Characterizations of Random Set-Walks

by

Joseph Hao Tan Wong

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF

THE REQUIREMENTS FOR THE DEGREE OF

Bachelor of Science with Honours

in

THE FACULTY OF SCIENCE

(Department of Computer Science)

I accept this thesis as conforming
to the required standard

The University of British Columbia

April 2001

Abstract

In this thesis, we introduce a new class of set-valued random processes called random set-walk, which is an extension of the classical random walk that takes into account both the nonhomogeneity of the walk's environment, and the additional factor of nondeterminism in the choices of such environments. We also lay down the basic framework for studying random set-walks.

We define the notion of a characteristic tuple as a 4-tuple of first-exit probabilities which characterizes the behaviour of a random walk in a nonhomogeneous environment, and a characteristic tuple set as its analogue for a random set-walk. We prove several properties of random set-walks and characteristic tuples, from which we derive our main result: the long-run behaviour of a sequence of random set-walks, relative to the endpoints of the walks, converges as the length of the walks tend to infinity.

Contents

Acknowledgements

I would like to express my sincerest gratitude to my supervisor, Professor Anne Condon. It was during my research assistantship with her that I was first introduced to the analysis of 2npfa's, the motivation behind the development of the random set-walk model. Throughout this past year, she has also provided guidance and motivation, shared her many insightful ideas, and gave up countless hours of her time to the discussion and refinement of the ideas which became this thesis. I also thank her for her patience while I stumbled over definitions and proofs, and for her optimism and her belief in me even at times when the tasks at hand seemed insurmountable. Her role as my supervisor, my role model, and my friend is deeply appreciated, and I am very fortunate to have the opportunity to work with such a wonderful person.

I would also like to thank Professor John Walsh, for his courses which formally introduced me to the study of probability and random processes. Without them, the results in this thesis could not have been proven in full rigour.

Finally, I am indebted to my family, for their constant encouragement and support throughout these years of my undergraduate career.

JOSEPH WONG

The University of British Columbia
April 2001

CHAPTER 1

Introduction

The theory of random walk has been studied extensively since the development of modern probability. Over the years, researchers have expanded the study of random walk to its various extensions, including random walk on random environment (RWRE), random walk with internal states, random walk on trees and groups, and others. However, much is unknown with regards to the properties and behaviour of a random walk over the set of all possible nonhomogeneous environments. In this thesis we set out to generalize the notion of a random walk and develop the theory of *random set-walk* to study the effects of *nondeterministic environments*.

1.1 A First Picture of a Random Set-Walk

Consider a particle inhabiting the set of integers which is at the initial position Z_0 at time $t = 0$. At subsequent discrete time-steps $1, 2, \ldots$, the particle will move either one step left or one step right from its previous position, with probabilities q and $p = 1 - q$ respectively. Moreover, each such move by the particle is independent of all other moves. Let η_i be the random variable associated with the particle making a right move at time i, i.e. $\mathsf{P}(\eta_i = 1) = p$ and $\mathsf{P}(\eta_i = -1) = q$. Then the position of the particle at time t can be expressed as $Z_t = Z_0 + \sum_{i=1}^{t} \eta_i$. The system described here is the “simple random walk”.

Imagine that the pair of probabilities (q, p) is associated with a biased coin, and that this coin is flipped to decide the direction of the particle’s movement. One particular extension to this coin-tossing game involves putting a different biased coin at each position on the integer line, and having the player pick up the coin at the particle’s current position, flip it to decide the next move, and put the coin down again. We shall limit the number of different biases for all the coins to be finite – one can think of this restriction as having only a finite number of mints, each producing coins with a different bias. Since all coins produced by the same mint are identically biased, we can refer to them collectively as (q_i, p_i); and hence, we can let $\mathcal{D} = \{(q_i, p_i)\}$ be the set of all possible coin biases. A particle moving according to these rules is an example of a random walk in a nonhomogeneous environment (with a single internal state).

Another extension to the random walk involves giving the particle an internal state $s \in \{1, \ldots, m\}$. Now, not only will the particle move back and forth on the integer line, but it will also change its internal state. The role of the coin is expanded to dictate the movement between the internal states as well as the left and right movement of the particle. In this scenario, $\alpha = (q_{ij})$ and $\gamma = (p_{ij})$ are non-negative matrices where q_{ij} is the probability of the particle moving into state j from state i and step towards the left, and p_{ij} is the corresponding right-move probability. In addition, the sum of these two matrices will be a stochastic $m \times m$ matrix. We now have a general random walk in a nonhomogeneous environment, with a finite internal state space.

Finally, suppose we have infinitely many particles, each involved in its own random walk on a different nonhomogeneous environment, but with the same set $\mathcal{D}$. This collection of random walks is referred to as a “random set-walk” in the theory developed in the subsequent chapters, and is the main topic of this thesis.

If we ignore left-right movement of all the particles, the random set-walk reduces to the Markov set-chain as introduced by D. Hartfiel [22], with the transition set $\Sigma = \{\alpha + \gamma \mid (\alpha, \gamma) \in \mathcal{D}\}$. On the other hand, if the set D has only one coin – that is, if spatial nonhomogeneity is taken out, the system is known as a random walk with internal states, as studied by D. Szász and A. Krámli [28, 29, 27]. This lineage is reflected in our choice of the term "random set-walk" as the name of the class of random processes of interest.

Remark. While we have been using coins with two possible outcomes (with probabilities α and γ) in the previous examples, the actual theory uses the more general notion of a "transition triple" $\mathsf{T} = (\alpha, \beta, \gamma)$, where α and γ are as before, and β is the probability (or the matrix of probabilities) that the particle stays at the same position. The analogue of the set $\mathcal{D}$ described above is the "transition set" $\mathcal{T}$. All these terms will be defined rigorously in Chapter 2.

1.2 Motivation

The development of the theory of random set-walk was motivated by the analysis of bounded-error two-way nondeterministic probabilistic finite automata (2npfa), as studied by Condon et al. [10]. A 2npfa is a finite automata with both nondeterministic and probabilistic states, and a two-way head. When the analysis was focused on the unary languages accepted by a bounded-error 2npfa, it was realized that the operation of a 2npfa exhibits a random walk behaviour when its input is a string 0^n from the alphabet $\{0\}$. This behaviour is captured by the theory of random set-walk, where the walk component corresponds to the head moving back and forth on the input tape, and where the internal state component corresponds the state of the automaton.

It is conjectured that all unary languages accepted by a 2npfa are regular. A key step in proving this conjecture is to establish the validity of the following:

The sets of first-exit probability matrices corresponding to a sequence of ran-

dom set-walks of lengths $d, 2d, 3d, \ldots$ converges in the Hausdorff metric, for some periodicity d.

A special case of this result, for the case when the internal state space contains only one state, is proved in Chapter 3.

1.3 Related Work

1.3.1 Markov Set-Chains

The theory of Markov set-chains has been studied for over twenty years by Hartfiel [22] and others. A Markov set-chain is a sequence $(M^i)_{i=0}^{\infty}$ of sets of finite, substochastic matrices, where $M^{k+1} = MM^k = \{A_1A_2 \mid A_1 \in M, A_2 \in M^k\}$ is the set of all possible products of matrices in the set M of length k. Hartfiel proved many results concerning Markov set-chains, including ones on the criteria for guaranteeing convergence of a Markov set-chain, and decomposition of the state space into ergodic and transient classes. In particular, he proved that a Markov set-chain is convergent if it is product-scrambling, i.e. if there exists an $r \in \mathbb{N}$ such that all products of matrices $A_1A_2 \ldots A_r$ in M of length r have a coefficient of ergodicity $\tau(A_1A_2 \ldots A_r) < 1$, where the coefficient of ergodicity τ is defined as $\tau(A) = \frac{1}{2} \max_{i,j} \sum_k |a_{ik} - a_{jk}|$. There is an emphasis in Hartfiel's work on Markov set-chains built on matrix intervals $M = [P, Q]$ where $P < Q$ componentwise. While these set-chains model the physical world quite well, their analysis do not lend themselves well to studying more general set-chains, where M is not an interval.

1.3.2 Product of Sets of Matrices

While Hartfiel used mostly matrix-theoretic techniques in the analysis of Markov set-chains, Saks and Condon [34] studied the products of sets of stochastic matrices using graph-theoretic techniques. They showed that for any finite set of stochastic

A, there exist a $k \in \mathbb{N}$, such that the sequence of sets $(A^{ki})_{i \in \mathbb{N}}$ converges in the Hausdorff metric. This result is quite strong in the sense that it does not rely on any spectral properties or the matrices, nor on any properties of the set A.

1.3.3 Random Walk with Internal States

The theory of random walk with internal states was first introduced by Sinai [37] and later studied in depth by Krámli and Szász [28, 29, 27]. The model was introduced as a tool in the study of the behaviour of a class of physical processes known as the Lorentz process. In their formulation, a random walk with internal states $X_t = (Z_t, Y_t)$ has both a random walk component $Z_t \in \mathbb{Z}$ and internal states $Y_t \in E$ where E is a finite state space. The process satisfies a translation invariance property, in the sense that the transition probabilities are not dependent on the current Z position of the walk.

In part 2 of their paper [29], Krámli and Szász added two absorbing barriers to the model, and analyzed the asymptotics of the first-hitting probabilities of the random walk as the distance between the barriers tends to infinity. Using a technique that involves analyzing the spectral properties of the defining matrices of the system, they were able to derive an explicit form for the first-hitting probabilities in the limit.

Although a random walk with internal states with the two absorbing barriers, is equivalent to a random set-walk with no nondeterminism, Krámli and Szász's analysis assumes the reversibility of the random walk. While reversibility of the random walk may hold in the case of physical processes such as the Lorentz process, it cannot be assumed in general.

1.3.4 Nonhomogeneous Markov Chains and Products of Non-negative Matrices

The generalization of Markov chains to nonhomogeneous products of matrices was first studied by Hajnal [18, 19, 20], Cohn, and Seneta. There are many known results regarding such chains and products, and in particular results on the weak and strong ergodicity of nonhomogeneous products. Hajnal [20] proved that a sequence of partial products in the form $H_r = H_{r-1}A_r$ is weakly ergodic if the nonzero components of each stochastic A_i are uniformly bounded away from zero. In another paper, he showed that the "backward" product $B_r = A_r B_{r-1}$ is strongly ergodic, which can be interpreted as the existence of a unique limit B_{lim} such that $|B_{\mathrm{lim}} - B_r| < \varepsilon$ for any $\varepsilon > 0$ and r sufficiently large.

1.4 Organization

This thesis is organized in the following way:

- Chapter 2 covers the background material, and rigorously defines random set-walk and related concepts.
- Chapter 3 looks at the case when the internal state space contains only a single internal state, and proves the main result of this thesis.
- Chapter 4 summarizes our work and discusses some open problems, in particular the general case of a random set-walk with a finite internal state space containing two or more states.

CHAPTER 2

Background and Definitions

2.1 Markov Chains

Let $(\Omega, \mathcal{F}, \mathsf{P})$ be a probability space. Then a random (or *stochastic*) process is a family $(X_t)_{t \in T}$ of random variables indexed by some index set T. The set S in which these random variables take values is referred to as the *state space* of the process. When $T \subseteq \mathbb{N}$, we call the process a *discrete-time* process.

This thesis is mainly concerned with the analysis of discrete-time processes; thus, $T = \mathbb{N}$ is assumed unless otherwise specified. The notation (X_t), or $(X_t)_t$ when there is ambiguity in the index variable, shall be a shorthand for $(X_t)_{t \in \mathbb{N}}$.

Definition 2.1 (Markov chain). A process (X_t), with state space S, is a *Markov chain* (or simply, a chain) if it satisfies the *Markov property*:

$$\mathsf{P}(X_t = x_t \mid X_i = x_i, \forall i = 0, 1, \ldots, t-1) = \mathsf{P}(X_t = x_t \mid X_{t-1} = x_{t-1}) \tag{2.1}$$

for all $t \geq 1$ and $x_0, x_1, \ldots, x_t \in S$.

The Markov property simply stipulates that, conditional on its value at the nth step, a process's future value does not depend on its previous values. This is a very useful property in the analysis of Markov chains, since it simplifies complex conditional probabilities to ones depending only on the last value of the process.

As the evolution of a chain is described by its defining transition probabilities, which depend upon the source state i, the destination state j, and the time step t, it can be quite complicated to analyze. Homogeneity is a requirement on a Markov chain that its transition probabilities are dependent only on i and j, and are independent of n.

Definition 2.2 (Homogeneity). A Markov chain (X_t) with state space S is called *homogeneous* if

$$\mathsf{P}(X_{t+1} = j \mid X_t = i) = \mathsf{P}(X_1 = j \mid X_0 = i) \tag{2.2}$$

for all $t \geq 0$ and $i, j \in S$.

Let (X_t) be a chain with an *at most countable* (i.e. finite or countable) state space S. The behaviour of such a chain can be characterized by its transition probabilities and its initial distribution.

Definition 2.3 (Transition matrix). The *transition matrix* $P = [p_{ij}]$ of a homogeneous chain (X_t) is the $|S| \times |S|$ matrix of *transition probabilities*

$$p_{ij} = \mathsf{P}(X_1 = j \mid X_0 = i)\ . \tag{2.3}$$

Similarly, we can define the *t-step transition matrix* $P^{(t)} = \left[p_{ij}^{(t)}\right]$ of a homogeneous chain as the matrix of *t-step transition probabilities*

$$p_{ij}^{(t)} = \mathsf{P}(X_t = j \mid X_0 = i) \tag{2.4}$$

with $P^{(1)} = P$. By the Chapman-Kolmogorov equations [16]

$$p_{ij}^{(m+n)} = \sum_k p_{ik}^{(m)} p_{kj}^{(n)}\ ,$$

we can obtain the identity

$$P^{(t)} = P^t\ . \tag{2.5}$$

Consequently, the probability $\mathsf{P}(X_t = i)$ satisfies

$$\mathsf{P}(X_t = i) = \mu P^t \tag{2.6}$$

where $\mu = (\mathsf{P}(X_0 = i))_{i \in S}$ is the *initial distribution vector*, which is the row vector of initial probabilities for the chain.

Definition 2.4 (Stochastic and substochastic matrices). Let $M = [m_{ij}]$ be a non-negative matrix. If M satisfies $\sum_j m_{ij} = 1$ (resp. ≤ 1), then M is a *stochastic* (resp. *substochastic*) matrix.

