Games, Walks and Grammars

Problems I've Worked On

ILLC Dissertation Series DS-2000-03

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

For further information about ILLC-publications, please contact

Institute for Logic, Language and Computation
Universiteit van Amsterdam
Plantage Muidergracht 24
1018 TV Amsterdam
phone: +31-20-525 6051
fax: +31-20-525 5206
e-mail: illc@wins.uva.nl
homepage: http://www.illc.uva.nl/

Games, Walks and Grammars

Problems I've Worked On

Academisch Proefschrift

ter verkrijging van de graad van doctor aan de
Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof.dr. J.J.M. Franse
ten overstaan van een door het college voor
promoties ingestelde commissie, in het openbaar
te verdedigen in de Aula der Universiteit
op dinsdag 19 September 2000, te 12.00 uur

door

Marco Robert Vervoort

geboren te Amsterdam.

Promotor: Prof.dr. M. Keane
Co-promotor: Dr. M. van Lambalgen
Faculteit der Natuurwetenschappen, Wiskunde en Informatica
Universiteit van Amsterdam
Plantage Muidergracht 24
1018 TV Amsterdam

This research was supported by NWO-PIONIER project PGS 22-262.

Cover design by Rob Molthoff.
Printed and bound by PrintPartners Ipskamp.

ISBN: 90–5776–055–X

Contents

Acknowledgments

First and foremost, I would like to thank Michiel van Lambalgen, who gave me a position as a Logic PhD student, kept on supporting me and being interested even when I wandered into completely different fields of mathematics, and when I went through a motivational crisis, found people I could work with on my own terms. Without him, this thesis would not have been written.
Tony Martin, Mike Keane and Pieter Adriaans each gave me questions to work on, and to a greater or lesser extent helped shape the answers. If not for them, this thesis might have had only one subject. :-)
In no particular order, I would also like to thank Alex Heinis, Robert Maas, Dick de Jongh, Nikos Massios, Carlos Areces, Wendy de Leeuw, Marco de Vries, Judith Keijsper, Rosalie Iemhoff, Annette Bleeker, Ingrid van Loon, Joost J. Joosten, Renata Wassermann, Gertjan van der Hoeven, and anyone else I might have forgotten, for making work interesting, and for being interested in my work.
Last but not least, I would like to thank my parents, both for giving me solid ground to stand on, and for allowing me to make my own mistakes.

Chapter 1

Introduction

This dissertation is about games, walks and grammars. That's what the title says, and it's completely correct. The second part of the title is 'problems I've worked on', and that too is correct. However, those of you that are expecting deep and interesting connections between these subjects, I have to give warning: you are going to be disappointed. I worked on these problems as separate problems in separate fields, and quite frankly, it seems almost hypocritical to try to establish a connection at this point. If you wish, you may see it as a demonstration that the techniques and the manner of thinking used in different fields of mathematics are not all that different.

The first part, titled 'Blackwell Games', is an extension of the thesis I wrote as a graduate student in 1995. It is about the problem of determinacy of Blackwell games, a class of infinite games of imperfect information, where both players simultaneously select moves from a finite set, infinitely many rounds are played, and payoff is determined by a Borel measurable function f on the set of possible resulting sequences of moves. In the original thesis I gave elementary proofs of determinacy for Blackwell games whose payoff function is an indicator function of a Borel set up to complexity $G_{\delta\sigma}$. D.A. Martin later found a reduction of the problem for general Borel payoff functions to the known result of determinacy of Borel perfect information games[16]. Both results are presented here, reworked to fit into a single framework (yielding some new proofs for auxiliary results). We also consider the Axiom of Blackwell Determinacy, an analogue for Blackwell games of the Axiom of Determinacy, and give some new results regarding the consequences of this axiom.

In the second part, titled 'Random Walks', we consider recurrence in reinforced random walks, where edges in a graph are traversed with probabilities that may be different (reinforced) at second, third etc. traversals. We focus on the case where the probability for any edge only changes once, after its first traversal. Thomas Sellke showed that in the case of the once-reinforced random walk on the infinite ladder, the walk is almost surely recurrent if reinforcement is small[31].

Here, we present some general tools which allow us to obtain the same result as a special case. After considering some other interesting cases, we combine these tools with an application of nonstandard analysis to graph theory, and show that the walk on the infinite ladder is also recurrent if reinforcement is sufficiently *large*. For readers who are not familiar with nonstandard analysis, a brief overview is provided.

The third part, titled 'The EMILE Grammar Inducer', is about the EMILE program, a program that I wrote for Pieter Adriaans this past year, and that I am still working on. The program reads in a text, and without prior knowledge attempts to determine the grammatical structure of the language, outputting the results in various ways. One purpose of the program is to verify whether natural languages satisfy a condition known as *shallowness*: if this is the case, the EMILE program should work well on natural languages. Here, we first look at the basic concepts and algorithms underlying the program. Then we consider the results of this approach, both in theory and in practice. In a separate appendix, explicit pseudo-code for each of the sub-algorithms of EMILE is given.

Part I

Blackwell Games

Overview

In this part of the dissertation we consider Blackwell games. Blackwell games are infinite games of imperfect information, where both players simultaneously make their moves, infinitely many rounds are played, and payoff is determined by a Borel measurable function f on the set of possible resulting sequences of moves. In particular, we consider the problem of determinacy for Blackwell games, and give elementary proofs of determinacy for Blackwell games whose payoff function is an indicator function of a Borel set up to complexity $G_{\delta\sigma}$. For general Borel payoff functions, we give a reduction, found by T. Martin[16], to the determinacy of Borel perfect information games.

In Chapter 2, we informally introduce the concepts behind Blackwell games, for those that are unfamiliar with Blackwell games or game theory in general.

In Chapter 3, we formally define Blackwell games and other concepts, and prove several basic results that are used in the other chapters.

In Chapter 4, we look at proofs of determinacy for Blackwell games with payoff functions of varying complexity.

In the last chapter of this part, Chapter 5, we consider Blackwell games whose payoff function is not Borel measurable, formulate an analogue of the Axiom of Determinacy for these games, and compare some of the consequences of this 'Axiom of Blackwell Determinacy' with those of the original Axiom of Determinacy.

Chapter 2

An Introduction to Game Theory

2.1 Games of Perfect Information

Imagine two players playing a game of blind chess. The only board they have is in their minds, and they make their moves merely by announcing them. Someone who doesn't know the rules would find a game like this difficult to follow. If that someone was of a literal bend, he might describe it like this:

> "There were two players, playing against each other. The first player said something, and I was told it was her move, and that she had made the move by saying it. The other player thought for a while, and then announced his own move. Then the first player made a move again, then the second player, and so forth. The moves always sounded similar, something like 'pawn from ee-four to ee-five'. So I think they couldn't just say anything, but had to select their moves from only a few possible options. And suddenly they stopped, and shook hands, and I was told that the first player had won, apparently because of the moves she and her opponent had played."

If no one gave the poor fellow a copy of the rules of chess, the way a sequence of moves determines which player wins would probably seem quite arbitrary. And our hypothetical observer might be quite impressed that apparently chess-players are able to memorize this long list of what the result is of each possible sequence of moves.

Of course, the game of chess is not really that arbitrary, and those of us who play chess only need to know a few simple rules to figure out which player has won. But we can use this concept of a game to construct a quite general mathematical game $\Gamma_{p.i.}(f)$.

> Let there be given two finite sets X and Y, an integer n, and a function f assigning to each sequence w of length n of pairs $(x_i, y_i) \in X \times Y$,

a *payoff* $f(w) \in \mathbb{R}$. Two players are playing against each other. First, player I makes a move by selecting an element $x_1 \in X$, and announcing his or her selection. Then player II makes a move by selecting an element $y_1 \in Y$ and announcing his or her selection. Then they each in turn make a second move in this fashion, and a third move, and continue making moves until n rounds have been played. This generates a sequence w of length n of pairs $(x_i, y_i) \in X \times Y$. Then they stop, and player II pays player I the amount $f(w)$.

With the right choices for X, Y, n and f, the game $\Gamma_{p.i.}(f)$ can 'emulate' the game of chess. For if we let X and Y be the set of all possible chess moves, and $n = 6350$[1] then a sequence w corresponds to a finished game of chess. We now set $f(w) = 1$, $f(w) = 0$, or $f(w) = \frac{1}{2}$, depending on whether the corresponding game is a win for White, a win for Black, or a draw.[3] And voilá, we have our chess emulator.
But chess is not the only game that can be 'emulated' in this manner. The same can be done with Noughts-and-Crosses, Connect-Four, Go and Checkers. In general, the games $\Gamma_{p.i.}(f)$ can emulate any game G that has the following properties:

- There are two players.
- There is no element of chance
- Moves are essentially made by selecting them and announcing them.
- There is no hidden information: a player knows all the moves made so far when making his or her current move, and there is nothing going on simultaneously either (*Perfect Information*).
- If one player loses (a certain amount) the other player wins (that same amount) (*Zero-Sum*).
- The game can last no more than a certain number of rounds (*Finite Duration*).

[1]There exists a rule in chess (the fifty-move rule) stating that if no piece has been captured and no pawn has been moved for fifty turns, the game is a draw. Since there are only 30 pieces that can be captured, and each of the 16 pawns can make at most 6 moves, it is easy to show that no game can last longer than $(30 + 16 \cdot 6) \cdot 50 + 50 = 6350$ moves[2].

[2]With a little more effort, this upper bound can be improved to 5950 by observing that pawns in the same column cannot pass one another without one of them making a capture. A game of length 5950 can be constructed, so this is also the maximum length of a game. Optionally the arbiter may allow for greater intervals between captures and/or pawn moves in certain endgames which are known to require this; the maximum length of games under this optional rule may be different.

[3]If w does not correspond to a *legal* chess game, we count it as a win for White if the first illegal move is made by Black, and vice versa.

- There is a maximum number of alternatives each player can select from (*Finite Choice-of-Moves*).

These games are called *games of perfect information*, and any results for the games $\Gamma_{p.i.}(f)$ apply to all the games with these properties.

2.2 Strategies and Values

D. Blackwell described the concept of a *strategy* as [4]:

> Imagine that you are to play the white pieces in a single game of chess, and that you discover you are unable to be present for the occasion. There is available a deputy, who will represent you on the occasion, and who will carry out your instructions exactly, but who is absolutely unable to make any decisions of his own volition. Thus, in order to guarantee that your deputy will be able to conduct the white pieces throughout the game, your instructions to him must envisage every possible circumstance in which he may be required to move, and must specify, for each such circumstance, what his choice is to be. Any such complete set of instructions constitutes what we shall call a *strategy*.

Thus, a strategy for a given player in a given game consists of a specification, for each position in which he or she is required to make a move, of the particular move to make in that position. In turn, a position can be specified by the moves made to get to that position. If we apply this to the game $\Gamma_{p.i.}(f)$, a strategy becomes a function from the set of sequences of length $< n$ of pairs $(x_i, y_i) \in X \times Y$ (followed by single elements $x_j \in X$ in case of a strategy for player II), to the set of possible selections X or Y.

Given strategies for each of the players, the outcome of the game is fixed: each move follows from the current position and the strategy of the player whose turn it is to move, and determines the next position. So, the totality of all the decisions to be made can be described by a single decision - the choice of a strategy. This is the *normal form* of a game: the two players independently make a single move, which consists of selecting a strategy, and then payoff is calculated and made.

Of course, there are good strategies and bad strategies. The *value* of a strategy for a given player is the result of that strategy against the best counter-strategy. The *value* of a game for a given player is the best result that that player can guarantee, i.e. the value of that player's best strategy. A game is called *determined* if its value is the same for both players. That value is the result that will occur if both players are playing perfectly.[4]

[4] In more general cases, we allow ϵ-approximation, i.e. a game is *determined* if and only if there exists a value v such that for any $\epsilon > 0$, the two players have strategies guaranteeing them a payoff of at least $v - \epsilon$ or at most $v + \epsilon$, respectively.

V. Allis [1] recently demonstrated that in a game of Connect-Four, the first player can win, i.e. has a strategy that wins against any counter-strategy. And countless persons throughout the ages have independently discovered that in the game of Noughts-and-Crosses, both players can force a draw. These are both examples of determinacy. It can be shown (using induction) that any game $\Gamma(f)$ as defined above is determined, and hence any game with all of the properties mentioned above is determined. In the case of Go, chess, and checkers, this means that either one of the players has a winning strategy, or both players have a drawing strategy.

2.3 Games of Imperfect Information

Now consider the game of Scissors-Paper-Stone. In this game, the two players simultaneously 'throw' one of three symbols: 'stone' (hand balled in a fist), 'paper' (hand flat with the palm down) or 'scissors' (middle and forefinger spread, pointing forwards). If both players throw the same symbols, the result is a draw; otherwise, 'paper' beats 'stone', 'stone' beats 'scissors', and 'scissors' beats 'paper' (the mnemonic being "paper wraps stone, stone blunts scissors, and scissors cut paper"). In this game, the players do not make moves in turn, but simultaneously. In other words, both players make moves, and neither player knows what move the other is making. This is an example of a game of *imperfect information*. In general, games of imperfect information have all the properties that games of perfect information have, except that players make moves at the same time instead of one after the other.

For games of imperfect information, we need to redefine the concept of strategy. If we keep to the existing definition, then the only possible strategies in (for example) the game of Scissors-Paper-Stone, are strategies of the type 'throw *this*'. However, any such strategy, for either player, is a losing strategy: for instance, the strategy 'throw stone' loses against the counter-strategy 'throw paper'. So in terms of the concept of strategy described above, this game is not determined. On the other hand, consider the 'strategy' 'throw scissors 1/3 of the time, throw paper 1/3 of the time, and throw stone the remaining 1/3 of the time'. Against any other strategy, this strategy loses, draws and wins 1/3 of the time each, for an 'average result' of a draw. This strategy does not fit in the concept of strategy given above, but it is clearly worth considering.

Strategies of this new type are called *mixed* strategies, as opposed to the old type of strategies, the *pure* strategies. A mixed strategy for a given player in a given game consists of a specification, for each position in which he or she is required to make a move, of the *probability distribution* to be used to determine what move to make in that position.[5] Given mixed strategies for each of the players,

[5] Standard game theory defines a mixed strategy as a probability distribution on pure strategies, but the above definition can be shown to be equivalent to that one.

the outcome of the game is not determined, but we can calculate the *probability* of each outcome. If we assign values to winning and losing ('the loser pays the winner one dollar'), then we can calculate the average profit/loss one player can expect to make from the other, playing those strategies.
The *value* of a mixed strategy is therefore the *expected average result* against the best counter-strategy. And a game is called determined if, for some value v, one of the players has a strategy with which she can always expect to make (on average) at least v, no matter what the other plays, while the other player has a strategy with which he can always expect to lose (on average) at most v, no matter what the other plays. As before, it can be shown (using induction and a theorem of von Neumann) that all finite two-person zero-sum games with Imperfect Information (i.e. the games with the properties mentioned above, except that players make moves at the same time instead of one after the other) are determined.

2.4 Infinite Games

All the games mentioned so far are of *finite* duration. Let, as before, X and Y be two finite sets, and let f be a function assigning to each countably infinite sequence w of pairs $(x_i, y_i) \in X \times Y$, a payoff $f(w) \in \mathbb{R}$. We first consider games of infinite duration and *perfect* information:

> Two players are playing against each other. Each player, in turn, makes a move by selecting an element $x_1 \in X$ or $y_1 \in Y$, respectively, and announcing his or her selection. Then they each in turn make a second move, and a third move, and continue making moves for a *countably infinite* number of rounds. This generates an infinite sequence w of pairs $(x_i, y_i) \in X \times Y$. 'Then' they stop, and player II pays player I the amount $f(w)$.

The problem with infinite games, of course, is that the outcome is only known after an infinite number of moves, and thus it is impractical to play the game as it is. But our concept of a strategy as a specification of which move to make in each position, is still valid in the case of games of infinite duration. And given strategies for both players we can construct the infinite sequence of moves that will be played (or the probability distribution thereof), and apply the payoff function to obtain our (expectation of the) outcome. Hence we can still play the game in a fashion, by using its *normal form*.
The concepts of values and determinacy carry over as well. But it is no longer provable that all such games are determined. For some payoff functions f, such as bounded Borel-measurable functions f, it has been proven that the infinite game of perfect information $\Gamma_{p.i.}(f)$ is determined. But using the Axiom of Choice, a non-measurable payoff function f can be constructed such that $\Gamma(f)$ is not

determined [6]. The *Axiom of Determinacy*, the axiom that all games $\Gamma(f)$ are determined, is a commonly used alternative to AC [8, 17].
A game of infinite duration and *imperfect* information is similar, except that both players make their n^{th} move at the same time. These games are called *Blackwell* games, named after D. Blackwell, the first person to describe and study these games [2]. For Blackwell games, it was quickly proven that the game $\Gamma(f)$ is determined if f is the indicator function of an open or G_δ set, [2, 3], and after a considerable period determinacy was also shown for the case where f is the indicator function of a $G_{\delta\sigma}$ set[20]. In 1996, D.A. Martin finally proved determinacy of $\Gamma(f)$ for the case that f is Borel[16].

Chapter 3

Definitions, Lemmas and Terminology

In this chapter, we define the basic concepts of Blackwell games, introduce some terminology, and derive some lemmas that we'll be using in the chapters to come.

3.1 Games, Strategies and Values

The definitions in this section are fairly standard, and merely formalize the intuitive concepts from the introduction. The lemmas are all basic properties of game-values.
First we will define the concept of Blackwell games itself, and the basic concepts of plays and positions. Fix two finite, nonempty sets X and Y, and put $Z = X \times Y$, $W = Z^{\mathbb{N}}$,

3.1.1. Definition. Let $f : W \to \mathbb{R}$ be a bounded Borel-measurable[6] Borel-measurable function. The *Blackwell game* $\Gamma(f)$ with *payoff function* f is the two-person zero-sum infinite game of imperfect information played as follows:

> Player I selects an element $x_1 \in X$ (we say that she *makes the move* x_1) and simultaneously player II selects an element $y_1 \in Y$. Then both players are told $z_1 = (x_1, y_1)$. Next, player I selects $x_2 \in X$, and simultaneously player II selects $y_2 \in Y$. Then both players are told $z_2 = (x_2, y_2)$. Next, both players simultaneously selects $x_3 \in X$ and $y_3 \in Y$, etc. In this manner they produce an infinite sequence $w = (z_1, z_2, \ldots)$. Finally, player II pays player I the amount $f(w)$, ending the game.

[6]These conditions on the payoff function ensure that expectations and values of strategies can be easily defined. In the next chapters, if a function is introduced as the payoff function of a Blackwell game, it may be implicitly assumed to be bounded and Borel-measurable unless explicitly stated otherwise.

If the payoff function is the indicator function I_S of a set $S \subseteq W$, then we will often write $\Gamma(S)$ for $\Gamma(I_S)$.

3.1.2. DEFINITION. Let $\Gamma(f)$ be a Blackwell game. A countably infinite sequence w of pairs $(x, y) \in Z$ is called an *(infinite) play* of $\Gamma(f)$. A finite sequence p of length k of pairs $(x, y) \in Z$ is called a *finite play* or *position* of $\Gamma(f)$, of length k. Note that a finite play of length k contains k moves of each of the players, not k moves total.
If an infinite play w starts with a finite play p of length k, we say that w *hits* p *on round* k, or that w *passes through* p *on round* k. Analogously for a finite play p' instead of w. If p' passes through p and $p' \neq p$, we also say that p' *follows* p, or that p *precedes* p'.

NOTATION. We use the following notational conventions:

- w and p are used to denote infinite plays and positions, respectively.
- W denotes the set of infinite plays, as previously defined.
- P denotes the set of positions $Z^{<\omega}$.
- For all $n \in \mathbb{N}$, W_n denotes the set Z^n of finite plays of length n.
- $\text{len}(p)$ denotes the length of a finite play p.
- $w_{|k}$ $[p_{|k}]$ denote the position w $[p]$ passes through on round k.
- $p \subset w$ $[p \subseteq p']$ denote that w $[p']$ passes through p on some round.
- $p \subset p'$ denotes that p precedes p'.
- e denotes the position of length 0, i.e. the empty sequence.
- $p\frown p'$ $[p\frown w]$ denotes the sequence consisting of the finite sequence p followed by the finite sequence p' [the infinite sequence w].
- $[p]$ denotes the set $\{w \in W \mid w \supset p\}$ of all plays hitting the position p.
- $[H]$ denotes the set $\{w \in W \mid \exists p \in H : w \supset p\}$ of all plays hitting any position in a set of positions H.
- We sometimes write $(x_1, y_1, x_2, y_2, \ldots)$ instead of $((x_1, y_1), (x_2, y_2), \ldots)$.

3.1.3. REMARK. We give W the usual topology by letting the basic open sets be the sets of the form $[H]$ for some some set $H \subseteq W_n$ of positions of fixed length n. Then the open subsets of W are exactly those of the form $[H]$ for some set H of positions. The G_δ subsets of W are exactly those of the form $\{w \in W \mid \#\{p \in H \mid w \text{ hits } p\} = \infty\}$ for some set H of positions. Note that under this topology, W is a compact space.

Next we will give the definitions of strategies and values, and some basic properties thereof. Proofs are omitted for reasons of conciseness.

3.1.4. Definition. A *strategy* for player I in a Blackwell game $\Gamma(f)$ is a function σ assigning to each position p a probability distribution on X. More formally, σ is a function $P \to [0,1]^X$ satisfying $\forall p \in P : \sum_{x \in X} \sigma(p)(x) = 1$.
Analogously, a *strategy* for player II is a function τ assigning to each position p a probability distribution on Y.

3.1.5. Definition. Let σ and τ be strategies for players I and II in a Blackwell game $\Gamma(f)$. σ and τ determine a probability measure $\mu_{\sigma,\tau}$ on W, induced by

$$\mu_{\sigma,\tau}[p] = P\{w \mid w \text{ hits } p\} = \prod_{i=1}^{n} \Big(\sigma(p_{|(i-1)})(x_i) \bullet \tau(p_{|(i-1)})(y_i)\Big) \qquad (3.1)$$

for any position $p = (x_1, y_1, \ldots, x_n, y_n) \in P$.
The *expected income* of player I in $\Gamma(f)$, if she plays according to σ and player II plays according to τ, is the expectation of $f(w)$ under this probability measure:

$$E(\sigma \text{ vs } \tau \text{ in } \Gamma(f)) = \int f(w) d\mu_{\sigma,\tau}(w) \qquad (3.2)$$

3.1.6. Definition. Let $\Gamma(f)$ be a Blackwell game. The *value* of a strategy σ for player I in $\Gamma(f)$ is the expected income player I can guarantee if she plays according to σ. Similarly, the *value* of a strategy τ for player II in $\Gamma(f)$ is the amount to which player II can restrict player I's income if he plays according to τ. I.e.

$$\text{val}(\sigma \text{ in } \Gamma(f)) = \inf_{\tau} E(\sigma \text{ vs } \tau \text{ in } \Gamma(f)) \qquad (3.3)$$

$$\text{val}(\tau \text{ in } \Gamma(f)) = \sup_{\sigma} E(\sigma \text{ vs } \tau \text{ in } \Gamma(f)) \qquad (3.4)$$

3.1.7. Definition. Let $\Gamma(f)$ be a Blackwell game. The *lower value* of $\Gamma(f)$ is the smallest upper bound on the income that player I can guarantee. Similarly, the *upper value* of $\Gamma(f)$ is the largest lower bound on the restrictions player II can put on player I's income. I.e.

$$\text{val}^{\downarrow}(\Gamma(f)) = \sup_{\sigma} \text{val}(\sigma \text{ in } \Gamma(f)) = \sup_{\sigma} \inf_{\tau} E(\sigma \text{ vs } \tau \text{ in } \Gamma(f)) \qquad (3.5)$$

$$\text{val}^{\uparrow}(\Gamma(f)) = \inf_{\tau} \text{val}(\tau \text{ in } \Gamma(f)) = \inf_{\tau} \sup_{\sigma} E(\sigma \text{ vs } \tau \text{ in } \Gamma(f)) \qquad (3.6)$$

Clearly, for all games $\Gamma(f)$, $\text{val}^{\downarrow}(\Gamma(f)) \leq \text{val}^{\uparrow}(\Gamma(f))$. If $\text{val}^{\uparrow}(\Gamma(f)) = \text{val}^{\downarrow}(\Gamma(f))$, then $\Gamma(f)$ is called *determined*, and we write $\text{val}(\Gamma(f)) = \text{val}^{\uparrow}(\Gamma(f)) = \text{val}^{\downarrow}(\Gamma(f))$.

3.1.8. DEFINITION. Let $\Gamma(f)$ be a Blackwell game, and let $\epsilon > 0$. A strategy σ for player I in $\Gamma(f)$ is *optimal* if $\mathrm{val}(\sigma \text{ in } \Gamma(f)) = \mathrm{val}^{\downarrow}(\Gamma(f))$. A strategy σ for player I in $\Gamma(f)$ is ϵ-*optimal* if $\mathrm{val}(\sigma \text{ in } \Gamma(f)) > \mathrm{val}^{\downarrow}(\Gamma(f)) - \epsilon$. Similarly, a strategy τ for player II in $\Gamma(f)$ is *optimal* if $\mathrm{val}(\tau \text{ in } \Gamma(f)) = \mathrm{val}^{\uparrow}(\Gamma(f))$, and ϵ-*optimal* if $\mathrm{val}(\tau \text{ in } \Gamma(f)) < \mathrm{val}^{\uparrow}(\Gamma(f)) + \epsilon$.

3.1.9. LEMMA. *Let f, g be two payoff functions such that for all $w \in W$, $f(w) \leq g(w)$. Then* $\mathrm{val}^{\downarrow}(\Gamma(f)) \leq \mathrm{val}^{\downarrow}(\Gamma(g))$ *and* $\mathrm{val}^{\uparrow}(\Gamma(f)) \leq \mathrm{val}^{\uparrow}(\Gamma(g))$.

3.1.10. LEMMA. *Let f be a payoff function, and let $a, c \in \mathbb{R}, a \geq 0$. Then* $\mathrm{val}^{\downarrow}(\Gamma(af + c)) = a\,\mathrm{val}^{\downarrow}(\Gamma(f)) + c$ *and* $\mathrm{val}^{\uparrow}(\Gamma(af + c)) = a\,\mathrm{val}^{\uparrow}(\Gamma(f)) + c$.

3.1.11. LEMMA. *Let f be a payoff function, and let $f_{sw} : (Y \times X)^{\mathbb{N}} \to \mathbb{R}$ be defined by*

$$f_{sw}((y_1, x_1), (y_2, x_2), \ldots) = f((x_1, y_1), (x_2, y_2), \ldots) \tag{3.7}$$

Then

$$\mathrm{val}^{\downarrow}(\Gamma(-f)) = -\mathrm{val}^{\uparrow}(\Gamma_{sw}(f_{sw})) \tag{3.8}$$

$$\mathrm{val}^{\uparrow}(\Gamma(-f)) = -\mathrm{val}^{\downarrow}(\Gamma_{sw}(f_{sw})) \tag{3.9}$$

where $\Gamma_{sw}(f_{sw})$ is the Blackwell game with payoff function f_{sw} in which player I selects moves from Y and player II selects moves from X.

3.1.12. LEMMA. *Let $(f_i)_i$ be a sequence of payoff functions $f_i : W \to [a, b]$ such that $(f_i)_i$ converges pointwise to a function $f : W \to [a, b]$. Then for any two strategies σ, τ,* $\lim_{i\to\infty} E(\sigma \text{ vs } \tau \text{ in } \Gamma(f_i)) = E(\sigma \text{ vs } \tau \text{ in } \Gamma(f))$

3.1.13. LEMMA. *Let $(f_i)_i$ be a sequence of payoff functions $f_i : W \to [a, b]$ such that $(f_i)_i$ converges uniformly to a function $f : W \to [a, b]$. Then* $\mathrm{val}^{\downarrow}(\Gamma(f)) = \lim_{i\to\infty} \mathrm{val}^{\downarrow}(\Gamma(f_i))$ *and* $\mathrm{val}^{\uparrow}(\Gamma(f)) = \lim_{i\to\infty} \mathrm{val}^{\uparrow}(\Gamma(f_i))$.

3.2 Starting and Stopping

On occasion, we would like to 'fix' a finite sequence of initial moves, and to consider the game starting from that position rather than the game starting from the empty position e. Similarly, we sometimes like to consider a game with positions such that if the game ever hits such a position, the outcome is known and the players may stop playing. This section will lay the basics for using starting and stopping positions. The next section will give some tools for using starting and stopping positions to 'combine' games.

3.2.1. DEFINITION. Let $f : W \to I\!R$ be a bounded Borel function, and $p = ((x_1, y_1), (x_2, y_2), \ldots, (x_n, y_n))$ a position. The *subgame* $\Gamma(f, p)$ *starting from position* p is the game played like $\Gamma(f)$, except that the players start at round $n+1$, and the first n moves are supposed to have been $x_1, y_1, x_2, y_2, \ldots, x_n, y_n$. The game $\Gamma(f, p)$ is played exactly the same as the game $\Gamma(g)$, where g is the payoff function defined by $g(w) = f(p^\frown w)$.
As before, strategies σ and τ determine a probability measure $\mu_{\sigma,\tau \text{ in } \Gamma(f,p)}$ on W. This measure is equal to the conditional probability measure obtained from $\mu_{\sigma,\tau}$ given $[p]$, i.e.

$$\mu_{\sigma,\tau \text{ in } \Gamma(f,p)}(S) = \frac{\mu_{\sigma,\tau}(S \cap [p])}{\mu_{\sigma,\tau}[p]} \quad \text{if } \mu_{\sigma,\tau}[p] = 0 \tag{3.10}$$

The expected income of player I, the value of a strategy σ, etc. are defined for the games $\Gamma(f, p)$ in the same manner as for the games $\Gamma(f)$.

3.2.2. DEFINITION. A *stopping position* in a Blackwell game $\Gamma(f)$ is a position p, such that for all plays $w, w' \in [p]$, $f(w) = f(w')$. We will denote this value by $f(p)$. A *stopset* in a Blackwell game $\Gamma(f)$ is a set H of stopping positions, such that no stopping position $p \in H$ precedes another stopping position $p' \in H$.

We will often define a payoff function f using the following format:

$$\begin{aligned} f(p) &= \text{some formula} && \text{for } p \in H \\ f(w) &= \text{some other formula} && \text{for } w \notin [H] \end{aligned}$$

where H is a set of positions such that no position $p \in H$ precedes another position $p' \in H$. In the game $\Gamma(f)$ constructed in this fashion, H is a stopset.

3.2.3. REMARK. If p is a stopping position, any moves made at or after p will not affect the outcome of the game. It is often convenient to assume that both players will stop playing if a stopping position is hit. If $\Gamma(f)$ is a Blackwell game, and H is a stopset, we write $\Gamma_H(f)$ to explicitly denote that players stop playing at the positions in H. In this case, we only require strategies to be defined on non-stopping positions. Similarly, with respect to a subgame $\Gamma(f, p)$, we only require strategies to be defined on positions that are following or equal to p.

Using stopsets, a finite game can be treated as a special type of infinite game.

3.2.4. DEFINITION. Let $\Gamma(f)$ be a Blackwell game. If, for some n, all positions in W_n are stopping positions, then $\Gamma(f)$ is called *finite (of length n)*. If $\Gamma(f)$ is finite, we can stop after playing n rounds, and we will denote this by writing $\Gamma_n(f)$.

3.3 Equivalent Truncated Subgames

In games like chess, Go, or even Risk or Monopoly, a player is usually allowed to give up if he has no hope of winning. He or she doesn't have to play it out in the hope that the other player will make a mistake. Two players can agree beforehand to stop in certain positions, and pay out the value of the game at that position rather than continue playing. Provided their assessment of that value is accurate, this does not change the value of the total game. We will call a game resulting from such an alteration a *truncated subgame.*

3.3.1. DEFINITION. Let f, g be two payoff functions, and H a stopset in $\Gamma(g)$. $\Gamma_H(g)$ is an *equivalent truncated subgame* of $\Gamma(f)$ (*truncated at* H), if for any play $w \notin [H]$, $f(w) = g(w)$, and for any $p \in H$, $g(p) = \mathrm{val}(\Gamma(f,p))$. $\Gamma_H(g)$ is a *truncated subgame, equivalent for player I [for player II]*, if for any play $w \notin H$, $f(w) = g(w)$, and for any $p \in H$, $g(p) = \mathrm{val}^{\downarrow}(\Gamma(f,p))$ $[g(p) = \mathrm{val}^{\uparrow}(\Gamma(f,p))]$. In all three cases, $\Gamma(f)$ is called an *extension* of $\Gamma_H(g)$.

Note that $\Gamma_H(g)$ is an equivalent truncated subgame of $\Gamma(f)$ if and only if it is a truncated subgame equivalent for both player I and player II.

3.3.2. LEMMA. *Let* $\Gamma(f)$ *be a Blackwell game, and let* $\Gamma_H(g)$ *be a truncated subgame of* $\Gamma(f)$*, truncated at a set of positions* H*, equivalent for player I [for player II]. Then* $\mathrm{val}^{\downarrow}(\Gamma(f)) = \mathrm{val}^{\downarrow}(\Gamma_H(g))$ $[\mathrm{val}^{\uparrow}(\Gamma(f)) = \mathrm{val}^{\uparrow}(\Gamma_H(g))]$*. Furthermore, for any* $\epsilon > 0$*, any* ϵ*-optimal strategy for player I [player II] in* $\Gamma_H(g)$ *(if it is undefined on all positions at or after positions in* H*) can be extended to an* ϵ*-optimal strategy for player I [player II] in* $\Gamma(f)$*, and conversely, any* ϵ*-optimal strategy for player I [player II] in* $\Gamma(f)$ *is also an* ϵ*-optimal strategy for that player in* $\Gamma_H(g)$*.*

Proof

Let σ_0 be an ϵ-optimal strategy for player I in $\Gamma_H(g)$. Let $\delta = \mathrm{val}^{\downarrow}(\Gamma_H(g)) - \mathrm{val}(\sigma \text{ in } \Gamma_H(g))$, and for each stopping position $p' \in H$, let $\sigma_{p'}$ be an $(\epsilon - \delta)$-optimal strategy for player I in $\Gamma(f,p')$. We can combine these strategies in a single strategy σ for player I in $\Gamma(f)$, by setting

$$\sigma(p) = \begin{cases} \sigma_{p'}(p) & \text{if for some } p' \in H,\, p' \subseteq p \\ \sigma_0(p) & \text{otherwise} \end{cases} \tag{3.11}$$

It is easy to verify that this is an ϵ-optimal strategy for player I in $\Gamma(f)$. The converse holds trivially.

□

3.3.3. COROLLARY. *Let* $\Gamma(f)$ *be a Blackwell game, and let* $\Gamma_H(g)$ *be an equivalent truncated subgame of* $\Gamma(f)$ *(truncated at* H*). If* $\Gamma_H(g)$ *is determined, then* $\Gamma(f)$ *is determined, and* $\mathrm{val}(\Gamma(f)) = \mathrm{val}(\Gamma_H(g))$*. Furthermore, for any* $\epsilon > 0$*, any* ϵ*-optimal strategy for player I or player II in* $\Gamma_H(g)$ *can be extended to an* ϵ*-optimal strategy for player I or player II in* $\Gamma(f)$*.*

3.3.4. Corollary. *Let $\Gamma(f), \Gamma_H(g)$ be Blackwell games. If for any $p \in H$, $g(p) \leq \mathrm{val}^{\downarrow}(\Gamma(f,p))$, and for any $w \notin [H]$, $g(w) \leq f(w)$, then $\mathrm{val}^{\downarrow}(\Gamma_H(g)) \leq \mathrm{val}^{\downarrow}(\Gamma(f))$. Similarly for the value and the upper value, and for $\geq$ instead of $\leq$.*

Truncated subgames may be nested. If we have a nested series of truncated subgames, then we may extend a strategy for the smallest subgame to a strategy for *all* subgames. This allows us to approximate complicated games with a series of simpler, truncated subgames, obtain a strategy that is (ϵ-)optimal in all the subgames. The final lemma in this section allows us to prove results for that strategy in the original game.

3.3.5. Definition. Let, for $n \in \mathbb{N}$, f_n be a payoff function, and H_n a set of stopping positions in $\Gamma(f_n)$. If for all $n \in \mathbb{N}$, $\Gamma_{H_n}(f_n)$ is a truncated subgame of $\Gamma_{H_{n+1}}(f_{n+1})$, and equivalent to $\Gamma_{H_{n+1}}(f_{n+1})$ [for player I, II], then the series of games $(\Gamma_{H_n}(f_n))_{n\in\mathbb{N}}$ is called a *nested series of equivalent truncated subgames* [equivalent for player I, II].

3.3.6. Lemma. *Let $(\Gamma_{H_i}(f_i))_{i\in\mathbb{N}}$ be a nested series of truncated games equivalent for player I [player II]. Then all the games have the same lower value [upper value]. Furthermore, we can find a strategy for player I [player II] that is ϵ-optimal in all the games $\Gamma_{H_i}(f_i)$.*

Proof

We may assume without loss of generality that $H_0 = \{e\}$ and that for any $i \in \mathbb{N}$, any stopping position in $\Gamma_{H_i}(f_i)$ is either equal to or preceded by a stopping position from the set H_i. Note that for $j > i$, all elements of H_j are stopping positions in $\Gamma_{H_i}(f_i)$.

For $i > 0$ and $p' \in H_{i-1}$, let $\sigma_{i,p'}$ be a $2^{-i}\epsilon$-optimal strategy for player I in $\Gamma_{H_i}(f_i, p')$. We can combine these strategies in a single strategy σ for player I in $\Gamma(f)$, by setting

$$\sigma(p) = \begin{cases} \sigma_{i,p'}(p) & \text{if } p' \in H_i,\ p' \subseteq p, \text{ and } \neg\exists p'' \in H_{i+1} : p'' \subseteq p \\ \text{arbitrary} & \text{if } \forall i \exists p' \in H_i : p' \subseteq p \end{cases} \tag{3.12}$$

It is easy to verify that for all i, this is an $(1-2^{-i})\epsilon$-optimal strategy for player I in the game $\Gamma_{H_i}(f_i)$. Conversely, any ϵ-optimal strategy for player I in a game $\Gamma_{H_i}(f_i)$ is also an ϵ-optimal strategy for player I in the games $\Gamma_{H_j}(f_j)$, for all $j < i$.

□

3.3.7. Lemma. *Let $(\Gamma_{H_i}(f_i))_{i\in\mathbb{N}}$ be a nested series of equivalent truncated subgames equivalent for player I [player II]. Let $f : W \to \mathbb{R}$ be a bounded Borel function such that*

$$\forall w \in W : \liminf_{i\to\infty} f_i(w) \leq f(w) \leq \limsup_{i\to\infty} f_i(w) \tag{3.13}$$

Then the game $\Gamma(f)$ has the same lower value [upper value] as the games $\Gamma_{H_i}(f_i)$.

Proof

The proof of this lemma is based on work of D.A. Martin[16, proof of Lemma 1.1]. We may assume without loss of generality that $H_0 = \{e\}$ and that for any $i \in I\!N$, any stopping position in $\Gamma_{H_i}(f_i)$ is either equal to or preceded by a stopping position from the set H_i. Note that for $j > i$, all elements of H_j are stopping positions in $\Gamma_{H_i}(f_i)$.

