 - (e) Multiply c_0 by $\alpha(z_0)$ and c_1 by $\alpha(z_1)$;
5. The limiting expected entropy H_∞ is estimated by

$$\widehat{H}_\infty(N) = \frac{\widehat{C}_1 \widehat{H}_0 + \widehat{C}_0 \widehat{H}_1}{\widehat{C}_1 \widehat{I}_0 + \widehat{C}_0 \widehat{I}_1}.$$

Proposition 3.2. *The estimate $\widehat{H}_\infty(N)$ satisfies the error bound*

$$|\widehat{H}_\infty(N) - H_\infty| \leq \frac{16\alpha^N}{(1 - \alpha)^4 Q(N)^2},$$

where $Q(N) = \widehat{C}_1 \widehat{I}_0 + \widehat{C}_0 \widehat{I}_1$ is the denominator of $\widehat{H}_\infty(N)$, and

$$\alpha = \sup_{i,z} \alpha_i(z) = 1 - \min \left\{ b(1 - p), (1 - a)(1 - p), (1 - b)(1 - q), a(1 - q) \right\}.$$

Proof. The formula for α follows from Definition 2.14, since the α -functions are linear and hence their maxima occur at the endpoints $z = 0$ or $z = 1$. Note that $\alpha < 1$ follows from the requirement that $a, b, p, q \in (0, 1)$. Since $Q(N)$ is monotonic increasing, this implies that the error bound is finite and vanishes as $N \rightarrow \infty$.

Since each series has non-negative summands, each truncated series is smaller than the untruncated series, hence

$$\begin{aligned} \left| H_\infty - \frac{C_1 H_0 + C_0 H_1}{\widehat{C}_1 \widehat{I}_0 + \widehat{C}_0 \widehat{I}_1} \right| &= \frac{(C_1 H_0 + C_0 H_1)(C_1 I_0 + C_0 I_1 - \widehat{C}_1 \widehat{I}_0 - \widehat{C}_0 \widehat{I}_1)}{(C_1 I_0 + C_0 I_1)(\widehat{C}_1 \widehat{I}_0 + \widehat{C}_0 \widehat{I}_1)} \\ &\leq \frac{C_1 H_0 + C_0 H_1}{(\widehat{C}_1 \widehat{I}_0 + \widehat{C}_0 \widehat{I}_1)^2} ((C_1 I_0 - \widehat{C}_1 \widehat{I}_0) + (C_0 I_1 - \widehat{C}_0 \widehat{I}_1)). \end{aligned}$$

The k th summand in each series is bounded by $c_k^{(i)} \leq \alpha^k$, hence

$$\begin{aligned} C_1 I_0 - \widehat{C}_1 \widehat{I}_0 &= I_0(C_1 - \widehat{C}_1) + C_1(I_0 - \widehat{I}_0) - (C_1 - \widehat{C}_1)(I_0 - \widehat{I}_0) \\ &\leq \frac{1}{1-\alpha} \frac{\alpha^N}{1-\alpha} + \frac{1}{1-\alpha} \frac{\alpha^N}{1-\alpha} = \frac{2\alpha^N}{(1-\alpha)^2}. \end{aligned}$$

The same bound holds for $C_0 I_1 - \widehat{C}_0 \widehat{I}_1$, hence

$$\left| H_\infty - \frac{C_1 H_0 + C_0 H_1}{\widehat{C}_1 \widehat{I}_0 + \widehat{C}_0 \widehat{I}_1} \right| \leq \frac{8\alpha^N}{(1-\alpha)^4 Q(N)^2}.$$

Similarly, since $Q(N) = \widehat{C}_1 \widehat{I}_0 + \widehat{C}_0 \widehat{I}_1 \leq 2/(1-\alpha)^2$,

$$\begin{aligned} \left| \widehat{H}_\infty(N) - \frac{C_1 H_0 + C_0 H_1}{\widehat{C}_1 \widehat{I}_0 + \widehat{C}_0 \widehat{I}_1} \right| &= \frac{(C_1 H_0 - \widehat{C}_1 \widehat{H}_0) + (C_0 H_1 - \widehat{C}_0 \widehat{H}_1)}{\widehat{C}_1 \widehat{I}_0 + \widehat{C}_0 \widehat{I}_1} \\ &\leq \frac{4\alpha^N}{(1-\alpha)^2 Q(N)} \leq \frac{8\alpha^N}{(1-\alpha)^4 Q(N)^2}. \end{aligned}$$

Combining via the triangle inequality gives the required bound. \square

Note that this bound depends only on the quantities α and $Q(N)$, which are easily calculated. In particular, this allows us to prescribe a precision and

calculate the limiting expected entropy to within that precision by running the algorithm with unbounded N and terminating when error bound reaches the desired precision. Furthermore, since $Q(N)$ appears only in the denominator and grows monotonically with N , we can replace it with $Q(N_0)$ for some small, fixed N_0 to calculate a priori a sufficient number of steps for any given precision.

Example 3.3. In later simulations, we will pick $a, b, p, q \in [0.025, 0.975]$, so that $\alpha \leq 0.999375$. This gives an error bound of

$$|\widehat{H}_\infty(N) - H_\infty| \leq (16 \times 400^4)Q(N)^{-2}(0.999375)^N.$$

While the constant appears daunting at first glance, solving for a prescribed error of 10^{-10} gives

$$\log 16 + 4 \log 400 - 2 \log Q(N) + N \log 0.999375 \leq -10 \log 10.$$

Hence, we require

$$N \geq 79598 + 3199 |\log Q(N)|.$$

For any realistic value of $Q(N)$, this is easily within computational bounds, as each iteration requires at most 36 arithmetic operations, 2 calls to the policy function, and 4 calls to the logarithm function.