For example, the transition probability matrix P of a homogeneous Markov chain is a stochastic matrix.

2.2 Random Walks

A classic example of a homogeneous Markov chain is the simple random walk, which one can picture as describing the motion of a particle travelling on the integers. At each step, this particle either moves one step to the right, with probability p, or one step to the left, with probability $q = 1 - p$. In addition, the steps made by the particle are independent of each other.

Definition 2.5 (Simple random walk). The *simple random walk* is a homogeneous chain (Z_t) taking values in $\mathbb{Z}$, where

$$\mathsf{P}(Z_{t+1} = \begin{Bmatrix} z_t + 1 \\ z_t \\ z_t - 1 \end{Bmatrix} \mid Z_t = z_t) = \begin{cases} p \\ 0 \\ q & = 1 - p \end{cases}. \tag{2.7}$$

When $p = q = 1/2$, we call the chain a *symmetric* simple random walk.

Lemma 2.6. *The simple random walk is* spatially homogeneous, *that is,* (Z_t) *satisfies the spatial homogeneity condition:*

$$\mathsf{P}(Z_t = n + d \mid Z_0 = m + d) = \mathsf{P}(Z_t = n \mid Z_0 = m) \tag{2.8}$$

for any t, m, n *and displacement* $d \in \mathbb{Z}$.

2.2.1 Gambler's Ruin and Stopping Time

While the simple random walk is a process that takes values on the all of the integers, we can restrict a random walk to a particular interval of integers, which traditionally is the interval $0, \dots, N$.

Notation. We shall denote the integer interval $a, \dots, b$ as $[a : b]$. More formally,

$$[a : b] \doteq \{i \in \mathbb{Z} \mid a \le i, i \le b\} \tag{2.9}$$

where $a, b \in \mathbb{Z}$. Thus $[a : b] = \phi$ if $a > b$.

One way of restricting a random walk to $[0 : N]$ is to put *absorbing barriers* at 0 and N, which means that we modify the defining probabilities (2.7) to:

$$\mathsf{P}(Z_{t+1} = z_t + 1 \mid Z_t = z_t) = \begin{cases} p & \text{if } z_t \in [1 : N-1] \\ 0 & \text{if } z_t = 0 \text{ or } z_t = N \end{cases} \tag{2.10a}$$

$$\mathsf{P}(Z_{t+1} = z_t \mid Z_t = z_t) = \begin{cases} 0 & \text{if } z_t \in [1 : N-1] \\ 1 & \text{if } z_t = 0 \text{ or } z_t = N \end{cases} \tag{2.10b}$$

$$\mathsf{P}(Z_{t+1} = z_t - 1 \mid Z_t = z_t) = \begin{cases} q = 1 - p & \text{if } z_t \in [1 : N-1] \\ 0 & \text{if } z_t = 0 \text{ or } z_t = N \end{cases} \tag{2.10c}$$

This random walk with absorbing barriers at 0 and N is more commonly known as *gambler's ruin*, which can be stated as follows: A gambler starts with an initial capital of k, and gambles with an adversary with initial capital $N - k$. On each bet, the gambler wins or loses a dollar with probabilities p and q, respectively. He plays this game repeatedly until one player wins the whole pot N (and the other player is ruined).

Definition 2.7 (Stopping time). A random variable T is a *stopping time* for a random process (X_t) if the event $\{T = n\}$ is dependent only on the values of $X_0, \dots, X_n$.

Let T_{ab} be a stopping time defined by

$$T_{ab} = \inf\{n : Z_n \notin [a+1 : b-1]\} \tag{2.11}$$

Then the gambler's probabilities of ultimately winning and losing the gamble are

$$\mathsf{P}(\text{win}) = \mathsf{P}(Z_{T_{1,N-1}} = N \mid Z_0 = k) = \begin{cases} \frac{1-(q/p)^k}{1-(q/p)^N} & \text{if } p \neq 1/2 \\ k/N & \text{if } p = q = 1/2 \end{cases} \tag{2.12}$$

$$\mathsf{P}(\text{loss}) = \mathsf{P}(Z_{T_{1,N-1}} = 0 \mid Z_0 = k) = \begin{cases} \frac{(q/p)^k-(q/p)^N}{1-(q/p)^N} & \text{if } p \neq 1/2 \\ 1-(k/N) & \text{if } p = q = 1/2 \end{cases}. \tag{2.13}$$

2.2.2 *Strong Markov Property*

Related to the notion of a stopping time is the strong Markov property (abbreviated SMP), which is the same result as the Markov property, but for a Markov chain that is restarted at a stopping time.

Theorem 2.8 (Strong Markov property). *Let T be a stopping time for the Markov chain X_t. Then, conditional on $\{T < \infty\}$, the process $(X_{T+t})_{t\in\mathbb{N}}$ is a Markov chain with the same transition probabilities as the original chain, conditionally independent of $X_0, \dots, X_T$ given X_T.*

Proof.

$$\begin{aligned}
&\mathsf{P}(X_T = i_0, X_{T+1} = i_1, \dots, X_{T+n} = i_n \mid T < \infty) \\
&= \frac{1}{\mathsf{P}(T<\infty)}\mathsf{P}(T<\infty, X_T = i_0, \dots, X_{T+n} = i_n) \\
&= \frac{1}{\mathsf{P}(T<\infty)}\sum_{k=0}^{\infty}\mathsf{P}(T = k, X_k = i_0, \dots, X_{k+n} = i_n) \\
&= \frac{\sum_{k=0}^{\infty}\mathsf{P}(X_{k+n} = i_n \mid X_{k+n-1} = i_{n-1}, \dots, X_k = i_0, T = k)\mathsf{P}(X_{k+n-1} = i_{n-1}, \dots, X_k = i_0, T = k)}{\mathsf{P}(T<\infty)}
\end{aligned}$$

by the Markov property applied at $k+n-1$, and the fact that $\{T = k\}$ is dependent only on $X_0, \dots, X_k$,

$$
\begin{aligned}
&= \frac{1}{\mathsf{P}(T < \infty)} \sum_{k=0}^{\infty} \mathsf{P}(X_{k+n} = i_n \mid X_{k+n-1} = i_{n-1}) \mathsf{P}(X_{k+n-1} = i_{n-1}, \dots, X_k = i_0, T = k) \\
&= \frac{1}{\mathsf{P}(T < \infty)} \sum_{k=0}^{\infty} p_{i_{n-1}, i_n} \mathsf{P}(X_{k+n-1} = i_{n-1}, \dots, X_k = i_0, T = k)
\end{aligned}
$$

which, by an induction argument, yields,

$$
\begin{aligned}
&= \frac{1}{\mathsf{P}(T < \infty)} \sum_{k=0}^{\infty} \mathsf{P}(X_k = i_0, T = k) p_{i_0, i_1} \cdots p_{i_{n-1}, i_n} \\
&= \mathsf{P}(X_T = i_0 \mid T < \infty) p_{i_0, i_1} \cdots p_{i_{n-1}, i_n} \\
&= \mathsf{P}(X_{T+n} = i_n \mid X_T = i_0) \mathsf{P}(X_T = i_0 \mid T < \infty) \ .
\end{aligned}
$$

Hence, conditional on the distribution of X_T and T being finite, X_{T+t} is another Markov chain with transition probabilities p_{ij}. $\square$

2.3 Nonhomogeneous Environments and Nondeterminism

From random walks and gambler's ruin, we turn our attention to more general classes of processes based upon the simple random walk. All of these "walks" share the common feature that within each, there is a notion of taking random steps to the left or right as the process evolves.

We begin our exploration with random walk with internal states, which has been studied by D. Szász and A. Krámli [28, 29, 27] and others. Whereas their analysis requires that the process is time reversible, we place no such restriction on our random walk with internal states.

Definition 2.9 (Transition triple). Let E be an at most countable set, and α, β, γ be three substochastic $|E| \times |E|$ matrices such that the sum $\alpha + \beta + \gamma$ is a stochastic matrix. We shall call the triple $\mathsf{T} = (\alpha, \beta, \gamma)$ a *transition triple* on E.

Also, suppose T' is an arbitrary transition triple on E. Let us denote its component matrices by $\alpha[\mathsf{T}']$, $\beta[\mathsf{T}']$ and $\gamma[\mathsf{T}']$. In other words, $\mathsf{T}' = (\alpha[\mathsf{T}'], \beta[\mathsf{T}'], \gamma[\mathsf{T}'])$.

The role of a transition triple in a random walk with internal states is analogous to that of the familiar left- and right-step probabilities q and p in a simple random walk, directing the process in its move left and right on the integers, as well as the transitions taken by its internal state.

Definition 2.10 (Random walk with internal states). Let (X_t) be a random process where each random variable $X_i = (Z_i, Y_i) \in \mathbb{Z} \times E$, and E is an at most countable set. If the process satisfies

$$\mathsf{P}(X_{t+1} = (\begin{Bmatrix} z_t + 1 \\ z_t \\ z_t - 1 \end{Bmatrix}, j) \mid X_t = (z_t, i)) = \begin{cases} \alpha[\mathsf{T}]_{ij} \\ \beta[\mathsf{T}]_{ij} \\ \gamma[\mathsf{T}]_{ij} \end{cases} \tag{2.14}$$

where T is a transition triple on E, then call this process a *random walk with internal states*. Here, Z_t is the external walk component of the process, while Y_t represents its internal state. We shall call the set E the *internal state space* of this process.

2.3.1 *Definition of Random Set-Walk*

So far we have dealt with random walks that are spatially homogeneous. While homogeneity in space makes the analysis of random walks relatively straightforward, it is sometimes too restrictive a condition to place upon a process. Hence, we shall look at extensions of random walk that are not homogeneous in space. One way to introduce spatial nonhomogeneity is to specify that each possible position of the walk is associated with a different transition triple. We call such a mapping of positions to transition triples a strategy.

Definition 2.11 (Transition Set). Let E be an at most countable set. If $\mathcal{T} = \{\mathsf{T}_i\}$ is a non-empty set of transition triples on E, then we shall call $\mathcal{T}$ a *transition set*.

The transition set is to be the set of all possible triples to which positions on the walk can be mapped.

Definition 2.12 (Strategy). A *strategy* is a mapping $s : [a : b] \to \mathcal{T}$. In other words, s maps each point $z \in [a : b]$ to a transition triple $\mathsf{T} \in \mathcal{T}$. The strategy is nonhomogeneous if $|s([a : b])| > 1$.

Denote the set of all strategies of the type $[a : b] \to \mathcal{T}$ by $\mathbb{S}_{[a:b]}(\mathcal{T})$, and the set of all strategies sharing the same range $\mathcal{T}$ by $\mathbb{S}(\mathcal{T})$. Thus,

$$\mathbb{S}(\mathcal{T}) = \bigcup_{a,b\in\mathbb{Z}} \mathbb{S}_{[a:b]}(\mathcal{T})$$

Remark. We shall say that a strategy s *covers* the point z if $z \in dom(s)$.

For the purpose of developing the theory of random set-walk, we focus our attention on nonhomogeneous random walks which are bounded, i.e. the walk never leaves a certain bounded integer interval. As such, we will need the notion of barriers that block the walk from heading out of this interval.

Definition 2.13 (Barrier and barrier set). Let T be a transition triple. If $\alpha[\mathsf{T}]$ (resp. $\gamma[\mathsf{T}]$) is the zero matrix, then T is a *left* (resp. *right*) *barrier*. A transition triple that is both a left barrier and a right barrier is an *absorbing* barrier.

We shall call a set $\mathcal{L}$ of left barriers a *left barrier set*, and a set $\mathcal{R}$ of right barriers a *right barrier set*.

Definition 2.14 (Closed strategy). A strategy $s \in \mathbb{S}_{[a:b]}(\mathcal{T})$ is *closed* if $s(a)$ is a left barrier, and $s(b)$ is a right barrier.

The significance of a closed strategy is that a random walk governed by such a strategy will never leave the interval $[a : b]$ if $z_0 \in [a : b]$.

Definition 2.15 (Random walk in a nonhomogeneous bounded environment). Let (X_t) be a random process where each random variable $X_i = (Z_i, Y_i) \in \mathbb{Z} \times E$, and E is an at most countable internal state space. If the process satisfies

$$\mathsf{P}(X_{t+1} = (\begin{Bmatrix} z_t + 1 \\ z_t \\ z_t - 1 \end{Bmatrix}, j) \mid X_t = (z_t, i)) = \begin{cases} \alpha[s(z_t)]_{ij} \\ \beta[s(z_t)]_{ij} \\ \gamma[s(z_t)]_{ij} \end{cases} \tag{2.15}$$

where $s \in \mathbb{S}_{[a:b]}(\mathcal{T})$ is a closed strategy covering Z_0 (i.e. $Z_0 \in [a : b]$), and $\mathcal{T}$ is a transition set on E, then this process is a *random walk in a nonhomogeneous bounded environment*.

We will denote this random walk by $w(s) \doteq (X_t^{(s)})_{t\in\mathbb{N}}$.

The last addition to the model is that of nondeterminism, which in this context refers to choosing all possible strategies *nondeterministically* and looking at the evolution of the resultant random walks in parallel. Here, we are interested in studying the properties of the set of these random walks as one combined, nondeterministic entity, which is what we call a random set-walk.

Definition 2.16 (Random set-walk). Let E be an internal state space, $\mathcal{T}$ be a transition set, and $\mathcal{S}$ be a set of closed strategies covering 0. Define a set-valued random process $\mathcal{W}(\mathcal{S}) = (\mathcal{W}_t)$, where $\mathcal{W}_t = \{X_t^{(s)} \mid s \in \mathcal{S}\}$, and $\mathcal{W}_0 = \{(0, Y_0)\}$ for some fixed $Y_0 \in E$. Such a set-valued process is a *random set-walk*.

Suppose $\mathcal{T}$ is a transition set, $\mathcal{L} \subset \mathcal{T}$ is a non-empty left barrier set, $\mathcal{R} \subset \mathcal{T}$ is a non-empty right barrier set, and $\mathcal{S} \subseteq \mathbb{S}_{[0:n+1]}(\mathcal{T})$ is a strategy set. If $s(0) \in \mathcal{L}$ and $s(n+1) \in \mathcal{R}$ for all $s \in \mathcal{S}$, then $W(\mathcal{S})$ is a *random set-walk with barriers at* 0 *and* $n + 1$.

2.3.2 *More on Strategies*

For analyzing the properties of random set-walk, we will need to develop some new tools in order to manage the complexity.