Let v be the value of the games $\Gamma_{H_i}(f_i)$, First we will show that the lower value of the game $\Gamma(f)$ is at least v. Let $\epsilon > 0$, and consider the strategy σ defined in Lemma 3.3.6. The value of this strategy in $\Gamma(f)$ is at least $v - 3\epsilon$. For let τ be an arbitrary counter-strategy for player II, and suppose that $E(\sigma \text{ vs } \tau \text{ in } \Gamma(f)) < v - 3\epsilon$. We can find a continuous bounded function $g : W \to I\!R$ such that $f \leq g$ and[7]

$$E(\sigma \text{ vs } \tau \text{ in } \Gamma(g)) = \int g(w) d\mu_{\sigma,\tau}(w) \leq \int f(w) d\mu_{\sigma,\tau}(w) + \epsilon < v - 2\epsilon \qquad (3.14)$$

Define the functions $g_i : W \to I\!R$, for $i \in I\!N$, by setting

$$\begin{aligned} g_i(p) &= E(\sigma \text{ vs } \tau \text{ in } \Gamma(g, p)) && \text{for } p \in H_i \\ g_i(w) &= g(w) && \text{for } w \notin [H_i] \end{aligned} \qquad (3.15)$$

Then we have for any $i \in I\!N$ and any position $p \in H_i$,

$$E(\sigma \text{ vs } \tau \text{ in } \Gamma_{H_{i+1}}(g_{i+1}, p)) = g_i(p) \qquad (3.16)$$

By the construction of σ, we know that for any $i \in I\!N$ and any position $p \in H_i$,

$$E(\sigma \text{ vs } \tau \text{ in } \Gamma_{H_{i+1}}(f_{i+1}, p)) > f_i(p) - 2^{-i-1}\epsilon \qquad (3.17)$$

Using induction, we can now construct a sequence of positions p_i such that for all i, $p_i \in H_i$, $p_i \subseteq p_{i+1}$, and

$$g_i(p_i) < f_i(p_i) - (1 + 2^{-i})\epsilon \qquad (3.18)$$

For $i = 0$, we can take $p_0 = e$, since by our assumption that $H_0 = \{e\}$ we can write

$$g_0(e) = E(\sigma \text{ vs } \tau \text{ in } \Gamma(g)) < v - 2\epsilon = f_0(e) - 2\epsilon \qquad (3.19)$$

For $i + 1$, we use the induction hypothesis and equations (3.16) and (3.17) to derive

$$E(\sigma \text{ vs } \tau \text{ in } \Gamma_{H_{i+1}}(g_{i+1})) < E(\sigma \text{ vs } \tau \text{ in } \Gamma_{H_{i+1}}(f_{i+1})) - (1 + 2^{-i-1})\epsilon \qquad (3.20)$$

which, together with the observation that for all $w \in W - [H_i]$, $g_i(w) = g(w) \geq f(w) = f_i(w)$, implies that there must be a position $p_{i+1} \in H_{i+1}$ extending p_i and satisfying (3.18).

[7] The existence of g is guaranteed by the observation that f is bounded, and that we can contain the discontinuities of f in an open set of arbitrary small $\mu_{\sigma,\tau}$-measure[9, Theorem 17.12].

If for some i, $p_j = p_i$ for all $j > i$, then we are done, for then p_i is a stopping position of f and $E(\sigma \text{ vs } \tau \text{ in } \Gamma(g, p_i)) < f(p) - \epsilon$, contradicting $g \geq f$. So assume that the sequence $(p_i)_{i \in \mathbb{N}}$ is not eventually constant. For all $i \in \mathbb{N}$ we can find $w_i \in [p_i]$ such that

$$g(w_i) \leq E(\sigma \text{ vs } \tau \text{ in } \Gamma(g, p_i)) = g_i(p_i) \tag{3.21}$$

Since $(p_i)_{i \in \mathbb{N}}$ is not eventually constant, the sequence $(w_i)_{i \in \mathbb{N}}$ converges to some $w \in W$. Since g is continuous, we now have

$$g(w) = \lim_{i \to \infty} g(w_i) \leq \liminf_{i \to \infty} g_i(p_i) \leq \liminf_{i \to \infty} f_i(p_i) - \epsilon \leq f(w) - \epsilon \tag{3.22}$$

again contradicting $g \geq f$. It follows that there exists no counter-strategy τ for player II with $E(\sigma \text{ vs } \tau \text{ in } \Gamma(f)) < v - 3\epsilon$, as required.
To show that the lower value of the game $\Gamma(f)$ is not greater than v, let σ be a strategy for player I in $\Gamma(f)$, and suppose that σ is of value $> v + 3\epsilon$. In the manner of the construction of Lemma 3.3.6, we can construct a counter-strategy τ for player II such that for any $i \in \mathbb{N}$ and any position $p \in H_i$,

$$E(\sigma \text{ vs } \tau \text{ in } \Gamma_{H_{i+1}}(f_{i+1}, p)) < f_i(p) + 2^{-i-1}\epsilon \tag{3.23}$$

From here on, we can continue in the manner of the first part of this proof, to again derive a contradiction.

□

3.3.8. Corollary. *Let $(\Gamma_{H_i}(f_i))_{i \in \mathbb{N}}$ be a nested series of equivalent truncated subgames. If $\Gamma_{H_0}(f_0)$ is determined, then all the games $\Gamma_{H_i}(f_i)$ are determined, as well as the game $\Gamma(f)$, where f is as in Lemma 3.3.7. Furthermore, all the games have the same value, and we can find strategies for player I and player II that are ϵ-optimal in all the games $\Gamma_{H_i}(f_i)$ and the game $\Gamma(f)$.*

3.3.9. Remark. If the component games involved all have optimal strategies, then we can extend optimal strategies with optimal strategies to optimal strategies, i.e. drop the ϵ in the above lemmas and corollaries.

Chapter 4

Determinacy Results

In this chapter, we first give some proofs of determinacy for finite Blackwell games, and for Blackwell games whose payoff function is the indicator function of an open or G_δ set. We show that open or G_δ games can be approximated with finite or open games. Then we use this to show that Blackwell games over $G_{\delta\sigma}$ sets are determined as well. Finally, we show that the determinacy of Blackwell games for general Borel payoff functions follows from the determinacy of Borel perfect information games.

4.1 Finite Games

The most basic game of imperfect information is where both players select a single move simultaneously, and then the game is finished and the payoff is determined. The aforementioned game of Scissors-Paper-Stone is an example of such a game. Of course these games can also be seen as one-round Blackwell games. The determinacy of these games is the foundation supporting all the other results in this chapter.

4.1.1. Theorem (Von Neumann's Minimax Theorem[18]). *Let $\Gamma_1(f)$ be a finite one-round Blackwell game (i.e. of length 1). Then $\Gamma_1(f)$ is determined, and both players have optimal strategies.*

Proof
f can be interpreted as a function $X \times Y \to \mathbb{R}$. Without loss of generality we may assume that $X = \{1, \ldots, n\}, Y = \{1, \ldots, m\}$. Strategies σ and τ for players I and II in $\Gamma_1(f)$ can be represented by nonnegative vectors $\vec{x}^\sigma = (x_1^\sigma, \ldots, x_n^\sigma)$ and $\vec{y}^\tau = (y_1^\tau, \ldots, y_m^\tau)$ satisfying $\sum_{i=1}^n x_i^\sigma = \sum_{j=1}^m y_j^\tau = 1$. It is easily seen that

$$E(\sigma \text{ vs } \tau \text{ in } \Gamma_1(f)) \quad = \quad \sum_{i=1}^{n}\sum_{j=1}^{m} x_i^\sigma y_j^\tau f(i,j) \tag{4.1}$$

$$\text{val}(\sigma \text{ in } \Gamma_1(f)) = \min_{1\le j\le m} \sum_{i=1}^{n} x_i^\sigma f(i,j) \tag{4.2}$$

$$\text{val}(\tau \text{ in } \Gamma_1(f)) = \max_{1\le i\le n} \sum_{j=1}^{m} y_j^\tau f(i,j) \tag{4.3}$$

For each strategy σ for player I, let $\vec{z^\sigma} \in I\!R^m$ be the vector with coordinates $z_j^\sigma = \sum_{i=1}^n x_i f(i,j)$. For strategies σ and τ for players I and II, $E(\sigma \text{ vs } \tau \text{ in } \Gamma_1(f)) = \langle \vec{y^\tau}, \vec{z^\sigma} \rangle$ (where $\langle \cdot, \cdot \rangle$ denotes the in-product on $I\!R^m$), and val(σ in $\Gamma(f)$) is equal to the lowest coordinate of $\vec{z^\sigma}$. Let the closed convex set $C \subseteq I\!R^m$ be defined by

$$C := \{\vec{z} \in I\!R^m \mid \vec{z} \le \vec{z^\sigma} \text{ for some strategy } \sigma \text{ for player I}\} \tag{4.4}$$

Let $v = \text{val}^\downarrow(\Gamma_1(f))$, $\epsilon > 0$. Then there exists no strategy σ for player I of value $v+\epsilon$, and it is easily seen that this implies that the vector $\vec{b} = (v+\epsilon, \ldots, v+\epsilon) \notin C$. Let $\vec{c} \in C$ be the vector of C closest to $\vec{b}$ (i.e. such that $|\vec{b} - \vec{c}|$ is minimal). For all $j \le m$, $c_j \le b_j$, because otherwise we would be able to move closer to $\vec{b}$ by reducing c_j to b_j. Since $\vec{b} \ne \vec{c}$, it follows that $b_1+\ldots+b_m > c_1+\ldots+c_m$, and we can set

$$\vec{y} = \frac{1}{(b_1+\ldots+b_m)-(c_1+\ldots+c_m)} \left(\vec{b} - \vec{c}\right) \tag{4.5}$$

Now $y_1 + \ldots + y_m = 1$ and $y_j \ge 0$ for all $j \le m$, so $\vec{y} = \vec{y^\tau}$ for some strategy τ for player II. Furthermore,

$$\langle \vec{y}, \vec{c} \rangle = \langle \vec{y}, \vec{b} \rangle - \langle \vec{y}, \vec{b} - \vec{c} \rangle = (v+\epsilon) - \frac{|\vec{b} - \vec{c}|^2}{(b_1+\ldots+b_m)-(c_1+\ldots+c_m)} < v + \epsilon \tag{4.6}$$

So for any $\vec{z} \in I\!R^m$, if $\langle \vec{y}, \vec{z} \rangle \ge v + \epsilon$, then $\langle \vec{y}, \vec{z} - \vec{c} \rangle > 0$, This means that for $\lambda > 0$ small enough, $\vec{c} + \lambda(\vec{z} - \vec{c})$ is closer to $\vec{b}$ than $\vec{c}$ is, and hence by our choice of $\vec{c}$, $\vec{c} + \lambda(\vec{z} - \vec{c}) \notin C$. From the convexity of C it follows that $\vec{z} \notin C$. It follows that if $\vec{z} \in C$, then $\langle \vec{y}, \vec{z} \rangle < v + \epsilon$. In particular, for any strategy σ for player I,

$$E(\sigma \text{ vs } \tau \text{ in } \Gamma_1(f)) = \langle \vec{y}, \vec{z^\sigma} \rangle < v + \epsilon \tag{4.7}$$

and hence the strategy τ is of value $\le v + \epsilon$. Since $\epsilon > 0$ was taken arbitrarily, it follows that $\text{val}^\uparrow(\Gamma(f)) = \text{val}^\downarrow(\Gamma(f))$, i.e. $\Gamma(f)$ is determined.
The existence of an optimal strategy for player I follows from the observation that the space of possible strategies for player I can be seen as a closed and bounded subset of $I\!R^n$, and the value-function is continuous on this set. The same argument holds for strategies for player II.

□

For finite games in general, we also have determinacy. Furthermore, the value-function is, in a sense, continuous.

4.1.2. THEOREM. *Let $\Gamma_n(f)$ be a finite Blackwell game of length n. Then $\Gamma_n(f)$ is determined, and both players have optimal strategies.*

Proof

We will proof this by induction on the length n of the games.

For $n = 0$, determinacy is trivial. So let $n > 0$, let $\Gamma_n(f)$ be a finite Blackwell game of length n, and suppose that all finite Blackwell games of length $m < n$ have already been shown to be determined, and to have optimal strategies. For each position $p \in W_1$, the game $\Gamma(f, p)$ is finite of length $\leq n - 1$. Hence, by the induction hypothesis, each of those games is determined and has a value $\mathrm{val}(\Gamma(f, p))$. Define the payoff function $g : W \to I\!R$ by

$$g(p) = \mathrm{val}(\Gamma(f, p)) \text{ for } p \in W_1 \tag{4.8}$$

Then by von Neumann's Minimax Theorem, the game $\Gamma_1(g)$ is determined. $\Gamma_1(g)$ is by its definition an equivalent truncated subgame of $\Gamma_n(f)$, so by Corollary 3.3.3, $\Gamma_n(f)$ is determined. Furthermore, by Remark 3.3.9 the optimal strategy produced by von Neumann's Mini-Max Theorem can be extended to an optimal strategy in $\Gamma_n(f)$.

□

4.1.3. LEMMA. *Let $n \in I\!N$. Let $(f_i)_i$ be a sequence of payoff functions $f_i : W_n \to [a, b]$ such that $(f_i)_i$ converges to a payoff function $f : W_n \to [a, b]$. Then* $\mathrm{val}(\Gamma_n(f)) = \lim_{i\to\infty} \mathrm{val}(\Gamma_n(f_i))$.

Proof

Directly from Lemma 3.1.13 and the observation that since W_n is finite, convergence is uniform.

□

4.2 Generalized Open Games

In this section we prove determinacy of a class of 'generalized open games', where payoff for a play is calculated as the supremum of values associated with the positions hit in the play. In addition we derive a result for these and open games comparable to the compactness of W.

4.2.1. THEOREM. *Let $u : P \to I\!R$ be a bounded function, for $n \in I\!N$ let $f_n : W_n \to I\!R$ be the payoff function defined by $f_n(p) = \sup_{j\leq n} u(p_{|j})$ for $p \in W_n$, and let $f : W \to I\!R$ be the payoff function defined by $f(w) = \sup_{j\in I\!N} u(w_{|j})$. Then $\Gamma(f)$ is determined, and*

$$\mathrm{val}(\Gamma(f)) = \lim_{n\to\infty} \mathrm{val}(\Gamma_n(f_n)) \tag{4.9}$$

Proof
For any $p \in P$, and any $n \in I\!N$, the game $\Gamma_n(f_n, p)$ is finite (of length $\leq n$), and hence determined. It is easily seen that $f_0 \leq f_1 \leq f_2 \leq \ldots \leq f \leq 1$. Consequently, for any $p \in P$,

$$\mathrm{val}(\Gamma_0(f_0, p)) \leq \mathrm{val}(\Gamma_1(f_1, p)) \leq \mathrm{val}(\Gamma_2(f_2, p)) \leq \ldots \leq \mathrm{val}^{\downarrow}(\Gamma(f, p)) \leq r \quad (4.10)$$

where r is an upper bound for u. We can approximate $\Gamma(f)$ with a collection of finite auxiliary games $\Gamma_n(g_n)$ by setting, for $n \in I\!N$,

$$g_n(p) := \lim_{i \to \infty} \mathrm{val}(\Gamma_i(f_i, p)) \text{ for } p \in W_n \quad (4.11)$$

The games $\Gamma_n(g_n)$ form a nested series of equivalent truncated subgames. For if we define the functions $h_n^k : W_n \to I\!R$ by setting $h_n^k = \mathrm{val}(\Gamma_k(f_k, p))$ for $k, n \in I\!N$, $p \in W_n$, then we have for all $n \in I\!N$ and $p \in W_n$:

$$\forall k \in I\!N : h_n^k(p) = \mathrm{val}(\Gamma_{n+1}(h_{n+1}^k, p)) \quad (4.12)$$

by Corollary 3.3.3, and then

$$g_n(p) = \lim_{k \to \infty} h_n^k(p) = \mathrm{val}(\Gamma_{n+1}(\lim_{k \to \infty} h_{n+1}^k, p)) = \mathrm{val}(\Gamma_{n+1}(g_{n+1}, p)) \quad (4.13)$$

by Lemma 4.1.3 and the fact that W_{n+1} is finite.
Furthermore, for all $p \in W_n$,

$$g_n(p) \geq \mathrm{val}(\Gamma_n(f_n, p)) = f_n(p) \quad (4.14)$$

and hence for all $w \in W$

$$\lim_{n \to \infty} g_n(w_{|n}) \geq \lim_{n \to \infty} f_n(w_{|n}) = f(w) \quad (4.15)$$

So by Lemma 3.3.7,

$$g_0(e) \geq \mathrm{val}^{\uparrow}(\Gamma(f)) \quad (4.16)$$

Since by definition

$$g_0(e) = \lim_{n \to \infty} \mathrm{val}(\Gamma_n(f_n)) \leq \mathrm{val}^{\downarrow}(\Gamma(f)) \quad (4.17)$$

it follows that the game $\Gamma(f)$ is determined, and has value equal to

$$\mathrm{val}(\Gamma(f)) = \lim_{n \to \infty} \mathrm{val}(\Gamma_n(f_n)) \quad (4.18)$$

□

4.2.2. Corollary. *Let O be an open set. Then $\Gamma(O)$ is determined.*

Proof
There exists a set of positions H such that $O = [H]$. Then for all $w \in W$, $I_O(w) = \sup_{n \in I\!N} I_H(w_{|n})$. Applying Theorem 4.2.1 yields the corollary.

□

4.2.3. COROLLARY. *Let* $O = \bigcup_i O_i$ *be the union of open sets. Then* $\mathrm{val}(\Gamma(O)) = \lim_{n\to\infty} \mathrm{val}(\Gamma(\bigcup_{i\leq n} O_i))$.

Proof
As the union of open sets, O is open, and hence there is a set of positions H such that $O = [H]$, i.e.

$$O = \bigcup \{[p] \mid p \in H\} \tag{4.19}$$

(where $[p]$ denotes the set $\{w \in W \mid p \subset w\}$ as usual). Define the basic open sets $B_m \subseteq O$ by

$$B_m = \bigcup \{[p] \mid p \in H, \mathrm{len}(p) \leq m\} \tag{4.20}$$

then applying Theorem 4.2.1 with $u = I_H$ yields

$$\mathrm{val}(\Gamma(O)) = \lim_{m\to\infty} \mathrm{val}(\Gamma_m(B_m)) \tag{4.21}$$

For each $m \in I\!N$, B_m is a closed set covered by the open sets $(O_i)_{i \in I\!N}$. So by the compactness of W there is for each $m \in I\!N$ a $n_m \in I\!N$ such that $B_m \subseteq \bigcup_{i=1}^{n_m} O_i$. Then for all $n \geq n_m$,

$$\mathrm{val}(\Gamma_m(B_m)) \leq \mathrm{val}(\Gamma(\bigcup_{i=1}^{n} O_i)) \leq \mathrm{val}(\Gamma(O)) \tag{4.22}$$

The corollary follows immediately.

□

4.2.4. COROLLARY. *Let* f *be a continuous function. Then* $\Gamma(f)$ *is determined.*

Proof
As W is compact, and f is continuous, $f[W]$ is compact, and hence bounded. Define $u : P \to I\!R$ by $u(p) := \inf_{w \in [p]} f(w)$. Then u is well-defined and bounded, and by the continuity of f, $f(w) = \sup_{n \in I\!N} u(w_{|n})$ for all $w \in W$. Applying Theorem 4.2.1 yields the corollary.

□

4.2.5. REMARK. If the games $\Gamma_n(g_n)$ are as in the proof of Theorem 4.2.1, then Lemma 3.3.6 yields a strategy τ for player II[8]. And since all the games involved are finite, by Remark 3.3.9 we can even take this strategy τ to be optimal in all

[8]This strategy basically consisting of playing, at every position, the one-round subgame whose payoff (after one move) is the expected payoff of the continuance of the game.

the games $\Gamma_n(g_n)$. This strategy can easily be shown to be optimal in the game $\Gamma(f)$ as well. For let σ be a counter-strategy for player I. Then

$$E(\sigma \text{ vs } \tau \text{ in } \Gamma(f)) = \lim_{n\to\infty} E(\sigma \text{ vs } \tau \text{ in } \Gamma_n(f_n)) = \lim_{n\to\infty} E(\sigma \text{ vs } \tau \text{ in } \Gamma_n(g_n)) = v \tag{4.23}$$

However, for player I there does not always exist an optimal strategy, as the following example shows.

4.2.6. Example. Consider the following Blackwell game. Each round, both players say either 'Stop' or 'Continue'. If both players say 'Continue', then play continues. Otherwise, the game halts: player II wins (payoff 0) if both players said 'Stop', while player I wins (payoff 1) if only one of the players said 'Stop'. If play continues indefinitely, and neither player ever says 'Stop', then payoff is 0, i.e. player II wins.
This is clearly an open game. An interpretation of this game is, that player II tries to guess on which round player I will say 'Stop', and tries to match her. If player II guesses wrong, i.e. says 'Stop' too soon or not soon enough, then player I wins, if player II guesses right, then he wins.
A strategy of value $1-\frac{1}{n}$ for player I is, to select at random a number i between 1 and n, and say 'Stop' on round i. Translated to the standard format for strategies, this becomes:
on round 1, say 'Stop' $\frac{1}{n}$ of the time,
on round 2, if not yet stopped, say 'Stop' $\frac{1}{n-1}$ of the time,
on round 3, if not yet stopped, say 'Stop' $\frac{1}{n-2}$ of the time,
$\vdots$
on round n, if not yet stopped, say 'Stop' $\frac{1}{1}$ of the time.
Hence, the value of this game is 1. In fact, the value of this game at any position in which game has not yet ended is 1. But there exists no optimal strategy of value 1. For suppose there exists such a strategy, of value 1. Then on any round (in which play has not yet ended), the chance that player I will say 'Stop' in that round is 0%. For otherwise, the strategy would not score 100% against the counter-strategy that player II says 'Stop' on that round. But then, player I will never say 'Stop', and this strategy will lose against the counter-strategy that player II never says 'Stop'. So any strategy for player I has value strictly less than 1, although there are strategies with values arbitrarily close to 1. This game is an example of a game in which one of the players has no optimal strategy.

4.3 G_δ-sets

M. Davis gave a proof of determinacy for $G_{\delta\sigma}$ games of perfect information [5] that was based upon the idea of 'imposing restrictions' on the range of moves player II can make. I.e. certain moves were declared 'forbidden', or a loss for player II,

in such a way that (a) player I did not get a win she did not have before, and (b) a particular G_δ subset of player I's winning set would be avoided with certainty. By applying this to all the G_δ subsets of a $G_{\delta\sigma}$ set, and using compactness, Davis showed that if player I could not force the resulting sequence to be in one of the G_δ sets, player II could force the resulting sequence to be outside all of them.
The union of all the sequences in which one of the 'forbidden' moves is played, formed an open set containing the G_δ set under consideration. One way of looking at Davis' proof is, that each of the G_δ sets was enlarged to an open set without increasing the (lower) value of the game, in order to be able to apply determinacy of open games.
In this section, we show that this holds (in a fashion) for Blackwell games, i.e. that a G_δ set can be 'enlarged' to an open set without increasing the lower value of the game by more than an arbitrarily small amount, even in the presence of a 'background function', a payoff function for those sequences that are not in the G_δ set.

4.3.1. THEOREM. *Let $f : W \to [0,1]$ be a measurable function and let D be a G_δ set. Then*

$$\mathrm{val}^{\downarrow}(\Gamma(\max(f, I_D))) = \inf_{O \supseteq D, O \text{ open}} \mathrm{val}^{\downarrow}(\Gamma(\max(f, I_O))) \tag{4.24}$$

Proof
For any G_δ set D we can find a set of positions H, such that $D = \{w \in W \mid \#\{p \in H \mid p \subset w\} = \infty\}$. We may assume that $e \in H$. We partition the set H into sets H_i by setting, for $i \in I\!N$,

$$H_i := \{p \in H \mid \text{there are exactly } i \text{ positions } p' \text{ in } H \text{ strictly preceding } p\} \tag{4.25}$$

Now we can approximate the game $\Gamma(\max(f, I_D))$ by a series of games $\Gamma_{H_i}(g_i)$ by setting, for $i \in I\!N$,

$$\begin{array}{lcll} g_i(p) & = & \inf_{O \supseteq D, O \text{ open}} \mathrm{val}^{\downarrow}(\Gamma(\max(f, I_O), p)) & \text{for } p \in H_i \\ g_i(w) & = & f(w) & \text{for } w \notin [H_i] \end{array} \tag{4.26}$$

The games $\Gamma_{H_i}(g_i)$ form a nested series of truncated subgames equivalent for player I. To see this, fix $i \in I\!N$ and $p \in H_i$. For any open set $O \supseteq D$ and any $p' \in H_{i+1}$, $\mathrm{val}^{\downarrow}(\Gamma(\max(f, I_O), p')) \geq g_{i+1}(p')$. Hence by Corollary 3.3.4, for any open set $O \supseteq D$,

$$\mathrm{val}^{\downarrow}(\Gamma(\max(f, I_O), p)) \geq \mathrm{val}^{\downarrow}(\Gamma_{H_{i+1}}(g_{i+1}, p)) \tag{4.27}$$

On the other hand, for any $\epsilon > 0$, if we select for each $p' \in H_{i+1}$, an open set $O_{p'} \supseteq D$ such that $\mathrm{val}^{\downarrow}(\Gamma(\max(f, I_{O_{p'}}), p')) \leq g_{i+1}(p') + \epsilon$, and set $O = \bigcup_{p' \in H_{i+1}} ([p'] \cap O_{p'})$, then by 3.3.4 this open set $O \supseteq D$ satisfies

$$\mathrm{val}^{\downarrow}(\Gamma(\max(f, I_O), p)) \leq \mathrm{val}^{\downarrow}(\Gamma_{H_{i+1}}(g_{i+1}, p)) + \epsilon \tag{4.28}$$

It follows that

$$g_i(p) = \inf_{O \supseteq D, O \text{ open}} \mathrm{val}^{\downarrow}(\Gamma(\max(f, I_O), p)) = \mathrm{val}^{\downarrow}(\Gamma_{H_{i+1}}(g_{i+1}, p)) \tag{4.29}$$

Hence for each $i \in I\!N$, $G_{H_i}(g_i)$ is a truncated subgame of $G_{H_{i+1}}(g_{i+1})$, equivalent for player I.
Now for all $p \in H_i$,

$$g_i(p) \leq \mathrm{val}^{\downarrow}(\Gamma(\max(f, 1))) = 1 \tag{4.30}$$

Hence for $w \in D$, $g_i(w) \leq 1$ for all $i \in I\!N$. Furthermore, for $w \in W - D$, $g_i(w) = f(w)$ for i large enough. So for all $w \in W$,

$$\liminf_{i \to \infty} g_i(w) \leq \max(f, I_D) \tag{4.31}$$

So by Lemma 3.3.7,

$$g_0(e) \leq \mathrm{val}^{\downarrow}(\Gamma(\max(f, I_D))) \tag{4.32}$$

Since by Lemma 3.1.9 and the definition of g_0,

$$g_0(e) = \inf_{O \supseteq D, O \text{open}} \mathrm{val}^{\downarrow}(\Gamma(\max(f, I_O))) \geq \mathrm{val}^{\downarrow}(\Gamma(\max(f, I_D))) \tag{4.33}$$

it follows that

$$\mathrm{val}^{\downarrow}(\Gamma(\max(f, I_D))) = g_0(e) = \inf_{O \supseteq D, O \text{open}} \mathrm{val}^{\downarrow}(\Gamma(\max(f, I_O))) \tag{4.34}$$

□

4.3.2. COROLLARY. *Let S be a measurable set, and let D be a G_δ set. Suppose that $\Gamma(S \cup D)$ has lower value v. Then for any $\epsilon > 0$, there exists an open set O, $D \subseteq O$, such that $\Gamma(S \cup O)$ has lower value at most $v + \epsilon$.*

Proof
Take $f \equiv I_S$ and apply the non-trivial part of Theorem 4.3.1.

□

4.3.3. THEOREM. *Let D be a G_δ set. Then $\Gamma(D)$ is determined, and*

$$\mathrm{val}(\Gamma(D)) = \inf_{O \supseteq D, O \textit{ open}} \mathrm{val}(\Gamma(O)) \tag{4.35}$$

Proof
For any open set $O \supseteq D$, $\Gamma(O)$ is determined and $\mathrm{val}^{\downarrow}(\Gamma(D)) \leq \mathrm{val}^{\uparrow}(\Gamma(D)) \leq \mathrm{val}(\Gamma(O))$. Applying Theorem 4.3.1 with $f \equiv 0$ yields the required result. Note that the determinacy of G_δ Blackwell games was already proven by Blackwell [2, 3].

□

4.4 $G_{\delta\sigma}$-sets

In this section, we prove the determinacy of $\Gamma(f)$ in the case that f is the indicator function of a $G_{\delta\sigma}$ set. Structurally, this proof is similar to the aforementioned proof by M. Davis for $G_{\delta\sigma}$ games of perfect information [5]. We enlarge each of the G_δ sets composing a $G_{\delta\sigma}$ set to an open set, in such a way that the lower value of the game does not increase by more than an arbitrarily small amount. Instead of the compactness used in Davis' proof, we use Corollary 4.3.2

4.4.1. THEOREM. *Let $S = \bigcup_i D_i$ be a $G_{\delta\sigma}$ set. Then $\Gamma(S)$ is determined.*

Proof

Let $\epsilon > 0$. Using Corollary 4.3.2, we can inductively find open sets $O_i \supseteq D_i$ such that for all $j \in I\!N$,

$$\text{val}^\downarrow(\Gamma(S \cup \bigcup_{i \le j+1} O_i)) \le \text{val}^\downarrow(\Gamma(S \cup \bigcup_{i \le j} O_i)) + \epsilon/2^j \tag{4.36}$$

Then for all $j \in I\!N$,

$$\text{val}^\downarrow(\Gamma(S \cup \bigcup_{i \le j} O_i)) \le \text{val}^\downarrow(\Gamma(S)) + \epsilon \tag{4.37}$$

Now, the sets O_i are open, so the games $\Gamma(\bigcup_{i \in I\!N} O_i)$ and $\Gamma(\bigcup_{i \le j} O_i)$ are determined, and by Corollary 4.2.3,

$$\text{val}(\Gamma(\bigcup_{i \in I\!N} O_i)) = \lim_{j \to \infty} \text{val}(\Gamma(\bigcup_{i \le j} O_i)) \le \lim_{j \to \infty} \text{val}^\downarrow(\Gamma(S \cup \bigcup_{i \le j} O_i)) \le \text{val}^\downarrow(\Gamma(S)) + \epsilon \tag{4.38}$$

On the other hand, $S = \bigcup_{i \in I\!N} D_i \subseteq \bigcup_{i \in I\!N} O_i$. Therefore

$$\text{val}^\uparrow(\Gamma(S)) \le \text{val}(\Gamma(\bigcup_{i \in I\!N} O_i)) \le \text{val}^\downarrow(\Gamma(S)) + \epsilon \tag{4.39}$$

Since ϵ was chosen arbitrary, we conclude

$$\text{val}^\uparrow(\Gamma(S)) = \text{val}^\downarrow(\Gamma(S)) \tag{4.40}$$

□

4.5 Borel sets

In this section, we will consider a proof of Borel determinacy given by D.A. Martin[16]. Unlike the previous proofs, this is not an direct proof. Martin's method is to define auxiliary *perfect information* games $G_v(f)$, and use winning strategies in these games to construct mixed strategies in the game $\Gamma(f)$. This yields a reduction to the problem of the determinacy of perfect information games, which has been solved for Borel games[7, 13, 14].

For the next proofs, we need the following lemma:

4.5.1. LEMMA. *Let the function $u : P \to \mathbb{R}$ be such that for all $p \in P$,*

$$u(p) \leq \mathrm{val}(\Gamma_\Delta(u,p)) \;\; [u(p) \geq \mathrm{val}(\Gamma_\Delta(u,p))] \tag{4.41}$$

where $\Gamma_\Delta(u,p)$ is the game starting at position p, and stopping after one round at some position $p' = p\frown z$ with payoff $u(p')$. If $f : W \to \mathbb{R}$ is a payoff function such that for all $w \in W$,

$$f(w) \geq \liminf_{n\to\infty} u(w_{|n}) \;\; [f(w) \leq \limsup_{n\to\infty} u(w_{|n})] \tag{4.42}$$

then $\mathrm{val}^{\downarrow}(f) \geq u(e)$ $[\mathrm{val}^{\uparrow}(f) \leq u(e)]$.

Proof
Let r be a lower bound for the payoff function f. It is easily seen that without invalidating the conditions of the Lemma, we can modify the function u in such a way that the equality holds in equation 4.41 and that for all $p \in P$, if $u(p) \leq r$, then $u(p) = r$ and $u(p') = r$ for all p' following p. Now for $i \in \mathbb{N}$, let $\Gamma_{H_i}(f_i)$ be the finite game which stops after i rounds with payoff $f_i(p) = u(p)$. Applying Lemma 3.3.7 yields the desired result.

□

4.5.2. DEFINITION. Let $\Gamma(f)$ be a Blackwell game. For each $v \in \mathbb{R}$ we define a perfect information game $G_v(f)$:

> Let $p_0 = e$, $h_0(e) = v$.
> Player I selects[9] a function $h_1 : \{p_0\frown z \mid z \in Z\} \to \mathbb{Q}$ satisfying
>
> $$\mathrm{val}(\Gamma_\Delta(h_1, p_0)) \geq h_0(p_0) \tag{4.43}$$
>
> Then player II selects a position $p_1 \in \{p_0\frown z \mid z \in Z\}$.
>
> Next, player I selects a function $h_2 : \{p_1\frown z \mid z \in Z\} \to \mathbb{Q}$ satisfying
>
> $$\mathrm{val}(\Gamma_\Delta(h_2, p_1)) \geq h_1(p_1) \tag{4.44}$$
>
> Then player II selects a position $p_2 \in \{p_1\frown z \mid z \in Z\}$.
>
> The players continue in this fashion, generating a sequence of moves $(h_0,p_0,h_1,p_1,h_2,p_2,\ldots)$. Finally, player I wins the game $G_v(f)$ if and only if $\limsup_{i\to\infty} h_i(p_i) \leq f(w)$, where w is the unique infinite play of $\Gamma(f)$ which passes through the positions p_i.

[9]Note that since Z is finite, player I selects moves from a countable set, and player II from a finite set.

4.5.3. REMARK. One way to think of the game $G_v(f)$ is to imagine that player I is trying to show that $\mathrm{val}^{\downarrow}(\Gamma(f)) \geq v$ by constructing a 'witness' of the type used in Lemma 4.5.1. Indeed, in this manner we will show, in the next lemma, that if player I has a winning strategy for the game $G_v(f)$, then $\mathrm{val}^{\downarrow}(\Gamma(f)) \geq v$. The converse does not hold: if $\Gamma(f)$ is the game of Example 4.2.6, then $\mathrm{val}(\Gamma(f)) = 1$, but player II has a winning strategy in the corresponding game $G_1(f)$. However, a weaker form of the converse does apply: from a later lemma it follows that if $\mathrm{val}^{\downarrow}(\Gamma(f)) > v$, then player II will not have a winning strategy for the game $G_v(f)$.

4.5.4. LEMMA. *Let $\Gamma(f)$ be a Blackwell game, and let $v \in \mathbb{R}$. If player I has a winning strategy in the game $G_v(f)$, then $\mathrm{val}^{\downarrow}(\Gamma(f)) \geq v$.*

Proof

Let S be a winning strategy for player I in the game $G_v(f)$. Given S, the moves of player II determine the course of the game $G_v(f)$. Set $u(e) = v$, and for any position $p \in P$ of length $i \geq 1$, set

$$u(p) := h_i(p) \tag{4.45}$$

where $(h_0, p_{|0}, h_1, p_{|1}, \ldots, h_i, p)$ is the unique sequence of moves of the game $G_v(f)$, that is consistent with S and ends with the move p by player II. It is easily seen that this is a proper definition, and that for all $p \in P$,

$$u(p) \leq \mathrm{val}(\Gamma_{\Delta}(u, p)) \tag{4.46}$$

Furthermore, for any $w \in W$, there exists a unique play $(h_0, w_{|0}, h_1, w_{|1}, h_2, \ldots)$ of $G_v(f)$ consistent with S, and since S is a winning strategy for player I we have

$$f(w) \geq \liminf_{i\to\infty} h^i(w_{|i}) = \liminf_{i\to\infty} u(w_{|i}) \tag{4.47}$$

So by Lemma 4.5.1, $\mathrm{val}^{\downarrow}(\Gamma(f)) \geq v$.

□

4.5.5. LEMMA. *Let $\Gamma(f)$ be a Blackwell game, and let $v \in \mathbb{R}$. If player II has a winning strategy in the game $G_v(f)$, then $\mathrm{val}^{\uparrow}(\Gamma(f)) \leq v$.*

Proof

Let T be a winning strategy for player II in the game $G_v(f)$, and let $\epsilon > 0$. By induction on i we will define, for any $p \in P$ of length $i \geq 1$, a value $u(p)$ and (possibly) a function h^p and a set S^p. Set $u(e) = v$, $h^e(e) = v$. Now suppose that for some $i > 0$ and some position p of length i, $h^e, u(e), h^{p_{|1}}, u(p_{|1}), \ldots, h^{p_{|i-1}}, u(p_{|i-1})$ are all defined. If the set

$$S^p := \left\{ h \,\middle|\, \begin{array}{l} (h^e, e, h^{p_{|1}}, p_{|1}, \ldots, p_{|i-1}, h, p) \text{ is a sequence of} \\ \text{moves consistent with } T \text{ and ending with } p \end{array} \right\} \tag{4.48}$$

is nonempty, set $u(p) := \inf\{h(p) \mid h \in S(p)\}$ and select $h^p \in S^p$ such that $h^p(p) < u(p) + \epsilon/2^{i+1}$. If S^p is empty, set $u(p) = r$, and set $u(p') = r$ for all p' following p, where $r \in \mathbb{Q}$ is an upper bound of f.
Now for all positions $p \in P$ of length i for which h^p is defined,

$$u(p) + \epsilon/2^{i+1} > h^p(p) \geq \text{val}(\Gamma_\Delta(u, p)) \tag{4.49}$$

For suppose that for some $\delta > 0$

$$h^p(p) < \text{val}(\Gamma_\Delta(u, p)) - \delta \tag{4.50}$$

Let $h : \{p\hat{\ }z \mid z \in Z\} \to \mathbb{R}$ be a function such that for all $z \in Z$,

$$u(p\hat{\ }z) - \delta \leq h'(p\hat{\ }z) < u(p\hat{\ }z) \tag{4.51}$$

The sequence of moves $(h^e, e, h^{p_{|1}}, p_{|1}, \ldots, p_{|i-1}, h^p, p)$ is consistent with T, and h is a valid move for player I in this situation, so the strategy T prescribes some response $p\hat{\ }z$, For this response, $h \in S^{p\hat{\ }z}$, contradicting the fact that $u(p\hat{\ }z) > h(p\hat{\ }z)$.
For those positions $p \in P$ of length i for which h^p is not defined, we also have

$$u(p) + \epsilon/2^{i+1} = r + \epsilon/2^{i+1} > r = \text{val}(\Gamma_\Delta(u, p)) \tag{4.52}$$

So if we set

$$u'(p) = u(p) + \epsilon/2^{\text{len}(p)} \tag{4.53}$$

then u' satisfies

$$u'(p) > \text{val}(\Gamma_\Delta(u', p)) \tag{4.54}$$

Furthermore, for any play $w \in W$ of $\Gamma(f)$, either $h^{p_{|i}}$ is not defined for some $i \in \mathbb{N}$, and then

$$\limsup_{i\to\infty} u'(w_{|i}) = r \geq f(w) \tag{4.55}$$

or $(h^e, e, h^{p_{|1}}, p_{|1}, h^{p_{|2}}, p_{|2}, \ldots)$ is a play of $G_v(f)$ consistent with T, and since T is a winning strategy for player II,

$$\limsup_{i\to\infty} u'(w_{|i}) \geq \liminf_{i\to\infty} u'(w_{|i}) \geq \liminf_{i\to\infty} h^{w_{|i}}(w_{|i}) \geq f(w) \tag{4.56}$$

So by Lemma 4.5.1, $\text{val}^{\uparrow}(\Gamma(f)) \leq u'(e) = v + \epsilon$.