An alternative approach to estimating limiting expected entropy would be to simulate observation of the hidden Markov model under the given policy. The major drawback of this method is that it requires working with the information state Z_t , which takes values in $\mathcal{P}(\mathcal{P}(S)) = \mathcal{P}([0, 1])$, the set of probability measures on the unit interval, which is an infinite dimensional space.

One possibility is to divide $[0, 1]$ into subintervals and treat each subinterval as a discrete atom, but this produces a very imprecise result. Even using an unrealistically small subinterval width of 10^{-6} , the entropy function has a variation of over 10^{-5} across the first subinterval, restricting the error bound to this order of magnitude regardless of the number of iterations. In comparison, Example 3.3 shows that the direct estimation method has greatly superior performance.

An improvement is to use the fact that the limiting distribution Z_∞ is discrete, and store a list of pairs containing the locations and masses of discrete points. Since any starting point moves to either 0 or 1 in one step, at the N th iteration, the list of points must contain at least the first N points in the orbit of either 0 or 1. Each such point requires a separate calculation at each iteration, and thus the number of computations is $O(N^2)$ rather than $O(N)$ as for Algorithm 3.1.

Since the number of iterations N corresponds to the last point in the orbit of 0 or 1 which is represented, for any given N , this method differs from the direct computation method only in the masses on these points, thus we would expect the relationship between precision and number of iterations to be similar. Since the simulation method has quadratically growing number of computations, this would suggest that it is slower than the direct computation method, and indeed, this is also indicated by empirical trials.

We will use the direct computation method of estimating limiting expected entropy for all of our results.

3.2 Optimal Threshold Policy

The problem of finding the policy which minimises limiting expected entropy is made much easier by restricting the search space to the set of threshold policies, as these can be represented by a single number representing the threshold, and a sign representing which observation process is used on which side of the threshold.

The simplest approach is to pick a collection of test thresholds uniformly in $[0, 1]$, either deterministically or randomly, and test the limiting expected entropy at these thresholds, picking the threshold with minimal entropy as the optimal threshold policy. However, this method is extremely inefficient. Proposition 2.21 shows that the policy only matters on the countable set $R \subset [0, 1]$, so moving the threshold does not change the policy as long as it does not move past a point in R .

As shown in Figures 3.2–3.7, points in R tend to be quite far apart, and thus the naive approach will cause a large number of equivalent policies to be tested. On the other hand, points in R close to the accumulation points are closely spaced, so even with a very fine uniform subset, some policies will be missed when the spacing between points in R becomes less than the spacing between test points.

A better way is to decide on a direction in which to move the threshold, and select the test point exactly at the next point in the previous realisation of R in the chosen direction, so that every threshold in between the previous point and the current point gives a policy equivalent to the previous policy. This ensures that each equivalence class of policies is tested exactly once,

thus avoiding the problem with the naive method.

However, a new problem is introduced in that the set of test points depends on the iteration number N , which determines the number of points of R that are considered. This creates a circular dependence, in that the choice of N depends on the desired error bound, the error bound depends on the policy, and the set of policies to be tested depends on N . We can avoid this problem by adapting Proposition 3.2 to a uniform error bound across all threshold policies.

Proposition 3.4. *For a threshold policy and $N \geq L$, the error is*

$$|\widehat{H}_\infty(N) - H_\infty| < \frac{16\alpha^N}{\bar{\alpha}^{2L}(1-\alpha)^6},$$

where L is the smallest integer such that $r_0^L(0) > r_1^L(1)$, and

$$\bar{\alpha} = \inf_{i,z} \alpha_i(z) = 1 - \max \left\{ b(1-p), (1-a)(1-p), (1-b)(1-q), a(1-q) \right\}.$$

Proof. First note that L exists, since by Lemma 2.32, iterations of r_0 and r_1 converge to $\eta_0 > \eta_1$ respectively.

Using Proposition 3.2, it suffices to prove that for $N \geq L$,

$$Q(N) = \widehat{C}_1 \widehat{I}_0 + \widehat{C}_0 \widehat{I}_1 \geq \bar{\alpha}^L (1 - \alpha).$$

It is not possible for $\{0, r(0), \dots, r^L(0)\} \subset A_0$ and $\{1, r(1), \dots, r^L(1)\} \subset A_1$, since this would mean $r^L(0) = r_0^L(0)$ and $r^L(1) = r_1^L(1)$, which gives the ordering $0 < r^L(1) < r^L(0) < 1$, but A_0 and A_1 are intervals for a threshold policy.

Hence, either $z_\ell^{(0)} = r^\ell(0) \in A_1$ or $z_\ell^{(1)} = r^\ell(1) \in A_0$ for some $\ell \leq L \leq N$. If $z_\ell^{(0)} \in A_1$, then $\widehat{C}_0 \geq c_\ell^{(0)}(1 - \alpha_1(z_\ell^{(0)})) \geq \bar{\alpha}^\ell (1 - \alpha)$. Since $\widehat{I}_1 \geq c_0^{(1)} = 1$, this

gives $Q(N) \geq \bar{\alpha}^L(1 - \alpha)$, as required. A similar argument holds in the case $z_\ell^{(1)} \in A_0$. \square

The existence of this uniform bound for $Q(N)$ in the threshold case is closely related to Proposition 2.33, which states that the exception Case 1, where $R_0 \subset A_0$ and $R_1 \subset A_1$, cannot occur in a threshold policy. In this exceptional case, Proposition 2.34 does not hold, as the denominator is zero, and hence $Q(N) = 0$ for all N . The fact that this cannot occur in a threshold policy is the key ingredient of this uniform bound.

Now that we have an error bound which does not depend on the policy, we can determine a uniform number of iterations N that will suffice for estimating the limiting expected entropy for all threshold policies. This reduces the search space to a finite one, as each point in the orbits of 0 and 1 must be in one of the two policy regions, hence, there are at most 2^{2N+2} policies. Most of these will not be threshold policies, but since orbit points need not be ordered, there is no obvious bound on the number of threshold policies that need to be checked. Simulation results later in this section will show that in most cases, the number of such policies is small enough to be computationally feasible.