Definition 2.17 (Matching strategies). Two strategies $s \in \mathbb{S}_{[a:b]}(\mathcal{T})$ and $s' \in \mathbb{S}_{[c:d]}(\mathcal{T})$

match, denoted $s \simeq s'$, if

$$\forall z \in [a : b] \cap [c : d] : s(z) = s'(z) \ . \tag{2.16}$$

Definition 2.18 (Monotone strategy sequence). $S = (s_i)_{i \in \mathbb{N}}$ is a *monotone strategy sequence* if the following holds for all i:

1. $s_i \in \mathbb{S}_{[a:b]}(\mathcal{T})$, $s_{i+1} \in \mathbb{S}_{[c:d]}(\mathcal{T})$,
2. $c < a$ or $b < d$ (or both), and
3. $s_i \simeq s_{i+1}$.

A monotone strategy sequence S is a *forward* sequence if $\forall i, s_i \in \mathbb{S}_{[1:i]}(\mathcal{T})$.

A monotone strategy sequence S is a *backward* sequence if $\forall i, s_i \in \mathbb{S}_{[-i:0]}(\mathcal{T})$.

Definition 2.19 (Uniform strategy sequence). Let $U(\mathsf{T}) = (s_i)$ be a forward monotone strategy sequence where

$$s_i(z) = \mathsf{T} \text{ for all } z \in [1 : i], i \in \mathbb{N}. \tag{2.17}$$

We shall call this a *uniform strategy sequence*.

Definition 2.20 (Closure of a strategy). Let $s \in \mathbb{S}_{[a:b]}(\mathcal{T})$. Define the *closure* of this strategy, $\overline{s} \in \mathbb{S}_{[a-1:b+1]}(\mathcal{T}')$, where $\mathcal{T}' = \mathcal{T} \cup \{I\}$, by

$$\overline{s}(z) = \begin{cases} s(z) & \text{, if } z \in [a : b] \\ I & \text{, if } z = a - 1 \text{ or } z = b + 1 \end{cases} . \tag{2.18}$$

Here, I is the identity matrix.

Definition 2.21 (Concatenation). Let $s \in \mathbb{S}_{[a:b]}(\mathcal{T})$, and $s' \in \mathbb{S}_{[c:d]}(\mathcal{T})$. Define the binary operation of *concatenation*, denoted $s \cdot s'$, by

$$(s \cdot s')(z) = \begin{cases} s(z) & \text{if } z \in [a : b] \\ s'(z - b + c - 1) & \text{if } z \in [b + 1 : b + d - c + 1] \end{cases} . \tag{2.19}$$

Therefore, $s \cdot s' \in \mathbb{S}_{[a:b+d-c+1]}(\mathcal{T})$.

Similarly, if $\mathsf{T} \in \mathcal{T}$, define $s \cdot \mathsf{T}$ and $\mathsf{T} \cdot s$ by

$$(s \cdot \mathsf{T})(z) = \begin{cases} s(z) & \text{if } z \in [a : b] \\ \mathsf{T} & \text{if } z = b + 1 \end{cases} \tag{2.20}$$

and

$$(\mathsf{T} \cdot s)(z) = \begin{cases} \mathsf{T} & \text{if } z = a \\ s(z-1) & \text{if } z \in [a+1 : b+1] \end{cases} \tag{2.21}$$

respectively. Thus both $s \cdot \mathsf{T} \in \mathbb{S}_{[a:b+1]}(\mathcal{T})$ and $\mathsf{T} \cdot s \in \mathbb{S}_{[a:b+1]}(\mathcal{T})$.

In addition, define $\$ \cdot s$ and $s \cdot \#$ by

$$(\$ \cdot s)(z) = s(z + a - 1) \tag{2.22}$$

$$(s \cdot \#)(z') = s(z + b) \tag{2.23}$$

for $z \in [1 : b - a + 1]$ and $z' \in [a - b : 0]$. Consequently, $\$ \cdot s \in \mathbb{S}_{[1:b-a+1]}(\mathcal{T})$ and $s \cdot \# \in \mathbb{S}_{[a-b:0]}(\mathcal{T})$. $\$$ can be thought of as the empty strategy at 1, while $\#$ flushes the right end of a strategy to 0.

Notation. When the context is clear, the $\cdot$ can be omitted, as in $s_1 s_2$, $s\mathsf{T}$, or $\$s$.

Remark. Concatenation is associative, and thus $(\mathbb{S}(\mathcal{T}), \cdot)$ is a semigroup.

2.3.3 Transition Probabilities

Definition 2.22 (Walk instance). Let $\mathcal{T}$ be a transition set and $s \in \mathbb{S}_{[1:n]}(\mathcal{T})$ be a strategy. Call the process $W(s) = w(\overline{s})$ the *walk instance* of s.

The closure of the strategy s is needed here since $\mathcal{T}$ is an arbitrary transition set, and consequently $s(1)$ and $s(n)$ may not be barriers.

Definition 2.23 (Characteristic probabilities). Given a transition set $\mathcal{T}$ and a strategy $s \in \mathbb{S}_{[1:n]}(\mathcal{T})$, define a matrix of transition probabilities $C_s(x,y)$ for the walk instance $(\xi_t) = W(s)$ by

$$C_s(x,y) = [c_{ij}(x,y)] \tag{2.24}$$

where $x \in [1:n]$ and $y = 0$ or $n+1$, and

$$\begin{aligned} c_{ij}(x,y) &= \mathsf{P}(\xi_T = (y,j) \mid \xi_0 = (x,i)) \\ &= \lim_{t\to\infty} \mathsf{P}(\xi_t = (y,j) \mid \xi_0 = (x,i)) \end{aligned} \tag{2.25}$$

where T is the stopping time

$$T = \inf\{t : \xi_t \notin [1:n] \times E\} .$$

We shall call $C_s(x,y)$ a *characteristic probability matrix* of s.

Definition 2.24 (Characteristic tuple). Define the 4-tuple $\varphi_s = \langle C_s(n,0), C_s(1,0), C_s(n,n+1), C_s(1,n+1)\rangle$ as the *characteristic tuple* of s.

We shall also use the notation $a[s], b[s], c[s], d[s]$ to refer to the component matrices $C_s(n,0), C_s(1,0), C_s(n,n+1), C_s(1,n+1)$, respectively.

Suppose s, s' are two arbitrary strategies defined on the same transition set $\mathcal{T}$. Then according to the previous definition, there are characteristic tuples φ_s and $\varphi_{s'}$ associated with the strategies s and s', respectively. Moreover, one can calculate the characteristic tuple $\varphi_{s's}$ of the concatenation of these strategies. In fact, $\varphi_{s's}$ can be derived directly from φ_s and $\varphi_{s'}$.

Theorem 2.25 (Concatenation of characteristic tuples). *Let $s' \in \mathbb{S}_{[1:m]}(\mathcal{T})$ and $s \in \mathbb{S}_{[1:n]}(\mathcal{T})$, and let $\varphi_{s'} = \langle a_l, b_l, c_l, d_l\rangle$ and $\varphi_s = \langle a, b, c, d\rangle$ be the two characteristic tuples associated with the strategies s and s'. Define the concatenation of the characteristic tuples $\varphi_{s'} \cdot \varphi_s$, as the characteristic tuple of the concatenation of the underlying strategies, $\varphi_{s's}$.*

Then, $\varphi_{s'} \cdot \varphi_s$ satisfies the following system of equations:

$$a' = a(1 - c_l b)^{-1} a_l \tag{2.26a}$$

$$b' = b_l + d_l(1 - bc_l)^{-1} b a_l \tag{2.26b}$$

$$c' = c + a(1 - c_l b)^{-1} c_l d \tag{2.26c}$$

$$d' = d_l(1 - bc_l)^{-1} d \tag{2.26d}$$

where $\varphi_{s'} \cdot \varphi_s = \langle a', b', c', d' \rangle$. Consequently, the set of characteristic tuples under the concatenation operation is a semigroup.

Proof. It suffices to show that Equations (2.26a) and (2.26c) hold, since the argument is symmetric.

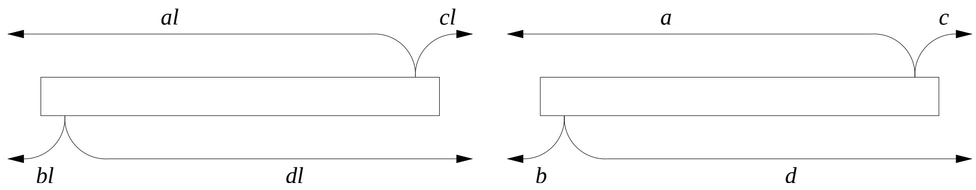

Figure 2.1: Concatenating two characteristic tuples

Let $(\xi_t) = W(s's)$ be the walk instance of $s's$, and the stopping times T_0, T_1 and T_2 be defined as follows:

$$T_0 = \inf\{t \geq 0 : \xi_t \notin [1 : m] \times E\}$$

$$T_1 = \inf\{t \geq 0 : \xi_t \notin [m+1 : m+n] \times E\}$$

$$T_2 = \inf\{t > 0 : \xi_t \in \{m\} \times E\}$$

Given $a, b, c, d, a_l, b_l, c_l, d_l, a', b', c', d'$ as defined in the theorem statement (see Figure 2.1), we have

$$\begin{aligned}[c_l b]_{ij} &= \sum_{k \in E} \mathsf{P}(\xi_{T_1} = (m, j) \mid \xi_0 = (m+1, k))\mathsf{P}(\xi_{T_0} = (m+1, k) \mid \xi_0 = (m, i)) \\ &= \mathsf{P}(\xi_{T_2} = (m, j) \mid T_2 > T_0, \xi_0 = (m, i))\end{aligned}$$

by the strong Markov property (SMP). Therefore,

$$\begin{aligned}\left[(1-c_l b)^{-1}\right]_{ij} &= \left[\sum_{N=0}^{\infty}(c_l b)^N\right]_{ij} \\ &= \mathsf{P}(\xi_T = (m,j) \mid \xi_0 = (m,i))\end{aligned}$$

where $T = \inf\{t \geq 0 : \xi_t \in \{m\} \times E\}$, once again by SMP. So multiplying this matrix on the left and right by a and a_l respectively yields

$$\begin{aligned}\left[a(1-c_l b)^{-1}a_l\right]_{ij} = \sum_{k\in E}\sum_{l\in E}(&\mathsf{P}(\xi_{T_1} = (m,k) \mid \xi_0 = (m+n,i)) \\ &\mathsf{P}(\xi_T = (m,l) \mid \xi_0 = (m,k)) \\ &\mathsf{P}(\xi_{T_0} = (0,j) \mid \xi_0 = (m,l)))\end{aligned}$$

where T, T_0, T_1 are as above

$$\begin{aligned}&= \mathsf{P}(\xi_{T'} = (0,j) \mid \xi_0 = (m+n,i)) \\ &= a'\end{aligned}$$

where $T' = \inf\{t : \xi_t \notin [1:m+n] \times E\}$, by SMP applied twice.

As for Equation (2.26c), we have

$$\left[(1-c_l b)^{-1}\right]_{ij} = \mathsf{P}(\xi_T = (m,j) \mid \xi_0 = (m,i))$$

from above. Hence multiplying on the left and right by a and $c_l d$ yields

$$\begin{aligned}\left[a(1-c_l b)^{-1}c_l d\right]_{ij} = \sum_{k\in E}\sum_{l\in E}\sum_{h\in E}(&\mathsf{P}(\xi_{T_1} = (m,k) \mid \xi_0 = (m+n,i)) \\ &\mathsf{P}(\xi_T = (m,l) \mid \xi_0 = (m,k)) \\ &\mathsf{P}(\xi_{T_0} = (m+1,h) \mid \xi_0 = (m,l)) \\ &\mathsf{P}(\xi_{T_1} = (m+n+1,j) \mid \xi_0 = (m+1,h)))\end{aligned}$$

$$= \mathsf{P}(\xi_{T'} = (m+n+1,j) \mid \xi_1 \in [1:m+n] \times E,\ \xi_0 = (m+n,i))\ .$$

By adding c to $a(1-c_l b)^{-1}c_l d$, we get

$$\begin{aligned}\left[c + a(1-c_l b)^{-1}c_l d\right]_{ij} &= \mathsf{P}(\xi_{T'} = (m+n+1, j) \mid \xi_0 = (m+n, i)) \\ &= c' \,.\end{aligned}$$

□

2.4 Sets and the Hausdorff Metric

When we say that a sequence of real numbers (x_i) converges, we mean

$$\exists L \in \mathbb{R}\, \forall \epsilon > 0\, \exists N \in \mathbb{N}\, \forall n > N : d(x_i, L) < \epsilon$$

where $d(\cdot,\cdot)$ is the regular Euclidean metric.

Now, suppose we are interested in looking at the potential convergence behaviour of a sequence of *sets* (Γ_i), what criteria shall we use? Specifically, how can we measure the distance between sets?

One solution is to define the Hausdorff metric on sets.

Definition 2.26 (Hausdorff metric). For any sets $\Gamma, \Lambda \in H(S)$, where S is a metric space and $H(S)$ denotes the set of compact sets in S, define

$$\delta(\Gamma, \Lambda) = \max_{x\in\Gamma} \min_{y\in\Lambda} d_S(x, y) \tag{2.27}$$

$$d(\Gamma, \Lambda) = \max(\delta(\Gamma, \Lambda), \delta(\Lambda, \Gamma)) \tag{2.28}$$

where d_S is the metric of the underlying space S.

The function d is called the *Hausdorff metric* on $H(S)$.

Equivalently, we can define the Hausdorff metric d by

$$d(\Gamma, \Lambda) = \min\{\epsilon \mid \Gamma \subseteq B_\epsilon(\Lambda) \text{ and } \Lambda \subseteq B_\epsilon(\Gamma)\} \tag{2.29}$$

where $B_\epsilon(X) = \{y \mid \exists x \in X\ d_S(x, y) < \epsilon\}$ denotes the *ϵ-envelope around X*.