□

Now for the main theorem

4.5.6. THEOREM. *Let f be a bounded Borel function. Then the Blackwell game $\Gamma(f)$ is determined.*

Proof

If f is Borel, then the games G_v are all Borel games of perfect information, with the players selecting their moves from countable sets. It is well-known that these games are determined[7, 13, 14]. Let v be the least upper bound of all the v such that player I has a winning strategy in the game G_v. Then by Lemma 4.5.4, $\mathrm{val}^{\downarrow}(\Gamma(f)) \geq v$, and by Lemma 4.5.5, $\mathrm{val}^{\uparrow}(\Gamma(f)) \leq v$. Hence $\Gamma(f)$ is determined, and $\mathrm{val}(\Gamma(f)) = v$.

□

The proof of Theorem 4.4.1 shows that any $G_{\delta\sigma}$ set (and a fortiori any set of lesser complexity) can be enlarged to an open set such that the value of the Blackwell game on that set is not increased by more than an arbitrarily small amount. For games of Perfect Information, this holds for all Borel sets, and given an optimal strategy T for player II, such an open set can be constructed as the set of plays that cannot occur if player II uses T. We can derive a similar result for Blackwell games.

4.5.7. Theorem. *Let S be a Borel set. Then*

$$\mathrm{val}(\Gamma(S)) = \inf\{\mathrm{val}(\Gamma(O)) \mid O \subseteq W \textit{ open}, O \supseteq S\} \tag{4.57}$$

Proof

Let $\epsilon > 0$, let $f : W \to \{0,1\}$ be the indicator function of S, and let $v = \mathrm{val}(\Gamma(S))$. Let $u : P \to I\!R$ be the function constructed in the proof of Lemma 4.5.5 as a witness that $\mathrm{val}^{\uparrow}(\Gamma(S)) \leq v+\epsilon$. Note that u is constructed in such a manner that $u(p) \geq \mathrm{val}^{\uparrow}(\Gamma(S,p))$ for all $p \in P$, and hence $u \geq 0$. Define the set of stopping positions H by setting

$$H := \{p \in P \mid u(p) > 1-\epsilon \wedge \neg\exists p' \in P : (p' \subset p \wedge u(p') > 1-\epsilon)\} \tag{4.58}$$

Now $[H] \subseteq W$ is an open set. If $w \in W - [H]$, then for any $i \in I\!N$, $u(w_{|i}) < 1-\epsilon$, and then by the construction of u, $f(w) < 1-\epsilon$ and $w \notin S$. It follows that $[H] \supseteq S$. Furthermore, if we define the payoff function $g : W \to I\!R$ by

$$\begin{aligned} g(p) &= u(p) && \text{for } p \in H \\ g(w) &= \liminf_{i\to\infty} u(w_{|i}) && \text{for } w \notin [H] \end{aligned} \tag{4.59}$$

then $\mathrm{val}(\Gamma(g)) = u(e) \leq v + \epsilon$, and $\chi_{[H]} \leq g + \epsilon$ (where $\chi_{[H]}$ is the indicator function of $[H]$), and hence

$$\mathrm{val}(\Gamma(H)) \leq \mathrm{val}(\Gamma(g)) + \epsilon \leq v + 2\epsilon \tag{4.60}$$

Note that this result was also obtained by D.A. Martin[16], using a slightly different proof. Furthermore, Maitra, Purves and Sudderth had already shown that this follows from Borel Blackwell determinacy[12].

□

Chapter 5

Non-Borel Blackwell Games

Up to this point, we have implicitly assumed that any payoff function f is Borel measurable. Indeed, properly speaking Blackwell games are (as yet) only defined for Borel measurable payoff functions. In this chapter we will look at non-Borel Blackwell games.

5.1 Definitions

If f is a bounded non-Borel function, then we can play the corresponding Blackwell game in exactly the same manner as before. However, when we try to define the value of a strategy, we run into a problem: f is not necessarily integrable under arbitrary measures $\mu_{\sigma,\tau}$. So instead of the value of the integral itself, we have to use approximations from above and below.

5.1.1. Definition. Let f be a bounded but not necessarily Borel measurable function. The *Blackwell game* $\Gamma(f)$ with *payoff function* f is played exactly as in Definition 3.1.1. For any strategies σ and τ for players I and II, we define the probability measure $\mu_{\sigma,\tau}$ as in Definition 3.1.5. However, instead of the expected income of player I (if she plays according to σ and player II plays according to τ) we now define the *lower* and *upper expected income* of player I:

$$E^-(\sigma \text{ vs } \tau \text{ in } \Gamma(f)) = \sup_{g\le f, g \text{ Borel measurable}} \int g(w)\,d\mu_{\sigma,\tau}(w) \tag{5.1}$$

$$E^+(\sigma \text{ vs } \tau \text{ in } \Gamma(f)) = \inf_{g\ge f, g \text{ Borel measurable}} \int g(w)\,d\mu_{\sigma,\tau}(w) \tag{5.2}$$

The definition of the value of a strategy σ or τ in $\Gamma(f)$ is slightly different as well:

$$\text{val}(\sigma \text{ in } \Gamma(f)) = \inf_{\tau} E^-(\sigma \text{ vs } \tau \text{ in } \Gamma(f)) \tag{5.3}$$

$$\text{val}(\tau \text{ in } \Gamma(f)) = \sup_{\sigma} E^+(\sigma \text{ vs } \tau \text{ in } \Gamma(f)) \tag{5.4}$$

The definitions of lower and upper values and of determinacy are the same as in Definition 3.1.7.

Note that when f is measurable, these definitions are all equivalent to the old definitions.
It is easy to see that Lemmas 4.5.4 and 4.5.5 also hold for non-Borel Blackwell games. This means that determinacy of a particular class of perfect information games implies determinacy of the corresponding class of Blackwell games. The determinacy of many classes of perfect information games can be deduced from so-called large cardinal axioms, and hence we have corresponding results for Blackwell games[15]. Martin gives the example that, for $n \geq 0$, determinacy of Σ^1_{n+1} Blackwell games follows from determinacy of Σ^1_{n+1} perfect information games, which in turn follows from the existence of n Woodin cardinals with a measurable cardinal above them.

5.2 The Axiom of Blackwell Determinacy

A well-known axiom in set theory is the Axiom of Determinacy[17].

5.2.1. Definition. The *Axiom of Determinacy* (AD) is the assertion that

> all game of Perfect Information and finite or countable choice of moves are determined.

Using the Axiom of Choice (AC), a non-measurable payoff function f can be constructed such that $\Gamma(f)$ is not determined [6], hence AD contradicts the Axiom of Choice. AD has many other interesting consequences, such as the existence of an ultrafilter on $\aleph_1$ and of a complete measure on $I\!R$. It is commonly used in large cardinal theory as an alternative to AC [8, 17]. We can formulate an analogue of this axiom for Blackwell games.

5.2.2. Definition. The *Axiom of Determinacy for Blackwell Games* (AD-Bl) is the assertion that

> all Blackwell games are determined

In this case as well, we can use Lemmas 4.5.4 and 4.5.5 to obtain

5.2.3. Theorem. *Working in ZF without the Axiom of Choice, AD implies AD-Bl.*

It is unknown whether the converse also holds, i.e. whether AD-Bl also implies AD. For a given game of Perfect Information, we can easily construct a Blackwell game that is 'equivalent', and assuming AD-Bl we can find an ϵ-optimal *mixed* strategy for that equivalent Blackwell-game. But to derive AD from AD-Bl, we

need to have a *pure* strategy, and even though we can interpret any mixed strategy as a probability distribution on pure strategies, there is no guarantee that any of these pure strategies by itself will do as well as the mixed strategy.
However, a number of consequences of AD *can* be derived from AD-Bl. In fact, one of them is almost trivial:

5.2.4. THEOREM. *Assuming AD-Bl, it follows that all sets of reals are Lebesgue measurable.*

Proof
It suffices to show that the Lebesgue measure on $[0,1]$ is complete. Set $X = Y = \{0,1\}$, and define $\phi : W \to [0,1]$ by

$$\phi((x_1, y_1, x_2, y_2, \ldots)) = \sum_{i=1}^{\infty} 2^{-i}(x_i \oplus y_i) \tag{5.5}$$

where $0 \oplus 0 = 1 \oplus 1 = 0$ and $0 \oplus 1 = 1 \oplus 0 = 1$. Now let σ, τ be strategies, and suppose that one of those strategies is the strategy that assigns the $\frac{1}{2}$-$\frac{1}{2}$ probability distribution on X or Y, respectively. Then for any $i \in I\!N$, $x_i \oplus y_i$ has equal chances of being 0 or 1, and the distribution of $\phi(w)$ on $[0,1]$ is the uniform distribution on $[0,1]$ under the Lebesgue measure on $[0,1]$. It follows that for any subset $S \subset [0,1]$,

$$\mu_{\sigma,\tau}^{\text{inner}}(\phi^{-1}[S]) = \mu_{\lambda}^{\text{inner}}(S) \tag{5.6}$$
$$\mu_{\sigma,\tau}^{\text{outer}}(\phi^{-1}[S]) = \mu_{\lambda}^{\text{outer}}(S) \tag{5.7}$$

where $\mu^{\text{inner}}(A) = \sup_{B \subseteq A \text{ measurable}} \mu(B)$, $\mu^{\text{outer}}(A) = \inf_{B \supseteq A \text{ measurable}} \mu(B)$.
Let $S \subset [0,1]$. No strategy for player I in the game $\Gamma(\phi^{-1}[S])$ can have value greater than $\mu_{\lambda}^{\text{inner}}(S)$, since this is the lower expected income of any strategy for player I against the $\frac{1}{2}$-$\frac{1}{2}$ strategy. Similarly, no strategy for player II in the game $\Gamma(\phi^{-1}[S])$ can have value less than $\mu_{\lambda}^{\text{outer}}(S)$. Therefore

$$\text{val}^{\downarrow}(\Gamma(\phi^{-1}[S])) \le \mu_{\lambda}^{\text{inner}}(S) \le \mu_{\lambda}^{\text{outer}}(S) \le \text{val}^{\uparrow}(\Gamma(\phi^{-1}[S])). \tag{5.8}$$

From the determinacy of the game $\Gamma(\phi^{-1}[S])$, it now follows that

$$\mu_{\lambda}^{\text{inner}}(S) = \mu_{\lambda}^{\text{outer}}(S) \tag{5.9}$$

Since this holds for arbitrary sets $S \subset [0,1]$, all subsets of $[0,1]$ are Lebesgue measurable.

□

5.2.5. COROLLARY. *AD-Bl is not consistent with AC, and the consistency of ZF + AD-Bl cannot be proven in ZFC.*

5.3 Mixed Strategies in Games of Perfect Information

There are a number of other consequences of AD that are also consequences of AD-Bl [11]. In some cases, it is possible to adapt the proofs for AD to use mixed strategies instead of pure strategies (an example will be given in the next section). In general, a requirement for such an adaption is that we may assume the mixed strategy is of value 0 or 1, rather than somewhere in between. The theorems in this section show that for games of perfect information, this is a valid assumption.

5.3.1. REMARK. The games of perfect information that fall under AD, and are used in this section and the next, are usually defined on a *game tree*. Aside from the fact that only one player moves at a time, games defined on game trees differ from the games as defined in definition 3.1.1, in that the sets from which the players select their moves are dependent on the current position, and may even be countably infinite. However, all such games are equivalent to games with binary choice-of-moves. For instance, a selection of a natural number can be emulated by letting the player repeatedly choose between 0 and 1 until the player chooses 1, and interpreting the number of times the player selected 0 as the natural number selected. In turn, a binary perfect information game can be simulated by a Blackwell game with $X = Y = \{0, 1\}$, simply by constructing a payoff function which only depends on the moves made in the Blackwell game by the player whose turn it was to move in the perfect information game (at that position).
It follows that all the definitions and statements relating to Blackwell games (including AD-Bl) also apply to mixed strategies in countable perfect information games. Therefore we will simply treat these games as Blackwell games, ignoring the formal differences.

We first need an auxiliary lemma:

5.3.2. LEMMA. *Let $\Gamma_n(f)$ be a finite perfect information game, and let σ, τ be mixed strategies for players I and II of values v and v'. Then*

$$\forall x > v' : \mu_{\sigma,\tau}\{w \in W_n \mid f(w) \geq x\} \leq \frac{v' - v}{x - v} \tag{5.10}$$

and

$$\forall x < v : \mu_{\sigma,\tau}\{w \in W_n \mid f(w) \leq x\} \leq \frac{v' - v}{v' - x} \tag{5.11}$$

Proof
We will prove this by induction on the length n of the game $\Gamma_n(f)$. For $n = 0$, it is trivial. For $n + 1$, we may assume without loss of generality that in the first round, player I is to move. Let $p_1, p_2, \ldots$ denote the positions of Γ_n that can be

reached from the starting position e in a single move. For $i \in I\!N$, let z_i denote the probability (according to player I's strategy σ) that player I will move to p_i, and let v_i and v_i' denote the values of the strategies σ and τ at that position. Then for all $i \in I\!N$, $v_i \leq v_i' \leq v'$, and $v = \sum_{i \in I\!N} z_i v_i$.
From the induction hypothesis it follows that for $x > v'$ and $i \in I\!N$,

$$\mu_{\sigma,\tau}(\{w \in [p_i] \mid f(w) \geq x\}) \leq z_i \frac{v_i' - v_i}{x - v_i} \leq z_i \frac{v' - v_i}{x - v_i} = z_i - z_i \frac{x - v'}{x - v_i} \tag{5.12}$$

and hence for $x > v'$

$$\mu_{\sigma,\tau}(\{w \in W_n \mid f(w) \geq x\}) = \sum_{i \in I\!N} \mu_{\sigma,\tau}(\{w \in [p_i] \mid f(w) \geq x\}) \tag{5.13}$$

$$\leq \sum_{i \in I\!N} \left(z_i - z_i \frac{x - v'}{x - v_i} \right) \leq 1 - \frac{x - v'}{x - v} = \frac{v' - v}{x - v} \tag{5.14}$$

Similarly it follows from the induction hypothesis that for $x < v$ and $i \in I\!N$,

$$\mu_{\sigma,\tau}(\{w \in [p_i] \mid f(w) \leq x\}) \leq z_i \cdot \begin{cases} \frac{v_i' - v_i}{v_i' - x} & \text{if } x < v_i \\ 1 & \text{if } v_i \leq x \end{cases} \leq z_i \frac{v' - v_i}{v' - x} \tag{5.15}$$

and hence for $x < v$

$$\mu_{\sigma,\tau}(\{w \in W_n \mid f(w) \leq x\}) = \sum_{i \in I\!N} \mu_{\sigma,\tau}(\{w \in [p_i] \mid f(w) \leq x\}) \tag{5.16}$$

$$\leq \sum_{i \in I\!N} z_i \frac{v' - v_i}{v' - x} = \frac{v' - v}{v' - x} \tag{5.17}$$

□

5.3.3. Theorem (0-1 Law for Mixed Strategies). *Let $\Gamma(S)$ be an countably infinite perfect information game, whose payoff function is the characteristic function of a set S. If $\Gamma(S)$ is determined (in the sense of mixed strategies), then its value is either 0 or 1.*

Proof

The following proof is based on communications with D.A. Martin. Let v be the value of $\Gamma(S)$, and suppose that $0 < v < 1$. Then we can select $\epsilon > 0$ such that $2\epsilon < v$, $2\epsilon < 1 - v$ and $v - 2\epsilon > 2\epsilon/(1 - v - 2\epsilon)$. Let σ and τ be ϵ-optimal mixed strategies for players I and II in $\Gamma(S)$, and assume that players I and II play according to σ and τ. Let H be the set of positions

$$H := \{p \in P \mid E(\sigma \text{ vs } \tau \text{ in } \Gamma(S, p)) \geq 1 - \epsilon\} \tag{5.18}$$

Now, we can find open sets $O \supseteq S$ with $\mu_{\sigma,\tau}(O)$ arbitrarily close to $\mu_{\sigma,\tau}(S)$, and for any such O, $(1-\epsilon)\mu_{\sigma,\tau}(O) + \epsilon\mu_{\sigma,\tau}([H]) \geq \mu_{\sigma,\tau}(S)$. It follows that

$$\mu_{\sigma,\tau}([H]) \geq \mu_{\sigma,\tau}(S) \geq \text{val}(\sigma \text{ in } \Gamma(S)) \geq v - \epsilon \tag{5.19}$$

i.e. with probability at least $v - \epsilon$, the game will eventually visit a position in H. Let $n \in \mathbb{N}$ be such that the probability of visiting H before time n is at least $v - 2\epsilon$, and let $\Gamma_{H'}(f)$ be the truncated subgame which stops at time n or when H is visited (whichever happens first), with payoff $f(p) = E(\sigma \text{ vs } \tau \text{ in } \Gamma(S, p))$. Now, since for all $p \in H'$, $f(p) \geq \text{val}(\sigma \text{ in } \Gamma(S, p))$, we have

$$\text{val}(\tau \text{ in } \Gamma(f)) \geq \text{val}(\sigma \text{ in } \Gamma(f)) \geq \text{val}(\sigma \text{ in } \Gamma(S)) > v - \epsilon \tag{5.20}$$

and similarly,

$$\text{val}(\sigma \text{ in } \Gamma(f)) \leq \text{val}(\tau \text{ in } \Gamma(f)) \leq \text{val}(\tau \text{ in } \Gamma(S)) < v + \epsilon \tag{5.21}$$

so by applying the previous Lemma we obtain,

$$\mu_{\sigma,\tau}\{w \in W_n \mid f(w) \geq 1-\epsilon\} \leq \frac{2\epsilon}{1 - v - 2\epsilon} < v - 2\epsilon \tag{5.22}$$

contradicting our choice of n.

□

5.3.4. THEOREM. *Assuming AD-Bl, in every countably infinite perfect information game $\Gamma(S)$ whose payoff function is the characteristic function of a set S, either player I has a strategy of value 1 or player II has a strategy of value 0*[10].

Proof

By AD-Bl and the 0-1 Law for Mixed Strategies, for any position $p \in P$ the value of the game $\Gamma(S, p)$ is either 0 or 1. Without loss of generality we may assume that $\Gamma(S)$ itself is of value 1. If we eliminate positions of value 0 (or equivalently, constrain player I to avoid those positions), the value of the game will not change, so we may also assume without loss of generality that for any position $p \in P$, $\text{val}(\Gamma(S, p)) = 1$.

Now, for any $p \in P$, let σ_p be a strategy of value $> 2/3$ in the game $\Gamma(S, p)$[11]. Set $H_0 = \{e\}$. We can inductively define nested stopping sets H_i such that, for all $i \geq 0$, H_{i+1} consists of all the first positions p' following a position p in H_i,

[10] To avoid confusion, remember: for player II, lower values are better. A strategy for player II of value 0 is a strategy such that the expected payoff is at most 0, i.e. such that player II wins almost surely.

[11] This does not require the Axiom of Choice. Consider an auxiliary game where player II first selects a position $p \in P$, and the players then play $\Gamma(S, p)$. Obviously player II cannot have a strategy of value < 1, so by AD-Bl, player I has a strategy of value $> 2/3$. This strategy contains all the strategies h_p.

such that $\mathrm{val}(\sigma_p \text{ in } \Gamma(S, p')) < 1/3$. Let σ be the strategy where player I starts out by playing according to σ_e, and whenever a position $p \in H_0 \cup H_1 \cup \ldots$ is hit, player I switches to playing according to σ_p. Then for any $i \geq 0$, and any $p \in P$, $\mathrm{val}(\sigma \text{ in } \Gamma_{H_i}(S \cup [H_i], p)) \geq 1/3$.

If player I uses this strategy, and player II uses some strategy τ, then for any $i \geq 0$, the probability (conditional on hitting H_i) of hitting H_{i+1} is at most $1/2$. Hence $\mu_{\sigma,\tau}(\bigcap_{i \geq 0} H_i) = 0$. It follows that for all $p \in P$,

$$\begin{aligned} E(\sigma \text{ vs } \tau \text{ in } \Gamma(S, p)) &= E(\sigma \text{ vs } \tau \text{ in } \Gamma(S \cup \bigcap_{i \geq 0} [H_i], p)) && (5.23) \\ &= \lim_{i \to \infty} E(\sigma \text{ vs } \tau \text{ in } \Gamma(S \cup H_i, p)) && (5.24) \\ &\geq 1/3 && (5.25) \end{aligned}$$

Now, for any strategy τ for player II, if $E(\sigma \text{ vs } \tau \text{ in } \Gamma(S)) < 1$, then for any $\epsilon > 0$ there would be positions $p \in P$ such that

$$E(\sigma \text{ vs } \tau \text{ in } \Gamma(S, p)) < \epsilon \tag{5.26}$$

We conclude that such a strategy τ does not exist, and that therefore σ is of value 1 in $\Gamma(S)$.

□

5.4 Constructing a Free and σ-Complete Ultrafilter on ω_1

A *σ-complete ultrafilter* on a set X is a collection U of subsets of X, closed under countable intersection and taking supersets, such that for each set $V \subseteq X$, exactly one of $V, X - V$ is in the ultrafilter. For any cardinal α, U is called α-complete if it is closed under the intersection of less than α sets. An ultrafilter is called *free* if it is not of the form $\{V \subseteq X \mid x \in V\}$ for some $x \in X$. If there exists a free, α-complete ultrafilter on a set of cardinality α, we say that α is *measurable.*

Under AC, $\aleph_1$ (the first uncountable cardinal) is not measurable, i.e. there exists no free, σ-complete ultrafilter on a set of cardinality $\aleph_1$. It is a well-known theorem of large cardinal theory that under AD, $\aleph_1$ *is* measurable[8]. In this section we take a construction of a free σ-complete ultrafilter on the set $\omega_1 = \{\alpha \in ORD \mid \alpha \text{ is finite or countable}\}$ (a set of cardinality $\aleph_1$) which uses the Axiom of Determinacy [19], and modify it to use AD-Bl instead.

Let $V \subseteq \omega_1$. We define auxiliary perfect information games $\Gamma^{\omega_1}(V)$ in which players I and II independently construct countably many countable ordinals, represented as subsets of $\mathbb{Q}$. The two players maintain separate (countable) collections of (initially empty) subsets of $\mathbb{Q}$, and each round adds finitely many elements to finitely many subsets. Player I wins at the 'end' of the game if the supremum of the ordinals represented by the constructed subsets is in V, otherwise player II wins. Formally:

5.4.1. Definition. Let $\mathcal{A}$ and $\mathcal{B}$ be two countably infinite, disjoint sets. For any subset $V \subseteq \omega_1$, the game $\Gamma^{\omega_1}_{\mathcal{A},\mathcal{B}}(V)$ is defined as follows:

- In any round $i \geq 1$, first player I selects a finite set a_i of pairs $(a, q) \in \mathcal{A} \times \mathbb{Q}$, and then player II selects a finite set b_i of pairs $(b, r) \in \mathcal{B} \times \mathbb{Q}$.
- Set $\mathcal{I} := \mathcal{A} \cup \mathcal{B}$, and define the *result* z by $z := a_1 \cup b_1 \cup a_2 \cup b_2 \cup \ldots \subseteq \mathcal{I} \times Q$.
- Let the function $\pi : \mathcal{P}(\mathbb{Q}) \to \omega_1$ be defined by

$$\pi(R) := \begin{cases} \text{the order type of } (R, <) & \text{if } (R, <) \text{ is well-ordered} \\ 0 & \text{otherwise} \end{cases} \tag{5.27}$$

  Set $\Pi_{\mathcal{A},\mathcal{B}}(z) := \sup_{i \in \mathcal{I}} \pi(\{q \in \mathbb{Q} \mid (i, q) \in z\})$.
  Player I wins if $\Pi_{\mathcal{A},\mathcal{B}}(z) \in V$, otherwise player II wins.

If we interpret $\Pi_{\mathcal{A},\mathcal{B}}$ as a function from $\mathcal{P}(\mathcal{I} \times \mathbb{Q})$ to ω_1, then we can set

$$\underline{V}_{\mathcal{A},\mathcal{B}} := \Pi^{-1}_{\mathcal{A},\mathcal{B}}[V] = \{z \subseteq \mathcal{I} \times \mathbb{Q} \mid \Pi_{\mathcal{A},\mathcal{B}}(z) \in V\} \tag{5.28}$$

and the winning condition of the game can be reformulated as "Player I wins if $z \in \underline{V}_{\mathcal{A},\mathcal{B}}$, otherwise player II wins".

5.4.2. Remark. For technical reasons, it is necessary at some places in the proof to be able to react to one's own moves as if they had been made by the opponent. This is done by temporarily considering some of one's own subsets-under-construction to belong to the other player, for the purpose of reacting to the moves made in them. The 'index-structure' $(\mathcal{A}, \mathcal{B})$ used in the definition above is used to facilitate this.
For any two index-structures $(\mathcal{A}, \mathcal{B})$, $(\mathcal{A}', \mathcal{B}')$, there exist bijective mappings $\mathcal{A} \leftrightarrow \mathcal{A}'$ and $\mathcal{B} \leftrightarrow \mathcal{B}'$, which in turn induce bijective mappings between the moves, games and strategies of $\Gamma^{\omega_1}_{\mathcal{A},\mathcal{B}}(V)$ and those of $\Gamma^{\omega_1}_{\mathcal{A}',\mathcal{B}'}(V)$. Therefore, when the distinction is not important, we write $\Gamma^{\omega_1}(V)$, $\Pi(z)$ and $\underline{V}$ for $\Gamma^{\omega_1}_{\mathcal{A},\mathcal{B}}(V)$, $\Pi_{\mathcal{A},\mathcal{B}}(z)$ and $\underline{V}_{\mathcal{A},\mathcal{B}}$. Note that $\Pi_{\mathcal{A},\mathcal{B}}$ and $\underline{V}_{\mathcal{A},\mathcal{B}}$ depend on $\mathcal{A} \cup \mathcal{B}$ only.

An ultrafilter U can be thought of as a partitioning of the subsets of X into 'large' subsets (those in U) and 'small' subsets (those not in U). The property 'player I can almost surely force the supremum to be in V' intuitively seems likely to be a 'largeness'-type property. And indeed, we will show that

5.4.3. Theorem. *Under AD-Bl, the set*

$$U := \{V \subseteq \omega_1 \mid \textit{player I has a strategy of value } 1 \textit{ in } \Gamma^{\omega_1}(V)\} \tag{5.29}$$

is a free and σ-complete ultrafilter on ω_1.

To prove this, first we give lemmas constructing strategies in $\Gamma^{\omega_1}(V)$, for several different constructions of V from other sets.

5.4.4. LEMMA. *If player I has a strategy of value* 1 *in the game* $\Gamma^{\omega_1}(V)$, *and* $V \subseteq W$, *then player I has a strategy of value* 1 *in* $\Gamma^{\omega_1}(W)$.

Proof
Any strategy of value 1 for player I in $\Gamma^{\omega_1}(V)$ is also of value 1 in $\Gamma^{\omega_1}(W)$.

□

5.4.5. LEMMA. *If V is a singleton, then player II has a strategy of value* 1 *in the game* $\Gamma^{\omega_1}(V)$.

Proof
If $V = \{\beta\}$, then player II can win by constructing $\beta + 1$.

□

5.4.6. LEMMA. *If player I has a strategy of value* 1 *in the game* $\Gamma^{\omega_1}(V)$, *then player II has a strategy of value* 0 *in* $\Gamma^{\omega_1}(\omega_1 - V)$, *and vica versa.*

Proof
Suppose that player I has a strategy f of value 1 in the game $G_{\mathcal{A},\mathcal{B}}(V)$. Then this is also a strategy for player II of value 0 in the game $G_{\mathcal{B},\mathcal{A}}(\omega_1 - V)$, except that since II does not have the first move, player II's response to any move is always 'delayed' by one round. Formally, we construct a strategy g for player II by setting

$$g(\langle b_1, a_1, b_2, a_2, \ldots, a_{k-1}, b_k \rangle) := f(\langle a_1, b_1, a_2, b_2, \ldots, a_{k-1}, b_{k-1} \rangle) \tag{5.30}$$

For any moves for player I in $G_{\mathcal{B},\mathcal{A}}(\omega_1 - V)$, if player II plays according to g, then the resulting sequence of moves $\langle b_1, a_1, b_2, a_2, \ldots \rangle$ corresponds to a sequence of moves $\langle a_1, b_1, a_2, b_2, \ldots \rangle$ in the game $G_{\mathcal{A},\mathcal{B}}(V)$, such that the probability distribution of player I's moves is according to the strategy f. Hence we have

$$z = b_1 \cup a_1 \cup b_2 \cup a_2 \cup \ldots = a_1 \cup b_1 \cup a_2 \cup b_2 \cup \ldots \in \underline{V} \text{ almost surely} \tag{5.31}$$

So g is a strategy for player II of value 0 in $G_{\mathcal{B},\mathcal{A}}(V)$.
Now suppose that player II has a strategy g of value 0 in the game $G_{\mathcal{B},\mathcal{A}}(\omega_1 - V)$. Then this is also a strategy for player I of value 1 in the game $G_{\mathcal{A},\mathcal{B}}(V)$, except that player I has a first move in which she does nothing. Formally, we construct a strategy f for player I by setting

$$f(\langle\rangle) := \text{'play } \emptyset\text{'}, f(\langle a_1, b_1, a_2, b_2, \ldots, a_{k-1}, b_{k-1} \rangle) := g(\langle b_1, a_2, b_2, \ldots, a_{k-1}, b_{k-1} \rangle) \tag{5.32}$$

For any moves for player II in $G_{\mathcal{A},\mathcal{B}}(V)$, if player I plays according to f, then the resulting sequence of moves $\langle\emptyset, b_1, a_2, b_2, a_3, \ldots\rangle$ corresponds to a sequence of moves $\langle b_1, a_2, b_2, a_3, \ldots\rangle$ in the game $G_{\mathcal{B},\mathcal{A}}(\omega_1 - V)$, such that the probability distribution of player II's moves is according to the strategy g. Hence we have

$$z = \emptyset \cup b_1 \cup a_2 \cup b_2 \cup a_3 \cup \ldots = b_1 \cup a_2 \cup b_2 \cup a_3 \cup \ldots \in \underline{V} \text{ almost surely} \quad (5.33)$$

So f is a strategy for player I of value 1 in $G_{\mathcal{A},\mathcal{B}}(V)$.

□

5.4.7. LEMMA. *Let $(V^i)_{i\geq 0}$ be a countable sequence of subsets of ω_1. If for all $i \geq 0$, player II has a strategy of value* 0 *in $\Gamma^{\omega_1}(V^i)$, then player II has a strategy of value* 0 *in $\Gamma^{\omega_1}(\bigcup_{i\geq 0} V^i)$.*

Proof

Let $(V^i)_{i\geq 0}$ be a countable sequence of subsets of ω_1, and suppose that for all $i \geq 0$, player II has a strategy of value 0 in $\Gamma^{\omega_1}(V^i)$. Let $(\mathcal{A}, \mathcal{B})$ be an index-structure for the game. We will construct a strategy g of value 0 for player II in $\Gamma^{\omega_1}_{\mathcal{A},\mathcal{B}}(\bigcup_{i\geq 0} V^i)$.

Let $(\mathcal{B}^i)_{i\geq 0}$ be a partitioning of $\mathcal{B}$ into a countably infinite number of disjoint countably infinite sets. Define $\mathcal{A}^i = (\mathcal{A} \cup \mathcal{B}) - \mathcal{B}^i$ for $i \geq 1$. Then for all $i \geq 1$, $(\mathcal{A}^i, \mathcal{B}^i)$ is an index-structure. By assumption, we can find[12] strategies g^i of value 0 for player II in each of the games $\Gamma^{\omega_1}_{\mathcal{A}^i,\mathcal{B}^i}(V^i)$.

Now let $w = (a_1, b_1, a_2, b_2, \ldots)$ be a play of $\Gamma^{\omega_1}_{\mathcal{A},\mathcal{B}}(V)$ such that for all $i \geq 0$ and $k \geq 1$, $b_{2^i(2k-1)} \in \mathcal{B}^i$. If we define for $i \geq 0$ and $k \geq 1$,

$$\begin{aligned} a^i_1 &= a_1 \cup b_1 \cup a_2 \cup \ldots \cup a_{2^i} \\ a^i_k &= a_{2^i(2k-3)+1} \cup b_{2^i(2k-3)+1} \cup a_{2^i(2k-3)+2} \cup \ldots \cup a_{2^i(2k-1)} \text{ for } k > 1 \\ b^i_k &= b_{2^i(2k-1)} \end{aligned}$$

then for all $i \geq 0$ and $k \geq 1$, a^i_k and b^i_k are finite subsets of $\mathcal{A}^i \times \mathbb{Q}$ and $\mathcal{B}^i \times \mathbb{Q}$, so $w^i = (a^i_1, b^i_1, a^i_2, b^i_2 \ldots)$ is a valid play of the game $\Gamma^{\omega_1}_{\mathcal{A}^i,\mathcal{B}^i}(V^i)$. So construct the strategy g for player II in $\Gamma^{\omega_1}_{\mathcal{A},\mathcal{B}}(V)$ by setting, for $i \geq 0$ and $k \geq 1$,

$$g((a_1, b_1, \ldots, a_{2^i(2k-1)})) := g^i((a^i_1, b^i_1, \ldots, a^i_k)) \quad (5.34)$$

It can easily be shown (inductively) that for all $i \geq 0$ and $k \geq 1$, $b_{2^i(2k-1)} \in \mathcal{B}^i$ always, so g is well-defined. Moreover, for all $i \geq 0$ and $k \geq 1$, the probability distribution of b^i_k is given by $g^i((a^i_1, b^i_1, a^i_2, \ldots, a^i_k))$, so for all $i \geq 0$, the probability distribution of w^i in the game $\Gamma^{\omega_1}_{\mathcal{A}^i,\mathcal{B}^i}(V^i)$ is consistent with player II's strategy g^i. It follows that for all $i \geq 0$,

$$z = a_1 \cup b_1 \cup a_2 \cup b_2 \cup \ldots = a^i_1 \cup b^i_1 \cup a^i_2 \cup b^i_2 \cup \ldots \in \underline{\omega_1 - V^i} \text{ almost surely} \quad (5.35)$$

[12]Again, we do not need the Axiom of Choice for this. See footnote 11.

and hence

$$z = a_1 \cup b_1 \cup a_2 \cup b_2 \cup \ldots \in \bigcap_{i \geq 0} \underline{\omega_1 - V^i} = \underline{\omega_1 - V} \text{ almost surely} \tag{5.36}$$

So g is a strategy of value 0 for player II in the game $\Gamma^{\omega_1}_{\mathcal{A},\mathcal{B}}(V)$.

□

Proof of Theorem 5.4.3

In any game $\Gamma^{\omega_1}_{\mathcal{A},\mathcal{B}}(V)$, there are only countably many possible moves each turn, since $\mathcal{A}$, $\mathcal{B}$ and $\mathbb{Q}$ are all countable, and therefore there are only countably many different finite collections of pairs $(a, q) \in \mathcal{A} \times \mathbb{Q}$ or $(b, q) \in \mathcal{B} \times \mathbb{Q}$. Hence Theorem 5.3.4 applies, and we have that for all $V \subseteq \omega_1$:

$$\begin{aligned} V \in U \quad &\Leftrightarrow \quad \text{player I has a strategy of value 1 in } \Gamma^{\omega_1}(V) \\ V \notin U \quad &\Leftrightarrow \quad \text{player II has a strategy of value 0 in } \Gamma^{\omega_1}(V) \end{aligned}$$

By the above equivalences, the previous lemmas correspond to the following properties of U:

1. For any $V, W \subseteq \omega_1$, if $V \in U$ and $V \subseteq W$, then $W \in U$.
2. For any $V \subseteq \omega_1$, if V is a singleton, then $V \notin U$.
3. For any $V \subseteq \omega_1$, $V \in U$ if and only if $\omega_1 - V \notin U$.
4. For any sequence $V_i \subseteq \omega_1$, if $V_i \notin U$ for all $i \geq 0$, then $\bigcup_{i \geq 0} V_i \notin U$.

and from the third and fourth property we can derive

5 For any sequence $V_i \subseteq \omega_1$, if $V_i \in U$ for all $i \geq 0$, then $\bigcap_{i \geq 0} V_i \in U$.

So U is a free and σ-complete ultrafilter on ω_1.

□

Part II

Random Walks

Overview

As the title indicates, this part of the dissertation is about random walks. Specifically, it is about recurrence in so-called *reinforced* random walks, where edges in a graph are traversed with probabilities that may be different (reinforced) at second, third etc. traversals.
We start by reviewing some general theory of random walks, in Chapter 6. Chapter 7 is a brief digression to a problem involving gambling and related to random walks, namely the problem of calculating the expected time until one of the players is broke. In Chapter 8, we introduce reinforced random walks. After some general results, we focus on the case where the probability for any edge only changes once, after its first traversal. As a special case, we show that the once-reinforced random walk on the infinite ladder is almost surely recurrent if reinforcement is small, extending a result by T. Sellke from an at this time unpublished article[31]. In Chapter 9, we briefly review the basics of nonstandard analysis and its application to graph theory. We use this in Chapter 10 to show that for a class of graphs which generalizes the infinite ladder, recurrence holds for sufficiently *large* reinforcements.

Chapter 6

Random Walks

In this chapter we will introduce the basic random walk. The stereotypical example of a random walk is that of the Drunkard's Walk, where a drunken vagrant, starting from Times Square, wanders the streets of Manhattan aimlessly and totally at random. The behavior of this famous alcoholic has been studied in many papers [25, 32]. One of the more interesting aspects of this stochastic process is that it has the property of *recurrence*. It can be shown that, given enough time, the drunkard is certain to return to Times Square eventually. Indeed, even if we extend Manhattan to some hypothetical infinite city (preserving the characteristic street pattern, of course), not only is the drunkard certain to return to Times Square, but eventually he will visit each and every corner of the city infinitely often. Or to be exact, the probability of his doing so is equal to 1.
We will start this chapter by reviewing notational conventions, defining the non-reinforced random walk and giving several characterizations of recurrence for this walk. Next we will introduce such basic concepts as martingales, stopping times and harmonic functions, and show how martingales naturally arise from random walks. Finally, we will characterize recurrence of non-reinforced random walks on graphs in terms of the existence of certain superharmonic functions on the vertices of these graphs, and give several examples of the application of these theorems to specific graphs. This chapter presumes some basic knowledge of graph theory and probability theory, but an effort has been made to make it as self-contained as possible.

6.1 The Non-Reinforced Random Walk

6.1.1. Remark. In this dissertation, random walks are always considered to be walks on the edges of weighted graphs with finitely or countably infinitely many vertices. We will assume that any given graph is connected, that there are no 'degenerate' edges of weight 0, and that each vertex has only finitely many neighbors. To avoid needless notational complications we will also assume that any given graph is countably infinite, and simple (i.e. without loops or parallel edges) unless explicitly stated otherwise.

The reader is invited to verify for him- or herself that all definitions, proofs and results in these chapters can easily be extended to non-simple graphs. Indeed, the generalization to non-simple graphs of Lemma 8.3.4 will be used in the proof of Theorem 8.3.9. However, since this extension does not add anything conceptually, and since it is convenient to be able to denote edges and arcs by their endpoints, we will concern ourselves with simple graphs, and postulate generalizations to non-simple graphs when necessary.