Definition 3.5. The **Orientation** of a threshold policy is the pair of whether A_0 is to the left or right of A_1 , and whether the threshold is included in the left or right interval. Let $[A_0][A_1]$, $[A_0](A_1)$, $[A_1][A_0]$ and $[A_1](A_0)$ denote the four possibilities, with the square bracket indicating inclusion of the threshold and round bracket indicating exclusion.

Our strategy for simplifying the space of threshold policies that need to

be considered is to note that the policy only matters on R , the support of the invariant measure. Although R depends on the policy, for a given orientation and threshold t , any policy with the same orientation and some other threshold t' such that no points of R lie between t and t' is an equivalent policy, in that sense that the invariant measure is the same, since no mass exists in the region where the two policies differ.

Thus, for each orientation, we can begin with $t = 1^+$, and at each iteration, move the threshold left past the next point in R , since every threshold in between is equivalent to the previous threshold. Although R changes at each step, this process must terminate in finite time since we already showed that there are only finitely many policies for any given N , and by testing equivalence classes of policies only once, it is likely that far fewer steps are required than the 2^{2N+2} bound.

Furthermore, since R is a discrete set, every threshold policy has an interval of equivalent threshold policies, so we can assume without loss of generality that the threshold is contained in the interval to the right, that is, only test the orientations $[A_0][A_1]$ and $[A_1][A_0]$.

Algorithm 3.6. Finding the optimal threshold policy.

1. Find L , the smallest integer such that $r_0^L(0) > r_1^L(1)$, by repeated application of r_0 and r_1 to 0 and 1 respectively;
2. Prescribe an error E and determine the number of iterations

$$N = \frac{\log E - \log 16 + 2L \log \bar{\alpha} + 6 \log(1 - \alpha)}{\log \alpha};$$

3. Start with the policy $A_0 = [0, t)$ and $A_1 = [t, 1]$ with $t = 1^+$, that is,

A_0 is the whole interval and A_1 is empty but considered as being to the right of A_0 , and loop until $t = 0$:

- (a) Run Algorithm 3.1, and let the next value of t be the greatest $z_k^{(i)}$ which is strictly less than the previous value of t ;
- (b) If entropy is less than the minimum entropy for any policy encountered so far, record it as the new minimum;

4. Repeat with A_0 considered to be on the right of A_1 .

We now calculate the location of the optimal threshold for a range of values of the parameters a , b , p and q . The results of this calculation, which will be presented in Figure 3.8, is the primary content of this section, as it will give an empirical description of the optimal threshold policy. Since we have not been able to prove optimality analytically, except in the symmetric case of Proposition 3.8, this empirical description will provide our most valuable insight into the problem.

In order to facilitate understanding, we will define six classes of threshold policies, depending on the location and orientation of the threshold. Diagrams of the invariant measure under each of these classes, presented in Figures 3.2–3.7, will demonstrate that these classes of thresholds are qualitatively different, which will further manifest itself in the next section.

We partition the interval into three subintervals with endpoints at the attractors η_0 and η_1 , noting that this produces three regions which consist entirely of a single equivalence class of policy. We also assign colours to these regions to accommodate the data presented in Figure 3.8.

Definition 3.7. The six **Threshold Regions** are defined by partitioning the set of all threshold policies first by whether the orientation is $[A_0][A_1]$ or $[A_1][A_0]$, then by the location of the threshold t in relation to the accumulation points η_1 and η_0 , with inclusion of endpoints defined more precisely overleaf. We number them using Roman numerals I through VI.

Note that if either $t \leq 0$ or $t > 1$, then either A_0 or A_1 occupies the whole interval, depending on the orientation. In particular, $[A_0][A_1]$ with $t = 0$ is equivalent to $[A_1][A_0]$ with $t = 1^+$, since in both cases A_1 is the whole interval, and similarly $[A_0][A_1]$ with $t = 1^+$ is equivalent to $[A_1][A_0]$ with $t = 0$.

Thus, the space of all possible threshold policies consists of two disjoint intervals $[0, 1^+]$, each of whose endpoints is identified with an endpoint of the other interval, which is topologically equivalent to a circle. To be technically correct, we note that identifying 0 and 1^+ does not present a problem, since we can simply extend the interval to $[0, 1 + \epsilon]$ for some small $\epsilon > 0$, and identify 0 and $1 + \epsilon$ instead. While this results in a subinterval of the circle corresponding to the same policy, this does not add additional complexity since every threshold policy has an interval of equivalent policies.

This topology is illustrated overleaf in Figure 3.1, followed by precise definitions and examples of the six threshold regions in Figures 3.2–3.7, and finally our computational results in Figure 3.8.

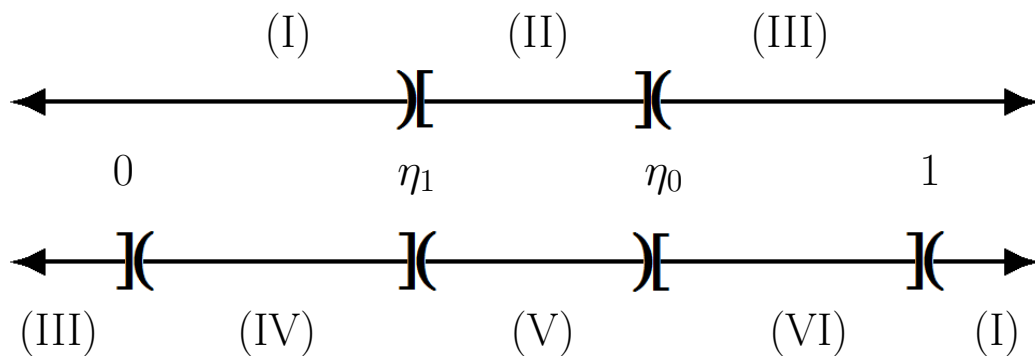


Figure 3.1: Space of all threshold policies. The top line represents the orientation $[A_0][A_1]$, while the bottom line represents the orientation $[A_1][A_0]$. The right end of the top line and the left end of the bottom line are both the policy $A_0 = [0, 1]$ and $A_1 = \emptyset$, so we can paste them together, and similarly for the left end of the top line and the right end of the bottom line. Hence, we see that the set of threshold policies is topologically a circle.