Thus, we can specify the convergence of the sequence of sets (Γ_i) as

$$\exists L \in H(S)\ \forall \epsilon > 0\ \exists N \in \mathbb{N}\ \forall n > N : d(\Gamma_i, L) < \epsilon \tag{2.30}$$

or as: there exists an $L \in H(S)$ such that both

$$\forall \epsilon > 0\ \exists N \in \mathbb{N}\ \forall n > N : \Gamma_n \subseteq B_\epsilon(L) \tag{2.31a}$$

and

$$\forall \epsilon > 0\ \exists N \in \mathbb{N}\ \forall n > N : B_\epsilon(\Gamma_n) \supseteq L\ . \tag{2.31b}$$

Chapter 3

Random Set-Walk with a Constant Internal State

Our development of the theory of random set-walk begins by exploring the case where the internal state space E satisfies $|E| = 1$, that is, $E = \{e\}$ where e is the sole element of the state space. When there is only a single internal state, all transition probability matrices can be reduced to scalar transition probabilities. Hence a transition triple T is in $[0,1]^3$, and a characteristic tuple φ is in $[0,1]^4$.

3.1 Main Result

The main result of this thesis concerns the long-run behaviour of random set-walks. Since, for the majority of transition sets $\mathcal{T}$, the external states $[1:n] \subset [0:n+1]$ of the walk instance $W(s)$ of $s \in \mathbb{S}_{[1:n]}(\mathcal{T})$ are transient, the stationary distribution is governed solely by the probabilities given in the characteristic tuple of the walk instance. Hence, looking at the set of characteristic probabilities over all $s \in \mathbb{S}_{[1:n]}(\mathcal{T})$ provides some insight into the stationary properties of the random set-walk $\mathcal{W}(\mathbb{S}_{[1:n]}(\mathcal{T}))$. We discovered that, in fact, these sets of characteristic probabilities converge as we lengthen the random set-walk by letting $n \to \infty$. In other words, the long-run behaviour of a random set-walk relative to its endpoints becomes increasingly independent of the length of the set-walk as it approaches infinity.

Definition 3.1 (Characteristic tuple set). Define the 4-tuple $\Phi_n(\mathcal{T}) = \langle \mathrm{A}_n, \mathrm{B}_n, \Gamma_n, \Delta_n \rangle$ as the *characteristic tuple set* of $\mathbb{S}_{[1:n]}(\mathcal{T})$, where

$$\begin{aligned}
\mathrm{A}_n &= a[\mathbb{S}_{[1:n]}(\mathcal{T})] &&= \{a[s] \mid s \in \mathbb{S}_{[1:n]}(\mathcal{T})\} \\
\mathrm{B}_n &= b[\mathbb{S}_{[1:n]}(\mathcal{T})] &&= \{b[s] \mid s \in \mathbb{S}_{[1:n]}(\mathcal{T})\} \\
\Gamma_n &= c[\mathbb{S}_{[1:n]}(\mathcal{T})] &&= \{c[s] \mid s \in \mathbb{S}_{[1:n]}(\mathcal{T})\} \\
\Delta_n &= d[\mathbb{S}_{[1:n]}(\mathcal{T})] &&= \{d[s] \mid s \in \mathbb{S}_{[1:n]}(\mathcal{T})\}.
\end{aligned}$$

Whereas a characteristic tuple describes a random walk in a nonhomogeneous bounded environment that is governed by a particular strategy s, a characteristic tuple set describes the random set-walk $\mathcal{W}(\mathbb{S}_{[1:n]}(\mathcal{T}))$ governed by the set of strategies $\mathbb{S}_{[1:n]}(\mathcal{T})$.

Theorem 3.2 (Main Theorem). *If the internal state space E contains only a single element, then, for any non-empty transition set $\mathcal{T}$, the sequence of characteristic tuple sets $(\Phi_n(\mathcal{T}))$ converges componentwise in the Hausdorff metric.*

This theorem statement is in fact slightly stronger than the claim made in the introductory paragraph of this section: we will prove the convergence of the characteristic tuple sets even when the individual walk instances $W(s)$ contain recurrent states other than the endpoints, which is the case when $(0, 1, 0) \in \mathcal{T}$.

3.1.1 Overview of the Proof

In the following sections, we present the proofs of the Main Theorem for various different cases, from which we will then be able to piece together the proof of the theorem in its most general context. The special cases that we will consider are as follows:

1. The case where all the transition triples in the transition set are "left-leaning", which is to say $\gamma[\mathsf{T}] \leq 1/2$ for each $\mathsf{T} \in \mathcal{T}$. The convergence proof for this case is presented in Section 3.3.

2. The case where all the triples in the transition set are "right-leaning", which is to say $\gamma[\mathsf{T}] > 1/2$ for each $\mathsf{T} \in \mathcal{T}$. The convergence proof for this case is presented in Section 3.4.

3. The case where there is a mix of left- and right-leaning transition triples in the transition set. The convergence proof for this case is presented in Section 3.5.

4. The border cases involving the triples $(1,0,0)$, $(0,0,1)$, and $(0,1,0)$ — the convergence proofs for which are presented in Sections 3.6, 3.7 and 3.8 respectively.

Note that, in these sections, the main focus will be establishing the convergence of the sets Γ_n. The convergence of Φ_n will in fact follow quite naturally, and is proved in full in Section 3.8.

3.2 Preliminaries

Here, we present some fundamental results concerning characteristic tuples and strategies that will be used in the proofs in upcoming sections.

3.2.1 *The Transition Triple* $\mathsf{T}(x)$

Definition 3.3. Define the transition triple $\mathsf{T}(x)$ by

$$\mathsf{T}(x) = (1-x, 0, x) \ .$$

In other words, $\alpha[\mathsf{T}(x)] = 1-x$, $\beta[\mathsf{T}(x)] = 0$, and $\gamma[\mathsf{T}(x)] = x$. Since we are concerned with random set-walks where $|E| = 1$, the function $\mathsf{T}(x)$ allows us to specify transition triples by their transition probabilities in a straightforward manner.

Moreover, this definition makes it simpler to divide the proof of the Main Theorem into special cases based on the "leanings" of the triples.

3.2.2 Basic Convergence Results

The following are three basic results regarding the convergence of a sequence of characteristic probabilities associated with a strategy sequence. These results are to be the building blocks with which we construct the subproofs of the Main Theorem.

Proposition 3.4. *If* $\mathsf{T}(x) \in \mathcal{T}$ *where* $x \leq 1/2$, *and* $U(\mathsf{T}(x)) = (s_n)$ *is a uniform strategy sequence, then there is a limit* c_∞ *such that* $c[s_n] \uparrow c_\infty$, *and* $d[s_n] \downarrow 0$.

Proof. By the gambler's ruin result (2.12),

$$c[s_n] = \begin{cases} \frac{\left(\frac{1-x}{x}\right)^n - 1}{\left(\frac{1-x}{x}\right)^{n+1} - 1} & \text{if } x < 1/2 \\ 1 - x/n & \text{if } x = 1/2 \end{cases}$$

$$\therefore \lim_{n\to\infty} c[s_n] = \frac{x}{1-x}$$

and

$$d[s_n] = \begin{cases} \frac{\frac{1-x}{x} - 1}{\left(\frac{1-x}{x}\right)^{n+1} - 1} & \text{if } x < 1/2 \\ x/n & \text{if } x = 1/2 \end{cases}$$

$$\therefore \lim_{n\to\infty} d[s_n] = 0, \text{ since } x \leq 1/2$$

□

Proposition 3.5. *Let* (s_i) *be a monotone strategy sequence on* $\mathcal{T}$ *where the sequence* $(c[s_i])$ *converges. Then for any* $s' \in \mathbb{S}(\mathcal{T})$, *the sequence* $(c[s_i s'])$ *also converges.*

Proof. Let $\varphi_{s'} = \langle a, b, c, d \rangle$. Then

$$\begin{aligned} c[s_i s'] &= c + \frac{a\, c[s_i]\, d}{1 - c[s_i] b} \\ &= f(c[s_i]) \qquad \text{where } f(x) = c + \frac{axd}{1 - xb}. \end{aligned}$$

By the continuity of f on $[0,1]$, if $c[s_i]$ converges as $i \to \infty$, then $f(c[s_i])$ converges as well. $\square$

Corollary 3.6. *Suppose $s \in \mathbb{S}(\mathcal{T})$, and $(l_i) = U(\mathsf{T}(l))$ is a uniform strategy sequence where $l < 1/2$ and $\mathsf{T}(l) \in \mathcal{T}$. Then the sequence $(c[l_i s])$ converges.*

Proof. By Proposition 3.4, $c[l_i]$ converges as $i \to \infty$, and hence it follows that $c[l_i s]$ converges as well, as $i \to \infty$, by Proposition 3.5. $\square$

3.2.3 The Comparison Lemma

Just as important as the previous results on convergence is the Comparison Lemma, so named because it allows us to state inequalities about two characteristic probabilities $c[s]$ and $c[s']$ based solely on the composition of the strategies s and s'. In essence, there is a partial ordering of the strategies in $\mathbb{S}(\mathcal{T})$ that is similar in nature to — but not the same as — a typographical ordering: for $s, s' \in \mathbb{S}_{[1:n]}(\mathcal{T})$, $c[s'] \leq c[s]$ if $\gamma[s'(z)] \leq \gamma[s(z)]$ for each $z \in [1:n]$.

Proposition 3.7. *Let $s_l, s_r, s \in \mathbb{S}(\mathcal{T})$. If $c[s_l] < c[s_r]$, then $c[s_l s] < c[s_r s]$.*

Proof. Let $\varphi_s = \langle a, b, c, d \rangle$. Then

$$c[s_l s] = c + \frac{a\, c[s_l] d}{1 - c[s_l] b}$$
$$c[s_r s] = c + \frac{a\, c[s_r] d}{1 - c[s_r] b}$$

and

$$\frac{a\, c[s_l] d}{1 - c[s_l] b} < \frac{a\, c[s_r] d}{1 - c[s_r] b} \qquad \text{since } c[s_l] < c[s_r],$$

thus $c[s_l s] < c[s_r s]$. $\square$

Proposition 3.8. *Let $s \in \mathbb{S}(\mathcal{T})$, and let $\mathsf{T}(l)$ and $\mathsf{T}(r)$ be two transition triples in $\mathcal{T}$. If $l < r$ and $c[s] < 1$, then $c[s\mathsf{T}(l)] < c[s\mathsf{T}(r)]$.*

Proof. First, if $l = 0$, then $c[s\mathsf{T}(l)] = 0 < c[s\mathsf{T}(r)]$, since $r > l = 0$.

Let $c = c[s]$. For the case $l > 0$, we have:

$$\begin{aligned} r - l &> 0 \\ 1 - c &> 0 \\ c(1-r) < c(1-l) &< 1 \\ 1 - c(1-r) > 1 - c(1-l) &> 0 \end{aligned}$$

So

$$\begin{aligned} \frac{(1-c)(r-l)}{(1-c(1-r))(1-c(1-l))} &> 0 \\ \frac{l}{1-c(1-l)} &< \frac{r}{1-c(1-r)} \\ \underbrace{l + \frac{(1-l)cl}{1-c(1-l)}}_{=c[s\mathsf{T}(l)]} &< \underbrace{r + \frac{(1-r)cr}{1-c(1-r)}}_{=c[s\mathsf{T}(r)]} \end{aligned}$$

Therefore, $c[s\mathsf{T}(l)] < c[s\mathsf{T}(r)]$. ☐

Proposition 3.9. *If $t, s \in \mathbb{S}(\mathcal{T})$, $c[t] < 1$, $c[s] < 1$, and $l < r$, then $c[t\mathsf{T}(l)s] < c[t\mathsf{T}(r)s]$.*

Proof. By Proposition 3.7, $c[\mathsf{T}(l)s] < c[\mathsf{T}(r)s]$. Thus it follows from Proposition 3.8 that $c[t\mathsf{T}(l)s] < c[t\mathsf{T}(r)s]$. ☐

Lemma 3.10 (Comparison Lemma). *For any $s, s' \in \mathbb{S}_{[a:b]}(\mathcal{T})$ such that $\gamma[s'(z)] \leq \gamma[s(z)]$ for each $z \in [a:b]$, the inequality $c[s'] \leq c[s]$ holds.*

Proof. Applying Proposition 3.9 to each position $z \in [a:b]$ where $\gamma[s'(z)] < \gamma[s(z)]$ yields this result. ☐

3.3 Left-leaning Transition Sets

This is the first in a series of special cases of the general result as stated in the Main Theorem. The intuition here for left-leaning transition sets is that the points in Γ_n are grouped together, and these groups refine and evolve in a self-similar fashion.

Let $0 < x_1 < x_2 < \cdots < x_d < 1/2$ be a non-empty, finite set of real numbers, and

$$\mathcal{T} = \{\mathsf{T}_i = \mathsf{T}(x_i) \mid i \in [1:d]\}$$

be a transition set.

For each x_i, define the uniform strategy sequence $(\mathsf{T}_i^j)_{j\in\mathbb{N}} = U(\mathsf{T}_i)$.

3.3.1 *Implicit Construction of the Limit Set*

While we may not be able to describe the exact shape of the limit set, we can construct sequences of approximations which themselves converge to the limit set. It is easier to work with these approximations than with the actual Γ_n, as they have relatively straightforward definitions compared to that of Γ_n.

Definition 3.11. Let $s \in \mathbb{S}(\mathcal{T})$. Define the sets $L(s)$ and $I(s)$ by

$$L(s) = \left\{ \lim_{m\to\infty} c[\mathsf{T}_i^m s] \mid i \in [1:d] \right\} \tag{3.1}$$

$$I(s) = \left[\lim_{m\to\infty} c[\mathsf{T}_1^m s], \lim_{m\to\infty} c[\mathsf{T}_d^m s] \right] . \tag{3.2}$$

Now define the sets L_n and I_n in terms of $L(s)$ and $I(s)$ by

$$L_n = \bigcup_{s\in\mathbb{S}_{[1:n]}(\mathcal{T})} L(s) \tag{3.3}$$

$$I_n = \bigcup_{s\in\mathbb{S}_{[1:n]}(\mathcal{T})} I(s) . \tag{3.4}$$

Note that

$$L_n \subseteq I_n . \tag{3.5}$$

The sets L_n and I_n represent, respectively, the "lower" and "upper" approximations of the limit set, given information about the strategies $\mathbb{S}_{[1:n]}(\mathcal{T})$. In particular, the set $L(s)$ contains a finite number of isolated points, two of which are the minimum and maximum value of $c[s's]$ as computed over all strategies $s' \in \mathbb{S}(\mathcal{T})$, while the set $I(s)$ is the interval between said minimum and maximum.

So intuitively, the limit set should "contain" L_n for all n, while at the same time be "contained" by all I_n for all n. Moreover, it will be shown that, for sufficiently large n, the sets L_n, Γ_n and I_n are all quite close to each other, with distances measured with the Hausdorff metric. Hence the limiting nature of L_n and I_n should give us the desired result, namely Γ_n converges to this elusive limit set Λ.