NOTATION. We denote a weighted graph G as $G = (V, E, w)$, where V and E are the sets of vertices and edges of G, and $w : E \to \mathbb{R}_{>0}$ is its weight function. Edges are denoted by their endpoints, as in 'the edge uv'. Note that uv and vu denote the same edge. Whenever the order of the vertices is important (for instance, when we want to indicate the direction in which an edge has been traversed), we use *arcs* (oriented edges), denoted as in 'the arc $\overrightarrow{uv}$', instead of edges. u and v are called the *tail* and *head* of $\overrightarrow{uv}$, respectively.
Some other notational conventions:

- v and u are used for vertices.
- $N_G(v)$ denotes the *neighbor set* of a vertex v in a graph $G = (V, E, w)$, i.e. the set of vertices u such that $uv \in E$.
- $w_G(v)$ denotes the total weight $\sum_{u \in N(v)} w(vu)$ of the edges adjacent to v.
- $\rho_G(v)$ denotes the *degree* of v in G, i.e. the number of adjacent edges.
- $d_G(v, u)$ denotes the *distance* in G between the vertex v and the vertex u (i.e. the number of edges contained in the shortest $v - u$ path in G).
- $d_G(v, F)$ denotes the distance in G between a vertex v and a vertex-set F.

The index G is omitted when no confusion is possible.

A random walk is a stochastic process of traversing the edges of a graph, where, each time a vertex is reached, the random walk continues over a randomly selected adjacent edge. Specifically, the *non-reinforced random walk* on a graph $G = (V, E, w)$ starting at a vertex $v_0 \in V$, is the following stochastic process:

- We start with the vertex v_0.
- Next, we randomly pick an edge $v_0 v_1 \in E$ that connects v_0 with some other vertex $v_1 \in V$. All candidate edges have a probability of being picked proportional to their weight. The random walk is said to *traverse* the edge $v_0 v_1$, and to *visit* the vertex v_1 at *time* 1.

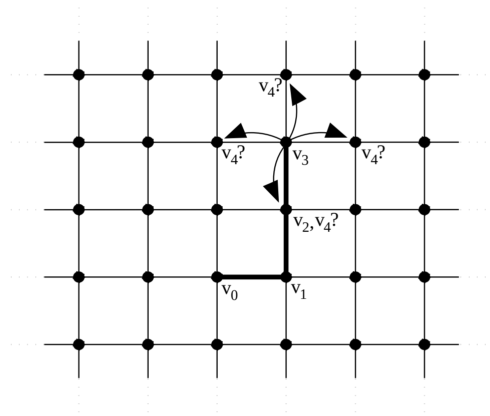

Figure 6.1: The first few steps of a random walk on the square lattice on $\mathbb{Z}^2$.

- Next, we randomly pick an edge $v_1v_2 \in E$ that connects v_1 with some other vertex $v_2 \in V$, in the same manner as in the previous step.
- Continuing in this manner, we obtain a path $v_0v_1v_2v_3\ldots$.

More formally,

6.1.2. DEFINITION. A *non-reinforced random walk* on a weighted graph $G = (V, E, w)$ is a series of stochastic variables $v_0, v_1, \ldots \in V$ such that for any *time* $t \in \mathbb{N}$,

$$P(v_{t+1} = u \mid \mathcal{F}_t) = \begin{cases} \frac{w(v_t u)}{w(v_t)} & \text{if } u \in N(v_t) \\ 0 & \text{otherwise} \end{cases} \tag{6.1}$$

where $\mathcal{F}_t$ denotes the σ-algebra of the history up to time t. Note that by our assumptions on graphs, $N(v) > 0$ for all $v \in V$.

NOTATION. v_t always denotes the location of the random walk at time t. Sometimes we write $v_0v_1v_2\ldots$ for the random walk itself. Throughout these chapters s and t are used for (integer) times, and the use of t as a subscript indicates a (stochastic) variable whose contents changes over time (such as v_t).

6.1.3. DEFINITION. A realization of a random walk is said to be *recurrent* if every vertex is visited infinitely often, and *transient* if every vertex is visited only finitely many times.

The question we are mainly concerned with in these chapters , is under what conditions a random walk is recurrent almost surely (i.e. with probability 1). For non-reinforced random walks we have the following observations:

6.1.4. LEMMA. *Let $G = (V, E, w)$ be a weighted graph, and consider the non-reinforced random walk on G starting in a vertex v_0. Then, depending on G, the random walk is either almost surely recurrent or almost surely transient.*

Proof

Whenever the random walk is at v_0, there is a probability p that v_0 is revisited at some later time. The probability that the random walk will revisit v_0 at least n times is exactly p^n. Hence, if $p = 1$ then almost surely v_0 will be visited infinitely often, and if $p < 1$ then almost surely v_0 will be visited only finitely often.
If two vertices $u, u' \in V$ are neighbors in G, then whenever u is visited, there is a probability $p' > 0$ that the next vertex visited will be u'. It follows that

$$P(u \text{ is visited infinitely often and } u' \text{ only finitely often}) = 0 \tag{6.2}$$

By induction on $d_G(u, u')$ we can show that the same holds for any two vertices $u, u' \in V$. The result follows.

□

6.1.5. LEMMA. *Let $G = (V, E, w)$ be a weighted graph, $F \subset V$ a finite set of vertices of G, and $v \in F$. Then the following are equivalent:*

(i) Any non-reinforced random walk on G is almost surely recurrent.

(ii) The non-reinforced random walk on G starting in v is almost surely recurrent.

(iii) The non-reinforced random walk on G starting in v returns to v almost surely.

(iv) Any non-reinforced random walk on G visits F almost surely.

Proof

$(i) \Rightarrow (ii)$ and $(ii) \Rightarrow (iii)$ are obvious.
If we assume (iv) does not hold, then for some $u \in V$, the non-reinforced random walk starting in u will not almost surely visit F. Since there is a path between v and u, the non-reinforced random walk starting in v will visit u with positive probability, and hence will not almost surely return to v, and (iii) fails.
Finally assume (iv) holds. Then starting at any vertex $u \in V$, F is visited almost surely. After this first visit F is almost surely visited again, and repeating this process we find that F is almost surely visited infinitely often. Therefore the non-reinforced random walk starting in u is almost surely not transient, and hence by Lemma 6.1.4 almost surely recurrent, proving (i).

□

6.2 Random Walks and Martingales

Our major tools for showing recurrence of random walks will be the concept of *martingales* and the Optional Stopping Theorem.

6.2.1. Definition. A series of stochastic variables $(M_t)_{t \in \mathbb{N}}$ is called a *martingale* if for all $t \in \mathbb{N}$,

$$M_t = E(M_{t+1} \mid \mathcal{F}_t) \tag{6.3}$$

where $E(M_{t+1} \mid \mathcal{F}_t)$ denotes the expectation, at time t, of the value of M_{t+1}.

6.2.2. Definition. A series of stochastic variables $(M_t)_{t \in \mathbb{N}}$ is called a *supermartingale* [*submartingale*] if for all $t \in \mathbb{N}$,

$$M_t \geq [\leq]\ E(M_{t+1} \mid \mathcal{F}_t) \tag{6.4}$$

It is easy to see that if M is a martingale, then $E(M_t) = M_0$ for any time $t \in \mathbb{N}$. The Optional Stopping Theorem for Martingales basically states that the same holds for the expectation of the value of the martingale at times which are defined in terms of states or conditions, such as the first time at which the value of the martingale is < 0 or > 100. To state the theorem, we need the concept of *stopping times.*

6.2.3. Definition. A *stopping time* is a stochastic variable τ, taking values in $\mathbb{N} \cup \{\infty\}$, such that for all $t \in \mathbb{N}$, $\{\tau = t\} \in \mathcal{F}_t$.

Stopping times are usually defined in the manner of 'let τ be the first time at which some condition holds'. Often we are only interested in the course of a random walk up to a certain event, such as its first visit to some given vertex. In that case we write 'the random walk which stops at time τ', or sometimes simply 'the random walk which stops as soon as some condition holds'.

6.2.4. Theorem (Optional Stopping Theorem for Martingales). *Let M_t be a martingale [supermartingale, submartingale] and τ a stopping time such that $\tau < \infty$ almost surely. If M_t is bounded [bounded from below, bounded from above] for $t < \tau$, then*

$$M_0 = [\geq, \leq] E(M_\tau) \tag{6.5}$$

and more generally

$$M_{t_0} = [\geq, \leq] E(M_\tau \mid \mathcal{F}_{t_0}) \tag{6.6}$$

if $t_0 \leq \tau$.

Kakutani[26] found that random walks give rise to martingales naturally, if we can find a function on the vertex-set of the graph with the property of *harmonicity*:

6.2.5. Definition. Let $G = (V, E, w)$ be a weighted graph, and let $h : V \to \mathbb{R}$ be a function. We say that h is *harmonic* [*superharmonic, subharmonic*] on a vertex-set $V' \subset V$ if for all $v \in V'$,

$$h(v) = [\geq, \leq] \sum_{u \in N(v)} h(u) \frac{w(vu)}{w(v)} \tag{6.7}$$

or, equivalently,

$$\sum_{u \in N(v)} w(vu)\Delta_h(\overrightarrow{vu}) = [\leq, \geq]\ 0 \tag{6.8}$$

where $\Delta_h(\overrightarrow{vu})$ denotes $h(u) - h(v)$.

6.2.6. LEMMA. *Let $G = (V, E, w)$ be a weighted graph, and let $h : V \to \mathbb{R}$ be a harmonic [superharmonic,subharmonic] function on a subset $V' \subset V$. Consider a non-reinforced random walk on G, and define*

$$M_t = \sum_{t'=0}^{t} \begin{cases} \Delta_h(v_{t'} v_{t'+1}) & \text{if } v_{t'} \in V' \\ 0 & \text{otherwise} \end{cases} \tag{6.9}$$

for $t \in \mathbb{N}$. Then M is a martingale [supermartingale, submartingale]. Furthermore, as long as $V - V'$ has not yet been visited,

$$M_t = h(v_t) - h(v_0) \tag{6.10}$$

Proof

If $v_t \in V - V'$ then $M_{t+1} = M_t$, otherwise

$$M_t \;=[\geq, \leq]\; M_t + \frac{1}{w(v_t)} \cdot \sum_{u \in N(v_t)} \Delta_h(v_t u) w(v_t u) \tag{6.11}$$

$$= M_t + \sum_{u \in N(v_t)} P(v_{t+1} = u \mid \mathcal{F}_t) \Delta_h(v_t u) \tag{6.12}$$

$$= E(M_{t+1} \mid \mathcal{F}_t) \tag{6.13}$$

The proof of the final statement is trivial.

□

More about martingales may be found in [24].

6.3 Recurrence and Superharmonic Functions

Now, if h is a superharmonic function on (a subset of) the vertex-set V of a graph G then the Optional Stopping Theorem for Martingales places bounds on the expected values of $h(v_t)$. We can use this to characterize recurrence of random walks in terms of the existence of superharmonic functions with certain properties.

6.3.1. DEFINITION. Let $G = (V, E, w)$ be a weighted graph, and let $h : V \to \mathbb{R}$ be a function. We say that $h(v)$ *goes to infinity if v goes to infinity* if

$$\forall r \in \mathbb{R}\ \exists n \in \mathbb{N}\ \forall v \in V\ (d_G(v_0, v) > n \ \Rightarrow\ h(v) > r) \tag{6.14}$$

For the graphs we are concerned about, in which no vertex has infinitely many neighbors, this is equivalent to the condition that

$$\text{for each } r \in \mathbb{R},\ \{v \in V \mid h(v) < r\} \text{ is finite} \tag{6.15}$$

6.3.2. THEOREM. *Let $G = (V, E, w)$ be a weighted graph. Then non-reinforced random walks on G are almost surely recurrent if there exists a function $h : V \to \mathbb{R}$ satisfying*

1. *h is superharmonic everywhere except on some finite set F.*
2. *$h(v)$ goes to infinity if v goes to infinity.*

Conversely, if non-reinforced random walks on G are almost surely recurrent, then a function h as above exists, and F may be chosen to be an arbitrary non-empty finite set.

Proof

First assume such a function h exists. Then h is bounded from below. We may assume without loss of generality that $h \geq 0$. So consider the random walk, starting at an arbitrary point $v_0 \in V$. By Lemma 6.1.5, it suffices to show that F will be visited almost surely. Let M_t be the martingale from Lemma 6.2.6, and let for $r > 0$ the stopping time τ_r be the first time at which the random walk leaves the finite set of vertices $\{v \in V \mid v \notin F \wedge h(v) < r\}$.

By Lemma 6.1.4, almost surely the random walk is either transient or recurrent, and in both cases the random walk visits infinitely many vertices. It follows that $\tau_r < \infty$ almost surely. Furthermore, $M_t = h(v_t) \geq 0$ for $t \leq \tau$. Hence we can use the Optional Stopping Time Theorem to obtain

$$M_0 \geq E(M_{\tau_r}) \geq (1 - P(v_{\tau_r} \in F))r \tag{6.16}$$

and hence $P(v_{\tau_r} \in F) \geq 1 - M_0/r$ for all $r > 0$. We conclude that $P(\exists t : v_t \in F) \geq 1 - \epsilon$ for arbitrarily small $\epsilon > 0$, and hence the random walk is almost surely recurrent.

For the converse implication, let $F \subset V$ be a non-empty finite set of vertices, and assume random walks on G almost surely visit F. Let, for any vertex $v \in V$, τ_v be the time that the random walk starting from v first visits F. Then for any vertex $v \in V$, $\tau_v < \infty$ almost surely. Now set $h(v) = E(f(\tau_v))$, where $f : \mathbb{N} \to \mathbb{R}$ is such that such that f monotonely diverges to infinity and $E(f(\tau_v))$ is finite for all $v \in V$.[13] Then h is well-defined, and by the monotonicity of f we have that for $v \in V - F$

$$\begin{aligned} h(v) &\geq E(f(\tau_v - 1)) && (6.17)\\ &= \sum_{u \in N(v)} E(f(\tau_v - 1) \mid v_1 = u)\frac{w(vu)}{w(v)} && (6.18)\\ &= \sum_{u \in N(v)} E(f(\tau_u))\frac{w(vu)}{w(v)} && (6.19)\end{aligned}$$

[13] For example, $f(n) = \min_{v \in V}\left(d(v, F) + (P(\tau_v \geq n))^{-1/2}\right)$ can be shown to have these properties. Unfortunately $E(\tau_v)$ is not generally finite, or we would not *need* f.

$$= \sum_{u \in N(v)} h(u) \frac{w(vu)}{w(v)} \tag{6.20}$$

So h is superharmonic on $V - F$, Furthermore, since starting from a vertex v F cannot be reached before time $t = d(v, F)$,

$$h(v) = E(f(\tau_v)) \geq f(d(v, F)) \to \infty \text{ if } v \to \infty \tag{6.21}$$

□

6.3.3. Example. The random walk on the square lattice graph on $\mathbb{Z}^2$ with unit weights is almost surely recurrent.

Proof
Let $h : \mathbb{Z}^2 \to \mathbb{R}$ be defined by

$$h(x,y) = \begin{cases} \log(1/12) & \text{if } (x,y) = (0,0) \\ \log(1/4) & \text{if } (x,y) = (0,\pm 1) \text{ or } (x,y) = (\pm 1, 0) \\ \log(x^2 + y^2 - 1) & \text{otherwise} \end{cases} \tag{6.22}$$

Then h satisfies the conditions of theorem 6.3.2, with $F = \{(0,0)\}$.

□

6.3.4. Example. For any $n \in \mathbb{N}_{>0}$, the random walks on the square lattice graphs on $\mathbb{Z} \times \{1, \ldots, n\}$ and $\mathbb{Z} \times (\mathbb{Z}/n\mathbb{Z})$ with unit weights are almost surely recurrent.

Proof
Let $h : (\mathbb{Z} \times \{1 \ldots n\}) \to \mathbb{R}$ be defined by

$$h(x, y) = |x| \tag{6.23}$$

Then h satisfies the conditions of theorem 6.3.2, with $F = \{(0, y) \mid 1 \leq y \leq n\}$. The proof for the cylinder lattice $\mathbb{Z} \times (\mathbb{Z}/n\mathbb{Z})$ is completely analogous.

□

Interestingly enough, the non-recurrence of random walks on a graph can *also* be characterized in terms of the existence of certain superharmonic functions.

6.3.5. Theorem. *Let $G = (V, E, w)$ be a weighted graph. Then non-reinforced random walks on G are not almost surely recurrent if and only if there exists a bounded non-constant function $h : V \to \mathbb{R}$ that is superharmonic on V.*

Proof

First assume that such a function h exists. Let $v_0, u \in V$ be vertices with $h(v_0) < h(u)$. By Lemma 6.1.5, it suffices to prove that the random walk starting at the vertex v_0 does not almost surely visit u. So consider the random walk, starting at the vertex v_0, which halts on visiting the vertex u, and assume that it does so almost surely. Then the stopping time $\tau = \min\{t \geq 1 \mid v_t = u\}$ is finite almost surely. By Lemma 6.2.6, the stochastic process $M_t = h(v_t)$ is a supermartingale, and by our initial assumption it is bounded. Hence we can use the Optional Stopping Times Theorem to obtain

$$h(v_0) = M_0 \geq E(M_\tau) = h(u) \tag{6.24}$$

contradicting our choice of v and u.

Now assume that random walks on G are not almost surely recurrent. Then there are vertices $v_0, u \in V$ such that starting at v_0, the random walk will not almost surely visit u. Define the function h by

$$h(v) = P(\text{the random walk starting at } v \text{ will reach } u) \tag{6.25}$$

Then $h : V \to [0,1]$ is bounded, $h(v_0) < h(u) = 1$, h is harmonic on $V - \{u\}$ and h is superharmonic on $\{u\}$.

□

6.3.6. Example. The random walk on the cubic lattice graph on $\mathbb{Z}^3$ is not almost surely recurrent.

Proof

Let $h : \mathbb{Z}^3 \to [0, 6^{-1/2}]$ be defined by

$$h(x,y,z) = \frac{1}{(x^2+y^2+z^2+6)^{1/2}} \tag{6.26}$$

Using a truncated Taylor series expansion of h, we can show that for all $x, y, z \in \mathbb{Z}$,

$$h(x+1,y,z) + h(x-1,y,z) \leq 2h(x,y,z) + \frac{2x^2 - y^2 - z^2}{(x^2+y^2+z^2+6)^{5/2}} \tag{6.27}$$

Analogous inequalities hold for $h(x, y+1, z) + h(x, y-1, z)$ and $h(x, y, z+1) + h(x, y, z-1)$. Taking the sum of these inequalities yields the superharmonicity inequality.

□

6.3.7. Example. Let $G = (V, E, w)$ be a weighted graph with $V = \{v^n \mid n \in \mathbb{Z}\}$, and $E = \{v^n v^{n+1} \mid n \in \mathbb{Z}\}$.[14] Then random walks on G are almost surely recurrent if and only if $\sum_{n=-\infty}^{-1}(1/w(v^n v^{n+1}))$ and $\sum_{n=0}^{\infty}(1/w(v^n v^{n+1}))$ both diverge.

[14] The superscript index v^n is used here to avoid confusion with the temporal index v_t.

Proof

If $\sum_{n=0}^{\infty}(1/w(v^n v^{n+1}))$ converges to $c \in \mathbb{R}$, then define $h : V \to [0, c]$ by setting $h(v^n) = \max(c, \sum_{k=n}^{\infty}(1/w(v^k v^{k+1})))$ for $n \in \mathbb{Z}$. It is easily verified that h is non-constant, harmonic on $V - \{v^0\}$ and superharmonic on $\{v^0\}$, fulfilling the conditions of Theorem 6.3.5. Likewise for the case that $\sum_{n=-\infty}^{-1}(1/w(v^n v^{n+1}))$ converges.

Now suppose $\sum_{n=-\infty}^{-1}(1/w(v^n v^{n+1}))$ and $\sum_{n=0}^{\infty}(1/w(v^n v^{n+1}))$ both diverge to ∞. Then define $h : V \to \mathbb{R}_{\geq 0}$ by setting $h(v^n) = \sum_{k=0}^{n-1}(1/w(v^k v^{k+1}))$ for $n \geq 0$ and $h(v^n) = \sum_{k=-n}^{-1}(1/w(v^k v^{k+1}))$ for $n < 0$. Again it is easily verified that h is non-constant, $h(v^n) \to \infty$ if $n \to \infty$ or $n \to -\infty$, and h is harmonic on $V - \{v^0\}$, fulfilling the conditions of Theorem 6.3.2

□

The next Theorem is included because it will be used in a later chapter. The proof, unfortunately, is beyond the scope of these pages. A beautiful proof was given by Doyle and Snell in [25].

6.3.8. THEOREM. *Let $G = (V, E, w)$ be a weighted graph such that non-reinforced random walks on G are almost surely recurrent. If $G' = (V', E', w')$ is a connected subgraph of G, possibly with lesser weights (i.e. $V' \subset V$, $E' \subset E$ and for all $e \in E'$, $w'(e) \leq w(e)$), then random walks on G' are almost surely recurrent.*

Chapter 7

Why The Gambler Came Home Late

Consider the following stochastic experiment. We have n players, each possessing a_1, a_2, ..., a_n coins. They play a gambling game, where each round, one of the players loses a coin to one of the other players (both selected randomly). The game continues until one of the players has no more coins. Let us denote the expected duration of the game (in rounds) by $T_n(a_1, a_2, \ldots, a_n)$.

7.0.9. Remark. This process of gambling can be viewed as a random walk on a finite graph. Consider the graph $G_n^s = (V_n^s, E_n^s)$ with vertices and edges

$$V_n^s = \{\vec{a} \in \mathbb{N}^n : a_1 + \ldots + a_n = s\} \tag{7.1}$$

$$E_n^s = \{\vec{a}\vec{b} \mid \vec{a}, \vec{b} \in V_n^s, \exists i \exists j : a_i = b_i + 1 \wedge a_j = b_j - 1 \wedge \forall k \neq i, j : a_k = b_k\} \tag{7.2}$$

Then the gambling process above corresponds to a random walk in G_n^s. where we take s to be the total number of coins the players possess. In particular, the expected duration of the gambling process corresponds to the expected time until the random walk reaches a vertex $\vec{a}$ with $\exists i : a_i = 0$. Since the graph on which the random walk takes place is finite, the expectation of this time is finite. Hence $T_n(a_1, \ldots, a_n)$ exists for all $n \geq 1$ and all $a_1, \ldots, a_n \geq 0$.

Question: What can be said about T_n as a function on $\mathbb{N}^n$?

7.1 Basic properties of T_n

To get an idea of the properties of the function T_n, let us consider first the case where $n = 2$. Here we have two players, and each round, one of the players loses one coin to the other player, until one of the players is broke. If the players currently possess a_1 and a_2 coins, respectively, then after one round, with probability 1/2 they will possess $a_1 - 1$ and $a_2 + 1$ coins, and otherwise they will possess $a_1 + 1$

and $a_2 - 1$ coins, respectively. It is easily seen that the function T_2 satisfies the equation

$$T_2(a_1, a_2) = \begin{cases} 0 & \text{if } a_1 = 0 \text{ or } a_2 = 0 \\ 1 + \frac{T_2(a_1-1,a_2+1)+T_2(a_1+1,a_2-1)}{2} & \text{otherwise} \end{cases} \tag{7.3}$$

Furthermore, for constant $s = a_1 + a_2$, this is a finite linear set of equations in the variables $T(0, s)$, $T(1, s-1)$, ..., $T(s, 0)$. It can be shown that this system of equations has a unique solution. Hence the function T_2 is uniquely determined by equation 7.3 Since the function $a_1 a_2$ satisfies this condition,

$$T_2(a_1 a_2) = a_1 a_2 \tag{7.4}$$

For $n = 3$, we can do something similar. In this case, each round there are 6 possibilities (one of three players loses a coin, and one of the remaining two players gains one), and we obtain the equation

$$T_3(a_1, a_2, a_3) = \begin{cases} 0 & \text{if } \exists i : a_i = 0 \\ 1 + \frac{T_3(a_1-1,a_2+1,a_3)+\ldots+T_3(a_1,a_2+1,a_3-1)}{6} & \text{otherwise} \end{cases} \tag{7.5}$$

Again, T_3 is uniquely defined by this equation (for a proof, see the general proof in the next lemma), and with some searching we can find the formula

$$T_3(a_1, a_2, a_3) = \frac{3a_1 a_2 a_3}{a_1 + a_2 + a_3} \tag{7.6}$$

In general

7.1.1. LEMMA. *If we define*

$$P_n = \{\vec{d} \in \mathbb{Z}^n \mid \exists i \exists j : d_i = 1 \wedge d_j = -1 \wedge \forall k \neq i, j : d_k = 0\} \tag{7.7}$$

then we have, for general n, that T_n satisfies and is uniquely defined by the equation

$$T_n(\vec{a}) = \begin{cases} 0 & \textit{if } \exists i : a_i = 0 \\ 1 + \frac{1}{n(n-1)} \sum_{\vec{d} \in P_n} T_n(\vec{a} + \vec{d}) & \textit{otherwise} \end{cases} \tag{7.8}$$

Proof

By Remark 7.0.9, T_n exists, and that T_n satisfies this equation follows directly from the definition of the gambling process, analogously to the cases with two and three players. To show that it is uniquely defined, fix $s \in \mathbb{N}$, and view the linear system of equations given by equation 7.8 for all $\vec{a} \in \mathbb{N}^n$ with $a_1 + a_2 + \ldots + a_n = s$, i.e. for all $\vec{a} \in V_n^s$. This system has exactly as many equations as it has variables, so to show that it has a unique solution, it suffices to show that the related system of equations

$$T'_n(\vec{a}) = \begin{cases} 0 & \text{if } \exists i : a_i = 0 \\ \frac{1}{n(n-1)} \sum_{\vec{d} \in P_n} T'_n(\vec{a} + \vec{d}) & \text{otherwise} \end{cases} \tag{7.9}$$

for all $\vec{a} \in V_n^s$, only has the solution with $T_n'(\vec{a}) = 0$ for all $\vec{a} \in V_n^s$. Now, if T_n' satisfies this system of equations, and if for some $a \in V_n^s$, $T_n'(\vec{a})$ is nonzero and either maximal or minimal in V_n^s, then we have that $T_n'(\vec{a}+\vec{d}) = T_n'(\vec{a})$ for all $d \in P_n$. Since G_s^n is a connected and finite graph, this would imply that $T_n'(\vec{a})$ is constant and non-zero on V_s^n, a contradiction.

□

Unfortunately, for $n > 3$ there is no formula known to satisfy this equation. Simply generalizing of the formulas for $n = 2$ and $n = 3$ to

$$T_n(\vec{a}) = \frac{C a_1 a_2 \dots a_n}{(a_1 + \dots + a_n)^m}, \text{ for some } C > 0,\ m \in I\!N \tag{7.10}$$

does not work: the differential

$$T_n(\vec{a}) - \frac{1}{n(n-1)} \sum_{\vec{d} \in P_n} T_n(\vec{a}+\vec{d}) \tag{7.11}$$

should be constant in order to satisfy equation (7.8), but if we write it out we get

$$\frac{C}{(a_1 + \dots + a_n)^m} \sum_{1 \le i < j \le n} \prod_{k \ne i,j} a_k \tag{7.12}$$

For $n > 3$, this expression is not constant for any $C > 0$ and $m \in I\!N$, and hence T_n cannot be expressed in this particular form. In Theorem 7.3.1, we will show that if such a formula exists, it must be considerably more complicated than the formulas for $n = 2$ and $n = 3$.

7.2 T_n, H_n and T_n^*

For each s, T_n can be thought of as a function on the graph G_n^s having constant curvature. There are known functions with constant curvature, such as $1/2|\vec{a}|^2$, which unfortunately do not have the right boundary values. However, the difference between such functions and T_n would be a function of *zero* curvature, i.e. a harmonic function.

7.2.1. LEMMA. *T_n can be written as*

$$T_n(\vec{a}) = \frac{1}{2}(H_n(\vec{a}) - |\vec{a}|^2) \tag{7.13}$$

where H_n is the unique function satisfying

$$H_n(\vec{a}) = |\vec{a}|^2 \quad \text{if } \exists i : a_i = 0 \tag{7.14}$$

$$\sum_{\vec{d} \in P_n} (H_n(\vec{a}+\vec{d}) - H_n(\vec{a})) = 0 \quad \text{otherwise} \tag{7.15}$$

Proof
It is straightforward to see that $H_n = 2T_n + |\vec{a}|^2$ satisfies the given equations. Uniqueness can be proven exactly as in the previous lemma.

□

7.2.2. REMARK. For any s, H_n is harmonic on G_n^s everywhere except on vertices $\vec{a}$ with $\exists i : a_i = 0$. As such, for $\vec{a} \in V_n^s$ $H_n(\vec{a})$ is equal to the *expected value* of $|\vec{b}|^2$, where $\vec{b}$ is the first vertex of G_n^s with $\exists i : b_i = 0$ which is visited by the random walk on G_n^s starting in $\vec{a}$.

Interestingly, if we look at a variation on the first game where, once one gambler is broke, the game continues with the remaining gamblers until one gambler has won all the money, we get an expected duration function T_n^* which *does* have a simple form:

7.2.3. THEOREM. *Let $T_n^*(\vec{a})$ denote the expected duration of the variation of the game where play continues until all but one gambler is broke. Then*

$$T_n^*(\vec{a}) = \frac{1}{2}((\sum_i a_i)^2 - |\vec{a}|^2) \tag{7.16}$$

Proof
We will prove this by induction on n. For $n = 1$, both sides of equation (7.16) are equal to 0, and the equation holds. For $n > 1$, using the methods of Lemma 7.1.1, it follows from the Induction Hypothesis that that T_n^* is the unique function satisfying

$$T_n(\vec{a}) = \begin{cases} \frac{1}{2}((\sum_i a_i)^2 - |\vec{a}|^2) & \text{if } \exists i : a_i = 0 \\ 1 + \frac{1}{n(n-1)} \sum_{\vec{d} \in P_n} T_n(\vec{a} + \vec{d}) & \text{otherwise} \end{cases} \tag{7.17}$$

and it is straightforward to check that this is satisfied by the formula of equation (7.16).

□

7.2.4. REMARK. Some calculation shows that

$$T_n^*(\vec{a}) = T_n(\vec{a}) - \frac{1}{2}H_n(\vec{a}) + \frac{1}{2}(\sum_i a_i)^2 \tag{7.18}$$

In other words, with respect to this other game, T_n is the expectation of the time *until* the first gambler goes broke, and $\frac{1}{2}((\sum_i a_i)^2 - H_n(\vec{a}))$ is the expectation of the time the game will last *after* the first gambler goes broke.

7.2.5. COROLLARY. *For all $\vec{a} \in I\!N^n$,*

$$0 \le T_n(\vec{a}) \le \frac{1}{2}((\sum_i a_i)^2 - |\vec{a}|^2) \tag{7.19}$$

7.2.6. COROLLARY. *For all $\vec{a} \in I\!N^n$,*

$$|\vec{a}|^2 \le H_n(\vec{a}) \le (a_1 + \ldots + a_n)^2 \tag{7.20}$$

7.3 Simple Rational Polynomials

Now, the formulas we have for T_2 and T_3 are rational polynomials. It would be nice to show that for $n \geq 4$, T_n *can't* be expressed as a rational polynomial. Although Lemma 7.4.1 is an effort in that direction, the proposition is still unproven. However, we can prove that T_n can't be expressed as a rational polynomial such that the denominator only depends on $a_1 + \ldots + a_n$:

7.3.1. THEOREM. *For all $n \geq 4$, T_n is not of the form*

$$T_n(\vec{a}) = \frac{P(a_1, \ldots, a_n)}{Q(a_1 + \ldots + a_n)} \tag{7.21}$$

for any two polynomials $P(a_1, \ldots, a_n)$ and $Q(s)$.

This theorem follows directly from the following two lemmas

7.3.2. LEMMA. *Suppose that, for some $n \geq 3$, T_n is as in equation (7.21), for some polynomials $P(a_1, \ldots, a_n)$ and $Q(s)$. Then there exists a nonzero polynomial $R(a_1, \ldots, a_{n-1})$ of degree $\geq n$, that is divisible by $a_1 \ldots a_{n-1}$ and $a_1 + \ldots + a_{n-1}$ and satisfies*

$$\sum_{\vec{d} \in P_n} D_{\vec{d}_{|n-1}} D_{\vec{d}_{|n-1}} R(\vec{a}) = 0 \tag{7.22}$$

where $D_{\vec{u}} f$ denotes the partial derivative of f in the direction $\vec{u}$, and $\vec{d}_{|n-1}$ denotes the vector obtained from $\vec{d}$ by discarding the final coordinate.

Proof
Without loss of generality we may assume that $P(a_1, \ldots, a_n)$ and $Q(s)$ have no common factors. Now, P and Q are polynomials, and $Q(a_1 + \ldots + a_n)$ only depends on the *total* amount of money, which does not change during the gambling. Combining these facts with the properties of T_n, we find that $P(a_1, \ldots, a_n)$ is divisible by $a_1 \ldots a_n$ and that for all $\vec{a} \in I\!R^n$

$$2n(n-1)Q(a_1 + \ldots + a_n) = \sum_{\vec{d} \in P_n} \left(P(\vec{a} + \vec{d}) + P(\vec{a} - \vec{d}) - 2P(\vec{a})\right) \tag{7.23}$$

We want to replace $Q(a_1 + \ldots + a_n)$ by 0, and $P(\vec{a} + \vec{d}) + P(\vec{a} - \vec{d}) - 2P(\vec{a})$ by the partial derivative $D_{\vec{d}} D_{\vec{d}} P(\vec{a})$. Both can be accomplished by taking the limit of an appropriate scaling, *provided* we can take $Q(a_1 + \ldots + a_n)$ to be constant. The easiest way to do this is to set $a_n = -(a_1 + \ldots + a_{n-1})$. So consider the polynomial

$$P^*(a_1, \ldots, a_{n-1}) = P(a_1, \ldots, a_{n-1}, -(a_1 + \ldots + a_{n-1})) \tag{7.24}$$

This polynomial is divisible by $a_1 \dots a_{n-1}$ and $a_1 + \dots + a_{n-1}$, and satisfies for $\vec{a} \in \mathbb{R}^{n-1}$

$$2n(n-1)Q(0) = \sum_{\vec{d} \in P_n} \left(P^*(\vec{a} + \vec{d}_{|n-1}) + P^*(\vec{a} - \vec{d}_{|n-1}) - 2P^*(\vec{a}) \right) \tag{7.25}$$

Now, if $P^*(a_1, \dots, a_{n-1})$ were the zero polynomial, then we would have $Q(0) = 0$, and then $P(a_1, \dots, a_n)$ and $Q(a_1 + \dots + a_n)$ would have a common factor $a_1 + \dots + a_n$, contradicting one of our starting assumptions. So $P^*(a_1, \dots, a_{n-1})$ is nonzero, and of total degree $d \geq n$. Let R be the uniformization of P, i.e. set

$$R(a_1, \dots, a_{n-1}) = \lim_{\delta \to 0} \delta^d P^*(\frac{a_1}{\delta}, \dots, \frac{a_{n-1}}{\delta}) \tag{7.26}$$

Then $R(a_1, \dots, a_{n-1})$ is a uniform polynomial[15] of total degree d, is divisible by $a_1 \dots a_{n-1}$ and $a_1 + \dots + a_{n-1}$, and satisfies

$$\begin{aligned}
&\sum_{\vec{d} \in P_n} D_{\vec{d}_{|n-1}} D_{\vec{d}_{|n-1}} R(\vec{a}) && (7.27)\\
&= \sum_{\vec{d} \in P_n} \lim_{\epsilon \to 0} \frac{1}{\epsilon^2} \left(R(\vec{a} + \epsilon\vec{d}_{|n-1}) + R(\vec{a} - \epsilon\vec{d}_{|n-1}) - 2S(\vec{a}) \right) && (7.28)\\
&= \sum_{\vec{d} \in P_n} \lim_{\epsilon \to 0} \lim_{\delta \to 0} \frac{\delta^d}{\epsilon^2} \left(P^*(\frac{\vec{a}}{\delta} + \frac{\epsilon}{\delta}\vec{d}_{|n-1}) + P^*(\frac{\vec{a}}{\delta} - \frac{\epsilon}{\delta}\vec{d}_{|n-1}) - 2P^*(\frac{\vec{a}}{\delta}) \right) && (7.29)\\
&= \lim_{\delta \to 0} \delta^{d-2} \sum_{\vec{d} \in P_n} \left(P^*(\frac{\vec{a}}{\delta} + \vec{d}_{|n-1}) + P^*(\frac{\vec{a}}{\delta} - \vec{d}_{|n-1}) - 2P^*(\frac{\vec{a}}{\delta}) \right) && (7.30)\\
&= \lim_{\delta \to 0} \delta^{d-2} 2n(n-1)Q(0) && (7.31)\\
&= 0 && (7.32)
\end{aligned}$$

Note that in the third equality, we are allowed to set $\epsilon = \delta$ and swap limits and sums, because the expression after the limit can be written as a polynomial in $a_1, \dots, a_{n-1}$, δ and ϵ.

□

7.3.3. Lemma. *For $n \geq 4$, there is no polynomial $R(a_1, \dots, a_{n-1})$, other than the zero polynomial, that is divisible by $a_1 \dots a_{n-1}$ and satisfies*

$$\sum_{\vec{d} \in P_n} D_{\vec{d}_{|n-1}} D_{\vec{d}_{|n-1}} R(\vec{a}) \text{ is constant} \tag{7.33}$$

Proof

Let $n \geq 4$ and let $R(a_1, \dots, a_{n-1})$ be a polynomial. Set

$$R^{(2)} = \frac{1}{2} \sum_{\vec{d} \in P_n} D_{\vec{d}_{|n-1}} D_{\vec{d}_{|n-1}} R(\vec{a}) \tag{7.34}$$

[15] R may be obtained from P^* by omitting all terms of total degree less than d

A bit of calculating shows that

$$R^{(2)} = (n-1) \sum_{1 \le i \le n-1} \frac{\partial^2 R}{\partial a_i^2} \quad - \sum_{1 \le i,j \le n-1, i \ne j} \frac{\partial^2 s}{\partial a_i \partial a_j} \tag{7.35}$$

Now suppose that R is divisible by $a_1, \ldots, a_{n-1}$, and that $R^{(2)}$ is constant. Since R contains no terms of total degree 2 or less, we immediately conclude that $R^{(2)} = 0$ everywhere. Now we shall prove by induction on m, that for all $m \in \mathbb{N}$ the following holds

> For all $i < j \le n-1$, R contains no terms such that the exponents of a_i and a_j in the term sum to m or less

Since R is divisible by $a_1, \ldots, a_{n-1}$, this is trivial for $m = 0, 1$. So let $m \ge 2$, let $1 \le i < j \le n-1$ and let P be any product of variables other than a_i and a_j. Set C to be the collection of terms of R such that the exponents of a_i and a_j sum to m and the remaining variables form P, i.e. C can be written as

$$\{c_k a_i^k a_j^{m-k} P \mid 0 \le k \le m\} \tag{7.36}$$

Each term of R contributes terms to $R^{(2)}$, as given by equation (7.35). Amongst the terms contributed by a term $c_k a_i^k a_j^{m-k} P$ of C are $(n-1)k(k-1)c_k a_i^{k-2} a_j^{m-k} P$, $-2k(m-k)c_k a_i^{k-1} a_j^{m-k-1} P$ and $(n-1)(m-k)(m-k-1)c_k a_i^k a_j^{m-k-2} P$. In these three terms of $R^{(2)}$, the exponents of a_i and a_j sum to $m-2$ and the remaining variables form P. It is easily seen that such terms are not contributed by any terms of R outside of C, since by the Induction Hypothesis, R does not contain any terms such that the exponents of a_i and a_j sum to $m-1$ or less. Since $R^{(2)} = 0$, this implies that for $k = 1, 2, \ldots, m-1$,

$$(n-1)(m-k+1)(m-k)c_{k-1} \; - \; 2k(m-k)c_k \; + \; (n-1)k(k+1)c_{k+1} \; = \; 0 \tag{7.37}$$

Multiplying by $(k-1)!(m-k-1)!$ and writing c_k^* for $(m-k)!k!c_k$ yields, for $k = 1, 2, \ldots, m-1$,

$$(n-1)c_{k-1}^* - 2c_k^* + (n-1)c_{k+1}^* = 0 \tag{7.38}$$

Furthermore, since a_i and a_j divide R, $c_0 = c_m = c_0^* = c_m^* = 0$. This system of linear equations has a nontrivial solution if and only if $\det B_{m-1}^{n-1} = 0$, where B_{m-1}^{n-1} is the $(m-1) \times (m-1)$ matrix

$$\begin{pmatrix} -2 & n-1 & 0 & 0 & \ldots & 0 & 0 \\ n-1 & -2 & n-1 & 0 & \ldots & 0 & 0 \\ 0 & n-1 & -2 & n-1 & \ldots & 0 & 0 \\ 0 & 0 & n-1 & -2 & \ldots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \ldots & -2 & n-1 \\ 0 & 0 & 0 & 0 & \ldots & n-1 & -2 \end{pmatrix} \tag{7.39}$$

Defining the matrices B_i^{n-1} analogously, we get the recursive equation

$$\det B_i^{n-1} = -2 \det B_{i-1}^{n-1} - (n-1)^2 \det B_{i-2}^{n-1} \tag{7.40}$$

with $\det B_0^{n-1} = 1$, $\det B_1^{n-1} = -2$. Now, if $n > 3$, then either $4|n-1$, or $p|n-1$ for some prime number $p > 2$. In the first case, it is easy to show by induction on i that for all i, $\det B_i^{n-1}$ is a multiple of 2^i but not of 2^{i+1}. In the second case, it is easy to show by induction on i that for all i, $\det B_i^{n-1}$ is not a multiple of p. In both cases, $\det B_i^{n-1} \neq 0$ for all $i \geq 0$. In particular, $\det B_{m-1}^{n-1} \neq 0$. [16]
It follows that all coefficients of terms of R in C are 0. We conclude that R contains no terms such that for some i and j, the exponents of a_i and a_j in the term sum to m or less. Since this holds for all $m \geq I\!N$, R is the zero polynomial.