Region I (represented by red in Figure 3.8): $[A_0][A_1]$ with $t < \eta_1$ or $[A_1][A_0]$ with $t > 1$. When $a + b > 1$, every policy here is equivalent to the all- A_1 policy, as the mass in the orbit of 0 approaches 0.




Figure 3.2: Invariant measure for the unique Region I policy with $a = 0.8$, $b = 0.3$, $p = 0.5$, $q = 0.3$. Entropy is 0.3145. Under the evolution function, any mass eventually enters the grey region since it contains both accumulation points in its interior, after which it cannot escape, hence in the limit, there is zero mass in the orbit of 0, and the policy is equivalent the all- A_1 policy.

Region II (yellow): $[A_0][A_1]$ with $\eta_1 \leq t \leq \eta_0$. This is the most difficult to understand of the threshold policies, as the orbits do not converge to the accumulation points η_0 and η_1 , but rather, oscillate around the threshold t .




Figure 3.3: Invariant measure for a typical Region II policy with $a = 0.8$, $b = 0.3$, $p = 0.5$, $q = 0.3$. Entropy is 0.3251. Note that the masses do not converge to the accumulation points.

Region III (green): $[A_0][A_1]$ with $t > \eta_0$ or $[A_1][A_0]$ with $t \leq 0$. When $a + b > 1$, every policy here is equivalent to the all- A_0 policy, since the mass in the orbit of 1 approaches 0.

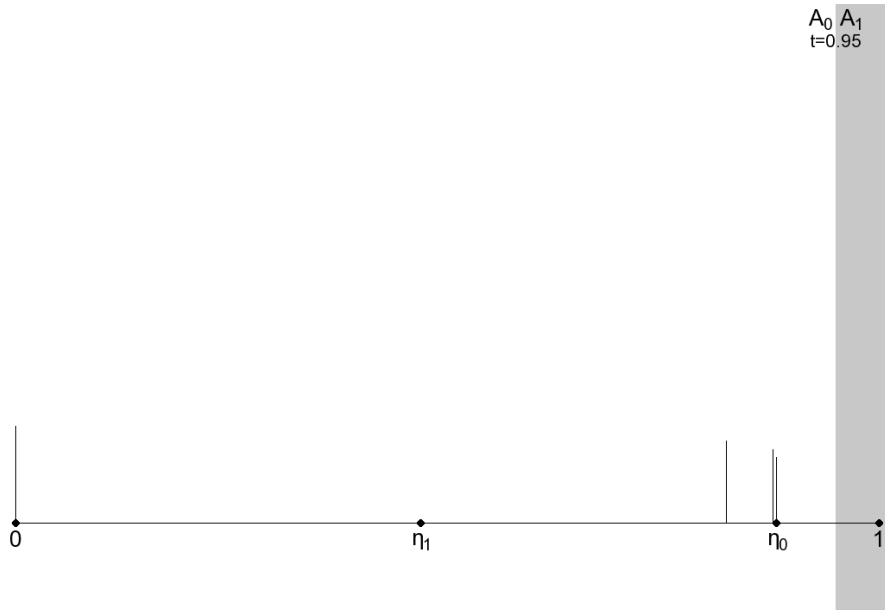


Figure 3.4: Invariant measure for the unique Region III policy with $a = 0.8$, $b = 0.3$, $p = 0.5$, $q = 0.3$. Entropy is 0.3337. Under the evolution function, any mass eventually enters the white region since it contains both accumulation points in its interior, after which it cannot escape, hence in the limit, there is zero mass in the orbit of 1, and the policy is equivalent to the all- A_0 policy.

Region IV (cyan): $[A_1][A_0]$ with $0 < t \leq \eta_1$. For a policy in this region, the first finitely many points of the orbit of 0 belongs to A_1 , while every other point lies in A_0 . Note that $t \leq 0$ is included in Region III.

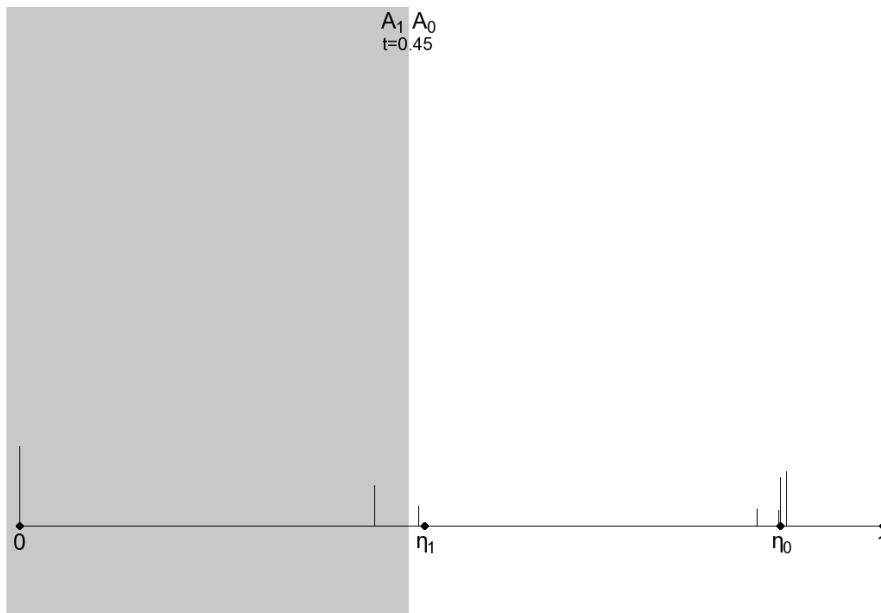


Figure 3.5: Invariant measure for a typical Region IV policy with $a = 0.8$, $b = 0.3$, $p = 0.5$, $q = 0.3$. Entropy is 0.3275. The orbit of 1 approaches η_0 , while the orbit of 0 follows an approach sequence to η_1 for a finite number of steps (in this case 2 steps) before also approaching η_0 .