In order to establish that the sets L_n and I_n do indeed converge, we will prove that the sequence L_n is monotonically increasing, while the sequence I_n is monotonically decreasing. This will imply that $L_n \to \bigcup_n L_n$, and $I_n \to \bigcap_n I_n$.

Proposition 3.12. $L_n \subseteq L_{n+1}$.

Proof. First, if $s' \in \mathbb{S}_{[1:n+1]}(\mathcal{T})$, then $s' = \mathsf{T}_i s$ for some $\mathsf{T}_i \in \mathcal{T}$ and $s \in \mathbb{S}_{[1:n]}(\mathcal{T})$. Hence, we begin with the proof of

$$L(s) \subseteq \bigcup_{j \in [1:d]} L(\mathsf{T}_j s) \tag{3.6}$$

with the observation that

$$\begin{aligned}
L(s) &= \left\{ \lim_{m\to\infty} c[\mathsf{T}_i^m s] \mid i \in [1:d] \right\} \\
&= \left\{ \lim_{m\to\infty} c[\mathsf{T}_i^m \mathsf{T}_i s] \mid i \in [1:d] \right\} \\
&\subseteq \left\{ \lim_{m\to\infty} c[\mathsf{T}_i^m \mathsf{T}_j s] \mid i \in [1:d], j \in [1:d] \right\} \\
&= \bigcup_{j \in [1:d]} \left\{ \lim_{m\to\infty} c[\mathsf{T}_i^m \mathsf{T}_j s] \mid i \in [1:d] \right\} \\
&= \bigcup_{j \in [1:d]} L(\mathsf{T}_j s) \ .
\end{aligned}$$

Now, since

$$L_n = \bigcup_{s \in \mathbb{S}_{[1:n]}(\mathcal{T})} L(s) \subseteq \bigcup_{s \in \mathbb{S}_{[1:n]}(\mathcal{T})} \bigcup_{j \in [1:d]} L(\mathsf{T}_j s) = \bigcup_{s' \in \mathbb{S}_{[1:n+1]}(\mathcal{T})} L(s') = L_{n+1} \ , \tag{3.7}$$

we can infer the result that $L_n \subseteq L_{n+1}$. $\square$

Proposition 3.13. $I_n \supseteq I_{n+1}$.

Proof. Again, we will begin with the proof of

$$I(s) \supseteq \bigcup_{j\in[1:d]} I(\mathsf{T}_j s) \tag{3.8}$$

by observing that

$$\begin{aligned}
I(s) &= \left[\lim_{m\to\infty} c[\mathsf{T}_1^m s], \lim_{m\to\infty} c[\mathsf{T}_d^m s]\right] \\
&= \left[\lim_{m\to\infty} c[\mathsf{T}_1^m \mathsf{T}_1 s], \lim_{m\to\infty} c[\mathsf{T}_d^m \mathsf{T}_d s]\right] \\
&\supseteq \bigcup_{j\in[1:d]} \left[\lim_{m\to\infty} c[\mathsf{T}_1^m \mathsf{T}_j s], \lim_{m\to\infty} c[\mathsf{T}_d^m \mathsf{T}_j s]\right] \quad \text{(by the Comparison Lemma)} \\
&= \bigcup_{j\in[1:d]} I(\mathsf{T}_j s)\ .
\end{aligned}$$

Therefore, since

$$I_n = \bigcup_{s\in\mathbb{S}_{[1:n]}(\mathcal{T})} I(s) \supseteq \bigcup_{s\in\mathbb{S}_{[1:n]}(\mathcal{T})} \bigcup_{j\in[1:d]} I(\mathsf{T}_j s) = \bigcup_{s'\in\mathbb{S}_{[1:n+1]}(\mathcal{T})} I(s') = I_{n+1}\ , \tag{3.9}$$

we can conclude that $I_n \supseteq I_{n+1}$. ☐

With the convergence of L_n and I_n proven, we proceed to show that the sets L_n, Γ_n, and I_n are close for sufficiently large values of n.

Lemma 3.14. *Given an arbitrary $\epsilon > 0$, there is an $N_0 \in \mathbb{N}$ such that for all $n > N_0$, $B_\epsilon(\Gamma_n) \supseteq I_n$.*

Proof. As $\lim_{n\to\infty} d[\mathsf{T}_d^n] = 0$, let us define N_0 to be such that $d[\mathsf{T}_d^n] < \epsilon/2$ for all $n > N_0$. Now, let $s \in \mathbb{S}_{[1:n]}(\mathcal{T})$, where $n > N_0$. Then, for all $i \in [1:d]$ and $m \in \mathbb{N}$,

$$\begin{aligned}
c[\mathsf{T}_i^m s] - c[s] &= a[\mathsf{T}_i^m](1 - c[\mathsf{T}_i^m]b[s])^{-1}c[\mathsf{T}_i^m]d[s] \\
&\le d[s] \qquad \because a[\mathsf{T}_i^m](1 - c[\mathsf{T}_i^m]b[s])^{-1}c[\mathsf{T}_i^m] \le 1 \\
&\le d[\mathsf{T}_d^n] \qquad \text{by the Comparison Lemma} \qquad < \epsilon/2\ .
\end{aligned}$$

Therefore, $\lim_{m\to\infty} c[\mathsf{T}_1^m s_n] - c[s_n] \leq \epsilon/2$ and $\lim_{m\to\infty} c[\mathsf{T}_d^m s_n] - c[s_n] \leq \epsilon/2$, leading us to conclude that

$$\begin{aligned} I(s) &= \left[\lim_{m\to\infty} c[\mathsf{T}_1^m s], \lim_{m\to\infty} c[\mathsf{T}_d^m s]\right] \\ &\subseteq B_\epsilon(c[s]) \end{aligned}$$

for all $s \in \mathbb{S}_{[1:n]}(\mathcal{T})$, and hence

$$\begin{aligned} I_n &= \bigcup_{s\in\mathbb{S}_{[1:n]}(\mathcal{T})} I(s) \\ &\subseteq \bigcup_{s\in\mathbb{S}_{[1:n]}(\mathcal{T})} B_\epsilon(c[s]) = B_\epsilon(\Gamma_n) \end{aligned}$$

for all $n > N_0$.

$\square$

Lemma 3.15. *Given an arbitrary $\epsilon > 0$, there is an $N_1 \in \mathbb{N}$ such that for all $n > N_1$, $\Gamma_n \subseteq B_\epsilon(L_n)$.*

Proof. As $\lim_{n\to\infty} d[\mathsf{T}_d^n] = 0$, let us define N_1 to be such that $d[\mathsf{T}_d^n] < \epsilon/2$ for all $n > N_1$. Now, let $s \in \mathbb{S}_{[1:n]}(\mathcal{T})$, where $n > N_1$. Then, by the reasoning employed in the previous proof, $c[\mathsf{T}_i^m s] - c[s] < \epsilon/2$ for all $i \in [1:d]$ and $m \in \mathbb{N}$.

Therefore, $\lim_{m\to\infty} c[\mathsf{T}_i^m s_n] - c[s_n] \leq \epsilon/2$ for all $i \in [1:d]$, leading us to conclude that

$$\begin{aligned} \delta(\{c[s]\}, L(s)) &< \epsilon \\ \delta\left(\bigcup_{s\in\mathbb{S}_{[1:n]}(\mathcal{T})} \{c[s]\}, \bigcup_{s\in\mathbb{S}_{[1:n]}(\mathcal{T})} L(s)\right) &< \epsilon \\ \delta(\Gamma_n, L_n) &< \epsilon \\ \Gamma_n &\subseteq B_\epsilon(L_n) \end{aligned}$$

for all $n > N_1$. $\square$

Having just proven the proximity of L_n, Γ_n and I_n to each other, we can finally identify the candidate for the limit set Λ. The set that we will consider is $\Lambda = \bigcap_n I_n$. Before proving the convergence result, we will need to establish that there is no gap between the limit $\overline{\bigcup_n L_n}$, and the limit $\bigcap_n I_n$, since otherwise there may be more than one choice for Λ, which would work against the proof that $\Gamma_n \to \Lambda$.

Note that we have to take the closure of $\bigcup_n L_n$, because while $\bigcup_n L_n$ contains a point $\lim_{m\to\infty} c[\mathsf{T}_i^m s]$ for each finite strategy $s \in \mathbb{S}(\mathcal{T})$ and $i \in [1 : d]$, it fails to capture monotone strategy sequences whose "limiting" sequences are "irrational", in the sense that they do not terminate with an infinite string of T_i. This is similar to the case of the Cantor set, where the majority of the points in the set are not the endpoints of the intervals removed (cf. $\bigcup_n L_n$), but rather their limit points (cf. $\overline{\bigcup_n L_n}$), which happen to be exactly the points in the intervals that were not removed (cf. $\bigcap_n I_n$).

Lemma 3.16. $\overline{\bigcup_n L_n} = \bigcap_n I_n$

Proof. First, we shall show $\overline{\bigcup_n L_n} \subseteq \bigcap_n I_n$.

Let $z \in \overline{\bigcup_n L_n}$. The definition of the closure of a set implies that there is a sequence lying strictly in $\bigcup_n L_n$ that converges to z. Thus, we may find a subsequence (z_i) of such a sequence, where $z_i \in L_{n_i}$, and $n_i < n_j$ for all $i < j$. Suppose $z \notin \bigcap_n I_n$. Then it must be that $z \notin I_N$ for some N. By Proposition 3.13, $I_n \supseteq I_m$ for all $n < m$, and hence $z \notin I_n$ for all $n > N$. Let $\epsilon = \delta(\{z\}, I_N) > 0$, where $\delta(\cdot,\cdot)$ is as defined in (2.27). So $\delta(z, I_n) \geq \delta(\{z\}, I_N) = \epsilon > 0$ for all $n > N$. Moreover, as $L_n \subset I_n$ for all n, $\delta(\{z\}, L_n) \geq \epsilon$ for all $n > N$. However, since $z_i \to z$, $\delta(\{z\}, L_n) \to 0$ as $n \to \infty$. This is a contradiction. Therefore we can conclude that if $z \in \overline{\bigcup_n L_n}$, then $z \in \bigcap_n I_n$.

Now, we will the show the reverse inclusion $\overline{\bigcup_n L_n} \supseteq \bigcap_n I_n$.

Let $z \in \bigcap_n I_n$. Suppose $z \notin \overline{\bigcup_n L_n}$. Then there exists an $\epsilon > 0$ such that $\delta(\{z\}, L_n) \geq \epsilon > 0$ for all large n. Now consider the ball $B_{\epsilon/2}(z)$. Let N be such that the intervals making up I_N have lengths less than $\epsilon/2$. Since $z \in \bigcap_n I_n \supseteq I_N$, $L_N \cap B_{\epsilon/2}(z)$ is non-empty, since the ball must cover at least one endpoint of the particular interval containing the point z. Thus $\delta(\{z\}, L_N) < \epsilon/2$, a contradiction. Consequently, we can infer that if $z \in \bigcap_n I_n$, then $z \in \overline{\bigcup_n L_n}$.

Thus, $\overline{\bigcup_n L_n} \subseteq \bigcap_n I_n \subseteq \overline{\bigcup_n L_n}$. In other words, $\overline{\bigcup_n L_n} = \bigcap_n I_n$. □

With the framework set up, we can finally prove the Main Theorem in the special case of a left-leaning transition set.

Theorem 3.17. $\lim_{n\to\infty} \Gamma_n = \Lambda = \bigcap_n I_n$ *in the Hausdorff metric.*

Proof. Let $\epsilon > 0$ be arbitrary. By Lemma 3.15, there is an N_1 such that for all $n > N_1$,

$$\begin{aligned}
\Gamma_n &\subseteq B_\epsilon(L_n) \\
&\subseteq B_\epsilon(\bigcup_m L_m) \\
&\subseteq B_\epsilon(\overline{\bigcup_m L_m}) \\
&= B_\epsilon(\Lambda) \qquad \text{by Lemma 3.16,}
\end{aligned}$$

so the requirement (2.31a) is satisfied. As well, by Lemma 3.14, there is an N_0 such that for all $n > N_0$,

$$\begin{aligned}
B_\epsilon(\Gamma_n) &\supseteq I_n \\
&\supseteq \bigcap_m I_m \\
&= \Lambda \qquad \text{by Lemma 3.16,}
\end{aligned}$$

so the requirement (2.31b) is also satisfied. This implies that $\Gamma_n \to \Lambda$ in the Hausdorff metric, as $n \to \infty$. □

3.4 Right-leaning Transition Sets

Compared to the previous case of left-leaning transition sets, this second special case is easier to deal with. Intuitively, starting a random set-walk near the right absorbing barrier, and then moving about as dictated by a set of right-leaning transition triples should result in a high probability in being ultimately absorbed by the right barrier. In fact, as n tends to infinity, all such probabilities should tend to 1. We show the correctness of this reasoning in this section.

Let $1/2 \leq y_1 < y_2 < \cdots < y_d < 1$ be a non-empty, finite set of real numbers, and

$$\mathcal{T} = \{\mathsf{T}_i = \mathsf{T}(y_i) \mid i \in [1:d]\}$$

be a transition set.

For each y_i, define the uniform strategy sequence $(\mathsf{T}_i^j)_{j\in\mathbb{N}} = U(\mathsf{T}_i)$.

Theorem 3.18. $\lim_{n\to\infty} \Gamma_n = \{1\}$.

Proof. For each strategy $s \in \mathbb{S}_{[1:n]}(\mathcal{T})$, $c[\mathsf{T}_1^n] \leq c[s] \leq c[\mathsf{T}_d^n]$ by the Comparison Lemma. Moreover, for all $i \in [1:d]$, $c[\mathsf{T}_i^n] \to 1$ as $n \to \infty$ by the gambler's ruin result (2.12) (for both the case $y_1 = 1/2$ and the case $y_1 > 1/2$). Thus $\Gamma_n \subseteq [c[\mathsf{T}_1^n], c[\mathsf{T}_d^n]]$ implies that $\lim_{n\to\infty} \Gamma_n \subseteq \{1\}$. As $\lim \Gamma_n \neq \phi$, $\lim \Gamma_n = \{1\}$. ☐

3.5 Mixed Transition Sets

In this third special case of the Main Theorem, we approach the problem of proving the convergence result from a different angle. We will actually specify at the outset the candidate for the limit set Λ, and then proceed to prove that the candidate is indeed the correct limit set. The method used to prove the correctness of the candidate mimics that of a greedy algorithm, where we show that given any point λ in the candidate set Λ, we can systematically construct a sequence of points $(c[s_n^\lambda])_{n\in\mathbb{N}}$, where $c[s_n^\lambda] \in \Gamma_n$, converging to λ. Hence we may conclude that the candidate limit is "close" to the sets Γ_n for sufficiently large n.