□

7.3.4. Remark. The proof of Lemma 7.3.3 also goes through, with only minor modifications, for the case where R can be written as a power series. However, scaling in the manner of Lemma 7.3.2 is not generally possible for analytic functions. When it is, i.e. when for some $d \geq 2$ and some $s \in I\!R$, the function

$$R(a_1, \ldots, a_{n-1}) := \lim_{\delta \to 0} \delta^d P(\frac{a_1}{\delta}, \ldots, \frac{a_{n-1}}{\delta}, \frac{-a_1 - \ldots - a_{n-1}}{\delta}) \tag{7.41}$$

can be written as a power series in $a_1, \ldots, a_{n-1}$ and equation (7.30) holds, then Lemma 7.3.3 can be applied. But aside from the polynomials there appear to be very few functions for which this is the case. For instance, if P is a rational polynomial, the denominator of R is 0 for $\vec{a} = \vec{0}$, and then R cannot be written as a power series in $a_1, \ldots, a_{n-1}$.

Unfortunately, when we try to adapt this proof to rational polynomials in general, we run into a number of problems. For instance, the formula for the second-order derivatives of the quotient of two functions involves taking products of the numerator, the denominator and their derivatives, and as a result the equations that the coefficients of the polynomials must satisfy are no longer *linear* equations.

7.4 General Rational Polynomials

Now, the following lemma *may* be useful in order to prove that T_n cannot be expressed as a rational polynomial, by giving a consequence of this premise which seems refutable. Unfortunately, there appear to be no general results in this area

[16] From the recursive equation we can derive an explicit formula, which turns out to be $B_i^{n-1} = (n-1)^i \sin(i+1)\alpha / \sin\alpha$, where $0 \leq \alpha \leq \pi$ and $\cos\alpha = -1/(n-1)$. Hence this result is distantly related to the fact that the only rational $c \in \mathbb{Q}$ such that $\cos\beta = c$ for some rational multiple β of π, are -1,-0.5,0,0.5 and 1.

that could be used to refute it. Please note that it is not a complete reduction of the problem: the implication is one-way. In particular, if a function h_n is found with the properties given below, it is not at all clear how to obtain T_n from h_n.

7.4.1. Lemma. *Suppose that $T_n(a_1, \ldots, a_n)$ can be expressed as a rational polynomial (the quotient of two polynomials). Then there exists a rational polynomial function h_n on an n-simplex $S \subset \mathbb{R}^{n-1}$, such that*

1. *h_n is harmonic on S (i.e. the Laplacian Δh_n is 0 on S).*
2. *for some point $\vec{m}$, $h_n(\vec{x}) = d^2(\vec{x}, \vec{m})$ for all $x \in \partial S$.*

Proof

Suppose that T_n can be expressed as a rational polynomial. Then so can H_n. Furthermore, whenever the denominator of T_n is non-zero, H_n is defined and satisfies

$$\sum_{\vec{d} \in P_n} (H_n(\vec{a} + \vec{d}) + H_n(\vec{a} - \vec{d}) - 2H_n(\vec{a})) = 0 \tag{7.42}$$

and

$$H_n(\vec{a}) = |\vec{a}|^2 \text{ if } \exists i : a_i = 0 \tag{7.43}$$

As in the proof of Theorem 7.3.1, we will substitute partial derivatives for their discrete counterparts, by taking the limit of an appropriate scaling. From Lemma 7.2.6 it follows that the total degree of the numerator is exactly 2 more than the total degree of the denominator. Let R be the uniformization of H_n, i.e. set

$$R(\vec{a}) = \lim_{\delta \to 0} \delta^2 H_n(\vec{a}/\delta) \tag{7.44}$$

on $\mathbb{R} \times \mathbb{R}^n$. R is a uniform rational polynomial[17] in $a_1, \ldots, a_n$ satisfying

$$R(\vec{a}) = |\vec{a}|^2 \text{ if } \exists i : a_i = 0 \tag{7.45}$$

and

$$|\vec{a}|^2 \leq R(\vec{a}) \leq (a_1 + \ldots + a_n)^2 \text{ if } \vec{a} \in R^n_{\geq 0} \tag{7.46}$$

In particular, R exists on $\mathbb{R}^n_{\geq 0}$. Furthermore, since R is a rational polynomial, the partial derivatives $D_d D_d \bar{R}(\vec{a})$ exist whenever $R(\vec{a})$ exists, and

$$\sum_{\vec{d} \in P_n} D_{\vec{d}} D_{\vec{d}} R(\vec{a}) \tag{7.47}$$

$$= \sum_{\vec{d} \in P_n} \lim_{\epsilon \to 0} \frac{1}{\epsilon^2} \left(R(\vec{a} + \epsilon\vec{d}) + R(\vec{a} - \epsilon\vec{d}) - 2R(\vec{a}) \right) \tag{7.48}$$

[17] R may be obtained from H_n by removing from the numerator and the denominator all terms of less than maximal total degree.

$$= \sum_{\vec{d}\in P_n} \lim_{\epsilon\to 0}\lim_{\delta\to 0}\frac{\delta^2}{\epsilon^2}\left(H_n(\frac{\vec{a}}{\delta}+\frac{\epsilon}{\delta}\vec{d}) + H_n(\frac{\vec{a}}{\delta}-\frac{\epsilon}{\delta}\vec{d}) - 2H_n(\frac{\vec{a}}{\delta})\right) \quad (7.49)$$

$$= \lim_{\delta\to 0}\sum_{\vec{d}\in P_n}\left(H_n(\frac{\vec{a}}{\delta}+\vec{d}) + H_n(\frac{\vec{a}}{\delta}-\vec{d}) - 2H_n(\frac{\vec{a}}{\delta})\right) \quad (7.50)$$

$$= 0 \quad (7.51)$$

Note that in the second-to-last equality, we are allowed to set $\epsilon = \delta$ and swap limits and sums, because the expression after the limit can be written as a rational polynomial in $a_1, \ldots, a_n$, δ and ϵ whose denominator is non-zero for $\delta = \epsilon = 0$ (provided the denominator of $R(\vec{a})$ is non-zero in the first place).
If we consider the simplex $S = \{\vec{a} \in I\!R^n : \sum_i a_i = 1, \forall i : a_i \geq 0\}$ in the hyperplane $V = \{\vec{a} \in I\!R^n : \sum_i a_i = 1\}$, then (after mapping V to $I\!R^{n-1}$) the function $h_n(\vec{a}) = R(\vec{a}) - 1/n$ satisfies the requirements, taking m to be the center of the simplex.

□

7.4.2. Remark. For $n = 2, 3$ we have

$$h_2(a_1, a_2) = 1 \text{ for } a_1 + a_2 = 1 \quad (7.52)$$

$$h_3(a_1, a_2, a_3) = a_1^2 + a_2^2 + a_3^2 + 6a_1a_2a_3 \text{ for } a_1 + a_2 + a_3 = 1 \quad (7.53)$$

Chapter 8

Reinforced Random Walks

In the orthodox random walk, the probability of traversing a specific street from a specific intersection is always the same, unaffected by anything that has gone before. In this chapter, we will study *reinforced* random walks, where the walk is given a particular kind of *feedback* such that edges already traversed are either more or less likely to be traversed in the future. In terms of the Drunkard's Walk example, the drunkard vaguely recognizes streets he has walked before, and is either more likely to traverse them (as he considers them safe) or less likely (as he considers them boring), depending on the conditions of the reinforcement.

Reinforced random walks were first introduced by Diaconis and Coppersmith[22], and generalized later by B. Davis[23] and Pemantle[29]. They were originally presented as an alternative to Pólya's urn[18]as a simplified model of a self-organizing system, i.e. a system whose basic parameters are very simple, and whose behavior 'evolves' to approach a (possibly random) limit. Such systems occur naturally, for instance in the formation of stalactites and stalagmites. For another example, consider a man who has just moved to a new city: if he does not know the shops, he will start out by visiting shops at random, but after a while he will develop preferences and habits.

8.1 General Reinforced Random Walks

First we will compare reinforced walks with non-reinforced random walks, give analogues of results and techniques from Chapter 6, and show that under some

[18] *Pólya's urn* is one of the simplest (and oldest) processes with reinforcement. In this model, there is an urn containing red and blue balls. At time $t = 0$, the urn contains one red and one ball. At each time $t > 0$, a ball is chosen uniformly from the contents of the urn, and is put back into the urn along with another ball of the same color. Eggenberger and Pólya[30] showed that the proportion of red balls converges almost surely, and that the limit is random with uniform distribution on $[0, 1]$.

very general conditions reinforced random walks on trees are almost surely recurrent. Then we will give a sufficient condition for recurrence of reinforced random walks on general graphs, which we will use in later sections.
In a reinforced random walk, when an edge has been traversed we change the probability that it will be traversed again, by increasing or decreasing the *weight* of the edge. In general reinforced random walks, the new weight may depend on many things, such as the edge in question, the number of times it has been traversed before, the time of traversal and the pattern formed by edges traversed at previous times, etc. etc. B. Davis [23] defines the category of reinforced random walks of *matrix* type, where for each edge vu, the current weight of vu is determined solely by the number of times $k_t(vu)$ it has been traversed up to then, and is not influenced by anything that has happened to any other edge. Note that in general walks of matrix type, the relationship between current weight and number of traversals may be different for each edge. In this chapter and the next we concern ourselves with a specific subclass of walks of matrix type, where a sequence $(\delta_k)_{k\in\mathbb{N}}$ is given which is the same for all edges, and the current weight of an edge at any given time is determined by multiplying its original weight by $\delta_{k_t(vu)}$. A formal definition:

8.1.1. Definition. Let $(\delta_k)_{k\in\mathbb{N}}$ be a sequence of strictly positive real numbers, Set the *weight of vu at time t* to

$$w_t(vu) = \delta_{k_t(vu)} w(vu) \tag{8.1}$$

where $k_t(vu)$ denotes the number of traversals of vu up to time t, i.e.

$$k_t(vu) = \#\{t' < t \mid v_{t'}v_{t'+1} = vu\} \tag{8.2}$$

A *reinforced random walk* on a graph $G = (V, E, w)$ with reinforcement sequence $(\delta_k)_{k\in\mathbb{N}}$, is a series of stochastic variables $v_0, v_1, \ldots \in V$ such that for all $t \in \mathbb{N}$,

$$P(v_{t+1} = u \mid \mathcal{F}_t) = \begin{cases} \frac{w_t(v_t u)}{w_t(v_t)} & \text{if } u \in N(v_t) \\ 0 & \text{otherwise} \end{cases} \tag{8.3}$$

Recurrence and transience are defined in the same manner as before.

8.1.2. Remark. The random walks defined above are similar but not identical to B. Davis' random walks of *sequence* type, where the current weight of an edge vu is defined as $w_t(vu) = w_0(vu) + \delta_{k_t(vu)}$ for some non-descending sequence $(\delta_k)_{k\in\mathbb{N}}$ [23]. Davis gave many results for walks of this type on the linear lattice $\mathbb{Z}$, most of which also hold for the random walks defined above. We will concern ourselves mainly with more general classes of graphs.

There are a number of differences between a reinforced and a non-reinforced random walk. For instance, a reinforced random walk is influenced by its history, and hence we might want to consider random walks with initial states in which some edges are considered to have been traversed already. Another difference is that if the reinforcement increases sharply enough, the random walk might get 'stuck' on an edge:

8.1.3. THEOREM. *Let $G = (V, E, w)$ be a weighted graph, and let $(\delta_k)_{k\in\mathbb{N}}$ be such that*

$$\sum_{k=0}^{\infty} \frac{1}{\delta_k} \text{ converges.} \tag{8.4}$$

Then for any edge $vu \in E$, and all random walks on G, there exists a $t_0 \geq 0$ such that

$$P(\forall t > t_0 : v_t \in \{v, u\}) > 0 \tag{8.5}$$

Proof

Since G is connected, every point is reachable, and hence there exists a $t_0 \in \mathbb{N}$ such that with non-zero probability v is visited at time t_0. Assume that it has. Then the probability that from time t_0 on, the random walk will keep traveling from v to u and back again, is

$$\begin{aligned}
& \prod_{i=0}^{\infty} \frac{w_{t_0+i}(vu)}{w_{t_0+i}(v_{t_0+i})} && (8.6)\\
\geq\ & \prod_{i=0}^{\infty} \frac{\delta_{k_{t_0}(vu)+i}\, w(vu)}{c + \delta_{k_{t_0}(vu)+i}\, w(vu)} && (8.7)\\
\geq\ & \prod_{k=k_{t_0}(vu)}^{\infty} e^{-c/(\delta_k w(vu))} && (8.8)\\
=\ & e^{-c/w(vu)\cdot\sum_{k=k_{t_0}(vu)}^{\infty}(1/\delta_k)} && (8.9)\\
>\ & 0 && (8.10)
\end{aligned}$$

where c is the total weight assigned at time t to edges other than vu that are incident with v or u.

□

The converse implication, that if $\sum_{k\in\mathbb{N}} \delta_k$ diverges, the random walk will almost surely not get 'stuck', does not hold in general.[19]However, it *does* hold for non-descending sequences, and for general sequences it is possible to come close, as the following analogues of Lemmas 6.1.4 and 6.1.5 show:

[19]For instance, if G is a tree with unit weights on which non-reinforced random walks are almost surely recurrent, then it can be shown that the reinforced random walk on G with reinforcement sequence $(\delta_k)_{k\in\mathbb{N}} = (1, 2, 1, 4, 1, 8, 1, 16, \ldots)$ starting from a vertex v_0 almost surely eventually stays within $\{v_0\} \cup N(v_0)$.

8.1.4. Lemma. *Let $G = (V, E, w)$ be a weighted graph, and let $(\delta_k)_{k\in \mathbb{N}}$ be such that*

$$\sum_{j=0}^{\infty} \frac{1}{\max(\delta_0, \delta_1, \ldots, \delta_j)} \text{ diverges.} \tag{8.11}$$

Then a reinforced random walk on G starting from any initial state will almost surely visit infinitely many vertices, and

$$P(\text{the walk is transient}) + P(\text{the walk is recurrent}) = 1 \tag{8.12}$$

Proof

The first assertion follows from the second, since both transient and recurrent walks visit infinitely many vertices. To prove the second assertion it suffices to show that for all $v, u \in V$

$$P(v \text{ is visited infinitely often and } u \text{ only finitely often}) = 0 \tag{8.13}$$

If we can show that the above holds for vertices $v, u \in V$ with $vu \in E$, then the general result follows by induction on the distance $d_G(v, u)$. So let $v, u \in V$ with $vu \in E$. Fix $t_0 \in \mathbb{N}$, and suppose that u has not been visited since time $t_0 \in \mathbb{N}$, and at some time $t > t_0$ v is visited again for the k-th time. Then $w_t(v) \leq w(v)\max\{\delta_0, \ldots, \delta_{2k}\}$, and since vu has been traversed at most t_0 times, $w_t(vu) \geq w(vu)\min\{\delta_0, \ldots, \delta_{t_0}\}$. Hence, the probability of *not* immediately traversing vu in this situation is at most

$$1 - c/\max\{\delta_0, \ldots, \delta_{2k}\} < e^{-c/\max\{\delta_0,\ldots,\delta_{2k}\}} \tag{8.14}$$

where $c = w(vu)\min\{\delta_0, \delta_1, \ldots, \delta_{t_0}\}/w(v)$.
Therefore, applying induction on k, we have that for all $k \geq 1$,

$$P(u \text{ is not visited between } t_0 \text{ and the } k+1\text{-th visit to } v) \tag{8.15}$$

$$\leq \prod_{k'=1}^{k} e^{-c/\max(\delta_0,\ldots,\delta_{2k'})} \tag{8.16}$$

$$= e^{-c\cdot\sum_{k'=1}^{k}(1/\max(\delta_0,\ldots,\delta_{2k'}))} \tag{8.17}$$

Consequently

$$P(v \text{ is visited infinitely often and } u \text{ never after time } t_0)$$

$$\leq e^{-c\sum_{k=1}^{\infty}(1/\max(\delta_0,\ldots,\delta_{2k}))} \tag{8.18}$$

$$= 0 \tag{8.19}$$

Summing over all times $t_0 \in \mathbb{N}$ gives the desired result.

□

8.1.5. LEMMA. *Let $G = (V, E, w)$ be a weighted graph, $F \subset V$ a finite set of vertices of G, and $v_0 \in V$. Let $(\delta_k)_{k \in I\!N}$ be such that equation (8.11) holds. Then for the reinforced random walk on G starting from v_0, the following are equivalent:*

(i) *The reinforced random walk on G starting from v_0 is almost surely recurrent.*

(ii) *For any $t_0 \in I\!N$, and any history up to time t_0, F will be (re)visited at some time at or after time t_0 almost surely.*

Proof

$(i) \Rightarrow (ii)$ is trivial. If (ii) holds, then by applying it repeatedly we find that the reinforced random walk on G starting from v_0 will almost surely visit F infinitely often. Then the random walk is almost surely not transient, and by the previous Lemma, this implies it is almost surely recurrent.

□

8.1.6. REMARK. In condition (ii) of Lemma 8.1.5, conceptually we *restart* the walk at time t_0. i.e. we look at a walk which starts at time t_0, with t_0 traversals part of a 'fixed' history up to time t_0 (as opposed to starting at time 0 with a blank initial state). If all such restarted walks can be shown to visit F almost surely, Lemma 8.1.5 states that the original reinforced random walk is almost surely recurrent.

8.1.7. LEMMA. *For random walks on weighted trees, the direction in which an edge is traversed is the same at all odd-numbered traversals (and opposite to the direction of traversal at all even-numbered traversals). This allows us to replace, for reinforced random walks on weighted trees, the condition of Lemmas 8.1.4 and 8.1.5 by the condition that*

$$\sum_{k=0}^{\infty}(1/\delta_{2k}) \text{ and } \sum_{k=0}^{\infty}(1/\delta_{2k+1}) \text{ both diverge.} \tag{8.20}$$

Proof

Consider a random walk on a weighted tree $G = (V, E, w)$, and assume that equation (8.20) holds. In order to show that the conclusions of Lemmas 8.1.4 and 8.1.5 hold, it suffices to show that for all vertices $v, u \in V$ with $vu \in E$,

$$P(v \text{ is visited infinitely often and } u \text{ only finitely often}) = 0 \tag{8.21}$$

So let $v \in V$, and let $u^0, u^1, \ldots, u^m$ be the neighbors of v in G, with u^0 being the unique neighbor of v that is on a path between v and v_0 if $v \neq v_0$. Set, for $i \leq m, k \in I\!N$,

$$R_k^i = \delta_{2k+1} w(vu^i) \text{ if } i = 0 \text{ and } v \neq v_0, \;\; R_k^i = \delta_{2k} w(vu^i) \text{ otherwise} \tag{8.22}$$

Then R_k^i is the weight of the edge vu^i if v is visited and the arc vu^i has been traversed (in that direction) k times before.
The next part of the proof is based on a proof of H. Rubin concerning a generalized Pólya Urn problem [23]. Let Y_k^i be independent exponential random variables such that $E(Y_k^i) = 1/R_k^i$,[20] and put

$$A^i = \left\{ \sum_{k'=0}^{k} Y_{k'}^i, k' \geq 0 \right\} \text{ for } i \leq m \tag{8.23}$$

□

Define a sequence of edges vu^i by making the k-th element of the sequence vu^i if the k-th smallest element of $A_0 \cup \ldots \cup A_m$ is from A_i. Now since by equation (8.20)

$$\text{for all } i \leq m, \sum_{k=0}^{\infty} \frac{1}{R_k^i} \text{ diverges.} \tag{8.24}$$

we have that almost surely

$$\text{for all } i \leq m, \sum_{k=0}^{\infty} Y_{k'}^i \text{ diverges.} \tag{8.25}$$

and hence almost surely vu^i will appear infinitely often in the sequence for all $i \leq m$.
As it turns out, this sequence has exactly the same probability distribution as the sequence of edges traversed from v in the reinforced random walk. In other words, we may decide that at visits to v we traverse successive arcs of the sequence, without changing any probabilities. The proof of this relies on properties of exponential random variables, and is straightforward but cumbersome. Interested readers are referred to Rubin's proof [23]. We conclude that equation (8.21) holds.

□

8.2 A Martingale for Reinforced Random Walks

Now let us consider recurrence for reinforced random walks. The proofs given in Chapter 6 used the fact that, if a function h on the vertex-set of a weighted graph G is harmonic, then in a *non-reinforced* random walk, $h(v_t)$ behaves like a martingale. This does not in general hold for *reinforced* random walks. If h is a harmonic function, then a vertex which has neighbors with higher h-values will also have neighbors with lower h-values, but the probabilities of the corresponding edges being traversed are not necessarily balanced, or even constant over time. In order to find an analogue of Lemma 6.2.6, we will need to compensate for the difference in probabilities.

[20] I.e. the probability distribution of Y_k^i is given by $P(Y_k^i > r) = e^{-rR_k^i}$ for all $r \in \mathbb{R}$

8.2.1. LEMMA. *Let $G = (V, E, w)$ be a weighted graph, and let $h : V \to \mathbb{R}$ be a harmonic [superharmonic, subharmonic] function on a subset $V' \subset V$. Consider the reinforced random walk with reinforcement sequence $(\delta_k)_{k\in\mathbb{N}}$ and define*

$$M_t = \sum_{t'=0}^{t} \begin{cases} \frac{\Delta_h(v_{t'}v_{t'+1})}{\delta_{k_{t'}(v_{t'}v_{t'+1})}} & \text{if } v_{t'} \in V' \\ 0 & \text{otherwise} \end{cases} \tag{8.26}$$

for $t \in \mathbb{N}$, where (as before) $\Delta_h(\overrightarrow{vu})$ denotes $h(u)-h(v)$. Then M is a martingale [supermartingale,submartingale].

Proof
If $v_t \in V - V'$, then $M_{t+1} = M_t$, otherwise

$$M_t = [\geq, \leq] \quad M_t + \frac{1}{w_t(v_t)} \cdot \sum_{u\in N(v_t)} w(v_t u)\Delta_h(\overrightarrow{v_t u}) \tag{8.27}$$

$$= \quad M_t + \sum_{u\in N(v_t)} \frac{w_t(v_t u)}{w_t(v_t)} \cdot \frac{\Delta_h(\overrightarrow{v_t u})}{\delta_{k_t(v_t u)}} \tag{8.28}$$

$$= \quad M_t + \sum_{u\in N(v_t)} P(v_{t+1} = u \mid \mathcal{F}_t) \frac{\Delta_h(\overrightarrow{v_t u})}{\delta_{k_t(v_t u)}} \tag{8.29}$$

$$= \quad E(M_{t+1} \mid \mathcal{F}_t) \tag{8.30}$$

□

As an application of the above martingale, we will show that if non-reinforced random walks on a weighted tree are almost surely recurrent, then for reinforced random walks on that tree, a very weak condition on the reinforcement sequence suffices to show recurrence.

8.2.2. THEOREM. *Let $G = (V, E, w)$ be a weighted tree, with the property that non-reinforced random walks on G are almost surely recurrent. Let $(\delta_k)_{k\in\mathbb{N}}$ be a non-descending reinforcement sequence that satisfies the condition of Lemma 8.1.5 (or that of Lemma 8.1.7). Furthermore, assume either that $(\delta_k)_{k\in\mathbb{N}}$ is bounded, or that $\delta_{k+1} > \delta_k$ for some even $k \in \mathbb{N}$. Then the reinforced random walk with reinforcement sequence $(\delta_k)_{k\in\mathbb{N}}$ is almost surely recurrent.*

Proof
Consider a reinforced random walk on G starting from some vertex $v_0 \in V$. By Lemma 8.1.5 (or Lemma 8.1.7), to show recurrence, it suffices to show for all $t_0 \in \mathbb{N}$, and any history up to time t_0, that v_0 will be revisited almost surely at some time at or after time t_0. So let $t_0 \in \mathbb{N}$, and fix the history up to time t_0.
First, we need a function h on V that is superharmonic on $V - \{v_0\}$. Since non-reinforced random walks on G are almost surely recurrent, such a function h exists by Theorem 6.3.2. For $r \in \mathbb{R}$, define the stopping time τ_r as the first

time $t \geq t_0$ at which $v_t = v_0$ or $h(v_t) > r$. By Lemma 8.1.4, the random walk will almost surely leave the finite set of vertices $\{v \in V \mid h(v) \leq r\}$. Hence $\tau_r < \infty$ almost surely.

Next, let M_t be the martingale of Lemma 8.2.1. For walks on weighted trees, the direction of traversal of an edge is the same for all odd-numbered traversals, and opposite to the direction for all even-numbered traversals. Furthermore, all odd-numbered traversals are traversals going from the lower to higher h-value, for otherwise it would be possible to construct an infinite sequence of vertices of decreasing h-value, which would contradict the fact that $h \to \infty$ if $v \to \infty$. Hence, an edge vu which has been traversed k times at time t contributes

$$|\Delta_h(\overrightarrow{vu})| \cdot \sum_{j=0}^{k-1} \begin{cases} 1/\delta_j & \text{if } j \text{ is even and } vu \text{ is not incident with } v_0 \\ 0 & \text{if } j \text{ is even and } vu \text{ is incident with } v_0 \\ -1/\delta_j & \text{if } j \text{ is odd} \end{cases} \tag{8.31}$$

to the value of the martingale. Now by the conditions on $(\delta_k)_{k \in \mathbb{N}}$, there exists a $c > 0$ such that either $1/\delta_k > c$ for all $k \in \mathbb{N}$, or $1/\delta_k - 1/\delta_{k+1} > c$ for some even $k \in \mathbb{N}$. We can use either property, together with the monotonicity of $(\delta_k)_{k \in \mathbb{N}}$, to obtain the following lower bound on the above contribution:

$$|\Delta_h(\overrightarrow{vu})| \cdot \left(\begin{cases} c & \text{if } k \text{ is odd} \\ 0 & \text{if } k \text{ is even} \end{cases} - \begin{cases} \lceil k/2 \rceil / \delta_1 & \text{if } vu \text{ is incident with } v_0 \\ 0 & \text{otherwise} \end{cases} \right) \tag{8.32}$$

At any time t, the edges of G that have been traversed an odd number of times are *exactly* the edges of the unique path in G between v_0 and v_t. Furthermore, between times t_0 and τ, there will be no traversals of edges incident with v_0, except for a possible traversal *to* v_0 at time τ. Hence the martingale M_t satisfies

$$M_t \geq c(h(v_t) - h(v_0)) - c' \tag{8.33}$$

where $c' = \sum_{u \in N(v_0)} \Delta_h(v_0 u) \lceil k_{t_0}(v_0 u)/2 \rceil / \delta_1$. Now we can apply the Optional Stopping Times Theorem to obtain

$$M_{t_0} \geq E(M_{\tau_r}) \geq (1 - P(v_{\tau_r} = v_0)) c (r - h(v_0)) - c' \tag{8.34}$$

We conclude that $P(v_{\tau_r} = v_0) \geq 1 - (M_{t_0} + c')/c(r - h(v_0))$ for all $r > h(v_0)$, and hence v_0 is almost surely revisited at some time after time t_0.

□

8.2.3. Remark. For the proof of the above theorem, we can weaken the conditions on the reinforcement sequence to the conditions of Lemma 8.1.7 and, for some $c > 0$,[21] the inequality

$$\sum_{j=0}^{k-1} (-1)^j/\delta_k > 0 \text{ for } k \text{ even }, > c \text{ for } k \text{ odd} \tag{8.35}$$

8.3 Once-Reinforced Random Walks

In this section we will consider the once-reinforced random walk, where the weight of an edge only changes the *first* time it is traversed, and afterwards remains constant. For this walk, the martingale M_t defined in the previous section can be expressed as $h(v_t)$ plus a certain (bounded) *bias*. If the expectation of the bias is small enough, we will be able to show recurrence in a similar manner as in Chapter 6.

8.3.1. DEFINITION. Let $\delta > 0$. The *once-reinforced* random walk with *reinforcement factor* δ is the reinforced random walk with reinforcement sequence

$$(\delta_k)_{k=0}^{\infty} = (1, \delta, \delta, \delta, \delta, \ldots) \tag{8.36}$$

8.3.2. DEFINITION. Define the stochastic variables E_t and A_t, for $t \in \mathbb{N}$, by setting

$$E_t = \{v_s v_{s+1} \mid s < t\} \tag{8.37}$$

$$A_t = \{\overrightarrow{vu} \mid vu \in E_t, \overrightarrow{vu} = \overrightarrow{v_s v_{s+1}} \text{ for } s = \min\{s' < t \mid v_{s'} v_{s'+1} = vu\}\} \tag{8.38}$$

i.e. E_t is an edge-set containing the edges that have been traversed up to time t, and A_t is an arc-set obtained from E_t by orienting each edge in the direction that it was first traversed.

8.3.3. LEMMA. *In a once-reinforced random walk with reinforcement factor $\delta > 0$, let $t_0 \in \mathbb{N}$, and let M_t be as in Lemma 8.2.1 for some function $h : V \to \mathbb{R}$ which is (super/sub)harmonic on $V' \subset V$. Then for $t \geq t_0$,*

$$\delta(M_t - M_{t_0}) = h(v_t) - h(v_{t_0}) + (\delta - 1) \sum_{\overrightarrow{vu} \in A_t - A_{t_0}} \Delta_h(\overrightarrow{vu}) \tag{8.39}$$

as long as $V - V'$ has not been visited at any time between t_0 and t (including t_0 and excluding t).

[21] Regrettably, the constant $c > 0$ cannot be replaced by 0. We can find a counterexample on the linear lattice graph $G = (V, E, w)$ with $V = \{v^n \mid n \in \mathbb{N}\}$, $E = \{v^n v^{n+1} \mid n \in \mathbb{N}\}$ and $w(v^n v^{n+1}) = n + 1$. Non-reinforced random walks on this graph are recurrent, but the reinforced random walk with reinforcement sequence $(\delta_k)_{k \in \mathbb{N}} = (1, 1, 2, 2, 3, 3, \ldots)$ starting in v_0 is almost surely transient.

Proof

At time $t = t_0$, the equality holds. If an arc $\overrightarrow{vu}$ is traversed that has been traversed before, then M_t changes by $\Delta_h(\overrightarrow{vu})/\delta$, $h(v_t)$ changes by $\Delta_h(\overrightarrow{vu})$, and A_t does not change, so equality is preserved. If an arc vu is traversed that has not been traversed before, then M_t changes by $\Delta_h(\overrightarrow{vu})$, $h(v_t)$ changes by $\Delta_h(\overrightarrow{vu})$, and $\overrightarrow{vu}$ is added to A_t, so equality is again preserved.

□

Now, in our proof of the recurrence of non-reinforced random walks, a key point was that when we moved farther away from F, the value of the martingale increased as well. Since the expectation of the martingale was bounded, this implied that the probability of reaching a border decreased if we moved the border further away. In order to use similar reasoning here, we will need the bias $(\delta-1)\sum_{\overrightarrow{vu}\in A_t}\Delta_h(\overrightarrow{vu})$ to be positive in the long run, or at least not *too* negative.

8.3.4. LEMMA. *Let $G = (V, E, w)$ be a weighted graph. Let $h : V \to \mathbb{R}$ be a function satisfying*

1. *h is superharmonic everywhere except on a finite subset $F \subset V$.*
2. *h goes to infinity if v goes to infinity.*

Consider the once-reinforced random walk on G with reinforcement factor δ starting at some vertex v_0. Suppose that for some $\epsilon > 0$, the following holds for any time t_0 and any history up to time t_0:

> *There exists a $c \in \mathbb{R}$ such that for all $r_0 \in \mathbb{R}$ we can find $r > r_0$ with*
>
> $$(\delta-1)E\left(\sum_{\overrightarrow{vu}\in A_{\tau_r}}\Delta_h(\overrightarrow{vu})\,\middle|\,\mathcal{F}_{t_0}\right) \geq -(1-\epsilon)r - c \qquad (8.40)$$
>
> *(where the stopping time τ_r is the first time at or after t_0 that F is visited or $h(v_t) \geq r$).*

Then the once-reinforced random walk on G with reinforcement factor δ starting at v_0 is almost surely recurrent.

Proof

Without loss of generality we may assume that $h \geq 0$. Note that the reinforcement sequence satisfies the condition of Lemmas 8.1.4 and 8.1.5. Therefore it suffices to show for all $t_0 \in \mathbb{N}$, and any history up to time t_0, that F will be revisited almost surely at some time at or after time t_0. So let $t_0 \in \mathbb{N}$, and fix the history up to time t_0.

Let M_t be the supermartingale of Lemma 8.2.1, and let $r \in \mathbb{R}$. For any $t \leq \tau_r$, the set $A_t - A_{t_0}$ is contained in the finite set $\{\overrightarrow{vu} \in V \mid v \notin F \wedge h(v) < r\}$. So

M_t is bounded for $t \le \tau_r$, and furthermore $\tau_r < \infty$ almost surely by Lemma 8.1.4. Hence we can apply the Optional Stopping Times Theorem to obtain $E(\delta M_{\tau_r}) \le \delta M_{t_0}$, which by Lemma 8.3.3 is equivalent to

$$E(h(v_{\tau_r})) \le h(v_{t_0}) + (\delta-1)\sum_{\overrightarrow{vu}\in A_{t_0}} \Delta_h(\overrightarrow{vu}) - (\delta-1)E\left(\sum_{\overrightarrow{vu}\in A_{\tau_r}} \Delta_h(\overrightarrow{vu})\right) \quad (8.41)$$

Combining this with the formula $E(h(v_{\tau_r})) \ge (1 - P(v_{\tau_r} \in F \mid \mathcal{F}_{t_0}))r$, we obtain

$$P(v_{\tau_r} \in F \mid \mathcal{F}_{t_0}) \ge 1 - \frac{h(v_{t_0})}{r} - \frac{\delta-1}{r}\sum_{\overrightarrow{vu}\in A_{t_0}} \Delta_h(\overrightarrow{vu}) + \frac{\delta-1}{r}E\left(\sum_{\overrightarrow{vu}\in A_{\tau_r}} \Delta_h(\overrightarrow{vu})\right) \quad (8.42)$$

By assumption we can find $c, r \in I\!R$ such that

$$\frac{\epsilon}{2}r > c + h(v_{t_0}) + (\delta-1)\sum_{\overrightarrow{vu}\in A_{t_0}} \Delta_h(\overrightarrow{vu}) \quad (8.43)$$

and (8.40) holds. Then

$$P(v_{\tau_r} \in F \mid \mathcal{F}_{t_0}) \ge 1 - \epsilon/2 + \frac{c}{r} - \frac{(1-\epsilon)r + c}{r} = \epsilon/2 \quad (8.44)$$

So there is at least a chance of $\epsilon/2$ of coming back to F at time $t = \tau_r$. In the event that this does not happen, we repeat the entire process starting at time $\tau_r + 1$, and each time we have a chance of $\epsilon/2$ of visiting F. It follows that the random walk will visit F almost surely.

□

Next are some applications of this lemma. We will write the bias as the sum of 'local' biases in order to estimate it. The first application demonstrates how to use absolute bounds on $\sum_{\overrightarrow{vu}\in A_{\tau_r}} \Delta_h(\overrightarrow{vu})$, to show recurrence for δ close to 1.

8.3.5. THEOREM. *Let $n \ge 1$, and let $G = (V, E, w)$ be the square lattice graph on $\mathbb{Z} \times \{1, \ldots, n\}$ or on $\mathbb{Z} \times (\mathbb{Z}/n\mathbb{Z})$. If $1 - \frac{1}{n} < \delta < 1 + \frac{1}{n-2}$ (for $n \ge 3$), or $1 - \frac{1}{n} < \delta$ (for $n = 1, 2$), then the once-reinforced random walk on G with reinforcement factor δ is almost surely recurrent* [22].

Proof

First assume that G is the square lattice graph on $\mathbb{Z} \times \{1, \ldots, n\}$. With each vertex v of G we can associate coordinates x_v, y_v with $x_v \in \mathbb{Z}, y_v \in \{1, \ldots, n\}$, in the obvious fashion. We may assume that the random walk starts at a point v_0 with $x_{v_0} = 0$. For our superharmonic function h we will use $h(v) = |x_v|$, which

[22] Recurrence for $1 \le \delta < 1 + \frac{1}{n-2}$ was first proven by Sellke in [31], using different methods.