Region V (blue): $[A_1][A_0]$ with $\eta_1 < t < \eta_0$. When $a + b > 1$, every policy here is equivalent to the policy $R_0 \subset A_1$ and $R_1 \subset A_0$.




Figure 3.6: Invariant measure for the unique Region V policy with $a = 0.8$, $b = 0.3$, $p = 0.5$, $q = 0.3$. Entropy is 0.3265. Note that the orbit of 0 converges to η_1 while the orbit of 1 converges to η_0 .

Region VI (magenta): $[A_1][A_0]$ with $\eta_0 \leq t \leq 1$. For a policy in this region, the first finitely many points of the orbit of 0 belongs to A_1 , while every other point lies in A_0 . Note that $t > 1$ is included in Region I. The lack of symmetry with Region VI in terms of strict and non-strict inequalities is due to the choice that the threshold itself be included in the right region.

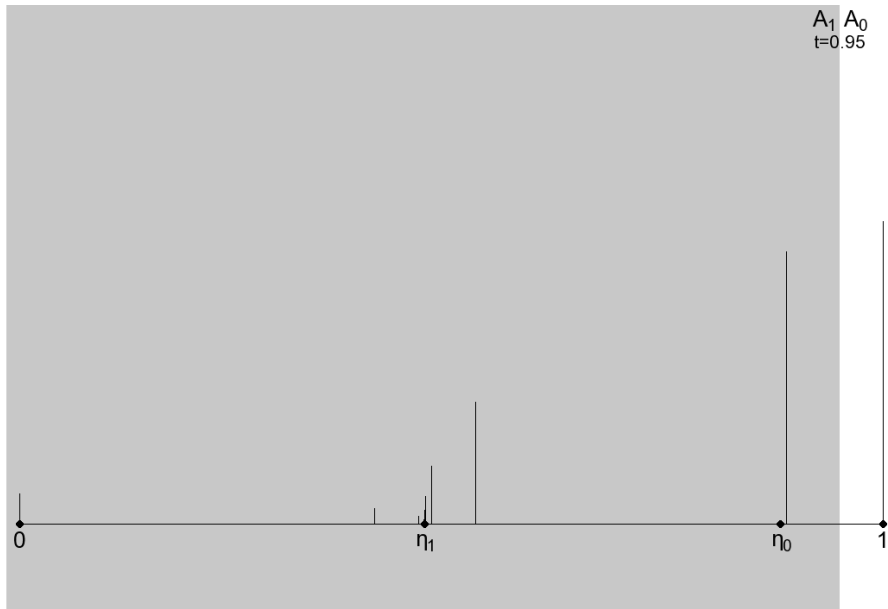


Figure 3.7: Invariant measure for a typical Region VI policy with $a = 0.8$, $b = 0.3$, $p = 0.5$, $q = 0.3$. Entropy is 0.3182. The orbit of 0 approaches η_1 , while the orbit of 1 follows an approach sequence to η_0 for a finite number of steps (in this case 1 step) before also approaching η_1 .




Figure 3.8: Location of the optimal threshold policy. The four-dimensional parameter space is represented with a on the major horizontal axis increasing rightwards, b on the major vertical axis increasing downwards, p on the minor horizontal axis increasing rightwards and q on the minor vertical axis increasing downwards. Each parameter is sampled by $\{0.025, 0.075, \dots, 0.975\}$, omitting the trivial cases when $a+b = 1$, resulting in 152000 sample points. Colours representing regions are defined in Figures 3.2–3.7.

The bound in Proposition 3.4 is consistent with the empirical running time for the calculations used to generate Figure 3.8. Our program experienced significant slowdowns when either a and p , or b and q were large, with the slowest occurring at the extreme point $a = b = p = q = 0.975$. This suggests that a uniform bound when $\alpha \rightarrow 1$ is impossible, and that Algorithm 3.6 will fail for sufficiently large values of α . On the other hand, either $a = 1$ or $b = 1$ results in trivial limiting behaviour of the underlying chain, so this situation is unlikely to occur in any practical application.

Note that in Figure 3.8, the half of the main diagonal where $a + b > 1$ is entirely blue. This can be proven analytically. We note that the condition $a + b > 1$ corresponds to the underlying Markov chain having positive one-step autocorrelation, which is a reasonable assumption when the frequency of observation is greater than the frequency of change in the observed system, since in this case, one would not expect the system to oscillate with each observation.

Proposition 3.8. *In the symmetric case $a = b$ and $p = q$, with $a + b > 1$, the optimal general policy is the unique Region V policy g_V given by $R_0 \subset A_1$ and $R_1 \subset A_0$.*

Proof. Proposition 2.34 gives the limiting expected entropy as a convex combination of two quantities, hence

$$H_\infty \geq \min \left\{ \frac{\sum_k c_k^{(0)} H(z_k^{(0)})}{\sum_k c_k^{(0)}}, \frac{\sum_k c_k^{(1)} H(z_k^{(1)})}{\sum_k c_k^{(1)}} \right\}. \quad (3.1)$$

Note that equality is realised when the two quantities above are equal, which occurs under g_V , since in that case $c_k^{(0)} = c_k^{(1)}$ and $z_k^{(0)} = 1 - z_k^{(1)}$.