Let $0 < x_1 < x_2 < \cdots < x_d < 1$, where $x_1 < 1/2 < x_d$ and $d > 1$, be a non-empty, finite set of real numbers, and

$$\mathcal{T} = \{\mathsf{T}_i = \mathsf{T}(x_i) \mid i \in [1:d]\}$$

be a transition set.

For each x_i, define the uniform strategy sequence $(\mathsf{T}_i^j)_{j\in\mathbb{N}} = U(\mathsf{T}(x_i))$.

Proposition 3.19. *The candidate for the limit set Λ is the closed interval*

$$\Lambda = \left[\lim_{m\to\infty} c[\mathsf{T}_1^m], 1\right] \ . \tag{3.10}$$

With the existing results proven in Section 3.2, we can already satisfy the condition (2.31a), which is half of the convergence argument.

Lemma 3.20. $\lim_{n\to\infty} \delta(\Gamma_n, \Lambda) = 0$.

Proof. By the Comparison Lemma, $c[s] \geq c[\mathsf{T}_1^n]$ for all $s \in \mathbb{S}_{[1:n]}(\mathcal{T})$, implying that $\Gamma_n \subseteq [c[\mathsf{T}_1^n], 1]$. By Proposition 3.4, $c[\mathsf{T}_1^n]$ converges to $\lim_{m\to\infty} c[\mathsf{T}_1^m] = x_1/(1-x_1)$. Hence $\delta(\Gamma_n, \Lambda) \leq \delta([c[\mathsf{T}_1^n], 1]\,, \Lambda) = (\lim_{m\to\infty} c[\mathsf{T}_1^m]) - c[\mathsf{T}_1^n]$, which approaches 0 as $n \to \infty$. In other words, $\lim_{n\to\infty} \delta(\Gamma_n, \Lambda) = 0$. ☐

Let us define the function

$$\begin{aligned} cl(s) &= \lim_{m\to\infty} c[\mathsf{T}_1^m s] \\ &= c[s] + \frac{a[s]\left(\frac{x_1}{1-x_1}\right)d[s]}{1 - b[s]\left(\frac{x_1}{1-x_1}\right)} \ . \end{aligned} \tag{3.11}$$

Intuitively, $cl(s)$ represents a lower bound for $c[s's]$ for any "infinite" strategy s'. It will be used extensively in the greedy algorithm.

Below, we prove three results about rates of convergence related to characteristic probabilities, to be used later in proving the correctness of the greedy algorithm.

Proposition 3.21. *For any strategy $s \in \mathbb{S}(\mathcal{T})$ where $c[s] < 1$,*

$$\lim_{m\to\infty} c[s\mathsf{T}_1^m] = \lim_{m\to\infty} c[\mathsf{T}_1^m] \ . \tag{3.12}$$

Proof. Let

$$\varphi_{\mathsf{T}_1^m} = \langle a, b, c, d \rangle$$
$$= \left\langle \frac{(\frac{1-x_1}{x_1})^{m+1} - (\frac{1-x_1}{x_1})^m}{(\frac{1-x_1}{x_1})^{m+1} - 1}, \frac{(\frac{1-x_1}{x_1})^{m+1} - \frac{1-x_1}{x_1}}{(\frac{1-x_1}{x_1})^{m+1} - 1}, \frac{(\frac{1-x_1}{x_1})^{m+1} - 1}{(\frac{1-x_1}{x_1})^{m+1} - 1}, \frac{\frac{1-x_1}{x_1} - 1}{(\frac{1-x_1}{x_1})^{m+1} - 1} \right\rangle .$$

Thus

$$\begin{aligned} c[s\mathsf{T}_1^m] &= c + \frac{a\, c[s]\, d}{1 - c[s]b} \\ &= c + a\, c[s] \frac{d}{1 - c[s](1-d)} \\ &= c + a\, c[s] \left(\frac{1 - c[s]}{d} + c[s] \right)^{-1} . \end{aligned}$$

As well, we have $\left(\frac{1-c[s]}{d} + c[s] \right)^{-1} \to 0$ as $m \to \infty$, since $c[s] < 1$, $x_1 < 1/2$, and $d \doteq d[\mathsf{T}_1^m] \to 0$ as $m \to \infty$. Hence

$$\lim_{m\to\infty} |c[s\mathsf{T}_1^m] - c| = \lim_{m\to\infty} |c[s\mathsf{T}_1^m] - c[\mathsf{T}_1^m]| = 0 ,$$

yielding

$$\lim_{m\to\infty} c[s\mathsf{T}_1^m] = \lim_{m\to\infty} c[\mathsf{T}_1^m] .$$

□

Proposition 3.22. *For any strategy $s \in \mathbb{S}(\mathcal{T})$,*

$$\lim_{m\to\infty} cl(\mathsf{T}_j \mathsf{T}_1^m s) = cl(s) . \tag{3.13}$$

Proof. First, let $s'_n = \mathsf{T}_1^n \mathsf{T}_j$, and notice that

$$\begin{aligned} cl(\mathsf{T}_j \mathsf{T}_1^m) &= \lim_{n\to\infty} c[\mathsf{T}_1^n \mathsf{T}_j \mathsf{T}_1^m] \\ &= \lim_{n\to\infty} c[\mathsf{T}_1^m] + \frac{a[\mathsf{T}_1^m] c[s'_n] d[\mathsf{T}_1^m]}{1 - c[s'_n] b[\mathsf{T}_1^m]} \\ &= c[\mathsf{T}_1^m] + \frac{a[\mathsf{T}_1^m] c_\infty d[\mathsf{T}_1^m]}{1 - c_\infty b[\mathsf{T}_1^m]} \qquad \text{where } c_\infty = \lim_{n\to\infty} c[s'_n] . \end{aligned}$$

Hence, since $c_\infty < 1$ by the nature of the sequence $s'_n = \mathsf{T}_1^n \mathsf{T}_j$, we can apply the result derived in the proof of the previous proposition and get

$$\lim_{m\to\infty} cl(\mathsf{T}_j \mathsf{T}_1^m) = \lim_{m\to\infty} c[\mathsf{T}_1^m] \,.$$

By Proposition 3.5, we can conclude

$$\lim_{m\to\infty} cl(\mathsf{T}_j \mathsf{T}_1^m s) = \lim_{m\to\infty} c[\mathsf{T}_1^m s] = cl(s) \,.$$

□

Proposition 3.23. *There is a $K \in (0,1]$ such that, for any strategy $s \in \mathbb{S}(\mathcal{T})$,*

$$\frac{cl(\mathsf{T}_d \mathsf{T}_1 s) - cl(s)}{cl(\mathsf{T}_d s) - cl(s)} \geq K > 0 \,. \tag{3.14}$$

Proof. Let $\varphi_s = \langle a, b, c, d\rangle$, then

$$\begin{aligned} cl(\mathsf{T}_d \mathsf{T}_1 s) - cl(s) &= \left(c + \frac{a\, cl(\mathsf{T}_d \mathsf{T}_1)\, d}{1 - cl(\mathsf{T}_d \mathsf{T}_1) b}\right) - \left(c + \frac{a\, cl(\mathsf{T}_1)\, d}{1 - cl(\mathsf{T}_1) b}\right) \\ &= \frac{a\, cl(\mathsf{T}_d \mathsf{T}_1)\, d}{1 - cl(\mathsf{T}_d \mathsf{T}_1) b} - \frac{a\, cl(\mathsf{T}_1)\, d}{1 - cl(\mathsf{T}_1) b} \\ &\geq \frac{ad(cl(\mathsf{T}_d \mathsf{T}_1) - cl(\mathsf{T} - 1))}{1 - cl(\mathsf{T}_d \mathsf{T}_1) b} \end{aligned}$$

(because $cl(\mathsf{T}_d \mathsf{T}_1) \geq cl(\mathsf{T}_1)$ by the Comparison Lemma)

$$\geq ad(cl(\mathsf{T}_d \mathsf{T}_1) - cl(\mathsf{T} - 1)) \,,$$

and

$$\begin{aligned} cl(\mathsf{T}_d s) - cl(s) &= \left(c + \frac{a\, cl(\mathsf{T}_d)\, d}{1 - cl(\mathsf{T}_d) b}\right) - \left(c + \frac{a\, cl(\mathsf{T}_1)\, d}{1 - cl(\mathsf{T}_1) b}\right) \\ &= \frac{a\, cl(\mathsf{T}_d)\, d}{1 - cl(\mathsf{T}_d) b} - \frac{a\, cl(\mathsf{T}_1)\, d}{1 - cl(\mathsf{T}_1) b} \\ &\leq \frac{ad\, cl(\mathsf{T}_d)}{1 - cl(\mathsf{T}_d) b} \\ &\leq \frac{ad\, cl(\mathsf{T}_d)}{1 - cl(\mathsf{T}_d)} \end{aligned}$$

So, taking the ratio of the two gives

$$\frac{cl(\mathsf{T}_d\mathsf{T}_1 s) - cl(s)}{cl(\mathsf{T}_d s) - cl(s)} \geq \frac{ad(cl(\mathsf{T}_d\mathsf{T}_1) - cl(\mathsf{T}-1))}{\frac{ad\ cl(\mathsf{T}_d)}{1-cl(\mathsf{T}_d)}}$$
$$= \frac{(cl(\mathsf{T}_d\mathsf{T}_1) - cl(\mathsf{T}-1))(1-cl(\mathsf{T}_d))}{cl(\mathsf{T}_d)}$$
$$= K > 0 \text{ , a constant.}$$

□

3.5.1 *Explicit Construction of the Limit Set*

Now, we need to prove $\lim_{n\to\infty}\delta(\Lambda,\Gamma_n) = 0$ so that, in conjunction with Lemma 3.20, we may establish the convergence result. We shall accomplish this through the following steps:

1. Given any $\epsilon > 0$, subdivide Λ into k intervals $[l_i, r_i]$ where the lengths $r_i - l_i < \epsilon/2$ for all i, and $\bigcup_{i\in[1:k]}[l_i, r_i] = \Lambda$.
2. For each interval $[l_i, r_i]$, pick the midpoint $p_i = (r_i - l_i)/2$, and find a backward monotone strategy sequence $(s_n^{p_i})_{n\in\mathbb{N}}$ such that $p_i = \lim_{n\to\infty} c[s_n^{p_i}]$.
3. Since there are only k strategy sequences, there exists an N such that for all $n > N$, $\lim_{m\to\infty} c[s_m^{p_i}] - c[s_n^{p_i}] < \epsilon/2$ *uniformly* in i.
4. Thus for each point $\lambda \in \Lambda$, $\lambda \in [l_i, r_i]$ for some $i \in [1:k]$. So, for all $n > N$,

$$|\lambda - c[s_n^{p_i}]| < |\lambda - p_i| + |p_i - \lim_{m\to\infty} c[s_m^{p_i}]| + |\lim_{m\to\infty} c[s_m^{p_i}] - c[s_n^{p_i}]|$$
$$< \frac{\epsilon}{2} + 0 + \frac{\epsilon}{2}$$
$$< \epsilon\ .$$

Hence, we have

$$\forall\epsilon > 0\ \exists N \in \mathbb{N}\ \forall n > N\ \Lambda \subseteq B_\epsilon(\Gamma_n) \tag{3.15}$$

and thus $\lim_{n\to\infty} \delta(\Lambda, \Gamma_n) = 0$ as required.

Consequently, our main task is to come up with a systematic way of identifying the backward monotone strategy sequences $(s_n^{p_i})_{n\in\mathbb{N}}$ satisfying the above property.

Lemma 3.24. *For each $\lambda \in \Lambda$, there exists a backward monotone strategy sequence, $(s_n^\lambda)_{n\in\mathbb{N}}$, such that*

$$\lim_{n\to\infty} c[s_n^\lambda] = \lambda \ . \tag{3.16}$$

Proof. Let $s_0^\lambda \in \mathbb{S}_{[0:0]}(\mathcal{T})$ be such that

$$s_0^\lambda(0) = \begin{cases} \mathsf{T}_d & \text{if } cl(\mathsf{T}_d) \leq \lambda \\ \mathsf{T}_1 & \text{otherwise,} \end{cases} \tag{3.17}$$

and let s_n^λ be defined recursively, for $n \geq 1$, as

$$s_n^\lambda = \begin{cases} \mathsf{T}_d s_{n-1}^\lambda \# & \text{if } cl(\mathsf{T}_d s_{n-1}^\lambda) \leq \lambda \\ \mathsf{T}_1 s_{n-1}^\lambda \# & \text{otherwise.} \end{cases} \tag{3.18}$$

Thus (s_n^λ) is a backward monotone sequence.

Since $c[s_n^\lambda] \leq c[s_{n+1}^\lambda] \leq 1$ holds for all n, the sequence $c[s_n^\lambda]$ converges as $n \to \infty$. Thus, we need to show that the limit of the above sequence is indeed λ.

First, we shall establish that $cl(s_n^\lambda) \leq \lambda$ for all n by induction. For the base case when $n = 0$, note that if $cl(\mathsf{T}_d) \leq \lambda$, then $s_0^\lambda = \mathsf{T}_d\#$, and so $cl(s_0^\lambda) = cl(\mathsf{T}_d\#) = cl(\mathsf{T}_d) \leq \lambda$. On the other hand, if $cl(\mathsf{T}_d) > \lambda$, then $s_0^\lambda = \mathsf{T}_1\#$, implying that $cl(s_0^\lambda) = cl(\mathsf{T}_1\#) = \lim_{m\to\infty} c[\mathsf{T}_1^m]$ by the definition of $cl(\cdot)$. However, as $\lambda \in \Lambda = [\lim_{m\to\infty} c[\mathsf{T}_1^m], 1]$, we can conclude that $cl(s_0^\lambda) = \lim_{m\to\infty} c[\mathsf{T}_1^m] \leq \lambda$.