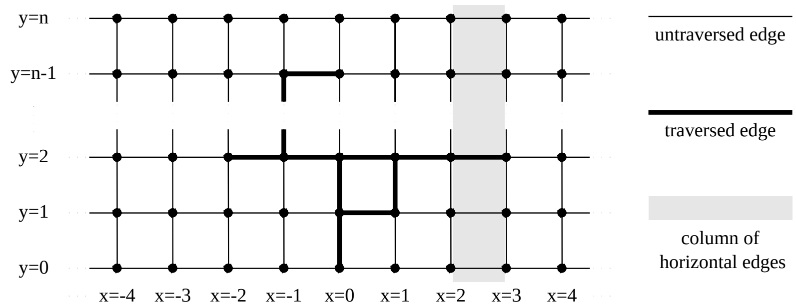

Figure 8.1: The square lattice graph on $\mathbb{Z} \times \{1, \ldots, n\}$.

is easily seen to be harmonic everywhere except on the finite set $F = \{v \in V \mid x_v = 0\}$.
With this function h, the only edges that contribute to the bias are horizontal edges. For any $c \in \mathbb{Z}$, consider the *column* C_c of n horizontal edges connecting points v with $x_v = c$ to points u with $x_u = c+1$. We need to estimate the number of edges of this column that, at first traversal, are traversed going from the lower to the higher h-value. This number is obviously at most n, and unless the column has not been traversed at all, it is at least 1 (since the random walk cannot reach the side of the column with higher h-values without crossing the column at least once). Similarly, the number of edges that, at first traversal, are traversed going from the higher to the lower h-value, is at least 0 and at most $n-1$. So the contribution of the column to the bias satisfies

$$(\delta-1) \sum_{\overrightarrow{vu} \in A_t, vu \in C_c} \Delta_h(\overrightarrow{vu}) \geq \begin{cases} (\delta-1)\max(0,n-2) & \text{if } \delta \geq 1 \\ (1-\delta)n & \text{if } \delta < 1 \end{cases} = -(1-\epsilon) \qquad (8.45)$$

where $\epsilon = 1 - (\delta-1)\max(0,n-2) > 0$ if $1 \leq \delta < 1 + 1/\max(0,n-2)$ and $\epsilon = 1 - (1-\delta)n > 0$ if $1 - 1/n < \delta < 1$.
Now for any t_0 and any $r > t_0/\epsilon$, if τ_r is the first time at or after t_0 that F is visited or $h(v_t) > r$, then the horizontal edges in A_{τ_r} are all contained in the $r + t_0$ columns with x-coordinates between $-t_0$ and r (in the case that $x_{v_{t_0}} > 0$) or between $-r$ and t_0 (in the case that $x_{v_{t_0}} > 0$). Summing all columns, we obtain

$$(\delta-1) \sum_{\overrightarrow{vu} \in A_\tau} \Delta_h(\overrightarrow{vu}) \geq -(1-\epsilon)r - (1-\epsilon)t_0 \qquad (8.46)$$

Hence the conditions of Lemma 8.3.4 are satisfied, and the reinforced random walk is almost surely recurrent.
The proof for the square lattice graph on the cylinder $\mathbb{Z} \times (\mathbb{Z}/n\mathbb{Z})$ is identical.
□

8.3.6. REMARK. Of course, using the absolute bound on $\sum\{\Delta_h(\overrightarrow{vu}) \mid \overrightarrow{vu} \in A_{\tau_r}\}$ is a very unsophisticated method of obtaining a bound on the *expected* value of

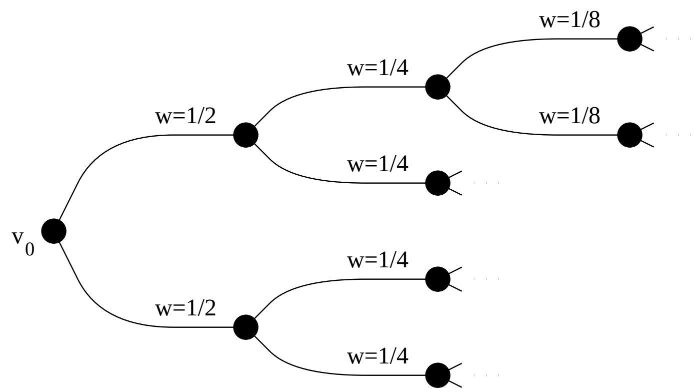

Figure 8.2: A tree on which the once-reinforced random walk with $\delta < 1/4$ is not almost surely recurrent.

the bias. In the above case, we could improve the bounds on the expected value of the bias with a few simple probabilistic calculations, resulting in slight improvements to our bounds on δ. This does not yield any substantial improvements, unfortunately.

In the next chapter a proof will be given of recurrence for *large* values of δ. No proof is yet known for intermediate values of δ. It is also not yet known whether the once-reinforced random walk on the square lattice graph on $\mathbb{Z}^2$ is recurrent for *any* reinforcement factor $\delta \neq 1$, although a related marginal result is given at the end of this chapter. The intuition, however, is that once-reinforced random walks on the square lattice graphs on $\mathbb{Z}^2$ and $\mathbb{Z} \times \{1, \ldots, n\}$ are recurrent for all $\delta \geq 1$.

For reinforced random walks on weighted trees, recurrence *can* be proven for all $\delta \geq 1$ (provided the tree is such that non-reinforced random walks are recurrent in the first place). This follows already from Theorem 8.2.2. To prove recurrence for *all* $\delta \geq 1$, in *any* graph on which non-reinforced random walks are recurrent, one could use something like

8.3.7. Proposition. For any edge vu that is 'far away' from all edges traversed so far, if v is closer than u to the origin of the walk (in the sense that $h(v) < h(u)$), then vu has at least as much chance of being traversed (the first time it is traversed) from v to u as it has of being traversed (the first time it is traversed) from u to v.

This means, very loosely formulated, that closer vertices are visited earlier. Alas, so far this proposition has neither been proved nor refuted.

8.3.8. Example. For negative reinforcements, recurrence is not necessarily preserved. There are examples of cases where the non-reinforced random walk on a graph is recurrent, but for certain $\delta < 1$, the once-reinforced random walk with

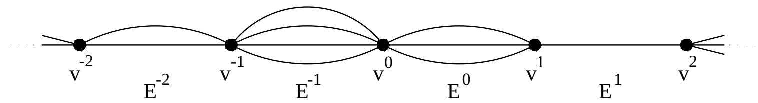

Figure 8.3: The graph of Theorem 8.3.9

reinforcement δ is not. Figure 8.2 is such a case: it can be shown that for $\delta < 1/4$, the once-reinforced random walk on this tree is not almost surely recurrent.

Next is an application of Lemma 8.3.4 that uses probabilistic methods rather than an absolute bound. Note that we actually use the generalization of Lemma 8.3.4 to graphs with parallel edges, rather than the lemma as written. As stated in Remark 6.1.1, we will simply postulate this generalization and proceed.

8.3.9. THEOREM. *Let $G = (V, E, w)$ be a weighted graph with vertices $V = \{v^i \mid i \in \mathbb{Z}\}$ and for any n, a finite non-zero number of parallel edges between the vertices v^n and v^{n+1}. If the non-reinforced random walk on G is almost surely recurrent, then the reinforced random walk on G is almost surely recurrent for any reinforcement factor $\delta > 0$.*

Proof

If the non-reinforced random walk on G is almost surely recurrent, then by Theorem 6.3.2 there exists a function $h : V \to \mathbb{R}$ such that h is superharmonic on $V - \{v^0\}$ and $h(v^n) \to \infty$ if $n \to \infty$ or $n \to -\infty$. It is easily seen that $h(v^n) > h(v^m)$ if $n > m > 0$ or $n < m < 0$.
Now for $n \in \mathbb{Z}$, let $G^n = (V^n, E^n, w^n)$ be the subgraph induced by $\{v^n, v^{n+1}\}$ (i.e. $V^n = \{v^n, v^{n+1}\}$, E^n is the set of edges between v^n and v^{n+1}, and $w^n = w_{|E^n}$). Although events outside G^n may effect *whether* and *when* an edge of G^n is traversed, *which* edge of G^n is traversed is only dependent on the relative current weights of the edges of G^n. So we can estimate the expected contribution to the bias of each set E^n separately, and take the sum to arrive at an estimate for the total expected bias. The possibility that at some point the walk in G will no longer return tot G_n can be simulated by a stopping time for the walk in G_n.
So fix $n \in \mathbb{Z}$ and consider the reinforced random walk on the finite graph G^n, starting in v^n if $n \geq 0$, and in v^{n+1} otherwise. Note that in both cases the random walk will start at the vertex with the lower h-value and then alternate between the two vertices. Set $c = |h(v^{n+1}) - h(v^n)|$. If for G^n we define A^n_t and E^n_t as usual, we have for any $t \in \mathbb{N}$

$$E\left(\sum_{\vec{a}\in A^n_{t+1}} \Delta_h(\vec{a}) \middle| \mathcal{F}_t\right) = \sum_{\vec{a}\in A^n_t} \Delta_h(\vec{a}) + (-1)^t \frac{w(E^n) - w(E^n_t)}{w(E^n) + (\delta-1)w(E^n_t)} \cdot c \tag{8.47}$$

where $w(X)$ denotes $\sum_{e\in X} w(e)$. This implies that for any stopping time τ

$$E\left(\sum_{\vec{a}\in A^n_\tau} \Delta_h(\vec{a})\right) = E\left(\sum_{t=0}^{\tau-1} (-1)^t \frac{w(E^n) - w(E^n_t)}{w(E^n) + (\delta-1)w(E^n_t)}\right) \cdot c \tag{8.48}$$

For all t, $w(E_t^n) \leq w(E_{t+1}^n) \leq w(E^n)$, and hence

$$\frac{w(E^n)-w(E_t^n)}{w(E^n)+(\delta-1)w(E_t^n)} \leq \frac{w(E^n)-w(E_{t+1}^n)}{w(E^n)+(\delta-1)w(E_{t+1}^n)} \tag{8.49}$$

We conclude that for any stopping time τ,

$$0 \leq E\left(\sum_{\vec{a}\in A_\tau^n} \Delta_h(\vec{a})\right) \leq c \tag{8.50}$$

Now let us return to the random walk on G. Fix $t_0 \in \mathbb{N}$ and the history up to time t_0. Then all the vertices that have been visited up to time t_0 have indices between $-t_0$ and t_0. Furthermore, all the vertices that can be visited after time t_0 are on the same side of v^0 until the first visit to v^0; without loss of generality we may assume that this is the side of the vertices with positive indices. If we transfer the results we obtained for the walks on the graphs G^n to the random walk on the graph G, and take the sum of the inequalities over all edge-sets E^n with $n \geq t_0$, then we obtain

$$-c' \leq E\left(\sum_{\vec{a}\in A_\tau} \Delta_h(\vec{a}) \middle| \mathcal{F}_{t_0}\right) \leq E(\max\{h(v_t) \mid t \leq \tau\}) - h(v^{t_0}) + c' \tag{8.51}$$

where τ is any stopping time such that the walk does not leave the set of vertices with positive indices, and $c' = \sum_{n=-t_0}^{t_0-1}(\#E_n)|h(v^{n+1}) - h(v^n)|$.
This implies the condition of (the generalization of) Lemma 8.3.4 for all $\delta > 0$.

□

The third application of Lemma 8.3.4 yields an (admittedly rather marginal) result for a *variant* once-reinforced random walk on the square lattice graph with unit weights on $\mathbb{Z}^2$. In this variant once-reinforced random walk, the reinforcement factor is not constant, but is allowed to be different for each edge. It is not difficult to modify Definition 8.1.1 to allow this type of once-reinforced random walks, although we will encounter some hidden complications in modifying some of the Lemmas given in this chapter.

8.3.10. Theorem. *Consider the variant once-reinforced random walk on the square lattice graph G with unit weights on $\mathbb{Z}^2$, where the reinforcement factor is not constant but is, for each edge, reciprocal to the Euclidean distance of the edge from the origin* $(0,0)$, *i.e. for some* $C > 0$,

$$w_t(vu) = \begin{cases} 1 & \textit{if } vu \textit{ has not yet been traversed} \\ 1 + C/(\max(|(x_v,y_v)|,|(x_u,y_u)|)) & \textit{if } vu \textit{ has been traversed} \end{cases} \tag{8.52}$$

This random walk is recurrent for $C < 1/(2\sqrt{2}\pi)$.

Proof

First we need to find analogues of Lemmas 8.2.1 and 8.3.3 for variant walks of this type. Unfortunately, the lack of a constant reinforcement factor δ makes it difficult to even formulate an analogue of Lemma 8.3.3, let alone prove it. This is not surprising: essential to the concept of Lemma 8.3.3 is that as long as the walk only traverses edges that have been traversed before, the change in the value of the martingale is reflected in the change of the current value of $h(v_t)$. This no longer holds if we use the definition of Lemma 8.2.1 and the reinforcement factor is not constant.

However, there is a solution to this dilemma: we can modify both the nominal weight of the edges and the *initial* reinforcement factors (the multipliers that are applied to edges that have not been traversed before), in such a way that the reinforcement factors that are applied to traversed edges become constant. I.e. without changing the *actual* walk, we consider it to be on the square lattice graph G' on $\mathbb{Z}^2$ with weights

$$w'(vu) = 1 + C/(\max(|(x_v, y_v)|, |(x_u, y_u)|)) \text{ for } vu \in E \tag{8.53}$$

and for each edge $vu \in E$ a reinforcement sequence

$$(\delta_k(vu))_{k=0}^{\infty} = (\frac{1}{1 + C/(\max(|(x_v, y_v)|, |(x_u, y_u)|))}, 1, 1, 1, \ldots) \tag{8.54}$$

In the actual walk, this yields the same weights as before. Of course, this means that we need to select h to be a function with the right properties on G' rather than on G. Fortunately, the weight of an edge of G' is never more than twice the weight of the corresponding edge of G, so by by Theorem 6.3.8 non-reinforced random walks are as recurrent on G' as they are on G. Hence a function h with the necessary properties exists. As it turns out, the function

$$h(x,y) = \begin{cases} \log(1/12) & \text{if } (x,y) = (0,0) \\ \log(1/4) & \text{if } (x,y) = (0,\pm 1) \text{ or } (x,y) = (\pm 1, 0) \\ \log(x^2 + y^2 - 1) & \text{otherwise} \end{cases} \tag{8.55}$$

which we used in Example 6.3.3 for G, also works for G'.

Now we can obtain, in sequence, analogues of Lemmas 8.2.1, 8.3.3 and 8.3.4. The proofs of the following statements is straightforward and quite similar to the proofs of the original lemmas, and therefore will be omitted. First, the stochastic process

$$M_t = \sum_{t'=0}^{t} \begin{cases} \frac{\Delta_h(v_{t'} v_{t'+1})}{\delta_{k_{t'}(v_{t'} v_{t'+1})}(v_{t'} v_{t'+1})} & \text{if } v_{t'} \in V' \\ 0 & \text{otherwise} \end{cases} \tag{8.56}$$

is an martingale. Next, for $t \geq t_0 \geq 0$ the equation

$$M_t - M_{t_0} = h(v_t) - h(v_{t_0}) + \sum_{\overrightarrow{vu} \in A_t - A_{t_0}} \left(\frac{1}{\delta_0(vu)} - 1 \right) \Delta_h(\overrightarrow{vu}) \tag{8.57}$$

holds, as long as $V - V'$ has not been visited at any time between t_0 and t. Finally, in order to show recurrence, it suffices to show that for some $\epsilon > 0$, for any time t_0, and for any history up to time t_0,

> There exists a $c \in I\!R$ such that for all $r_0 \in I\!R$ we can find $r > r_0$ with
>
> $$E\left(\sum_{\overrightarrow{vu} \in A_{\tau_r}} \left(\frac{1}{\delta_0(vu)} - 1\right) \Delta_h(\overrightarrow{vu}) \,\middle|\, \mathcal{F}_{t_0}\right) \geq -(1-\epsilon)r - c \tag{8.58}$$
>
> where the stopping time τ_r is the first time at or after t_0 that $(0,0)$ is visited or $h(v_t) \geq r$.

So now the only thing left to do are a few calculations in the manner of Theorem 8.3.5. First note that for an edge $vu \in E$ with $1 < |(x_v, y_v)| \leq |(x_u, y_u)|$,

$$\left|\left(\frac{1}{\delta_0(vu)} - 1\right) \Delta_h(\overrightarrow{vu})\right| = C \frac{\log(x_u^2 + y_u^2 - 1) - \log(x_v^2 + y_v^2 - 1)}{|(x_u, y_u)|} \tag{8.59}$$

$$\leq C \frac{x_u^2 + y_u^2 - x_v^2 - y_v^2}{(x_v^2 + y_v^2 - 1)|(x_v, y_v)|} \tag{8.60}$$

Now for most $v \in V$, there are two vertices $u \in N(v)$ with $|(x_u, y_u)| > |(x_v, y_v)|$, and

$$\sum_{u \in N(v), |(x_u, y_u)| > |(x_v, y_v)|} x_u^2 + y_u^2 - x_v^2 - y_v^2 \leq 2|x_v| + 2|y_v| + 2 \leq 2\sqrt{2}|(x_v, y_v)| + 2 \tag{8.61}$$

The exceptions are the vertices $v \in V$ with $x_v = 0$ or $y_v = 0$, and for those

$$\sum_{u \in N(v), |(x_u, y_u)| > |(x_v, y_v)|} x_u^2 + y_u^2 - x_v^2 - y_v^2 \leq 2|(x_v, y_v)| + 3 \leq 2\sqrt{2}|(x_v, y_v)| + 3 \tag{8.62}$$

Hence, for any $C' > 2\sqrt{2}C$, we can find an $R > 0$ such that for all $v \in V$ with $|(x_v, y_v)| > R$,

$$\sum_{u \in N(v), |(x_u, y_u)| > |(x_v, y_v)|} \left|\left(\frac{1}{\delta_0(vu)} - 1\right) \Delta_h(\overrightarrow{vu})\right| \leq \frac{C'}{(x_v^2 + y_v^2)} \tag{8.63}$$

Furthermore, for any $C'' > C'$, we can find $r_0 \geq \log(R'^2 - 1)$ such that for any $r > r_0$,

$$\sum_{\overrightarrow{vu} \in A_{\tau_r}} \left|\left(\frac{1}{\delta_0(vu)} - 1\right) \Delta_h(\overrightarrow{vu})\right|$$

$$\leq \sum_{vu \in E, |(x_v, y_v)| < \sqrt{e^r + 1}} \left|\left(\frac{1}{\delta_0(vu)} - 1\right) \Delta_h(\overrightarrow{vu})\right| \tag{8.64}$$

$$\leq \sum_{vu\in E,|(x_v,y_v)|\leq R} \left|\left(\frac{1}{\delta_0(vu)}-1\right)\Delta_h(\overrightarrow{vu})\right| + \sum_{v\in V,R<|(x_v,y_v)|<\sqrt{e^r+1}} \frac{C'}{(x_v^2+y_v^2)} \quad (8.65)$$

$$\leq \sum_{vu\in E,|(x_v,y_v)|\leq R} \left|\left(\frac{1}{\delta_0(vu)}-1\right)\Delta_h(\overrightarrow{vu})\right| + \oint_{D(O,R,e^r+1)} \frac{C''}{(x_v^2+y_v^2)} \quad (8.66)$$

$$\leq \sum_{vu\in E,|(x_v,y_v)|\leq R} \left|\left(\frac{1}{\delta_0(vu)}-1\right)\Delta_h(\overrightarrow{vu})\right| + 2C''\pi(\log(\sqrt{e^r+1})-\log(R)) \quad (8.67)$$

$$\leq \sum_{vu\in E,|(x_v,y_v)|\leq R} \left|\left(\frac{1}{\delta_0(vu)}-1\right)\Delta_h(\overrightarrow{vu})\right| + C''\pi r+1-2C''\pi\log(R) \quad (8.68)$$

where $D(O,R,e^r+1) = \{(x,y)\in I\!R^2 \mid R<|(x,y)|<\sqrt{e^r+1}\}$. For $C<1/(2\sqrt{2}\pi)$, we can select C' and C'' such that $C''<1$, and hence such that equation (8.58) is satisfied and the random walk is shown to be recurrent.

□

Chapter 9

Nonstandard Analysis

In Chapter 10, we will consider once-reinforced random walks whose reinforcement factor is near-infinite. To do that, we make extensive use of nonstandard analysis (NSA), the extension of real analysis with infinitesimals. Although it is known that anything that can be proven with NSA can also be proven without it, NSA allows for a much more intuitive treatment of concepts such as 'sufficiently large' and 'may be safely ignored', and hence is a very useful tool in this context.

This chapter aims to give a brief introduction to NSA, and an overview of its basic principles and techniques. A full treatment of NSA falls outside the scope of this chapter, but interested readers can find more material in [21], [27] and [28].

9.1 Introduction

The idea of using infinitesimals is nothing new. Newton used infinitesimals to define the derivative of a function, and in an old proof relating the area of a circle to its circumference the circle is treated as an infinity-sided polygon, and as a composition of triangles with infinitesimal bases. However, careless usage of infinitesimals and infinities can easily lead to contradictions, and hence the technique always was considered to be suspect. Eventually, the use of infinitesimals was discarded in favor of limit constructions.

The disadvantage of using limit constructions is that they are considerably less intuitive than infinitesimals. When we see the expression $\delta y/\delta x$, we may *define* it in terms of limits, but we *visualize* it as 'the rate of change over an infinitesimal interval'. For this reason, the study of infinitesimals was never wholly abandoned. And in 1961, these efforts finally bore fruit, as A. Robinson developed a consistent formalism for using infinitesimals, and founded the field of nonstandard analysis, abbreviated as NSA.

The nonstandard approach can also be applied to other fields of mathematics,

yielding concepts such as nonstandard ordinals, hyperfinitely-dimensional manifolds etc. In 1977, Nelson invented Internal Set Theory (abbreviated as IST) in an attempt to give a unified axiomatical background for nonstandard mathematics. IST extends ZFC, Zermelo-Fraenkel set theory, by adding a 'standardness' predicate, and three axioms, Idealization, Standardization and Transfer. In the next sections we will introduce and use this formalism.
Anything which holds in 'orthodox' mathematics also holds in nonstandard analysis. In a sense, the reverse holds as well: any statement that does not refer to any nonstandard concepts or constants, and that can be proven in nonstandard mathematics, can be proven in orthodox (non-nonstandard) mathematics. More formally:

9.1.1. THEOREM. *IST is a conservative extension of ZFC.*

Note that this implies that if ZFC is consistent, then so is IST. Nelson [28] gives an explicit algorithm to translate proofs to orthodox mathematics. So in a sense nonstandard analysis doesn't add anything new. However, proofs in nonstandard analysis are often much simpler.

9.2 The standardness predicate

Amusingly enough, the most important concept in nonstandard analysis is the concept of '*standard*'. The easiest way to introduce this concept is probably to consider infinitesimals, and what properties we *desire* them to have. For example, we want to be able to calculate with them: if δ is an infinitesimal, we want to be able to talk about 2δ, $1+\delta$, $1/\delta$ etcetera. Furthermore, we want those 'nonstandard' numbers to obey the same rules that 'normal' numbers do.
Robinson's original approach was to take a model of the real line, and construct a *new* model of the axioms of the real numbers, by adding the infinitesimals (and related numbers) in such a manner that everything that held in the old model also held in the new one. The drawback of this approach is that it is not possible to talk about infinitesimals without in some way referring to the old model and how it differs from the new one. For instance, it is common practice to refer to those real numbers that already exist in the old model as being '*standard*'. Using this concept of standardness we can define infinitesimals as

> An infinitesimal is a real number whose absolute value is smaller than every standard positive real

However, it is *impossible* to define infinitesimals without either using the concept of standardness or referring to the old model in some other manner. The reason for this is that in the old model it was impossible to 'access' infinitesimals because they didn't exist, and (by our design) everything that held in the old model also holds in the new model.

Nelson's approach made it possible to bypass this need to refer to different models of the reals. Since from 'within' a model of the reals it is impossible to see in which model you are, it is consistent to assume that the real line that you are using already contains infinitesimals. So Nelson simply *postulated* that they exist, and introduced a standardness predicate 'st(x)', and some axioms to describe its properties, that allowed us to access them. In this perspective, rather than adding or creating infinitesimals and other 'nonstandard' numbers, we *discover* them using the new predicate and axioms.

The two approaches are basically two different perspectives on the same concept. In the first perspective there are two worlds, an 'old' one and a 'new' one containing extra elements. In this second perspective there is no new 'world', merely aspects of the old 'world' that existed before but couldn't be seen. The difference between the two perspectives is more or less a matter of personal taste. In these pages we will keep to the second perspective.

So, nonstandard analysis contains all axioms of 'orthodox' real analysis, and hence all the usual theorems and tools are still valid. All the real numbers we knew before, such as 0, 1, e and π, still exist and have the same properties as before. But in addition they satisfy the standardness predicate, and we now can see that inside the gaps between standard reals exist *nonstandard* reals. However, the only abnormal property these reals have is that they do not satisfy the standardness predicate: they are not noticeable except for that.

9.2.1. Remark. As a rule of thumb, anything that can be defined using only standard parameters, is itself standard. Conversely, anything that can be used to define something that is known to be nonstandard, is itself nonstandard.

The standardness predicate can also be applied to sets. Again, sets such as $\{0\}$, $\mathbb{R}$ and $[0, e]$ still exist, have the same properties as before, and additionally satisfy the standardness predicate. But now we have new sets, such as $\{\delta\}$ and $[0, \delta]$, which are nonstandard (if $\delta \neq 0$ is an infinitesimal). Note that it is not true that standard sets are sets containing standard elements. The correspondence *does* hold for finite sets. But the set $\mathbb{R}$, for example, is a standard set containing nonstandard elements, since by definition $\mathbb{R}$ contains *all* the reals, including the nonstandard ones.

In fact, it can be shown that every infinite set, whether standard or nonstandard, contains a nonstandard element. Infinite collections containing only standard elements are undefinable as a set. The reason for this is that such sets would contradict some of the laws for sets that already hold in orthodox real analysis. For example, since in orthodox real analysis every bound set has a greatest lower bound, the same should hold in nonstandard analysis. However, if we consider the 'set' of positive standard reals, every infinitesimal is a lower bound of this set. So the greatest lower bound cannot be 0, it obviously cannot be a positive standard real, and if it were an infinitesimal then twice that value would be a greater lower bound. Hence this 'set' would not have a greatest lower bound. The solution to this apparent paradox is to disallow the use of formulas containing the standardness predicate when defining a subset, i.e. when using the Separation Axiom. The same applies when using the Replacement Axiom. This ensures that collections such as the collection of standard positive reals are undefinable *as a set.*[23]

9.3 Basics of Nonstandard Real Analysis

In a sense, nonstandard real analysis is about infinitesimals. To prove the existence of infinitesimals, we need the axioms of Internal Set Theory. However, rather than immediately reviewing these axioms, we will first look at the infinitesimals themselves, consider some related concepts and their properties, and give some examples of their use in mathematics.
First, we will formalize the definition of infinitesimals given before:

9.3.1. Definition. A real number $x \in I\!R$ is called *infinitesimal* if it satisfies

$$\forall^{\text{st}} \epsilon > 0 : |x| < \epsilon \tag{9.1}$$

Here the quantifier $\forall^{\text{st}} x$ is an abbreviation for $\forall x(\text{st}(x) \to \ldots)$, or 'for all standard $x, \ldots$'. Similarly, the quantifier $\exists^{\text{st}} \epsilon$ would be an abbreviation for $\exists x(\text{st}(x) \wedge \ldots)$, or 'there exists a standard x such that ...'. This definition properly captures the notion that infinitesimals are 'very, very small'. However, we also want to capture the notion that if δ is infinitesimal, then $1 + \delta$ is 'very, very close to 1':

9.3.2. Definition. $x, y \in I\!R$ are called *infinitesimally close* (denoted $x \approx y$) if

$$\forall^{\text{st}} \epsilon > 0 : |x - y| < \epsilon \tag{9.2}$$

$x \in I\!R$ is called *nearstandard* if x is infinitesimally close to some standard real:

$$\exists^{\text{st}} y \in I\!R : y \approx x \tag{9.3}$$

This standard real is called the *standard part* of x, denoted $°x$.

Now, for all $\delta > 0$, $1/\delta$ is a positive real. However, if δ is infinitesimal, then $1/\delta$ is very, very big. In fact, it is larger than every standard real. This is an example of a *hyperfinite* real:

[23] One of the few differences between the two perspectives we described before, is that in the first perspective, these collections are considered to exist as sets of elements of the model, or 'external sets'. What we consider to be a set, is called an 'internal set' in the first perspective, being a set of elements of the model that is also a set *in* the model. Since we keep to the second perspective, we only consider internal sets to be sets.

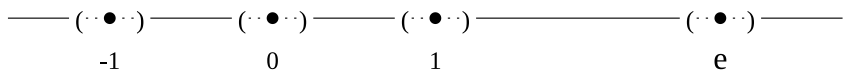

Figure 9.1: The standard reals -1, 0, 1 and e, each surrounded by a 'cloud' of infinitesimally close nonstandard reals.

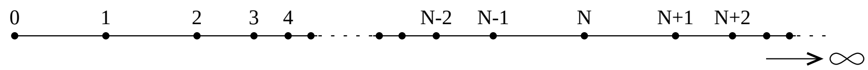

Figure 9.2: A hyperfinite number N

9.3.3. Definition. $x \in I\!R$ is called *hyperfinite* if it satisfies

$$\forall^{\text{st}} r \in I\!R : |x| > r \tag{9.4}$$

The reciprocals of hyperfinite numbers are non-zero infinitesimals, and vica versa. Hyperfinite numbers can also be characterized as those numbers that are not nearstandard. Note that although a hyperfinite number is larger (in absolute value) than all standard real numbers, it is still a member of $I\!R$, and as such *not* infinite. Figure 9.2 sketches the position of a hyperfinite number N, relative to the standard natural numbers and to infinity.
Now with these concepts, we will give some examples of their use.

9.3.4. Example. Let us imagine some object, say a bowling ball, falling from a large height with a constant acceleration $g = 9.8$. If the bowling ball starts out with speed $v(0) = 0$, its speed at time t satisfies

$$v(t) = gt \tag{9.5}$$

Now, let T be a standard time, and let N be a hyperfinite natural number. Then for any time t, $[t, t+T/N]$ is an infinitesimal time interval. The distance the ball travels in such an interval satisfies

$$gt(T/N) = v(t)(T/N) \leq s(t{+}T/N){-}s(t) \leq v(t{+}T/N)(T/N) = g(t{+}T/N)(T/N) \tag{9.6}$$

We can divide the time interval $[0, T]$ into N such intervals and take the sum over these intervals. Since the speed of the ball increases steadily, this yields

$$\begin{aligned}
\frac{1}{2}gT^2 &\approx \frac{1}{2}g(T/N)^2(N-1)N && (9.7)\\
&= g(T/N)\sum_{i=0}^{N-1}(T/N)i && (9.8)\\
&\leq s(T) - s(0) && (9.9)\\
&\leq g(T/N)\sum_{i=0}^{N-1}(T/N)(i+1) && (9.10)
\end{aligned}$$

$$= \frac{1}{2}g(T/N)^2N(N+1) \tag{9.11}$$

$$\approx \frac{1}{2}gT^2 \tag{9.12}$$

If g, $v(0)$ and T are all standard, then so is $s(T) - s(0)$, and hence

$$s(T) - s(0) = \frac{1}{2}gT^2 \tag{9.13}$$

9.3.5. Example. Consider a three-dimensional sphere S with center M and some standard radius $r > 0$. Let A denote the surface area of S, and V its volume. We can approximate this sphere by a polyhedron with hyperfinitely many faces, each of which is a triangle of infinitesimal dimensions. The surface area A' and volume V' of this polyhedron are infinitesimally close to A and V. Now, for each face DEF of the polyhedron, we can construct a tetrahedron $DEFM$. The volume of this tetrahedron is $r/3$ times the area of DEF. The polyhedron can be thought of as being composed of hyperfinitely many of these tetrahedrons, one for each face. Taking the sum over all these tetrahedrons, we get

$$V' = r/3 \cdot A' \tag{9.14}$$

Hence $V \approx r/3 \cdot A$. Since we are dealing with a sphere of standard radius, V and A are both standard, and hence

$$V = r/3 \cdot A \tag{9.15}$$

It can be shown (using the method of the previous example) that $V = 4/3\pi r^3$. Hence

$$A = 4\pi r^2 \tag{9.16}$$

9.4 Idealization, Standardization and Transfer

Nelson's Internal Set Theory extends ZFC with the standardness predicate and three axioms: Idealization, Standardization and Transfer. As stated before, the ZFC axiom schemas of Separation and Replacements are *not* extended to include formulae that use the standardness predicate. In this section we will review the three axioms and their common usage. We will formulate the axioms in terms of objects and sets rather than real numbers, in order to pave the way for the application of nonstandard analysis to graph theory. At the end of the section, we will give some examples of how the three axioms work together.

> **The Axiom of Idealization**: For there to exist an object which has a particular property relative to *all* standard objects, it suffices that there exist objects having that property relative to finitely many

standard objects at a time. Formally, for any formula $\phi(x, y)$ not containing the predicate st,

$$\left(\forall^{\text{st fin}} F\, \exists x\, \forall y \in F\, \phi(x, y)\right) \quad \leftrightarrow \quad \left(\exists x \forall^{\text{st}} y\, \phi(x, y)\right) \tag{9.17}$$

$\phi(x, y)$ may contain standard or nonstandard constants, or free variables other than x and y.

The right-to-left implication of this axiom is only used to show that finite sets are standard if and only if their elements are standard. The left-to-right implication is used to obtain nonstandard, 'idealized' objects. For instance, it is obvious that for any finite standard set of positive numbers F, there exists an $x \in \mathbb{R}$ such that $\forall y \in F : |x| < y$. Hence, if we apply Idealization to the formula $\phi(x, y) \equiv (y > 0 \to 0 < |x| < y)$, the left side of the equivalence holds, and we obtain the existence of $x \neq 0$ such that $\forall^{\text{st}} y > 0 : |x| < y$, i.e. x is infinitesimal. Another form of (the left-to-right-implication of) the Idealization Axiom is that of the *principle of Overflow*, which states that any (definable) set containing arbitrarily large standard reals also contains a hyperfinite real, and the related *principle of Underflow*, which states that any (definable) set containing arbitrarily small positive hyperfinite reals also contains a nearstandard real.

The Axiom of Standardization: If S is an arbitrary set, then we can obtain a (unique) standard set $^\circ S$, the *standardization* of S, by changing just the nonstandard elements. Formally, for any sets S there exists a standard set $^\circ S$ such that

$$\forall^{\text{st}} x (x \in {}^\circ S \leftrightarrow x \in S) \tag{9.18}$$

Standardization is often used to allow us to 'ignore' infinitesimal discrepancies. Note that the Standardization Axiom does not necessarily remove the nonstandard elements of a set: it makes the set as a whole standard by making arbitrary changes in its nonstandard elements. Standardizing a standard set such as $\mathbb{R}$ will have no effect, and standardizing the set $\{0, 1, \ldots, N\}$, where N is a hyperfinite natural number, will actually *add* nonstandard numbers (resulting in the standard set $\mathbb{N}$).

Furthermore, although we can represent objects such as functions as sets in order to apply the Standardization Axiom, this will result in the standardization of *all* aspects of the object. In the case of a function, the domain may well change, for instance. And the Standardization Axiom cannot be applied to infinitely many objects at the same time: to standardize each object in a collection, we have to standardize the collection itself, including its index set. For these reasons, one has to take care to set up the right conditions before using Standardization.

The Axiom of Transfer: The Transfer axiom states, that if something holds for all standard objects, it holds for all objects, and conversely if there exists an object satisfying some condition, there exists a standard object satisfying that condition. The formula involved may not refer to standardness or to nonstandard constants. Formally, for each formula $\phi(x, \bar{y})$ not containing 'st', nonstandard constants or free variables other than x and $\bar{y}$,

$$\forall^{\mathrm{st}} \bar{y}\, (\exists x \phi(x, \bar{y}) \leftrightarrow \exists^{\mathrm{st}} x \phi(x, \bar{y})\,) \tag{9.19}$$

The Transfer Axiom can be used to drop a condition of the form 'let x be standard', and to translate results back into the language of Real Analysis. It can also be used to show that we may take some entity to be standard. For instance, if $F(\bar{y})$ is a function definable without using nonstandard constants or the predicate st, then by applying Transfer to the formula $\phi(\bar{y}) \equiv x = F(\bar{y})$ we obtain that for all standard parameters $\bar{y}$, if $F(\bar{y})$ exists it is standard. Hence objects that can be uniquely defined using only standard constants are standard (as we already stated in Remark 9.2.1).

9.4.1. Example. The Idealization Axiom can be used to show that there exists a finite set containing all the standard elements. For it is obviously true that for any standard finite set F there exists a finite set x satisfying $y \in x$ for all $y \in F$: simply take F for x. Applying Idealization with $\phi(x, y) \equiv (x \text{ is finite}) \wedge (y \in x)$ yields the desired result. Note that the resulting set cannot be standard (else it would contain itself), and therefore must contain some nonstandard elements as well.

9.4.2. Example. Standardization and Transfer can be used to obtain the standard part of a real $x \in I\!R$, provided x is not hyperfinite. Let $S = \{z \in I\!R \mid z \leq x\}$. Since x is not hyperfinite, $-C < x < C$ for some standard $C > 0$, so S has a standard element and a standard upper bound. Now consider the least upper bound of its standardization ${}^{\circ}S$. If y is a standard upper bound of S, then $y \geq z$ for all standard $z \in {}^{\circ}S$, and by Transfer $y \geq z$ for all $z \in {}^{\circ}S$. So S and ${}^{\circ}S$ have the same standard upper bounds. By the definition of ${}^{\circ}S$, they also have the same standard elements. It follows that the least upper bound of ${}^{\circ}S$ exists and is infinitesimally close to the least upper bound of S, i.e. to x. Furthermore, by Transfer the least upper bound of ${}^{\circ}S$ is standard. So the least upper bound is equal to the standard part ${}^{\circ}x$ of x.
Note that if we try to use this approach with a hyperfinite real $x \in I\!R$, then the resulting set ${}^{\circ}S$ turns out to be equal to $\emptyset$ or $I\!R$, depending on whether x is negative or positive. Hence the nearstandard reals are exactly those reals that are not hyperfinite.

Chapter 10

Reinforced Random Walks with Large Reinforcements

In this chapter we will consider recurrence of once-reinforced random walks with very large reinforcement factors. To do this, we make extensive use of techniques from Nonstandard Analysis. Readers who are unfamiliar with NSA are advised to first read Chapter 9, which gives an overview of the basics.
The first section of this chapter elaborates on the application of NSA to graph theory, which is not always straightforward. After that, we consider once-reinforced random walks with hyperfinite reinforcement. We first show that if the reinforcement factor δ is hyperfinite, then the 'growth process' of the subgraph of traversed edges will not only be very slow, but also very uniform. Later in this chapter we use this to show that on some graphs, the estimated bias of the random walk is nonnegative, enabling us to use Lemma 8.3.4 to show recurrence. Finally we translate the result back to Real Analysis, and show recurrence for δ *large enough*.

10.0.3. REMARK. A word of caution: in the context of this chapter, 'almost surely' is taken to mean 'with probability 1', which is *not* the same as 'with probability infinitesimally close to 1'. We make similar distinctions between 'hyperfinitely often' and 'infinitely often', and between 'graphs of hyperfinite size' and 'graphs of infinite size'. The rule of thumb is, as always, that hyperfinite counts as finite, and infinitesimal is not the same as zero.

10.1 Nonstandard Analysis and Graph Theory

In the next section we will use nonstandard analysis applied to graph theory, with graphs that may be of hyperfinite size. Unfortunately there is a slight complication. We commonly consider isomorphic graphs to be essentially the same, merely different representations of a single object, and consider this single object to be the 'true' graph[24]. The problem is that in this case, the representation of a graph *does* matter: standardness doesn't respect graph isomorphisms, as the following example shows.

10.1.1. Example. Consider the graphs $G_n = (V, E, w_n)$ with $V = \mathbb{Z}$, $E = \{v_i v_{i+1} \mid i \in Z\}$, $w_n(v_n v_{n+1}) = 2$, $w_n(v_i v_{i+1}) = 1$ for all $i \neq n$. It is clear that all the graphs G_n are isomorphic. Given n, the graph G_n is uniquely defined, and conversely we can obtain n from G_n, both without using any (other) nonstandard constants. It follows from the Transfer Axiom that G_n is standard if and only if n is standard.

To solve this problem, we could try to find some kind of 'canonic' representation of a graph[25]. Unfortunately, for an infinite graph it is not possible to choose a new representation such that every vertex is standard, and there are few other notions on which to base canonicity. However, note that if G is a standard graph and v is a standard vertex of G, then the standard vertices of G are exactly those vertices that are within standard G-distance of v, i.e. if we try to choose a different representation of a graph, and we limit ourselves to standard representations, the choice of a single vertex to be standard determines the standardness of the remaining vertices. This gives rise to the following notion of *relative standardness*:

10.1.2. Definition. A graph $G = (V, E, w)$ is *standard relative to a vertex* $v \in V$ if there is a standard graph isomorphic to G, such that v corresponds to a standard vertex of that graph.