Let $z_k = z_k^{(0)}$, $c_k = c_k^{(0)} / \sum c_k^{(0)}$ and $H_k = H(z_k^{(0)})$. In order to prove that g_V is optimal, by symmetry, it suffices to prove that it minimises $\sum_k c_k H_k$.

First, we show that for each k and any other policy g , $H_k(g_V) \leq H_k(g)$, where the notation $H_k(g)$ makes the dependence on the policy explicit.

By Lemma 2.29, $z_k(g) = r^k(0) \geq r_1^k(0) = z_k(g_V)$. By Lemma 2.30 and our assumption that $a + b > 1$, iterations of r_0 and r_1 approach their limits monotonically, hence $z_k(g) < \eta_0$ and $z_k(g_V) < \eta_1$. These inequalities are illustrated in Figure 3.9.

Combining these inequalities gives $z_k(g_V) \leq z_k(g) < \eta_0$. Entropy is concave on the interval $[0, 1]$, and therefore on the subinterval $[z_k(g_V), \eta_0]$, hence

$$H(z_k(g)) \geq \min \{H(z_k(g_V)), H(\eta_0)\}. \quad (3.2)$$

Since $z_k(g_V) < \eta_1 < \frac{1}{2}$ and entropy is increasing on $[0, \frac{1}{2}]$, $H(z_k(g_V)) < H(\eta_1)$, which is equal to $H(\eta_0)$ by symmetry. Hence, the inequality (3.2) reduces to $H(z_k(g)) \geq H(z_k(g_V))$, that is, $H_k(g) \geq H_k(g_V)$.

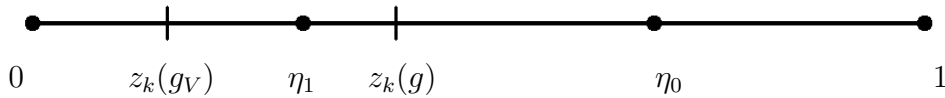


Figure 3.9: Diagram showing $z_k(g_V)$ and $z_k(g)$ in relation to η_1 , η_0 and 1. All positions are fixed except that $z_k(g)$ may be to the left of η_1 . Since $z_k(g_V)$ lies to the left of η_1 and the diagram is symmetric, $z_k(g_V)$ has lower entropy than η_0 . Since $z_k(g)$ lies between $z_k(g_V)$ and η_0 and entropy is concave, $z_k(g)$ has higher entropy than $z_k(g_V)$. Hence the g_V minimises the entropy at z_k .

Next, we show that for each k and any other policy g ,

$$c_0(g_V) + c_1(g_V) + \cdots + c_k(g_V) \geq c_0(g) + c_1(g) + \cdots + c_k(g). \quad (3.3)$$

Let $a_k = \alpha(z_k^{(0)})$. Since $a + b > 1$, by (2.36), α_1 is decreasing, while $\alpha_0(z) = \alpha_1(1 - z)$ is increasing. We have already established that $z_k(g_V) \leq z_k(g) < \eta_0$ and $z_k(g_V) < \eta_1 = 1 - \eta_0 < 1 - z_k(g)$, which implies $\alpha_1(z_k(g_V)) \geq \alpha_1(z_k(g))$ and $\alpha_1(z_k(g_V)) \geq \alpha_0(z_k(g))$ respectively. This shows that $a_k(g_V) \geq a_k(g)$.

For $\ell < k$, write

$$c_0 + \cdots + c_k = \frac{1 + a_0 + a_0 a_1 + \cdots + a_0 a_1 \cdots a_{k-1}}{1 + a_0 + a_0 a_1 + \cdots + a_0 a_1 \cdots a_{k-1} + \cdots} \equiv \frac{X + \alpha_\ell Y}{X + \alpha_\ell Z}.$$

Since $X > 0$ and $0 < Y < Z$, decreasing α_ℓ increases $c_0 + \cdots + c_k$, and the same is true for $\ell \geq k$, since in that case we can write the expression in the same way with $Y = 0$. This proves (3.3).

Using (3.3) and $H_0(g_V) < H_1(g_V) < H_2(g_V) < \cdots$, for any $\ell \in \mathbb{N}$,

$$\begin{aligned} S_\ell &\equiv \sum_{k \leq \ell} (c_k(g_V) - c_k(g)) H_\ell(g_V) + \sum_{k > \ell} (c_k(g_V) - c_k(g)) H_k(g_V) \\ &\leq \sum_{k \leq \ell} (c_k(g_V) - c_k(g)) H_{\ell+1}(g_V) + \sum_{k > \ell} (c_k(g_V) - c_k(g)) H_k(g_V) = S_{\ell+1}. \end{aligned}$$

Since $\sum_k c_k = 1$ identically, the second series vanishes as $\ell \rightarrow \infty$, while the first series is always non-negative by (3.3), hence $S_0 \leq 0$. Thus,

$$\sum_k c_k(g_V) H_k(g_V) \leq \sum_k c_k(g) H_k(g_V) \leq \sum_k c_k(g) H_k(g).$$

This proves the required minimisation. \square

Note that the proof above relies heavily on the fact that equality is attained in (3.1). This occurs only in the symmetric case, and thus this approach does

not generalise readily. The complexity of the proof in the symmetric case is indicative of the difficulty of the problem in general, and thus highlights the importance of the empirical description provided by Figure 3.8.

In the course of performing the computations to generate Figure 3.8, we noticed that entropy is unimodal with respect to threshold, with threshold considered as a circle as in Figure 3.1. While we cannot prove this analytically, it is true for each of the 152000 points in the parameter space considered.

This allows some simplification in finding the optimal threshold policy, since finding a local minimum is sufficient. Thus, we can alter Algorithm 3.6 to begin by testing only two policies, then testing policies in the direction of entropy decrease until a local minimum is found. However, the running time improvement is only a constant factor; if we model entropy as a standard sinusoid with respect to threshold, then the running time decreases by a factor of 3 on average.