For the inductive step, let us assume the hypothesis that $cl(s_n^\lambda) \leq \lambda$, and look at $cl(s_{n+1}^\lambda)$. By definition, if $cl(\mathsf{T}_d s_n^\lambda) \leq \lambda$, then $s_{n+1}^\lambda = \mathsf{T}_d s_n^\lambda$. Hence for this case, $cl(s_{n+1}^\lambda) = cl(\mathsf{T}_d s_n^\lambda) \leq \lambda$. If, on the other hand, $cl(\mathsf{T}_d s_n^\lambda) > \lambda$, then $s_{n+1}^\lambda = \mathsf{T}_1 s_n^\lambda$. By the definition of $cl(\cdot)$, $cl(s_{n+1}^\lambda) = cl(\mathsf{T}_1 s_n^\lambda) = cl(s_n^\lambda)$. However, as $cl(s_n^\lambda) \leq \lambda$ by the induction hypothesis, we have $cl(s_{n+1}^\lambda) \leq \lambda$ for this case as well.

Thus, by induction, we have shown that $cl(s_n^\lambda) \leq \lambda$ for all n.

Next, notice that, for any strategy s, $c[\mathsf{T}_d^m s] \to 1$ as $m \to \infty$, and so, $cl(\mathsf{T}_d^m s)$ converges to 1 as well. Thus, for each n, there exists an $M = M(n)$ such that

$$cl(\mathsf{T}_d^{M-1} s_n^\lambda) \leq \lambda \text{ and } cl(\mathsf{T}_d^{M} s_n^\lambda) > \lambda \ .$$

Hence by (3.18), $s_{n+M-1}^\lambda = \mathsf{T}_d^{M-1 s_n^\lambda}$.

The idea here is to group the triples assigned by the strategy $s_{n_w}^\lambda$ into w blocks of the form:

$$\cdots [\underbrace{\mathsf{T}_d \cdots \mathsf{T}_d}_{m_w \text{ copies}} \underbrace{\mathsf{T}_1 \cdots \mathsf{T}_1}_{m'_w \text{ copies}}] \cdots$$

where $n_w = \sum_{i=1}^{w}(m_w + m'_w)$. If the number of positions $z \in [-n, 0]$ mapped to T_d by the strategy s_n^λ approaches infinity as $n \to \infty$, one will be able to find an infinite subsequence (n_w) such that this grouping holds, i.e. the number of such blocks grows unboundedly as the strategy sequence progresses. If, on the other hand, the triple T_1 occurs in all but finitely many positions in the limit as $n \to \infty$, then the subsequence (n_w) is necessarily finite.

Let (n_w) be a maximal subsequence of $\mathbb{N}$ such that

$$s_{n_w}^\lambda = \mathsf{T}_d^{m_w} \mathsf{T}_1^{m'_w} s_{n_{w-1}}^\lambda$$

where $n_w = m_w + m'_w + n_{w-1}$, and $m_w > 0$, $m'_w > 0$ for all w.

If this subsequence is finite, then there is a W such that

$$s_n^\lambda = \mathsf{T}_1^{n-n_W} s_{n_W}^\lambda$$

for all $n > n_W$. In other words, $cl(\mathsf{T}_d \mathsf{T}_1^{n'-n_W} s_{n_W}^\lambda) > \lambda$ for all $n' \geq n_W$.

However, by Proposition 3.22, $\lim_{n' \to \infty} cl(\mathsf{T}_d \mathsf{T}_1^{n'-n_W} s_{n_W}^\lambda) = cl(s_{n_W}^\lambda)$. Thus $cl(s_{n_W}^\lambda) \geq \lambda$. Yet $cl(s_n^\lambda) \leq \lambda$ for all n by construction. So we have $cl(s_{n_W}^\lambda) = \lim_{n \to \infty} c[\mathsf{T}_1^{n-n_W} s_{n_W}^\lambda] = \lambda$, implying that $c[s_n^\lambda] \to \lambda$ as $n \to \infty$ — which is our intended result.

Now, suppose the maximal subsequence (n_w) is infinite. Then, for each w, we have

$$cl(\mathsf{T}_d \mathsf{T}_1^{m'_w - 1} s^\lambda_{n_{w-1}}) > \lambda$$
$$cl(\mathsf{T}_d \mathsf{T}_1^{m'_w} s^\lambda_{n_{w-1}}) \leq \lambda$$

and hence

$$\begin{aligned}
\frac{\lambda - cl(\mathsf{T}_d \mathsf{T}_1^{m'_w} s^\lambda_{n_{w-1}})}{\lambda - cl(s^\lambda_{n_{w-1}})} &= 1 - \frac{cl(\mathsf{T}_d \mathsf{T}_1^{m'_w} s^\lambda_{n_{w-1}}) - cl(s^\lambda_{n_{w-1}})}{\lambda - cl(s^\lambda_{n_{w-1}})} \\
&\leq 1 - \frac{cl(\mathsf{T}_d \mathsf{T}_1^{m'_w} s^\lambda_{n_{w-1}}) - cl(s^\lambda_{n_{w-1}})}{cl(\mathsf{T}_d \mathsf{T}_1^{m'_w - 1} s^\lambda_{n_{w-1}}) - cl(s^\lambda_{n_{w-1}})} \\
&\leq 1 - K < 1 \text{ by Proposition 3.23.}
\end{aligned}$$

Moreover, since

$$\lambda \geq cl(s^\lambda_{n_w}) = cl(\mathsf{T}_d^{m_w} \mathsf{T}_1^{m'_w} s^\lambda_{n_{w-1}}) \geq cl(\mathsf{T}_d \mathsf{T}_1^{m'_w} s^\lambda_{n_{w-1}}) \ ,$$

we have

$$\frac{\lambda - cl(s^\lambda_{n_w})}{\lambda - cl(s^\lambda_{n_{w-1}})} \leq 1 - K < 1 \ .$$

As this bound is independent of w, we can conclude that

$$\lambda - cl(s^\lambda_{n_w}) \leq (1-K)^w (\lambda - cl(s^\lambda_{n_0})) \ .$$

Consequently, $|\lambda - cl(s^\lambda_{n_w})| \to 0$ as $w \to \infty$, implying that $cl(s^\lambda_{n_w}) \to \lambda$. Since $(s^\lambda_{n_w})$ is a subsequence of (s^λ_n) and $(cl(s^\lambda_n))$ is monotone,

$$\lim_{n \to \infty} cl(s^\lambda_n) = \lambda \ .$$

In order to prove that the sequence $(c[s_n^\lambda])$ converges to λ as well, note that

$$\begin{aligned}
c[s_n^\lambda] \leq & \quad c[\mathsf{T}_1^m s_n^\lambda] \leq & c[s_{n+m}^\lambda] \\
\lim_{m\to\infty} c[s_n^\lambda] \leq & \lim_{m\to\infty} c[\mathsf{T}_1^m s_n^\lambda] \leq & \lim_{m\to\infty} c[s_{n+m}^\lambda] \\
c[s_n^\lambda] \leq & \quad cl(s_n^\lambda) \leq & \lim_{m\to\infty} c[s_m^\lambda] \\
\lim_{n\to\infty} c[s_n^\lambda] \leq & \lim_{n\to\infty} \quad cl(s_n^\lambda) \leq & \lim_{n\to\infty} \lim_{m\to\infty} c[s_m^\lambda] \\
\lim_{n\to\infty} c[s_n^\lambda] \leq & \lim_{n\to\infty} \quad cl(s_n^\lambda) \leq & \lim_{m\to\infty} c[s_m^\lambda] \ .
\end{aligned}$$

Hence

$$\lim_{n\to\infty} c[s_n^\lambda] = \lim_{n\to\infty} cl(s_n^\lambda) = \lambda \ .$$

□

As a consequence of the above lemma, we can finally conclude that

Theorem 3.25.

$$\lim_{n\to\infty} \Gamma_n = \Lambda = \left[\lim_{m\to\infty} c[\mathsf{T}_1^m], 1\right] \ .$$

3.6 The Transition Triple $\mathsf{T}(0)$

With the major special cases of the Main Theorem now resolved, there remains the border cases involving the barrier triples $\mathsf{T}(0) = (1,0,0)$, $\mathsf{T}(1) = (0,0,1)$, and $\mathsf{T}_\rho = (0,1,0)$. We begin with the discussion of $\mathsf{T}(0)$ in this section. It turns out that the sequence Γ_n is monotonically increasing when the transition set contains $\mathsf{T}(0)$, which naturally entails the convergence of the sequence.

Let $\mathcal{T}$ be a transition set where $\mathsf{T}(0) \in \mathcal{T}$. We can establish the convergence of Γ_n via the following result.

Lemma 3.26. *For* $\mathsf{T} = \mathsf{T}(0)$*, and any* $s, s' \in \mathbb{S}(\mathcal{T})$*,* $\varphi_{s'\mathsf{T}s} = \langle a[s], 1, c[s], 0\rangle$.

Proof. It is easy to see that the quantity $\varphi_{s'\mathsf{T}} = \langle 1, 1, 0, 0\rangle$, since T is a right barrier. Therefore

$$\begin{aligned}
\varphi_{s'\mathsf{T}s} &= \varphi_{s'\mathsf{T}} \cdot \varphi_s \\
&= \langle a[s](1 - 0 \cdot b[s])^{-1} \cdot 1 \qquad\qquad , 1 + 0 \cdot (1 - b[s] \cdot 0)^{-1} b[s] \cdot 1, \\
&\quad\ c[s] + a[s](1 - 0 \cdot b[s])^{-1} \cdot 0 \cdot d[s], 0 \cdot (1 - b[s] \cdot 0)^{-1} \cdot d[s]\rangle \\
&= \langle a[s], 1, c[s], 0\rangle \ .
\end{aligned}$$

□

In particular, we note that $c[s] = c[\mathsf{T}(0)s]$ for any $s \in \mathbb{S}(\mathcal{T})$. This is the basis of the monotonicity result for Γ_n which implies its convergence.

Theorem 3.27. $\lim_{n\to\infty} \Gamma_n = \bigcup_n \Gamma_n$ *in the Hausdorff metric.*

Proof. First, we note that

$$\begin{aligned}
\Gamma_n = \bigcup_{s\in\mathbb{S}_{[1:n]}(\mathcal{T})} c[s] &= \bigcup_{s\in\mathbb{S}_{[1:n]}(\mathcal{T})} c[\mathsf{T}(0)s] \\
&\subseteq \bigcup_{s'\in\mathbb{S}_{[1:n+1]}(\mathcal{T})} c[s'] = \Gamma_{n+1} \ ,
\end{aligned}$$

and so Γ_n increases monotonically to the limit $\bigcup_n \Gamma_n$. Since $\bigcup_n \Gamma_n$ is contained within the closed interval $[0, 1]$, it follows that the convergence holds also in the Hausdorff metric. □

3.7 The Transition Triple $\mathsf{T}(1)$

Whereas the inclusion of $\mathsf{T}(0)$ in a transition set causes the sequence Γ_n to become monotonic, the inclusion of $\mathsf{T}(1)$ has the effect of making the probability value of 1 a member of every set Γ_n. This corresponds to the intuition that, if $s(z) = \mathsf{T}(1)$ for some $z \in dom(s)$, then there is zero probability that the walk instance $W(s)$ will exit through the left endpoint, and hence $c[s] = 1$.

If $\mathcal{T} = \mathsf{T}(1) \cup \mathcal{T}'$ and $\mathsf{T}(1) \notin \mathcal{T}'$, then we can show the following convergence result:

Theorem 3.28. $\lim_{n\to\infty} \Gamma_n = \{1\} \cup \lim_{n\to\infty} \Gamma'_n$ *in the Hausdorff metric, where* $\Gamma'_n = c[\mathbb{S}_{[1:n]}(\mathcal{T}')]$.

Proof. We shall prove this by induction on n:

Base case: $n = 1$

$$\begin{aligned}\Gamma_1 = c[\mathbb{S}_{[1:1]}(\mathcal{T})] &= c[\mathbb{S}_{[1:1]}(\mathcal{T}')] \cup c[\mathsf{T}(1)] \\ &= c[\mathbb{S}_{[1:1]}(\mathcal{T}')] \cup \{1\} \ = \Gamma'_1 \cup \{1\}.\end{aligned}$$

Inductive step: Assume the hypothesis $\Gamma_n = \{1\} \cup \Gamma'_n$ holds for n, and check the case for $n+1$

$$\begin{aligned}\Gamma_{n+1} &= c[\mathbb{S}_{[1:n+1]}(\mathcal{T})] \\ &= \bigcup_{\mathsf{T}\in\mathcal{T}} \bigcup_{s\in\mathbb{S}_{[1:n]}(\mathcal{T})} \{c[\mathsf{T}s]\} \\ &= \left(\bigcup_{\mathsf{T}\in\mathcal{T}'} \bigcup_{s\in\mathbb{S}_{[1:n]}(\mathcal{T})} \{c[\mathsf{T}s]\} \right) \cup \bigcup_{s\in\mathbb{S}_{[1:n]}(\mathcal{T})} \{c[\mathsf{T}(1)s]\} \\ &= \left(\bigcup_{\mathsf{T}\in\mathcal{T}'} \left(\{1\} \cup \bigcup_{s\in\mathbb{S}_{[1:n]}(\mathcal{T}')} \{c[\mathsf{T}s]\} \right) \right) \cup \{1\} \qquad \text{by the induction hypothesis} \\ &= \Gamma'_{n+1} \cup \{1\}.\end{aligned}$$

Therefore, by induction, $\Gamma_n = \{1\} \cup \Gamma'_n$ holds for all n. Moreover, since Γ'_n converges by previous results on transition sets that do not contain $\mathsf{T}(1)$, $\lim_{n\to\infty} \Gamma_n = \{1\} \cup \lim_{n\to\infty} \Gamma'_n$. □

3.8 Proof of the Main Theorem

Proof. Let us now consider the convergence of the sequence Γ_n in general. In the previous sections, we have shown its convergence in the Hausdorff metric for the different compositions of the transition set $\mathcal{T}$. We will summarize the various scenarios into one consolidated result.

First, we note that for any arbitrary transition triple $\mathsf{T} = (\alpha, \beta, \gamma)$ where $\beta < 1$, there is a triple $\mathsf{T}' = (\alpha', 0, \gamma')$ such that $\varphi_{\mathsf{T}} = \varphi_{\mathsf{T}'}$, namely:

$$\alpha' = (1-\beta)^{-1}\alpha$$
$$\gamma' = (1-\beta)^{-1}\gamma \, .$$

So, for the cases where $(0,1,0) \notin \mathcal{T}$, we can reduce the set $\mathcal{T}$ to a new set $\mathcal{T}'$ where each triple in $\mathcal{T}'$ is of the form $\mathsf{T}(x) = (1-x, 0, x)$. Therefore $\Phi_n(\mathcal{T}) = \Phi_n(\mathcal{T}')$ for all n.