We can also define a notion of uniform standardness:

10.1.3. Definition. A graph $G = (V, E, w)$ is *uniformly standard* if it is standard relative to any vertex $v \in V$.

For example, by Transfer the graph G_0 of Example 10.1.1 is standard relative to v_0, but not relative to v_N for any nonstandard N, and therefore it is not uniformly standard. Examples of uniformly standard graphs are the square lattice graphs on $\mathbb{Z}^2$ and $\mathbb{Z} \times (\mathbb{Z}/n\mathbb{Z})$ with unit weights.

10.1.4. Remark. If a graph $G = (V, E, w)$ is standard relative to a vertex v, then it is not possible to define any hyperfinite number using G and v. If there is some anomalous feature of G, for instance a vertex $u \in V$ which is the only vertex of degree 1, then $d(v, u)$ would be definable using only G and v as parameters, and hence $d(v, u)$ would have to be standard. It follows that all definable features of G are somewhere within standard distance of v. Outside of these anomalies,

[24] Although technically it is not a graph but an equivalence class of graphs under graph isomorphism.

[25] I.e. a canonically chosen graph isomorphic with the original graph

and hence everywhere outside a 'standard-distance neighborhood' of v, G will follow some kind of regular pattern.
By the same token, if G is uniformly standard, then G has no anomalies *anywhere*. Hence the standard-distance neighborhood of any one vertex of G will look the same as that of any other.

The difficulties we encountered with standardness also apply to the notion of taking the standard part of a graph. The naive approach to taking the standard part would be to simply standardize the vertex-set, the edge-set and the weight function of the graph. Applying this notion of standardization to the graphs G_n of Example 10.1.1 yields G_n if n is standard, since then V, E and w_n are standard themselves. But if n is nonstandard, then the result is the graph $G_* = (V, E, w_*)$, where w_* is the weight function with $w_*(v_i v_{i+1}) = 1$ for all $i \in \mathbb{Z}$. Hence for this naive notion of taking the standard part, the result is dependent on the representation of the graph: not a desirable state of affairs. Again, we can solve this by defining a notion of *relative standard parts*:

10.1.5. DEFINITION. Let $G = (V, E, w)$ be a weighted graph, and let $v \in V$ be a vertex. The *standard part* $G^{(v)}$ *of* G *relative to* v is a graph $G^{(v)} = (V^{(v)}, E^{(v)}, w^{(v)})$ containing the vertex v, such that

(a) $G^{(v)}$ is standard relative to v.

(b) The vertices and edges of G within standard G-distance of v are the same as the vertices and edges of $G^{(v)}$ within standard $G^{(v)}$-distance of v.

(c) For any edge e within standard G-distance of v, $w^{(v)}(e) = {}^{\circ}(w(e))$.

For example, the standard part of the graph G_0 of Example 10.1.1 relative to v_0 is G_0 itself, but the standard part relative to v_N for some hyperfinite N is the graph G_* from the remark preceding the definition. Obviously, if G is uniformly standard, then the standard part of G relative to any vertex v is G itself.
Note that if G is a graph and v is a vertex, and within standard distance of v G has a vertex of nonstandard degree or an edge of hyperfinite weight, then it is easily seen that the standard part of G relative to v cannot exist. However, if this isn't the case, then the standard part of G relative to v does exist, and is unique up to isomorphism.

10.1.6. LEMMA. *Let $G = (V, E, w)$ be a weighted graph, and let $v \in V$ be a vertex, such that within standard distance of v, all vertices are of standard degree and all edges are of nearstandard weight. Then the standard part $G^{(v)}$ of G relative to v is well-defined and unique (up to isomorphism).*

Proof

Let G and v be as given. We will construct the relative standard part explicitly and show uniqueness afterwards. Our basic strategy will be to first change the representation, then standardize the vertex-set, edge-set and weight function, and finally change the representation back again.

We can arrange the vertices of G in order of ascending distance to v (resolving ties arbitrarily), thus creating a mapping $m : V \to \mathbb{N}$. Let $G^v = (\mathbb{N}, E^w, w^v)$ be the image of G under this mapping. By assumption, all vertices within standard distance of v are of standard degree. Hence the vertices within standard distance of v are exactly those that are mapped to standard numbers. It follows that the edges and vertices of G within standard distance of v correspond exactly to the standard edges and vertices of G^v, with v itself corresponding to 0.

Since all edges of G within standard distance of v have nearstandard weight, the standard part ${}^\circ(w^v)$ of the weight function of G^v exists. Now consider the standard graph ${}^\circ(G^v) = (\mathbb{N}, {}^\circ(E^v), {}^\circ(w^v))$. The standard vertices and edges of ${}^\circ(G^v)$ and G^v are the same. Moreover, both in G^v and ${}^\circ(G^v)$, no paths of standard length exist between standard and nonstandard vertices, so the standard vertices and edges are exactly those within standard distance of 0. So the edges and vertices of G within standard distance of v correspond exactly to the edges and vertices of ${}^\circ(G^v)$ within standard distance of 0. This correspondence is again given by the mapping m. Furthermore, for any edge $e \in E$ within standard distance of v, $({}^\circ(w^v))(m(e)) \approx w(e)$.

Now all we need to do to finish the construction is to represent each standard vertex of ${}^\circ(G^v)$ by the corresponding vertex of G. This can be done by taking the image of ${}^\circ(G^v)$ under the inverse mapping m^{-1}. The resulting graph $G^{(v)} = (V^{(v)}, E^{(v)}, w^{(v)})$ satisfies all the conditions of Definition 10.1.5.

It remains to show that the standard part of G relative to v is unique up to isomorphism. So let $G^{(v)}$, $G^{(v)'}$ be two graphs that satisfy the conditions of Definition 10.1.5. Then they are isomorphic with standard graphs H and H', with v corresponding to standard vertices u, u' of those graphs. For all standard $n \in \mathbb{N}$, the subgraphs of H and H' induced by the vertices within distance n of u or u' are isomorphic. Since the two constants involved, H and H', are standard, we can apply the Transfer Axiom to obtain that the above holds for *all* $n \in \mathbb{N}$. We conclude that H and H' are isomorphic.

□

10.1.7. REMARK. If we study what happens when we take the standard part $G^{(v)}$ of a graph G relative to a vertex v, we find that all the features found within standard distance of v are preserved, and all the features found outside standard distance of v are ignored. Speaking informally, $G^{(v)}$ can be viewed as the standard-distance neighborhood of v in G, extended to an entire graph. Within standard distance of v, $G^{(v)}$ approximates G. Outside of this area, $G^{(v)}$ continues the pattern found inside. For instance, if within standard distance of v, G looks like a square lattice, then $G^{(v)}$ looks like a square lattice *everywhere*.

10.1.8. Remark. By Remark 6.1.1, when we consider a graph G we assume that G is a connected countably infinite simple graph, without vertices with infinitely many neighbors. It is easy to see that the same then holds for $G^{(v)}$. For example, since all standard vertices of G^v are connected to v, so are all the standard vertices of ${}^{\circ}(G^v)$, and since ${}^{\circ}(G^v)$ is a standard graph, by Transfer this holds for *all* vertices of ${}^{\circ}(G^v)$. However, by Remark 6.1.1 we also assume that there are no 'degenerate' edges of weight 0, and this property is *not* preserved. In particular, if G contains edges of infinitesimal weight, $G^{(v)}$ may contain edges of weight 0, and deleting those edges might cause $G^{(v)}$ to be disconnected.
In using the graphs $G^{(v)}$ we will state preconditions to prevent this.

10.2 Large Reinforcements and Uniform growth

So let us start by considering what happens if the reinforcement δ is hyperfinite. Then an edge that has never been traversed before has an infinitesimally small chance of being selected, compared to an edge that *has* been walked before. As a result, the once-reinforced random walk will remain in the subgraph of traversed edges for (on average) hyperfinitely long periods. Each such period ends when the walk leaves the subgraph (thereby extending it with a new edge) and during each period the walk behaves like a non-reinforced random walk.
Now consider this non-reinforced random walk. Generalizing the situation, we have a graph and a subgraph, with the property that the weight of edges outside the subgraph is infinitesimal compared to the weight of edges in the subgraph, and we are considering the non-reinforced random walk which stops when it leaves this subgraph. If we compare this walk to the non-reinforced random walk in the standard part of the subgraph relative to some arbitrary vertex v_c, then we see that if the latter is almost surely recurrent, then the former will visit each vertex of the subgraph within standard distance of v_c, on average, a hyperfinite number of times. Furthermore, this number will be proportional to the total weight of edges adjacent to that vertex. This is enough to show that the probability that the walk ends by leaving the subgraph by a particular edge (and hence the probability that in the original once-reinforced random walk that edge will be added to the subgraph of traversed edges), is nearly the same for edges that are close together, relative to their weight.

10.2.1. Lemma. *Let $G = (V, E, w)$ be a weighted graph, and let $G' = (V', E', w')$ be a finite*[26] *connected subgraph of G with edges of non-infinitesimal weight. Consider the random walk starting from a vertex $v_0 \in V'$ and walking randomly in G*

[26] Again, this includes hyperfinite

until an edge in $E - E'$ is traversed. Now, if v^1u^1 and v^2u^2 are edges in $E - E'$, then

$$\frac{P(\text{the random walk ends by walking from } v^1 \text{ to } u^1)/w(v^1u^1)}{P(\text{the random walk ends by walking from } v^2 \text{ to } u^2)/w(v^2u^2)} \approx 1 \tag{10.1}$$

provided that there exists a vertex v^c, with $d_{G'}(v^c, v^1)$ and $d_{G'}(v^c, v^2)$ standard, such that the following holds

(a) If $e \in E - E'$ and $d_{G'}(v^c, e)$ is standard, then $w(e)$ is infinitesimal.

(b) The standard part $G'^{(v^c)}$ of G' relative to v^c exists and has no degenerate edges, i.e. if $e \in E'$ is an edge with $d_{G'}(v^c, e)$ standard, then its weight $w'(e)$ is nearstandard and non-infinitesimal, and if $u \in V'$ is a vertex with $d_{G'}(v^c, u)$ standard, then its degree $\rho_{G'}(y)$ is standard.

(c) Non-reinforced random walks on $G'^{(v^c)}$ are almost surely recurrent[27]*.*

Proof

Let G, G', v_0 and v^c be as stated, and let $G^{(v^c)} = (V'^{(v^c)}, E'^{(v^c)}, w'^{(v^c)})$ be the standard part of G relative to v^c. As before, we denote the combined weights of all the edges of G, G' and $G'^{(v^c)}$ adjacent to the vertex v by $w(v)$, $w'(v)$ and $w'^{(v^c)}(v)$, respectively. Note that for all $v \in V$ within standard G'-distance of v^c, $w(v) \approx w'(v) \approx w'^{(v^c)}(v)$.

Note that the random walk cannot end by traversing $\overrightarrow{v^1u^1}$ or $\overrightarrow{v^2u^2}$ without first visiting v^1 or v^2. So if the random walk does not start at v^1 or v^2, then we can walk until the first visit to v^1 or v^2 (or the walk ends by traversing some other edge $E - E'$), and therefore we can write the probabilities to be calculated as linear combinations of the probabilities for the cases where the random walk *does* start at v^1 or v^2. Hence it suffices to prove the Lemma for the case where $v_0 = v^1 = v^c$.

For any $v \in V$, there is some path from v to an edge of $E - E'$, and some $p_v > 0$, such that at every visit to v the walk will proceed with probability $> p_v$ by traversing the route to the edge of $E - E'$ and then traversing that edge. Then the expected number of visits to v before the random walk leaves G' is at most $1/p_v$, and therefore finite. Denote this number by $F(v)$. For any edge $vu \in E$, the expected number of traversals from v to u is equal to $F(v)w(vu)/w(v)$. Obviously, except for the initial visit to v_0, a vertex v is visited once for each time an edge is traversed to v. Hence F satisfies, for all $v \in V'$,

$$F(v) = \sum_{u \in N_{G'}(v)} \frac{F(u)w(vu)}{w(u)} + \begin{cases} 1 & \text{if } v = v_0 \\ 0 & \text{if } v \neq v_0 \end{cases} \tag{10.2}$$

[27] If non-reinforced random walks on $G'^{(v^c)}$ are almost surely recurrent, then each vertex of G' is going to be visited a hyperfinite number of times, and therefore the location of the starting vertex is of negligible importance. Otherwise the location of the starting vertex affects the probabilities, and the conclusion of the lemma does not hold. Hence this is a necessary requirement.

Setting $H(v) = \frac{F(v)}{w(v)}$, we obtain

$$H(v) \sum_{u \in N_G(v)} w(vu) = \sum_{u \in N_{G'}(v)} H(u)w(vu) + \begin{cases} 1 & \text{if } v = v_0 \\ 0 & \text{if } v \neq v_0 \end{cases} \tag{10.3}$$

which can be rewritten as

$$\sum_{u \in N_{G'}(v)} w(vu)(H(v) - H(u)) = -H(v)(w(v) - w'(v)) + \begin{cases} 1 & \text{if } v = v_0 \\ 0 & \text{if } v \neq v_0 \end{cases} \tag{10.4}$$

Now we can write down the following inequality:

$$\begin{aligned} 0 \ &\geq \ \sum \{w(vu)(H(v)-H(u)) \mid v,u \in V', vu \in E', H(v) > H(v_0)\} && (10.5) \\ &= \ \sum \{w(vu)(H(v)-H(u)) \mid v,u \in V', vu \in E', H(v) > H(v_0) \geq H(u)\} && (10.6) \\ &\geq \ 0 && (10.7) \end{aligned}$$

All sums in this inequality are defined, since V' is finite. It follows that all sums are equal to 0, and hence for all $v \in V'$, $H(v) \leq H(v_0)$. So setting $H'(v) = H(v)/H(v_0)$, we have that for all $v \in V'$, $H'(v)$ is nearstandard (and ≤ 1), and H' satisfies

$$\sum_{u \in N_{G'}(v)} w(vu)(H'(v) - H'(u)) \approx \begin{cases} 1/H(v_0) & \text{if } v = v_0 \\ 0 & \text{if } v \neq v_0 \end{cases} \tag{10.8}$$

for all $v \in V'$ within standard G'-distance of v^c.

By Standardization, there exists a function $H'^{(v^c)}$ on $V^{(v^c)}$ such that for all $v \in V'$ within standard G'-distance of v^c, $H'^{(v^c)}(v) = {}^\circ H'(v)$.[28] Then the following equation holds for all $v \in V^{(v^c)}$ that are within standard G'-distance of v^c, and hence by Transfer for all $v \in V^{(v^c)}$:

$$\sum_{u \in N_{G'^{(v^c)}}(v)} w'^{(v^c)}(vu)(H'^{(v^c)}(v) - H'^{(v^c)}(u)) = \begin{cases} {}^\circ(1/H(v_0)) & \text{if } v = v_0 \\ 0 & \text{if } v \neq v_0 \end{cases} \tag{10.9}$$

Clearly $H'^{(v^c)}$ is a well-defined bounded superharmonic function on $G'^{(v^c)}$. By assumption random walks on $G'^{(v^c)}$ are almost surely recurrent, hence by Theorem 6.3.5, there exist no *non-constant* bounded superharmonic functions on $V^{(v^c)}$. Therefore $H'^{(v^c)}$ must be a constant function on $V^{(v^c)}$, i.e. for all $v \in V^{(v^c)}$, $H'^{(v^c)}(v) = 1$. It follows that $H(v_0)$ is hyperfinite, and for all $v \in V'$ within standard G'-distance of v^c, $H(v)/H(v_0) \approx 1$.

[28]We obtain H' by mapping the domain of H to the graph G^{v^c} of Lemma 10.1.6, taking the standard part, and mapping the domain back to G^{v^c}. It is not generally possible to standardize the values of a function without changing its domain. In a sense, the whole purpose of the graph $G'^{(v^c)}$ is to provide the new domain.

Now we have that

$$P(\text{the random walk ends with } vu) = \frac{F(v)w(vu)}{w(v)} = w(vu)H(v) \tag{10.10}$$

and hence

$$\frac{P(\text{the random walk ends by walking from } v \text{ to } u)/w(vu)}{P(\text{the random walk ends by walking from } v' \text{ to } u')/w(v'u')} = \frac{H(v)}{H(v')} \approx 1 \tag{10.11}$$

□

Now let us apply this lemma to the reinforced walk with hyperfinite reinforcement factor δ. As before, let E_t denote the set of edges already-traversed at time t, and A_t the set of arcs (oriented edges) obtained by orienting all edges of E_t in the direction of first traversal. We say that an arc $\overrightarrow{vu}$ *extends* A_t if it could be added to A_t, i.e. if A_t does not contain $\overrightarrow{vu}$ but does contain an edge incident with v. Then the above lemma can be used to show that all edges that can extend A_t and are standardly close together, have approximately the same probability of being added to A_t, relative to their weight.
To apply the above lemma, we need to ensure that non-reinforced random walks on the standard part of the subgraph of traversed edges are always recurrent. We do this by requiring the existence of a uniformly standard graph on which non-reinforced random walks are recurrent, of which the original graph is required to be a subgraph. Note that, in order to be able to apply the result later, we will allow stopping times satisfying certain conditions.

10.2.2. Corollary. *Let $G = (V, E, w)$ be a weighted graph. Assume that all edges have non-infinitesimal, non-hyperfinite weights, and that G is a subgraph (possibly with lesser weights) of a uniformly standard graph G^* on which non-reinforced random walks are recurrent.*
Consider the once-reinforced random walk on G with hyperfinite reinforcement factor δ, let $t_0 \in I\!N$ and fix the history up to time t_0. Let B be a set of arcs extending E_{t_0}, all of which have tails within standard E_{t_0}-distance of one another. Then for any arc $\overrightarrow{vu} \in B$,

$$\frac{P(\textit{the first arc added to } A_t \textit{ is } \overrightarrow{vu})/w(vu)}{P(\textit{the first arc added to } A_t \textit{ is from } B)/w(B)} \approx 1 \tag{10.12}$$

where $w(B)$ denotes the sum of the weights of all arcs of B.
This also holds if the walk stops at a stopping time τ, provided the walk is certain not to stop within standard E_{t_0} distance of an arc of B.

Proof

Given the history up to time t_0, set G' to the graph of vertices and edges that have been traversed. For all v, the graph $G'^{(v)}$ is isomorphic to a subgraph of G^*

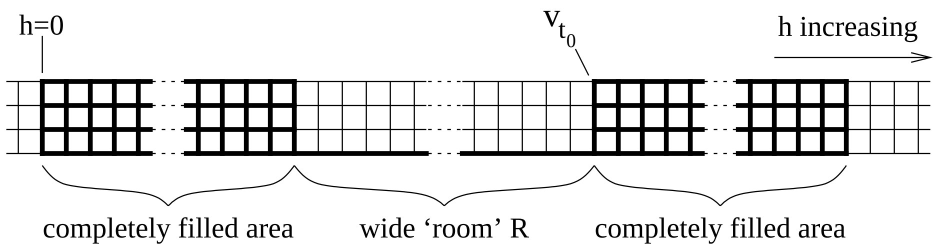

Figure 10.1: An anomalous situation

(possibly with lesser weights) by the uniform standardness of G^*, so by Theorem 6.3.8, random walks on $G'^{(v)}$ are almost surely recurrent for all v. If we divide all weights by δ, then all edges of E' have non-infinitesimal weight and all edges of $E - E'$ have infinitesimal weight. Hence, if we set v^c to be any vertex within standard E_{t_0}-distance of B, then the conditions of Lemma 10.2.1 are satisfied for any pair of arcs $\overrightarrow{vu}, \overrightarrow{v'u'} \in B$. The result follows.
If a stopping time τ is given, then by the given condition on τ and the Overflow Principle, there exists a hyperfinite N such that the random walk will not stop within E_{t_0}-distance N of an arc of B. By limiting G' to arcs and edges within E_{t_0}-distance $N - 1$ of arcs of B, there will be some edges of $E - E'$ with non-infinitesimal weight, but these will not be within standard E_{t_0}-distance of any arc of B, so the conditions of Lemma 10.2.1 will still be satisfied. If the walk leaves G' and then returns without stopping or adding an edge to A_t, we apply Lemma 10.2.1 again until the walk has stopped or an edge has been added to A_t.

□

10.3 Recurrence of Random Walks with Large Reinforcements

If the previous corollary would hold for arcs whose tails were arbitrarily far apart, then we would be done. For then the growth of A_t would always be uniform, and the expected change to the bias would be proportional to the weighted average of $\Delta_h(\overrightarrow{vu})$ over all arcs $\overrightarrow{vu}$ extending A_t. For any (super)harmonic function h this can be shown to be positive, and hence we would then be able to apply Lemma 8.3.4 to show recurrence on any graph.
Unfortunately, Corollary 10.2.2 is not that strong. In order to be able to use the corollary, we need to divide the graph into *smaller* areas, such that the arcs in that area extending A_t are standardly close together and have, on average, nonnegative $\Delta_h(\overrightarrow{vu})$. Then we can take the sum over the areas, and arrive at a nonnegative expected change to the total bias.
As it turns out, this does not always work. Figure 10.1 shows a situation where it doesn't, arising in the random walk on the ladder $\mathbb{Z} \times \{1, 2, 3, 4\}$ with $h(v) =$

$|x(v)|$. Thick lines in the figure denote edges that have been traversed (and hence are in A_t), and thin lines denote edges that are in the graph but have not yet been traversed. In this particular situation, all edges extending A_t that are not far away from v_{t_0} are contained in an extremely-wide 'room' R. If the next horizontal edge added to A_t is from the right side of the room, then it will be traversed right-to-left, and will contribute a negative Δ_h to the bias. If it is from the left side of the room, then it will be traversed left-to-right, and will contribute a positive Δ_h. $h(v) = |x(v)|$ is a harmonic function, and the harmonicity equations ensure that these potential contributions have a zero sum. So all would be well if the edges would have equal probabilities of being added to A_t. But if the width of the room is large enough (relative to δ), then in the reinforced random walk starting at v_{t_0}, the next arc added to A_t will be much more likely to come from the right side of the room than from the left side. In the given situation therefore, the expectation of the next change to the bias is negative.

So we need to make sure that such anomalous situations occur so rarely that they will have a negligible effect on the expectation of the change to the bias. In the case of the next theorem, this problem is addressed by considering only graphs that can be viewed as rows of connected vertex-sets V_i.[29] At any time there will be sets V_i such that all edges in the subgraph V_i have been traversed, which will act as 'walls', dividing the graph into 'rooms' similar to the one shown in figure 10.1. If the size of the sets V_i is standardly bounded, these walls will be only standardly far apart (on average). Anomalous situations can only occur in rooms with extremely large width, much greater than average, and therefore they will be extremely rare (outside a given initial situation) and will have little effect on the expectation of the bias.

This structure also solves an additional problem, namely that in general edges that are close together in G are not guaranteed to be close together in the subgraph of traversed edges, and if they are not, A_t has no uniform growth there. This cannot happen in rooms of small width, since the walls 'short-circuit' long paths.

10.3.1. THEOREM. *Let $G = (V, E, w)$ be a weighted graph, such that w has standard upper and lower bounds w_{max} and $w_{min} > 0$, there is a standard bound ρ_{max} to the degree of the vertices, and there is a partition $(V^i)_{i \in \mathbb{Z}}$ of V satisfying the following:*

1. *$|V^i|$ is bounded by some standard $k \in \mathbb{Z}$.*

2. *All edges in G are between vertices of V^i, or between vertices of V^i and V^{i+1}, for some $i \in \mathbb{Z}$.*

3. *For all i, $G_{|V^i}$ is connected.*

[29] This class of graphs may be viewed as a generalization of the infinite ladders $\mathbb{Z} \times \{1, \ldots, n\}$ and cylinders $\mathbb{Z} \times (\mathbb{Z}/n\mathbb{Z})$ of Theorem 8.3.5.

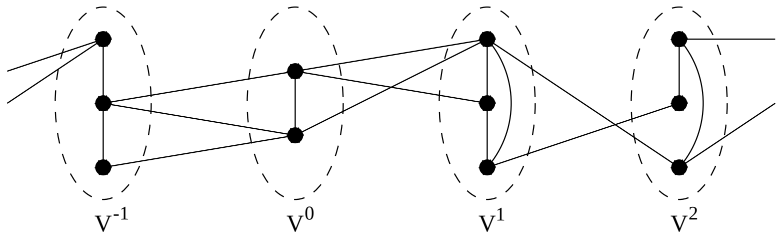

Figure 10.2: An example of a graph for Theorem 10.3.1, with $k = 3$.

Then a once-reinforced random walk on G with hyperfinite reinforcement factor δ is almost surely recurrent.

To show recurrence of the random walk, it suffices to show that at any time t_0, the probability of returning to V^0 is 1. Hence we may assume, without loss of generalization, that $v_{t_0} \in V^{\geq 0} = V^0 \cup V^1 \cup \ldots$, and that for purposes of this proof, $V^{-1} \cup V^{-2} \cup \ldots$ can be safely ignored. Unless stated otherwise, we assume $i \geq 0$ whenever we write V^i.
Before we start with the proof of this theorem, we need to a little more work on our tools. In particular we need to show that Corollary 10.2.2 applies, and since we are going to have to calculate with the results, it seems desirable to find explicit limits to replace 'standardly close' and '≈ 1'.

10.3.2. LEMMA. *Let G be a graph as in Theorem 10.3.1, and consider the once-reinforced random walk on G with hyperfinite reinforcement factor δ. Then there exists a hyperfinite $N \in \mathbb{N}$ and an infinitesimal $\eta > 0$ such that*

> *at any time t_0, if B is a set of arcs extending A_{t_0}, within E_{t_0}-distance N of one another, then for any arc $\overrightarrow{vu} \in B$*

$$\frac{1}{1+\eta} \leq \frac{P(\textit{the first arc added to } A_t \textit{ is } \overrightarrow{vu})/w(vu)}{P(\textit{the first arc added to } A_t \textit{ is from } B)/w(B)} \leq \frac{1}{1-\eta} \tag{10.13}$$

This also holds for the random walk that stops at V^0 if we fix a hyperfinite $c_{dist} \in \mathbb{R}$ and add the condition that

> *no arc of B is within G-distance c_{dist} of V^0*

Proof
Let k and w_{max} be as in Theorem 10.3.1, and define the graph $G^* = (V^*, E^*, w^*)$ by setting $V^* = \mathbb{Z} \times \{1, \ldots, k\}$, $E^* = \{((i_1, j_1), (i_2, j_2)) \mid |i_1 - i_2| \leq 1\}$ and $w^*(e) = w_{max}$. Then G is a subgraph of G^* (possibly with lesser weights). Moreover, it is easily seen that G^* is uniformly standard, and that non-reinforced random walks on G^* are almost surely recurrent. Hence the conditions of Corollary 10.2.2 are satisfied. If $c_{dist} \in R$ is hyperfinite and $N \in \mathbb{N}$ and $\eta > 0$ are standard, then the

conclusion of Corollary 10.2.2 is stronger than the statement we want to prove, so that statement holds. The Overflow Principle states that if it holds for all standard N and $\eta > 0$, then it also holds for some hyperfinite N and infinitesimal $\eta > 0$.

□

Since G is a subgraph of the graph G^* from the above lemma, non-reinforced random walks on G are almost surely recurrent, and hence there exists a superharmonic function on $V - V^0$ witnessing this, which we can use when applying Lemma 8.3.4. But using the given restrictions on G, we can construct a function h with some additional nice properties, which will simplify later calculations considerably:

10.3.3. Lemma. *Let G and δ be as above. Then there exists a function $h : V \to \mathbb{R}$ satisfying*

1. *h is harmonic on $V - V_0$.*
2. *$h \to \infty$ if $v \to \infty$.*
3. *For all vertices $vu \in E$, $|\Delta_h(\overrightarrow{vu})| \leq 1/w(vu) \leq 1/w_{min}$.*
4. *For all $i \in \mathbb{N}$, $\sum\{w(vu)\Delta_h(\overrightarrow{vu}) \mid v \in V^i, u \in V^{i+1}, vu \in E\} = 1$.*

Proof

Begin by defining, for $i > 0, v \in V^{\geq 0}$,

$$h_i(v) = P(\text{starting in } v,\, V^i \text{ is reached before } V^0) \qquad (10.14)$$

$$q_i = \sum\{w(vu)\Delta_{h_i}(\overrightarrow{vu}) \mid v \in V^0, u \in V^1, vu \in E\} \qquad (10.15)$$

It is easily seen that h_i is harmonic on $V^1 \cup V^2 \cup \ldots \cup V^{i-1}$, $h_i(v) = 0$ for $v \in V^0$, and $q_i > 0$. There exists a function $h : V^{\geq 0} \to \mathbb{R}_{\geq 0}$ that is a limit of $(h_i)_{i=1}^{\infty}$ in the sense that, for any finite set $F \subset V^{\geq 0}$, $(h(v))_{v \in F}$ is a limit point of $((h_i(v)/q_i)_{v \in F})_{i=1}^{\infty}$. It follows that this function h is harmonic on $V^1 \cup V^2 \cup \ldots$, $h(v) = 0$ for $v \in V^0$, and

$$\sum\{w(vu)\Delta_h(\overrightarrow{vu}) \mid v \in V^0, u \in V^1, vu \in E\} = 1 \qquad (10.16)$$

By taking linear combinations of the harmonicity equality, we obtain

$$\sum\{w(vu)\Delta_h(\overrightarrow{vu}) \mid v \in V^i, u \in V^{i+1}, vu \in E\} = 1 \text{ for } i \in \mathbb{N} \qquad (10.17)$$

$$\sum\{w(vu)\Delta_h(\overrightarrow{vu}) \mid v, u \in V, h(v) \leq r < h(u), vu \in E\} = 1 \text{ for } r \geq 0 \qquad (10.18)$$

From (10.18) we see that for all $vu \in E$ $|\Delta_h(\overrightarrow{vu})| \leq 1/w(vu) \leq 1/w_{min}$.
Finally, let $r \geq 0$, and set $U_r = \{v \in V^{\geq 0} \mid h(v) \leq r\}$. If U_r were to have an infinite connected component, then for some $i_0 > 1$, U_r would intersect all V^i with

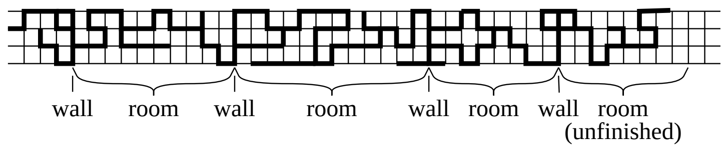

Figure 10.3: Traversed edges forming walls and rooms in the square lattice graph on $\mathbb{Z} \times \{1, 2, 3, 4\}$

$i \geq i_0$, and hence by the connectedness of the sets V_i, $h(v) < r + kw_{min}$ for all $v \in V^{\geq i_0}$. Then $-h$ would be a bounded non-constant superharmonic function, contradicting the fact that non-reinforced random walks on G are almost surely recurrent. On the other hand, if U_r were to have a finite non-empty component F disjoint from V^0, we could take linear combinations of the harmonicity equality to obtain

$$\sum \{w(vu)\Delta_h(\overrightarrow{vu}) \mid v \in F, u \in V - F, vu \in E\} = 0 \tag{10.19}$$

contradicting the fact that if $v \in F, u \in V - F$, then $h(u) \geq r > h(v)$. So for all $r \geq 0$, U_r consists only of a finite component containing V^0. We conclude that $h(v) \to \infty$ if $v \to \infty$, and that h is a proper witness to the recurrence of non-reinforced random walks in G.

□

Proof of Theorem 10.3.1

Let h be the function of Lemma 10.3.3. We will use Lemma 10.3.2 to show that the condition of Lemma 8.3.4 is satisfied.

We say that there is a *wall* at V^i, when all edges of E with both vertices in V^i are in E_t. The area between two successive walls we will call a *room*. Given that there are at most $k\rho_{max}$ edges involved, it can easily be seen that there is a non-infinitesimal $p_{wall} > 0$ such that the probability that a wall appears at V^i before a single vertex of V^{i+2} is visited is at least p_{wall}. It follows that, not taking into account rooms already present in any given initial situation, the expected average width of a room is at most p_{wall}^{-1}.

Let $t > 0$, and fix the history up to time t. Let $\overrightarrow{v^*u^*}$ denote the next arc (at or after time t) that is added to A_t. Now first consider a room R with walls at V^i and V^j and width $\leq N/k$, at distance $\geq c_{dist}$ from the origin. It is easily seen that any non-self-intersecting E_t-path that does not extend beyond the room is of length at most N, and since walls connect all their vertices, paths that leave the room and return can be 'short-circuited'. Hence all arcs in the room extending A_t have tails with E_t-distance N or less to one another, and for any such arc $\overrightarrow{vu}$,

$$P(\overrightarrow{v^*u^*} = \overrightarrow{vu} \mid \mathcal{F}_t)(1-\eta) \leq P(v^*u^* \text{ in } R \mid \mathcal{F}_t)\frac{w(vu)}{w_{ext}(R)} \leq P(\overrightarrow{v^*u^*} = \overrightarrow{vu} \mid \mathcal{F}_t)(1+\eta) \tag{10.20}$$

where $w_{ext}(R)$ denotes the total weight of all arcs in R that extend A_t. It follows that

$$P(\overrightarrow{v^*u^*}=\overrightarrow{vu}\,|\,\mathcal{F}_t)(\Delta_h(\overrightarrow{vu})+\eta|\Delta_h(\overrightarrow{vu})|)\geq\frac{P(v^*u^*\text{ in }R\,|\,\mathcal{F}_t)}{w_{ext}(R)}w(vu)\Delta_h(\overrightarrow{vu}) \quad (10.21)$$

By taking linear combinations of the harmonicity equation, we can show that

$$\sum_{\substack{vu\text{ in }R\\ \overrightarrow{vu}\text{ extends }A_t}} w(vu)\Delta_h(\overrightarrow{vu})=0 \quad (10.22)$$

and hence

$$\sum_{\substack{vu\text{ in }R\\ \overrightarrow{vu}\text{ extends }A_t}} P(\overrightarrow{v^*u^*}=\overrightarrow{vu}\mid\mathcal{F}_t)\,(\Delta_h(\overrightarrow{vu})+\eta|\Delta_h(\overrightarrow{vu})|)\geq 0 \quad (10.23)$$

Similarly, if we let i_t^{max} denote the largest index $i>0$ such that a vertex of V^i has been visited at time t, then the area between $V^{i_t^{max}+1}$ and the wall V^{i_0} with the largest index i_0 can be considered to be an 'unfinished room' R_{last}. Here, taking linear combinations of the harmonicity equation yields

$$\sum_{\substack{vu\text{ in }R_{last}\\ \overrightarrow{vu}\text{ extends }A_t}} w(vu)\Delta_h(\overrightarrow{vu})=1 \quad (10.24)$$

and if R_{last} has width $\leq N/k$, we derive in the same manner as before

$$\sum_{\substack{vu\text{ in }R_{last}\\ \overrightarrow{vu}\text{ extends }A_t}} P(\overrightarrow{v^*u^*}=\overrightarrow{vu}\mid\mathcal{F}_t)\,(\Delta_h(\overrightarrow{vu})+\eta|\Delta_h(\overrightarrow{vu})|)\geq\frac{P(v^*u^*\text{ in }R_{last}\mid\mathcal{F}_t)}{w_{ext}(R_{last})} \quad (10.25)$$

where $w_{ext}(R_{last})$ denotes the total weight of all arcs in R_{last} that extend A_t. Hence, if we pretend for the moment that we do not need to take into consideration edges within distance c_{dist} of the V^0, rooms of width $>N/k$, or the fact that to apply Lemma 8.3.4 we need to show that its condition holds *given* an arbitrary fixed history up to some time t_0, we obtain

$$E\left(\Delta_h(\overrightarrow{v^*u^*})+\eta|\Delta_h(\overrightarrow{v^*u^*})|\,\middle|\,\mathcal{F}_t\right)\geq\frac{P(v^*u^*\text{ in }R_{last}\mid\mathcal{F}_t)}{w_{ext}(R_{last})} \quad (10.26)$$

Combined with the inequality

$$P(i_{t+1}^{max}=i_t^{max}+1\mid\mathcal{F}_t)\leq\frac{k\rho_{max}w_{max}}{w_{ext}(R_{last})}\cdot P(v^*u^*\text{ in }R_{last}\mid\mathcal{F}_t)\cdot\frac{1}{1-\eta} \quad (10.27)$$

this implies

$$E\left(\Delta_h(\overrightarrow{v^*u^*}) + \eta|\Delta_h(\overrightarrow{v^*u^*})| - (i_{t+1}^{max} - i_t^{max})\frac{1-\eta}{k\rho_{max}w_{max}}\,\middle|\,\mathcal{F}_t\right) \geq 0 \qquad (10.28)$$

Now, if τ is a stopping time, then we can take the sum of the above equation over all $t < \tau$, to obtain

$$E\left(\sum_{\overrightarrow{vu}\in A_\tau} \Delta_h(\overrightarrow{vu}) + \eta \sum_{\overrightarrow{vu}\in A_\tau} |\Delta_h(\overrightarrow{vu})| - i_\tau^{max}\frac{1-\eta}{k\rho_{max}w_{max}}\right) \geq 0 \qquad (10.29)$$

Now let us consider those aspects that we chose to ignore before. We will show that the number of edges involved in those aspects is relatively small. One of the aspects we ignored was the existence of rooms of width N/k or greater. But if a new room is 'growing', then the probability is less than $(1-p_{wall})^{N/k}$ that the room will reach a width of N/k or greater, and even in that case the expected width of the room is at most $N/k + p_{wall}^{-1}$. Hence the expectation of the number of sets V^i contained in rooms of width $> N/k$ at any time τ is at most

$$E(i_\tau^{max})(1-p_{wall})^{N/k}(N/k + p_{wall}^{-1}) \qquad (10.30)$$

Another aspect we ignored was that, in order to apply Lemma 8.3.4 we need to show that its condition holds *given* an arbitrary fixed history up to some time t_0. But any such initial situation is contained within the area defined by $V^0 \cup \ldots \cup V^{i_{t_0}^{max}}$. The same holds for edges within distance c_{dist} of the origin: all such edges are contained in $V^0 \cup \ldots \cup V^{c_{dist}}$. Edges in rooms extending from one of these areas may be affected, but the expected number of sets V^i involved in such an 'extension' is at most p_{wall}^{-1}. Hence the expected number of sets V^i for which the previous calculations do not apply is at most

$$\max(i_{t_0}^{max}, c_{dist}) + p_{wall}^{-1} + E(i_\tau^{max})(1-p_{wall})^{N/k}(N/k + p_{wall}^{-1}) \qquad (10.31)$$

Each set V^i contains at most k vertices, each of which is adjacent to at most ρ_{max} edges. Each edge vu that we 'counted wrongly' before might have caused the expectation above to be higher than it should have been, but the difference will be at most $2\Delta_h(vu) \leq 2/w_{min}$. So the expectation above may be higher than it should have been, but not by more than

$$\frac{2k\rho_{max}}{w_{min}}\Big(\max(i_{t_0}^{max}, c_{dist}) + p_{wall}^{-1} + E(i_\tau^{max}\,|\,\mathcal{F}_{t_0})(1-p_{wall})^{N/k}(N/k + p_{wall}^{-1})\Big) \qquad (10.32)$$

Combining everything, and taking into account that

$$\eta \sum_{\overrightarrow{vu}\in A_\tau} |\Delta_h(\overrightarrow{vu})| \leq \eta\frac{k\rho_{max}i_t^{max}}{w_{min}} \qquad (10.33)$$

we derive that for any stopping time τ

$$E\left(\sum_{\overrightarrow{vu}\in A_\tau} \Delta_h(\overrightarrow{vu}) \,\middle|\, \mathcal{F}_{t_0}\right) \geq c \cdot E(i_\tau^{max} \mid \mathcal{F}_{t_0}) - c' \tag{10.34}$$

with

$$c = \frac{1-\eta}{k\rho_{max}w_{max}} - \frac{k\rho_{max}(2(1-p_{wall})^{N/k}(N/k + p_{wall}^{-1}) + \eta)}{w_{min}} \tag{10.35}$$

$$c' = \frac{2k\rho_{max}(\max(i_{t_0}^{max}, c_{dist}) + p_{wall}^{-1})}{w_{min}} \tag{10.36}$$

Since η is infinitesimal and N is hyperfinite, $c \approx 1/k\rho_{max}w_{max} > 0$, and we conclude that the condition of Lemma 8.3.4 is satisfied.