3.3 General Policies

The problem of determining the optimal general policy is much more difficult, due to the complexity of the space of general policies. Since a policy is uniquely determined by the value of the policy function at the orbit points, this space can be viewed as a hypercube of countably infinite dimension, which is much more difficult to study than the space of threshold policies, which is a circle.

One strategy is to truncate the orbit and consider a finite dimensional hypercube, justified by the fact that orbit points have masses which decay geometrically, and thus the tail contributes very little. However, a truncation at N (that is, force the policy to be constant on $\{r^N(0), r^{N+1}(0), \dots\}$, and similarly for the orbit of 1) gives 2^{2N+2} possible policies, which is still far too large to determine optimality by checking the entropy of each policy.

The next approximation is to only look for locally optimal policies, in the sense that changing the policy at each of the $2N + 2$ truncated orbit points increases entropy, and hope that by finding enough such locally optimal policies, the globally optimal policy will be among them. Since a hypercube has very high connectivity, regions of attraction tend to be large, which heuristically suggests that this strategy will be effective.

Algorithm 3.9. Finding a locally optimal truncated policy.

1. Pick N , and a starting policy, expressed as a pair of sequences of binary digits $g_k^{(i)} = g(z_k^{(i)})$, with $k = 0, 1, \dots, N$;
2. Cycle through the digits $g_k^{(i)}$, flipping the digit if it gives a policy with

lower entropy, otherwise leaving it unchanged;

3. If the previous step required any changes, repeat it, otherwise a locally optimal truncated policy has been found.

We picked $N = 63$ since this allows a policy to be easily expressed as two unsigned 64-bit integers, and for each of the 152000 uniformly spaced parameters of Figure 3.8, we generated 10 policies uniformly on the hypercube and applied Algorithm 3.9.

None of the locally optimal policies for any of the parameter values had lower entropy than the optimal threshold policy from Figure 3.8, and on average 98.3% of them were equivalent to the optimal threshold policy, up to a precision of 0.1%, indicating that the optimal threshold policy is locally optimal with a very large basin of attraction, which strongly suggests that it is also the globally optimal policy.

Conjecture 3.10. *In the special case, the infimum of entropy attainable under threshold policies is the same as that under general policies.*

The fact that a large proportion of locally optimal policies have globally optimal entropy gives a new method for finding the optimal policy. By picking 10 random truncated policies and running Algorithm 3.9, at least one of them will yield an optimal policy with very high probability. Empirical observations suggest that this method is slower than Algorithm 3.6 on average, but since the success rate remains high while Algorithm 3.6 becomes significantly slower as α approaches 1, this method is a better alternative for some parameter values.




Figure 3.10: Locally optimal policies. Axes are as in Figure 3.8. Darkness increases with the proportion of simulated locally optimal policies which have the same entropy as the optimal threshold policy, up to a precision of 0.1%. The average is 9.83 out of 10, but the distribution is far from uniform—local optima are exceedingly likely to be the same as the threshold optimum for some parameter values and exceedingly unlikely for others. The boundaries are approximately those of the threshold regions (see Figure 3.8), with some imprecision due to the non-deterministic nature of the simulation data.

One last policy of interest is the greedy policy. In the previous sections, we considered a long term optimality criterion in the minimisation of limiting expected entropy, but in some cases, it may be more appropriate to set a short term goal. In particular, one may instead desire to minimise expected entropy at the next step, in an attempt to realise maximal immediate gain while ignoring future implications.

Definition 3.11. The **greedy** policy is the policy such that the expected entropy after one observation is minimised. Up to an exchange of strict and non-strict inequalities, this is given by

$$A_0 = \{z \in [0, 1] : \alpha_0(z)H(r_0(z)) < \alpha_1(z)H(r_1(z))\},$$

$$A_1 = \{z \in [0, 1] : \alpha_0(z)H(r_0(z)) \geq \alpha_1(z)H(r_1(z))\}.$$

The greedy policy has the benefit of being extremely easy to use, as it only requires a comparison of two functions at the current information state. Since these functions are smooth, efficient numerical methods such as Newton-Raphson can be used to determine the intersection points, thus allowing the policy to be described by a small number of thresholds.

In fact, only one threshold is required, as computational results show the greedy policy always a threshold policy. Using the 152000 uniformly distributed data points from before, in each case the two functions defining the greedy policy crossed over at most once.

Conjecture 3.12. *The greedy policy is always a threshold policy.*

Idea of Proof. Note that $q\alpha_0(z)r_0(z) = r_1(z)\alpha_1(z)$. It may appear at first

glance that the factor of q violates symmetry, but recall that z maps to $1 - z$ under relabelling.

Using this identity, the intersection points that define the greedy policy satisfy $G(r_0(z)) = qG(r_1(z))$, where $G(z) = H(z)/z$. It is easy to see that G is monotonic decreasing on $[0, 1]$ with range $[0, \infty]$, hence $F(z) = G^{-1}(qG(z))$ is a well-defined one-parameter family of functions mapping $[0, 1]$ to itself with fixed points at 0 and 1.

On the other hand, $f(y) = r_0(r_1^{-1}(y)) = y/((1 - pq)y + pq)$ is also a one-parameter family of functions mapping $[0, 1]$ to itself with fixed points at 0 and 1. Since the range of r_0 is contained in $(0, 1)$, we can discount the endpoints $y = 0$ and $y = 1$, hence it suffices to show that the equation $F(y) = f(y)$ has at most one solution for $y \in (0, 1)$.

Convexity methods may help in this last step but we have not been able to complete the proof. □

Even when the greedy policy is not optimal, it is very close to optimal. Of the 152000 uniformly distributed data points in Figure 3.11 below, the greedy policy is non-optimal at only 6698 points, or 4.41%, up to an error tolerance of 10^{-12} . On average, the greedy policy has entropy 0.0155% higher than the optimal threshold policy, with a maximum error of 5.15% occurring at the sample point $a = 0.975$, $b = 0.925$, $p = 0.675$ and $q = 0.375$. Thus the greedy polices provides an alternative suboptimal policy which is very easy to calculate and very close to optimal.