Although not explicitly covered, the case when the transition set contains only one transition triple (i.e. $|\mathcal{T}| = 1$) reduces to an instance of the gambler's ruin problem, for which we know $\Gamma_n = \{c[s_n]\}$ converges as $n \to \infty$.

For the case $|\mathcal{T}| \geq 2$, let us enumerate the transition triples in the set $\mathcal{T}$ by $\mathsf{T}(x_1), \mathsf{T}(x_2), \ldots, \mathsf{T}(x_d)$, where $x_1 < x_2 < \cdots < x_d$. We can partition the problem space by the values of x_1 and x_d, and show that, for each partition, the convergence of Γ_n is proven by the conjunction of one or more previous theorems.

	$x_d = 0$	$0 < x_d < 1/2$	$x_d = 1/2$	$1/2 < x_d < 1$	$x_d = 1$
$x_1 = 0$	Case $\lvert\mathcal{T}\rvert = 1$	Thm 3.27	Thm 3.27	Thm 3.27	Thm 3.27
$0 < x_1 < 1/2$	—	Thm 3.17	Thm 3.17	Thm 3.25	Thms 3.17, 3.25, 3.28
$x_1 = 1/2$	—	—	Thm 3.18	Thm 3.18	Thms 3.18, 3.28
$1/2 < x_1 < 1$	—	—	—	Thm 3.18	Thms 3.18, 3.28
$x_1 = 1$	—	—	—	—	Case $\lvert\mathcal{T}\rvert = 1$

Table 3.1: Partition of the problem space, and the theorems proving the convergence of Γ_n for each case

Consequently, we have now proven that, for all non-empty, finite transition sets $\mathcal{T}$ that do not contain the triple $(0,1,0)$, the sequence of sets of first-exit probabilities Γ_n converges as $n \to \infty$.

As for the sets A_n, note that for the special case where $|E| = 1$, $a[s] = 1 - c[s]$ for all $s \in \mathbb{S}(\mathcal{T})$. Hence $\mathrm{A}_n = \{1 - c \mid c \in \Gamma_n\}$, and the convergence of Γ_n in the

Hausdorff metric implies the convergence of A_n in the same metric.

By symmetry, the arguments for the convergence of A_n and Γ_n can be applied to establish the convergence of Δ_n and B_n, respectively. As $\Phi_n(\mathcal{T}) = \langle \mathrm{A}_n, \mathrm{B}_n, \Gamma_n, \Delta_n \rangle$, we have now shown that $\Phi_n(\mathcal{T})$ converges componentwise in the Hausdorff metric to some $\Phi_\infty(\mathcal{T})$, as $n \to \infty$.

Returning to the case where the transition set $\mathcal{T}$ contains the triple $\mathsf{T}_\rho = (0, 1, 0)$, we note that the property $a[s] + c[s] = b[s] + d[s] = 1$ is no longer universally valid, since there is a non-zero probability that the walk will never reach the endpoints, but rather loop forever at a position z where $s(z) = \mathsf{T}_\rho$. However, since T_ρ is both a left barrier and a right barrier, we have the following identities:

- $\varphi_{\mathsf{T}_\rho} = \langle 0, 0, 0, 0 \rangle$
- For any $s, s' \in \mathbb{S}(\mathcal{T})$ where $\mathsf{T}_\rho \in \mathcal{T}$,

$$\begin{aligned} \varphi_{s'\mathsf{T}_\rho s} &= \langle a[s'\mathsf{T}_\rho s], b[s'\mathsf{T}_\rho s], c[s'\mathsf{T}_\rho s], d[s'\mathsf{T}_\rho s] \rangle \\ &= \langle 0, b[s'], c[s], 0 \rangle \end{aligned}$$

It follows from the above identities that, if $\mathcal{T} = \{\mathsf{T}_\rho\} \cup \mathcal{T}'$ where $\mathsf{T}_\rho \notin \mathcal{T}'$, and if $\Phi_n(\mathcal{T}') = \langle \mathrm{A}_n, \mathrm{B}_n, \Gamma_n, \Delta_n \rangle$, then

$$\Phi_n(\mathcal{T}) = \langle \mathrm{A}_n \cup \{0\}, \bigcup_{i \le n} \mathrm{B}_i, \bigcup_{i \le n} \Gamma_i, \Delta_n \cup \{0\} \rangle \ . \tag{3.19}$$

Therefore, if $\Phi_n(\mathcal{T}') \to \Phi = \langle \mathrm{A}, \mathrm{B}, \Gamma, \Delta \rangle$ as $n \to \infty$, then

$$\Phi_n(\mathcal{T}) \to \langle \mathrm{A} \cup \{0\}, \mathrm{B} \cup \bigcup_{i \in \mathbb{N}} \mathrm{B}_i, \Gamma \cup \bigcup_{i \in \mathbb{N}} \Gamma_i, \Delta \cup \{0\} \rangle \tag{3.20}$$

as $n \to \infty$.

We have now exhaustively proven that $(\Phi_n(\mathcal{T}))_{n \in \mathbb{N}}$ converges for any non-empty $\mathcal{T}$, and thus the Main Theorem is valid in its most general form.

□

CHAPTER 4

Conclusions

We have introduced an extension of random walk known as the random set-walk. Its purpose is to add the notion of nondeterminism to the existing models of random walk in a random environment (RWRE) and random walk with internal states. A random set-walk is nondeterministic in the sense that it executes all possible walk instances in parallel, where a walk instance is a nonhomogeneous random walk governed by a particular strategy. We showed that the space of strategies is itself a semigroup under the operation of concatenation. Furthermore, we defined, for each walk instance of a random set-walk, a 4-tuple of first-exit probabilities known as its characteristic tuple, which characterizes the long-run behaviour of the particular walk instance. The function φ_s mapping strategies to characteristic tuples was then shown to be a homomorphism under concatenation.

The fundamental result proven in this thesis is that, given an internal state space that contains only one distinguished state, and a sequence of sets of strategies of increasing length, the characteristic tuple sets of the corresponding random set-walks converge, on a term-by-term basis, in the Hausdorff metric on sets. We proved this Main Theorem in several parts, for different compositions of the transition set. For the case with purely left-leaning transition triples, we utilized the idea of bounding the sets Γ_n by approximations L_n and I_n, and proceeded to prove that

I_n and L_n converge to the same limit, entailing the convergence of Γ_n itself. With mixed transition sets, we specified a scheme – a greedy algorithm – with which we can systematically identify that each point λ of the candidate limit set Λ can be approached by a sequence of points $(c[s_n^\lambda])$ taken from the sets Γ_n. For the other cases, we employed results from the gambler's ruin problem to derive the desired result. Finally, we consolidated these subproofs and established the validity of the Main Theorem in its most general form.

4.1 Future Work

While the Main Theorem is an important result concerning the long-run behaviour of random set-walks, it is limited in scope, since it relies on the singleton nature of the internal state space. In particular, it does not solve the problem of establishing the regularity of the language class **2NPFA-unary**, as described in the Motivation section. We conjecture that a generalization of the Main Theorem to arbitrary finite internal state spaces is also valid.

Conjecture 4.1. *For any finite internal state space E, and any non-empty transition set $\mathcal{T}$, the sequence of characteristic tuple sets $(\Phi_n(\mathcal{T}))$ converges componentwise in the Hausdorff metric.*

During preliminary investigation, it was noted that the sequence of sets of matrices Γ_n approaches a "limit set" which looks macroscopically like a collection of matrix pencils. Moreover, the mapping $c[\mathbb{S}_{[1:n]}(\mathcal{T})] \mapsto c[\mathbb{S}_{[1:n]}(\mathcal{T}) \cdot s']$ appears to be a combination of rotation, translation, and scaling applied non-uniformly to the points in $\Gamma_n = c[\mathbb{S}_{[1:n]}(\mathcal{T})]$. Hence, it may be fruitful to develop an analogue of the method presented in Section 3.3, bounding each pencil individually by approximating sets, and then proving that the sequences of these approximations converge.

Bibliography

[1] Marc Artzrouni, *A theorem on products of matrices*, Linear Algebra Appl. **49** (1983), 153–159.

[2] ______, *On the convergence of infinite products of matrices*, Linear Algebra Appl. **74** (1986), 11–21.

[3] ______, *On the growth of infinite products of slowly varying primitive matrices*, Linear Algebra Appl. **145** (1991), 33–57.

[4] ______, *On the dynamics of the linear process $Y(k) = A(k)Y(k-1)$ with irreducible matrices $A(k)$*, SIAM J. Matrix Anal. Appl. **17** (1996), no. 4, 822–833.

[5] Michael N. Barber and B. W. Ninham, *Random and restricted walks. Theory and applications*, Gordon and Breach, New York-London-Paris, 1970.

[6] Abraham Berman and Robert J. Plemmons, *Nonnegative matrices in the mathematical sciences*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1994, Revised reprint of the 1979 original.

[7] S. Chatterjee and E. Seneta, *Towards consensus: some convergence theorems on repeated averaging*, J. Appl. Probability **14** (1977), no. 1, 89–97.

[8] Joel E. Cohen, John Hajnal, and Charles M. Newman, *Approaching consensus can be delicate when positions harden*, Stochastic Process. Appl. **22** (1986), no. 2, 315–322.

[9] Harry Cohn (ed.), *Doeblin and modern probability*, Providence, RI, American Mathematical Society, 1993.

[10] Anne Condon, Lisa Hellerstein, Samuel Pottle, and Avi Wigderson, *On the power of finite automata with both nondeterministic and probabilistic states*, SIAM J. Comput. **27** (1998), no. 3, 739–762 (electronic).

[11] J. L. Doob, *Stochastic processes*, John Wiley & Sons Inc., New York, 1953.

[12] Jewgeni H. Dshalalow (ed.), *Advances in queueing. Theory, methods, and open problems*, CRC Press, Boca Raton, FL, 1995.

[13] Ioannis I. Gerontidis, *Periodic strong ergodicity in nonhomogeneous Markov systems*, J. Appl. Probab. **28** (1991), no. 1, 58–73.

[14] M. Girault, *Stochastic processes*, Ökonometrie und Unternehmensforschung, III, Springer-Verlag New York, Inc., New York, 1966.

[15] Roe Goodman, *Introduction to stochastic models*, The Benjamin/Cummings Publishing Co. Inc., Menlo Park, CA, 1988.

[16] G. R. Grimmett and D. R. Stirzaker, *Probability and random processes*, second ed., Oxford Science Publications, 1993.

[17] Allan Gut, *Stopped random walks. Limit theorems and applications*, Springer-Verlag, New York, 1988.

[18] J. Hajnal, *The ergodic properties of non-homogeneous finite Markov chains*, Proc. Cambridge Philos. Soc. **52** (1956), 67–77.

[19] ———, *Weak ergodicity in non-homogeneous Markov chains*, Proc. Cambridge Philos. Soc. **54** (1958), 233–246.

[20] ———, *On products of non-negative matrices*, Math. Proc. Cambridge Philos. Soc. **79** (1976), no. 3, 521–530.

[21] ———, *Shuffling with two matrices*, Doeblin and modern probability (Blaubeuren, 1991) (Providence, RI), Amer. Math. Soc., 1993, pp. 271–287.

[22] Darald J. Hartfiel, *Markov set-chains*, Springer-Verlag, Berlin, 1998.

[23] John E. Hopcroft and Jeffrey D. Ullman, *Introduction to automata theory, languages, and computation*, Addison-Wesley Series in Computer Science, Addison-Wesley Publishing Co., Reading, Mass., 1979.

[24] Dean L. Isaacson and Richard W. Madsen, *Markov chains. Theory and applications*, Robert E. Krieger Publishing Co. Inc., Melbourne, Fla., 1985, Reprint of the 1976 original.

[25] Jānis Kaņeps, *Regularity of one-letter languages acceptable by 2-way finite probabilistic automata*, Fundamentals of computation theory (Gosen, 1991), Springer, Berlin, 1991, pp. 287–296.

[26] John G. Kemeny, J. Laurie Snell, and Anthony W. Knapp, *Denumerable Markov chains*, second ed., Springer-Verlag, New York, 1976, With a chapter on Markov random fields, by David Griffeath, Graduate Texts in Mathematics, No. 40.

[27] András Krámli, Nándor Simányi, and Domokos Szász, *Random walks with internal degrees of freedom. III. Stationary probabilities*, Probab. Theory Relat. Fields **72** (1986), no. 4, 603–617.

[28] András Krámli and Domokos Szász, *Random walks with internal degrees of freedom. I. Local limit theorems*, Z. Wahrsch. Verw. Gebiete **63** (1983), no. 1, 85–95.

[29] ______, *Random walks with internal degrees of freedom. II. First-hitting probabilities*, Z. Wahrsch. Verw. Gebiete **68** (1984), no. 1, 53–64.

[30] F. Thomson Leighton and Ronald L. Rivest, *The Markov chain tree theorem*, MIT Laboratory for Computer Science TM-249, Massachusetts Institute of Technology, Cambridge, Mass., November 1983.

[31] Günter Mayer, *On the convergence of the Neumann series in interval analysis*, Linear Algebra Appl. **65** (1985), 63–70.

[32] Michel Morvan, Christoph Meinel, and Daniel Krob (eds.), *STACS 98*, Berlin, Springer-Verlag, 1998.

[33] V. I. Romanovsky and E. Seneta, *Discrete Markov chains*, Wolters-Noordhoff Publishing, Groningen, The Netherlands, 1970, Translated from the Russian by E. Seneta.

[34] Michael Saks and Anne Condon, *A limit theorem for sets of stochastic matrices*, manuscript, 12 pages.

[35] Arto Salomaa and Matti Soittola, *Automata-theoretic aspects of formal power series*, Texts and Monographs in Computer Science, Springer-Verlag, New York, 1978.

[36] E. Seneta, *Non-negative matrices and Markov chains*, second ed., Springer-Verlag, New York, 1981.

[37] Ya. G. Sinai, *Random walks and some problems concerning lorentz gas*, Proceedings of the Kyoto Conference, 1981, pp. 6–17.

[38] Frank Spitzer, *Principles of random walk*, The University Series in Higher Mathematics, D. Van Nostrand Co., Inc., Princeton, N.J.-Toronto-London, 1964.

[39] George H. Weiss, *Aspects and applications of the random walk*, North-Holland Publishing Co., Amsterdam, 1994.