Figure 3.11: Optimality of the greedy policy. Axes are as in Figure 3.8. Light grey indicates that the greedy policy is the optimal threshold policy; darker points indicate suboptimality with darkness proportional to error. Similarly to Figure 3.10, the suboptimal points lie on the boundaries of the threshold regions.

We make a final remark that the likelihood of a locally optimal policy being globally optimal as shown in Figure 3.10, and the closeness of the greedy policy to the optimal threshold policy as shown in Figure 3.11, both exhibit a change in behaviour at the boundaries of the threshold regions as shown in Figure 3.8. This suggests that these regions are indeed qualitatively different, and are likely to be interesting objects for further study.

4 Conclusion and Future Work

This thesis presents initial groundwork in the theory of hidden Markov models with multiple observation processes. We prove a condition under which the information state converges in distribution, and give algorithms for finding the optimal policy in a special case, which provides strong evidence that the optimal policy is a threshold policy.

Future work will aim to prove the conjectures that we were not able to prove:

- The information state converges for an anchored hidden Markov model with only ergodicity rather than positivity;
- The greedy policy is always a threshold policy;
- Among threshold policies, the limiting expected entropy is unimodal with respect to threshold; and
- The optimal threshold policy is also optimal among general policies.

Possible approaches to these problems are likely to be found in [9], [10] and [11]. The author was not aware of these papers until after the date of submission, and thus was unable to incorporate their methods into this thesis.

Better algorithms and error bounds for finding the optimal policy are also a worthwhile goal. Although our algorithms are computationally feasible with reasonable prescribed errors, our focus was on finding workable rather than optimal algorithms, and thus there is plenty of room for improvement.

Another direction for future work would be to extend our results into more general cases.

References

- [1] L.E. Baum and T. Pétrie. Statistical inference for probabilistic functions of finite state Markov chains. *Annals of Mathematical Statistics*. Volume 6, Number 37, Pages 1554–1563, 1966.
- [2] O. Cappé, E. Moulines and T. Rydén. *Inference in Hidden Markov Models*. Springer, New York, 2005.
- [3] G. Casella and R.L. Berger. *Statistical Inference*, Second Edition. Thomson, Pacific Grove, 2002.
- [4] T.M. Cover and J.A. Thomas. *Elements of Information Theory*. Wiley, New York, 1991.
- [5] R.J. Elliott, L. Aggoun and J.B. Moore. *Hidden Markov Models: Estimation and Control*. Springer-Verlag, New York, 1995.
- [6] J.S. Evans and V. Krishnamurthy. Optimal sensor scheduling for hidden Markov model state estimation. *International Journal of Control*. Volume 18, Number 74, Pages 1737–1742, 2001.
- [7] L.A. Johnston and V. Krishnamurthy. Opportunistic file transfer over a fading channel: a POMDP search theory formulation with optimal threshold policies. *IEEE Transactions on Wireless Communications*. Volume 5, Number 2, Pages 394–205, 2006.
- [8] V. Krishnamurthy. Algorithms for optimal scheduling and management of hidden Markov model sensors. *IEEE Transactions on Signals Processing*. Volume 6, Number 50, Pages 1382–1397, 2002.

- [9] V. Krishnamurthy and D. Djonin. Structured threshold policies for dynamic sensor scheduling—a POMDP approach. *IEEE Transactions on Signals Processing*. Volume 5, Number 10, Pages 4938–4957, 2007.
- [10] F. Le Gland and L. Mevel. Exponential forgetting and geometric ergodicity in hidden Markov models. *Mathematics of Control, Signals and Systems*. Volume 13, Pages 63–93, 2000.
- [11] I.M. MacPhee and B.P. Jordan. Optimal search for a moving target. *Probability in the Engineering and Informational Sciences*. Volume 9, Number 2, Pages 159–182, 1995.
- [12] S.P. Meyn and R.L. Tweedie. *Markov Chains and Stochastic Stability*. Springer-Verlag, London, 1993.
- [13] L.R. Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*. Volume 2, Number 77, Pages 257–286, 1989.
- [14] M.J. Rezaeian. Sensor Scheduling for Optimal Observability Using Estimation Entropy. *Proceedings of the Fifth Annual IEEE International Conference on Pervasive Computing and Communications Workshops*. March 2007.
- [15] D. Sinno and D. Cochran. Dynamic Estimation with Selectable Linear Measurements. *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing*. May 1998.
- [16] D.M. Topkis. *Supermodularity and Complementarity*. Princeton University Press, 1998.

- [17] H.S. Witsenhausen. Separation of Estimation and Control for Discrete Time Systems. *Proceedings of the IEEE*. Volume 59, Number 11, 1971.
- [18] J.Y. Zhao. *Hidden Markov Models with Multiple Observation Processes*. Honours Thesis, University of Melbourne, 2007.



Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Zhao, James Yuanjie

Title:

Hidden Markov models with multiple observation processes

Date:

2010

Citation:

Zhao, J. Y. (2010). Hidden Markov models with multiple observation processes. Masters Research thesis, Department of Mathematics and Statistics and Department of Electrical and Electronic Engineering, The University of Melbourne.

Persistent Link:

<http://hdl.handle.net/11343/35789>

File Description:

Hidden Markov models with multiple observation processes

Terms and Conditions:

Terms and Conditions: Copyright in works deposited in Minerva Access is retained by the copyright owner. The work may not be altered without permission from the copyright owner. Readers may only download, print and save electronic copies of whole works for their own personal non-commercial use. Any use that exceeds these limits requires permission from the copyright owner. Attribution is essential when quoting or paraphrasing from these works.