□

10.3.4. COROLLARY. *Let n be standard. The once-reinforced random walks with hyperfinite reinforcement factor δ on the square lattice graphs on $\mathbb{Z} \times \{1, \ldots, n\}$ and $\mathbb{Z} \times (\mathbb{Z}/n\mathbb{Z})$ are almost surely recurrent.*

10.3.5. COROLLARY. *Let n be a natural number. The once-reinforced random walks with reinforcement factor δ on the square lattice graphs on $\mathbb{Z} \times \{1, \ldots, n\}$ and $\mathbb{Z} \times (\mathbb{Z}/n\mathbb{Z})$ are almost surely recurrent if δ is large enough, i.e. if $\delta > \delta_0$ for some $\delta_0 \in \mathbb{R}$.*

Proof

If n is standard, then the above holds for any hyperfinite δ_0. The property of 'for some δ_0, for all $\delta > \delta_0$, the random walk on the infinite ladder or cylinder with width n is almost surely recurrent' can be expressed without using nonstandard constants or the predicate $\mathrm{st}(x)$. Hence by Transfer the above holds for any n, standard or not. Note that if n is standard, it follows from the principle of Underflow that δ_0 may be taken standard. Also note that this Corollary is formulated without using nonstandard concepts.

□

10.3.6. REMARK. The proof above *can* be translated back to orthodox mathematics, at the expense of (even more) pages of computations. The adapted proof may be reduced in size somewhat by applying the techniques used by Doyle and Snell [25] to the proof of Lemma 10.2.1. The main contributor to the size of the lower bound δ_0 obtained by this proof, is the requirement that $N/k \gg p_{wall}^{-1}$, which unfortunately is hyperexponential in k. The resulting bound δ_0 is of order $(k\rho_{max}w_{max}/w_{min})^{2k\rho_{max}+6}$.

10.3.7. Remark. The function of walls in the above proof is, as was stated before, to limit the occurrence of anomalous situations, where the edges that are close together in the graph of traversed edges do not form areas balanced by the harmonicity equations and vica versa. This approach fails when considering the random walk on the square lattice on $\mathbb{Z}^2$. Although it is possible to view $\mathbb{Z}^2$ as a row of vertex-sets V_i, as i increases the size of the vertex-sets would also increase, the probability of a wall forming would converge to 0, and the average width of a room would diverge to ∞.
Although presumably the reasoning above would hold up in a variant random walk, where the value of δ increases with the distance from some arbitrary origin, this is a marginal result at best. A better avenue of investigation for this problem is likely to find some other way of limiting the expected occurrence of anomalous situations.

Part III

The EMILE Grammar Inducer

Overview

In this section of the dissertation we consider the EMILE program, a program that reads a text, and without prior knowledge, attempts to determine the grammatical structure of the language.

In chapter 11, we consider the problem of grammar inference, and introduce some of the basic concepts of EMILE.

In chapter 12, we study the algorithm underlying EMILE, starting with a very simple version of the basic algorithm, and changing it to the full algorithm in several steps, elaborating on the motivations for the change at each step.

In chapter 13, we consider the results of the EMILE program, both in theory and in practice. It is conjectured that natural languages satisfy the condition of *shallowness*, and that this implies that the EMILE program will work well for natural languages.

Finally, appendix A lists the sub-algorithms used in EMILE, giving both a synopsis and explicit pseudo-code for each sub-algorithm.

Chapter 11

The Basics of EMILE

11.1 Introduction

Human being are remarkably good in working with natural languages. Even if someone has no knowledge of the formal structure of a language, he or she will be able to tell when 'something' is like 'something else'. For instance, Lewis Caroll's famous poem 'Jabberwocky' starts with

> 'Twas brillig, and the slithy toves
> Did gyre and gimble in the wabe;
> All mimsy were the borogoves
> and the mome raths outgrabe.

Even without Humpty Dumpty's annotations, it is immediately obvious what the *syntactic structure* of the first sentence is: 'brillig' and 'slithy' are adjectives, 'toves' is a noun, 'gyre' and 'gimble' are verbs, etcetera.
So how do we know such things? The short answer is 'from context'. When a sentence starts with ' 'Twas', we are not surprised if the next word is an adjective. Similarly, if a sentence has the pattern 'the (.) did (.) and (.) in the (.)', we expect the missing phrases to be a noun-phrase, two verb-phrases and another noun-phrase, respectively.
The notion of *grammatical type* has many possible definitions. For instance, if we had a *context-free* grammar of a language, we can view each non-terminal symbol as a grammatical type. In general, one of the properties of a grammatical type is, that wherever some expression is used as an expression of that type, other expressions of that type can be substituted without making the sentence ungrammatical. This gives rise to a notion of a grammatical type as a set of expressions together with a set of contexts. For instance, the type 'noun-phrase' could be represented by the set of all noun-phrases, together with the set of all contexts in which a noun-phrase can appear. Combining any of the expressions of

a type with any of the contexts will yield a grammatical sentence. Many of these combinations might appear in actual texts, especially the short ones (in terms of number of words).

In this terminology, we can describe the above phenomenon, as the existence of contexts which are *characteristic* for a type, meaning that whenever something appears in that context, we assume it also belongs to that type. Some types may also have *characteristic expressions*, with the analogous property.

EMILE[30] 4.1 is a program based on the above concepts. It attempts to learn the grammatical structure of a language from sentences of that language, without being given any prior knowledge of the grammar. For any type in any valid grammar for the language, we can expect context/expression combinations to show up in a sufficiently large sample of sentences of the language. EMILE searches for such clusters of expressions and contexts in the sample, and interprets them as grammatical types. It then tries to find characteristic contexts and expressions, and uses them to extend the types. Finally, it formulates derivation rules based on the types found, in the manner of the rules of a context-free grammar. The program can present the grammatical structure found in several ways, as well as use it to parse other sentences or generate new ones.

The theoretical concepts used in EMILE 4.1 are elaborated on in P. Adriaans articles on EMILE 1.0/2.0 [33] and EMILE 3.0 [34]. in these chapters we will focus on the practical aspects. Note that although EMILE 4.1 is based on the same theoretical concepts as EMILE 3.0, it is not based on the same algorithm. More information on the precursors of EMILE 4.1 may be found in the above articles, as well as in the E. Dörnenburg's Master's Thesis[36].

11.2 Definitions

The three most basic concepts in EMILE are *contexts*, *expressions* and *context/expression pairs.*

11.2.1. DEFINITION. A *context/expression pair* is a sentence split into three parts, for instance

John (makes) tea

Here, 'makes' is called an *expression*, and 'John (.) tea' is called a *context* (with *left-hand side* 'John' and *right-hand side* 'tea').

11.2.2. DEFINITION. We say that an expression e *appears with* a context c, or that the context/expression pair (c, e) has been *encountered*, if $c_l\hat{}e\hat{}c_r$ appears

[30] EMILE 4.1 is a successor to EMILE 3.0, written by P. Adriaans. The original acronym stands for Entity Modeling Intelligent Learning Engine. It refers to earlier versions of EMILE that also had semantic capacities. The name EMILE is also motivated by the book on education by J.-J. Rousseau.

as a sentence in a text, where c_l and c_r are the left-hand side and the right-hand side of c, respectively, and $a^\frown b$ denotes the concatenation of a and b.

11.2.3. Remark. Context/expression pairs are not always sensible, as for instance in the sentence

John (drinks coffee, and Mary drinks) tea

where the expression 'drinks coffee, and Mary drinks' appears in the context 'John (.) tea'. EMILE will find such context/expression pairs and attempt to use them in the grammar induction process, But such pairs are usually isolated, i.e. they are not part of any significant clusters. So EMILE will fail to make use of them, and they will be effectively ignored.

As stated before, we view grammatical types in terms of the expressions that belong to that type, and the contexts in which they can appear (as expressions of that type). As such, we define grammatical types as follows:

11.2.4. Definition. In the context of this paper, a grammatical type T is defined as a pair (T_C, T_E), where T_C is a set of contexts, and T_E is a set of expressions. Elements of T_C and T_E are called *primary* contexts and expressions for T.

The intended meaning of this definition is, that all expressions of a type can appear with all of its the contexts.
In natural languages, the type of an expression is not always unambiguous. For instance, the word 'walk' can be both a noun and a verb. Hence 'walk' will not only appear in contexts for noun-phrases, but also in contexts for verb-phrases. The same does not hold for the phrase 'thing': 'thing' only appears in contexts for noun-phrases, and in any such context, any noun can be substituted for 'thing' without making the sentence ungrammatical. We say that 'thing' is *characteristic* for the type 'noun'. Formally,

11.2.5. Definition. An expression of a type T is *characteristic* for T if it only appears with contexts of type T.[31] Similarly, a context of a type T is *characteristic* for T if it only appears with expressions of type T.

In these chapters, we will also use *characteristic** and *secondary* expressions and contexts. However, as these definitions are rather dependent on the algorithms, they will be delayed until section 12.4. That section also has several examples of characteristic expressions and contexts.

[31] EMILE changes this definition slightly in implementation, in that contexts which have been assigned no type at all are completely ignored, i.e. an expression is characteristic if all contexts with which it appears are of type T, or untyped.

NOTATION. For any type T, T_E, T_E^{ch}, T_E^* and T_C^{se} denote the sets of primary, characteristic, characteristic* and secondary expressions of T, and T_C, T_C^{ch}, T_C^* and T_C^{se} denote the corresponding sets of contexts.

The EMILE program also attempts to transform the collection of grammatical types found into a context-free grammar consisting of derivation rules. Such rules generally are of the form

$$[T] \Rightarrow s_0[T_1]s_1[T_2]\dots[T_k]s_k$$

where $T, T_1, T_2, \dots, T_k$ are grammatical types, and $s_0, s_1, \dots, s_k$ are (possibly empty) sequences of words. Given a rule with left-hand side $[T]$, and a sequence of word-sequences and grammatical types containing $[T]$, that appearance of $[T]$ can be replaced by the right-hand side of the rule, (concatenating adjacent word-sequences as necessary). Any sequence which can be obtained from another sequence by such rule applications, is said to be derivable from that sequence. The language of a context-free grammar consists of those word-sequences e such that $[0] \Rightarrow e$ is derivable, where $[0]$ denotes the type of whole sentences.

Chapter 12

The algorithms of EMILE

This chapter attempts to give some insight into the reasoning underlying the algorithms of EMILE. We will start with a very simple version of the basic algorithm, and in several steps change it to the full algorithm, at each step elaborating on the motivations for the change.

12.1 1-dimensional clustering

Given a sample of sentences, we want to obtain sets of expressions and contexts that correspond to grammatical types. A simple clustering technique is to extract all possible context/expression combinations from a given sample of sentences, and group together expressions that appear with the same context.

12.1.1. Example. If we take the sample sentences 'John makes tea' and 'John likes tea', we get the following context/expression *matrix*:

	(.) makes tea	John (.) tea	John makes (.)	(.) tea	John (.)	(.)	(,) likes tea	John likes (.)
John	x						x	
makes		x						
tea			x					x
John makes				x				
makes tea					x			
John makes tea						x		
likes		x						
John likes				x				
likes tea					x			
John likes tea						x		

from which we can obtain the clusters

[‘John (.) tea’, {‘makes’, ‘likes’}]
[‘(.) tea’, {‘John makes’, ‘John likes’}]
[‘John (.)’, {‘makes tea’, ‘likes tea’}]
[‘(.)’, {‘John makes tea’, ‘John likes tea’}]

Next, we can group contexts together if they appear with exactly the same expressions.

12.1.2. Example. If we add the sentences ‘John makes coffee’, ‘John likes coffee’ to the previous sample, the relevant part of the context/expression matrix looks like

	John (.) tea	John (.) coffee	John makes (.)	John likes (.)
makes	x	x		
likes	x	x		
tea			x	x
coffee			x	x

which yields the clusters

[{‘John (.) tea’, ‘John (.) coffee’}, {‘makes’, ‘likes’}]
[{‘John makes (.)’, ‘John likes (.)’}, {‘tea’, ‘coffee’}]

As stated before, a grammatical type can be characterized by the expressions that are of that type, and the contexts in which expressions of that type appear. Hence the clusters we find here can be interpreted as grammatical types. For instance, the clusters in the above example could be said to correspond to the grammatical types of ‘verbs’ and ‘nouns’, respectively.

12.2 2-dimensional clustering

One of the flaws in this technique is that it doesn’t properly handle contexts whose type is ambiguous.

12.2.1. Example. If we add the sentences ‘John likes eating’ and ‘John is eating’ to the previous example, the relevant part of the context/expression matrix will look like this:

	John (.) tea	John (.) coffee	John (.) eating	John makes (.)	John likes (.)	John is (.)
makes	x	x				
likes	x	x	x			
is			x			
tea				x	x	
coffee				x	x	
eating					x	x

Here we can intuitively identify four grammatical types: noun-phrases, verb-phrases, 'ing'-phrases, and 'verbs-appearing-with-ing-phrases'-phrases. The context 'John likes (.)' is ambiguous, in the sense that it appears with both noun-phrases and 'ing'-phrases. If we proceed as before, we get the following clusters

[{'John (.) tea', 'John (.) coffee'}, {'makes', 'likes'}]
[{'John (.) eating'}, {'likes', 'is'}]
[{'John makes (.)'}, {'tea', 'coffee'}]
[{'John likes (.)'}, {'tea', 'coffee', 'eating'}]
[{'John is (.)'}, {'eating'}]

i.e. the context 'John likes (.)' is assigned a separate type.

Assigning ambiguous contexts a separate type not only results in a less natural representation, in a later step it will prevent us from correctly identifying the characteristic expressions of a type (as will be demonstrated in Example 12.4.2). A more natural representation would be to allow ambiguous contexts and expressions to belong to multiple types. For this, we need to use a different clustering method. The clustering method EMILE uses is to search for maximum-sized blocks in the matrix. This could be termed *2-dimensional clustering*.

12.2.2. Example. The following picture shows the matrix of the previous example, with the maximum-sized blocks indicated by rectangles.[32]

	John (.) tea	John (.) coffee	John (.) eating	John makes (.)	John likes (.)	John is (.)
makes	x	x				
likes	x	x	x			
is			x			
eating					x	x
tea				x	x	
coffee				x	x	

[32]Please note that the expressions and contexts have been arranged to allow the blocks to be easily indicated: in general, blocks will *not* consist of adjacent context/expression pairs.

These blocks correspond to the clusters

[{'John (.) tea', 'John (.) coffee'}, {'makes', 'likes'}]
[{'John (.) eating'}, {'likes', 'is'}]
[{'John makes (.)', 'John likes (.)'}, {'tea', 'coffee'}]
[{'John is (.)', 'John likes (.)'}, {'eating'}]
[{'John (.) tea', 'John (.) coffee', 'John (.) eating'}, {'likes'}]
[{'John likes (.)'}, {'eating', 'tea', 'coffee'}]

The last two clusters correspond to sets of context/expression pairs which are already 'covered' by the other blocks. In a sense these blocks are superfluous.

The algorithm to find these blocks is very simple: starting from a single context/expression pair, EMILE randomly adds contexts and expressions while ensuring that the resulting block is still contained in the matrix, and keeps adding contexts and expressions until the block can no longer be enlarged. This is done for each context/expression pair that is not already contained in some block. Once all context/expression pairs have been 'covered', the superfluous blocks (those completely covered by other blocks) are discarded.

12.3 Allowing for imperfect data

In the previous section, the requirement for a block was that it was entirely contained within the matrix, i.e. the clustering algorithm did not find a type unless every possible combination of contexts and expressions of that type had actually been encountered and stored in the matrix. This only works if a perfect sample has been provided. In practical use, we need to allow for imperfect samples. There are many context/expression combinations, such as for instance 'John likes evaporating', which are grammatical but nevertheless will appear infrequently, if ever.

To allow EMILE to be used with imperfect samples, two enhancements have been made to the algorithm. First, the requirement that the block is completely contained in the matrix, is weakened to a requirement that the block is *mostly* contained in the matrix. Specifically, a certain percentage of the context/expression pairs of the block as a whole should be contained in the matrix, as well as a certain percentage of the context/expression pairs in each individual row or column. We can express this as

$$\begin{array}{lll} & \#(M \cap (T_C \times T_E)) \geq \#(T_C \times T_E) & \cdot\ \texttt{total\_support\%} \\ \forall c \in T_C: & \#(M \cap (\{c\} \times T_E)) \geq \#T_E & \cdot\ \texttt{context\_support\%} \\ \forall e \in T_E: & \#(M \cap (T_C \times \{e\})) \geq \#T_C & \cdot\ \texttt{expression\_support\%} \end{array}$$

where M is the set of all encountered context/expression pairs, and the values `XXX_support%` are constants that can be set by the user.

12.3.1. Example. Suppose that the matrix of context/expression pairs EMILE has encountered has the following sub-matrix:

	John makes (.)	John likes (.)	John drinks (.)	John buys (.)
tea	x	x	x	x
coffee	x	x	x	x
lemonade	x	x	x	
soup	x	x	x	x
apples				x

If the settings `context_support%` and `expression_support%` have been set to 75%, and `total_support%` has been set to 80%, then the type represented by the cluster

$$\left[\begin{array}{c}\{\text{‘John makes (.)’, ‘John likes (.)’, ‘John drinks (.)’, ‘John buys (.)’}\},\\ \{\text{‘tea’, ‘coffee’, ‘lemonade’, ‘soup’}\}\end{array}\right]$$

will be identified, in spite of the fact that one of the context/expression pair of the block, (‘John buys (.)’, ‘lemonade’), does not appear in the matrix. However, the expression ‘apples’ will not be added to the above type, since it appears with less than `expression_support%` of the contexts.

Secondly, note that of the different expressions and contexts belonging to a grammatical type, it can be expected that the short and medium-length ones (in terms of number of words) will be encountered more often than the long ones. In other words, if we restrict the sample to short and medium-length contexts and expressions, it will be closer to a perfect sample. Implementing this notion, EMILE uses only short and medium-length contexts and expressions when searching for grammatical types.

12.4 Characteristic and secondary expressions and contexts

To search for longer expressions and contexts associated with types, EMILE uses characteristic expressions and contexts. As defined in Definition 11.2.5, a *characteristic* expression of a type T only appears with contexts that are of type T.[33] Since the types involved usually have not been fully identified yet, EMILE relaxes this requirement to also allow untyped contexts.

[33] Note that a context may have more than one type, so a context appearing with a expression characteristic for a type T may be of other types in addition to being of type T.

Occasionally, a type has no characteristic expressions (due to imperfections in the sample or the inherent ambiguity of the type): in such cases, the primary expressions of the type are used in place of the characteristic expressions. We call these the *characteristic** expressions of T, i.e. the *characteristic** expressions of T are defined as the characteristic expressions of T if there are any, and as the primary expressions of T otherwise.
The definitions of *characteristic* and *characteristic** contexts of a type T are analogous.
Any untyped context appearing with an characteristic expression of a type T is likely to belong to T as well. Contexts which appear with (a certain percentage of the) characteristic* expressions of T are called *secondary* contexts of T, as opposed to the *primary* contexts found by the clustering algorithm. Analogous for *secondary* expressions. Note that the constraint on the length of primary contexts and expressions does not apply to secondary contexts and expressions, and hence this allows for long contexts and expressions to be associated with types.

12.4.1. EXAMPLE. In the previous example, for the type represented by the cluster

[{'John likes (.)'}, {'eating', 'tea', 'coffee'}]

'John likes (.)' only appears with 'eating', 'tea' and 'coffee', so it is a characteristic (and hence characteristic*) context for this type. The expression 'eating' also appears with the context 'John is (.)', so it is not a characteristic expression. A similar condition obtains for 'tea' and 'coffee', so the type has no characteristic expressions at all. Consequentially, its primary expressions 'eating', 'tea' and 'coffee' are also its characteristic* expressions.
For the type represented by the cluster

[{'John makes (.)', 'John likes (.)'}, {'tea', 'coffee'}]

all its expressions and contexts are characteristic.

12.4.2. EXAMPLE. In Example 12.2.1, we used 1-dimensional clustering to obtain the cluster

[{'John makes (.)'}, {'tea', 'coffee'}]

Here, 'tea' and 'coffee' are not characteristic expressions, since they appear with the context 'John likes (.)', which here is not a context belonging to the type. So the type has no characteristic expressions. It is easy to see that when using 1-dimensional clustering, whenever a context is ambiguous[34], all types involved will lack characteristic expressions.

[34] 'Ambiguous' in the sense that the set of expressions it appears with is the union of several smaller sets associated with other contexts

12.4.3. Example. Assume that primary expressions are constrained to be at most 5 words long. If we add the sentence 'John makes really really really really strong coffee' to the sample of the previous example, then the expression 'really really really really strong coffee' will not be added as a primary expression to the type represented by the cluster

[{'John makes (.)', 'John likes (.)'}, {'tea', 'coffee'}]

However, since 'John makes (.)' is a characteristic expression of this type, the expression 'really really really really strong coffee' will be associated with the type as a secondary expression.

12.5 Finding rules

The EMILE program also transforms the grammatical types found into derivation rules. For reasons of simplicity, EMILE constructs a context-free grammar rather than a context-sensitive grammar. For this construction, only the sets of expressions associated with the types are needed: the sets of contexts associated with the types are not used in creating the derivation rules.
First, EMILE searches for rules that are *supported*. Obviously, if an expression e belongs to a type T (as a secondary expression), the rule

$$[T] \Rightarrow e$$

is supported. EMILE finds more complex rules, by searching for characteristic* expressions of one type that appear in the secondary expressions of another (or the same) type. For example, if the characteristic* expressions of a type T are

$$\{\text{dog}, \text{cat}, \text{gerbil}\}$$

and the type [0] contains the secondary expressions

$$\{\text{I feed my dog}, \text{I feed my cat}, \text{I feed my gerbil}\}$$

then EMILE will find the rule

$$[0] \Rightarrow \text{I feed my } [T]$$

This process of abstraction is repeated to obtain more abstract rules. Formally, a rule R is considered to be *supported* if it is of the form $[T] \Rightarrow e$ (with e being a secondary expression of T), or if it is of the form $[T] \Rightarrow s_0[T_1]s_1[T_2]\ldots s_k$, $k \geq 1$, and for some $i \in \{1,\ldots,k\}$,

$$\#\{e \in T_E^* \mid R \text{ with } [T_i] \text{ replaced by } e \text{ is supported}\} \geq \#T_E^* \cdot \texttt{rule\_support\%} \tag{12.1}$$

In certain cases, using characteristic* and secondary expressions in this manner allows EMILE to find recursive rules. For instance, a characteristic* expression of the type of sentences S might be

Mary drinks tea

If the maximum length for primary expressions is set to 4 or 5, the sentence

John observes that Mary drinks tea

will be a secondary expression of S, but not a primary or characteristic one. So if there are no other expressions involved, EMILE would derive the rules

$$\begin{array}{lcl} [S] & \Rightarrow & \text{Mary drinks tea} \\ [S] & \Rightarrow & \text{John observes that } [S] \end{array}$$

which would allow the resulting grammar to generate, for instance,

John observes that John observes that John observes that Mary drinks tea

EMILE creates a set of supported rules capable of generating all sentences in the original sample. To reduce the size of this grammar, the program discards from the final output rules which are superfluous, such as rules which are instantiations of other rules[35], and rules for types which aren't referred to in other rules.
Experiments showed that often, EMILE finds several types which where only slight variations of one another. If all these types are referred to in the rules, this results in a much larger rule-set than is necessary. The most recent incarnation of EMILE tries to prevent this by being actively conservative in the number of types used: a set of *used types* is maintained, and only rules using those types are considered for inclusion. This set initially contains only the whole-sentence type [0], and types are added only if this would result in a decrease in the size of the total rule-set.[36]

12.6 Future Developments

There is still a lot of room for improvement. The clustering algorithm could be extended to use negative samples (i.e. sentences which should not be constructible) as well as positive ones. Furthermore, a module can be added to EMILE which allows it to identify those sentences whose grammaticality is the most uncertain (from those sentences which EMILE considers grammatical but which are not in the original sample), which would allow it to query an oracle in a directed fashion.

[35] I.e. which can be obtained from other rules by replacing a type reference by a secondary expression of that type

[36] EMILE can also be set to allow a small increase: this often results in a more meaningful grammar at the expense of a slightly larger rule-set.

Another possible extension is to the algorithm constructing the derivation-rule grammars. Currently EMILE constructs a context-free grammar. It may be possible to adapt EMILE to produce a more sensible context-sensitive grammars, using the sets of contexts produced by the clustering algorithm.

Chapter 13

Results: Theory and Practice

13.1 Theoretical Results

First, we give some definitions and a conjecture adapted from P. Adriaans[34]:

13.1.1. DEFINITION. Let G be a grammar (context-free or otherwise) of a language L. G has *context separability* if each type of G has a characteristic context, and *expression separability* if each type of G has a characteristic expression.

13.1.2. DEFINITION. A class of languages[37] C is *shallow* if for each language L it is possible to find a context- and expression-separable grammar G, and a set of sentences S inducing characteristic contexts and expressions for all the types of G, such that the size of S and the length of the sentences of S are logarithmic in the descriptive length of L (relative to C).

Natural languages seem to be context- and expression-separable for the most part, i.e. if there are any types lacking characteristic contexts or expressions[38], these types are few in number, and rarely used. Furthermore, there is no known example of a syntactical construction in a natural language that cannot be expressed in a short sentence[39]. Hence the following conjecture seems tenable:

13.1.3. CONJECTURE. *Natural languages are (mostly) context- and expression-separable, and shallow.*

[37] Strictly speaking the shallowness property cannot be applied to single languages. However, when a language L has a grammar G and a set of sentences S as described in the definition, then the size of S relative to the logarithm of the descriptive length of L can be taken as a measure of the (un-)shallowness of S, so we can (imprecisely) speak of a language being 'very shallow' or 'not so shallow'.

[38] After rewriting types such as 'verbs that are also nouns' as composites of basic types.

[39] At the 1997 CSLI workshop, P. Adriaans offered a thousand dollars for a syntactical construction in any known language, that cannot be expressed in 16 words or less. Nobody has claimed this money yet. The offer is still open as of this writing.

Now, if a grammar is context- and expression-separable, then EMILE will be able to find its types given the proper settings and a sufficiently complete sample, as the following lemma shows

13.1.4. LEMMA. *Let T be a type with a characteristic context c^{ch} and a characteristic expression e^{ch}. Suppose that the maximum lengths for primary contexts and expressions are set to at least* $\text{len}(c^{ch})$ *and* $\text{len}(e^{ch})$ *and suppose that the* **total_support%**, **expression_support%** *and* **context_support%** *settings are all set to* 100%. *Let $T_C^{\leq max}$ and $T_E^{\leq max}$ be the sets of contexts and expressions of T that are small enough to be used as primary contexts and expressions. If EMILE is given a sample containing all combinations of contexts from $T_C^{\leq max}$ and expressions from $T_E^{\leq max}$, then EMILE will find type T.*

Proof

For any type U, if c^{ch} belongs to U, then all expressions of U appear with c^{ch}, and hence are also expressions of T. Similarly, for any type U, if e^{ch} belongs to U, then all contexts of U are also contexts of T. It follows that for any type U, if U covers the context/expression pair (c^{ch}, e^{ch}), then $U_C \times U_E \subseteq T_C \times T_E$. Conversely, $T_C \times T_E$ is a type covering (c^{ch}, e^{ch}). We conclude that $T_C \times T_E$ is the unique maximal type covering (c^{ch}, e^{ch}), and hence will appear in the grammar output by EMILE.

□

Given this result, if the conjecture that natural languages are (mostly) context- and expression-separable holds, then EMILE should have the potential to learn natural languages. If natural languages are also shallow, then the required sample can be relatively small. The question whether EMILE works in practice, and what constitutes a 'small' sample, will be considered in the next two sections.

13.2 Results for a Generated Sample

The EMILE program was given 100,000 different sentences generated by the following context-free grammar:

$$
\begin{aligned}
[S] &\Rightarrow [NP]\ [V_i]\ [ADV] \ \mid\ [NP_a]\ [VP_a] \ \mid\ [NP_a]\ [V_s]\ \text{that}\ [S] \\
[NP] &\Rightarrow [NP_a] \ \mid\ [NP_p] \\
[VP_a] &\Rightarrow [V_t]\ [NP] \ \mid\ [V_t]\ [NP]\ [P]\ [NP_p] \\
[NP_a] &\Rightarrow \text{John} \mid \text{Mary} \mid \text{the man} \mid \text{the child} \\
[NP_p] &\Rightarrow \text{the car} \mid \text{the city} \mid \text{the house} \mid \text{the shop} \\
[P] &\Rightarrow \text{with} \mid \text{near} \mid \text{in} \mid \text{from} \\
[V_i] &\Rightarrow \text{appears} \mid \text{is} \mid \text{seems} \mid \text{looks} \\
[V_s] &\Rightarrow \text{thinks} \mid \text{hopes} \mid \text{tells} \mid \text{says} \\
[V_t] &\Rightarrow \text{knows} \mid \text{likes} \mid \text{misses} \mid \text{sees} \\
[ADV] &\Rightarrow \text{large} \mid \text{small} \mid \text{ugly} \mid \text{beautiful}
\end{aligned}
$$

(where the '|' symbol is used to separate alternatives). The EMILE program used the following settings (see appendix A for the exact meaning of all the settings):

`maximum_sentence_length` = 14
`maximum_primary_expr_length` = 4
`maximum_primary_context_length` = 5
`minimum_contexts_per_type` = 3
`minimum_expressions_per_type` = 4
`type_usefulness_required` = 1
`ruleset_increase_disallowed` = 1
`total_support_percentage` = 44
`expression_support_percentage` = 25
`context_support_percentage` = 25
`secondary_expression_support%` = 25
`secondary_context_support%` = 25
`rule_support_percentage` = 25
`sesp_for_no_characteristics` = 34
`scsp_for_no_characteristics` = 26
`rsp_for_no_characteristics` = 26

After processing 100,000 sentences, EMILE generated the following grammar:

[0] $\Rightarrow$ [17] [6]
[0] $\Rightarrow$ [17] [22] [17] [6]
[0] $\Rightarrow$ [17] [22] [17] [22] [17] [22] [17] [6]
[6] $\Rightarrow$ misses [17]
[6] $\Rightarrow$ likes [17]
[6] $\Rightarrow$ knows [17]
[6] $\Rightarrow$ sees [17]
[6] $\Rightarrow$ [22] [17] [6]
[6] $\Rightarrow$ appears [34]
[6] $\Rightarrow$ looks [34]
[6] $\Rightarrow$ is [34]
[6] $\Rightarrow$ seems [34]
[6] $\Rightarrow$ [6] near [17]
[6] $\Rightarrow$ [6] from [17]
[6] $\Rightarrow$ [6] in [17]
[6] $\Rightarrow$ [6] with [17]
[17] $\Rightarrow$ the child
[17] $\Rightarrow$ Mary
[17] $\Rightarrow$ the city
[17] $\Rightarrow$ the man
[17] $\Rightarrow$ John
[17] $\Rightarrow$ the car
[17] $\Rightarrow$ the house
[17] $\Rightarrow$ the shop
[22] $\Rightarrow$ tells that
[22] $\Rightarrow$ thinks that
[22] $\Rightarrow$ hopes that
[22] $\Rightarrow$ says that
[22] $\Rightarrow$ [22] [17] [22]
[34] $\Rightarrow$ small
[34] $\Rightarrow$ beautiful
[34] $\Rightarrow$ large
[34] $\Rightarrow$ ugly

As can be seen, EMILE identifies most of the structures of the original grammar, and even manages to capture its recursive structure. Furthermore, the resulting grammar is not much larger than the original grammar. This gives hope that EMILE, or a program based on EMILE, could be used as a tool to find meaningful patterns in languages.

However, it should be noted that the grammar found by EMILE is weaker than the original grammar. For one, it does not differentiate between types of nouns, making possible sentences such as 'the car says that…'. Furthermore, in the original grammar, phrases such as 'with the car' were optional additions to certain sentences, and at most one such phrase could be appended. In the grammar found by EMILE, a second recursive structure allows any sentence to be followed by an arbitrary number of these phrases. Very likely, a higher value for 'minimum_expression_length' and higher values for the support settings will result in a grammar weakly equivalent to the original one, but for these settings, a larger sample will probably be required to achieve meaningful results.

The grammar found by EMILE contains a few superfluous rules, such as '[0] ⇒ [17] [22] [17] [22] [17] [22] [17] [6]'. This is caused by the fact that when checking which rules have been made superfluous, EMILE only checks one-step instantiations, i.e. those expressions which can be obtained from a rule by directly replacing type references with secondary expressions. To check more thoroughly, it is necessary to consider those expressions which can be obtained using successive rule substitutions, which is very expensive (in terms of computation time).[40]
To study the program's behavior as a function of the sample-size, the CFG was used to generate 1000 sentences at a time. This produced the following statistics:

number of sentences	number of types found	size of rule-set	number of types used	time used (minutes)[41]
1000	421	269	12	0
2000	647	288	19	3
3000	657	79	11	11
4000	643	82	14	30
5000	638	60	10	78
6000	566	72	12	78
7000	468	60	11	101
8000	283	98	12	98
9000	217	57	12	143
10000	195	87	12	112
11000	196	40	8	61
12000	211	68	8	45
13000	202	100	11	87
14000	214	84	11	50
15000	215	46	9	109
16000	202	41	8	157
17000	214	41	8	199
18000	211	63	10	107
19000	201	56	13	180
20000	199	66	12	169
30000	205	42	9	387
40000	180	40	6	521
50000	154	63	7	462
60000	168	38	7	773
70000	125	38	7	939
80000	112	34	6	838
90000	89	51	6	2806
100000	61	33	5	3598

[40]In fact this requires the program to solve the problem of whether two grammars are weakly equivalent, which is undecidable in general. However, in this case we can limit ourselves to expressions encountered by the program, which makes it decidable.

As can be seen, these statistics do not yield a smooth curve, although the deviations are not extravagantly large. This is probably caused by the randomizer, which is used at a number of points in the grammar deduction process, to implement nondeterministic selection. As can be seen, the number of rules found briefly increases, then drops, increases again, and then slowly drops to slightly more than 30. Something similar happens to the number of types found and the number of types used. Presumably, the increase around the 10,000 sentence-point is caused by a shift in probability distributions around then: the sentence generator is prohibited from repeating sentences, so around that point the proportion of long sentences vs. short sentences will start to change. Another observation is that, taking into account the large variations in used time caused by the changing CPU loads, the time used by EMILE does not seem to be exponential in the number of sentences, or even high-order polynomial.
We can conclude that in this experiment, the output of EMILE converges to a concise grammar, and that a sample of 30,000 sentences suffices to get good results.

Results for large real-world datasets

EMILE was given the text of the Bible (King James edition) to see if it could derive a grammar for the English language. Using the following settings, EMILE processed the 21070 different sentences of the Bible that were of length ≤ 14:

`maximum_sentence_length` = 14	`expression_support_percentage` = 40
`maximum_primary_expr_length` = 4	`context_support_percentage` = 40
`maximum_primary_context_length` = 5	`secondary_expression_support%` = 20
`minimum_contexts_per_type` = 2	`secondary_context_support%` = 20
`minimum_expressions_per_type` = 3	`rule_support_percentage` = 20
`type_usefulness_required` = 1	`sesp_for_no_characteristics` = 51
`ruleset_increase_disallowed` = 1	`scsp_for_no_characteristics` = 34
`total_support_percentage` = 64	`rsp_for_no_characteristics` = 34

The result was a grammar containing 20858 rules, only 212 less than the trivial grammar containing only the literal sentences. In fact most (20441) of the rules in the generated grammar are rules for literal sentences, such as

[0] ⇒ And the flood was forty days upon the earth ;

This indicates that for most sentences, EMILE could not discern a pattern, or at least not a pattern that could be used to reduce the size of the grammar. Amongst the patterns which EMILE did manage to discover. are some which might be significant:

[41] This time should be taken as a *very* rough indication, as it is strongly influenced by the load caused by other programs running on the same computer

[0] ⇒ Thou shall not [582]	[582] ⇒ eat it ;
[0] ⇒ Neither shalt thou [582]	[582] ⇒ kill .
	[582] ⇒ commit adultery .
	[582] ⇒ steal .
	[582] ⇒ bear false witness against thy neighbour .
	[582] ⇒ abhor an Edomite ;

and some which are probably mere accidents:

[0] ⇒ and [72]	[72] ⇒ Er and Onan died in the land of Canaan .
[0] ⇒ but [72]	[72] ⇒ let me not fall into the hand of man .
	[72] ⇒ they could not .
	[72] ⇒ he saw :
	[72] ⇒ now murderers .
	[72] ⇒ they shall know that I am the Lord GOD .

EMILE was also run with successively larger subsets of this sample, with the following results:

number of sentences	number of types found	size of rule-set	number of types used	time used (seconds)
2107	63	2097	7	4
4214	151	4172	11	9
6321	290	6255	13	16
8428	396	8337	18	23
10535	420	10410	18	29
12642	487	12506	20	37
14749	579	14582	22	46
16856	653	16670	28	62
18963	800	18753	29	79
21070	840	20858	33	102

There is no sign of the convergence that characterized the previous experiment: presumably, the Bible simply isn't big enough as a sample of the English language. A problem with using larger samples is that the EMILE program uses a lot of memory. To analyze the Bible, EMILE needs between 100 and 250 megabytes of memory (depending on the settings used): larger samples have proportionally larger memory requirements. It may be possible to design a version of EMILE which allows for the data to be distributed over several machines. However, even if a distributed version turns out to be impractical, this is a temporary problem, given the exponential growth of available computer memory of the last few years.

Appendix A

The inner workings of EMILE

At this moment of writing, the EMILE program consists of about 5500 lines of C++-code. However, most of that is code for data type representation, user interface, utility functions, various optimizations, etcetera. The algorithms themselves are fairly simple. Each of the sections of this chapter focuses on a different algorithm used in EMILE. For each algorithm, a synopsis is given, as well as explicit pseudo-code, and a summary of the constants controlling the algorithm that can be set by the user.

A.1 Gathering context/expression pairs.

Synopsis EMILE maintains a matrix M of the context/expression pairs it has encountered. This routine updates this matrix, given text from some input I.

Algorithm

```
sub learn_sentences(I)
    while (there is input to be read) do
        read the input I up to the next end-of-sentence marker;
        set s := the sentence read, converted to a sequence of words;
        if (length(s) ≤ max_sentence_length) then
            for each triple (c_l, e, c_r) with c_l^e^c_r = s do
                insert (c_l^"(.)"^c_r, e) into M;
            end for
        end if
    end while
end sub
```

User definable settings

- `max_sentence_length`: sentences longer than this are ignored.

- `end_of_sentence_markers`: a set of characters that EMILE interprets as marking the end of a sentence.

- `allow_multi_line_sentences`: a boolean variable. If this variable is set to `false`, the end of an input line is considered to indicate the end of the current sentence. If this variable is set to `true`, a sentence can span multiple lines. In either case, the `end_of_sentence_markers` also indicate the end of the sentence, as does an empty line (i.e. two end-of-lines separated by nothing but whitespace).

- `ignore_abbreviation_periods`: if the period symbol '.' is used as an end-of-sentence-marker, and this boolean variable is set to `true`, periods following a single letter are not considered to indicate the end of a sentence.

- `regular_expression_as_marker`: if this boolean variable is set to `true`, the content of the setting `end_of_sentence_markers` is interpreted as a regular expression, using the syntax of Extended Regular Expressions as defined in the `regex`(5) Unix man page. The standard C++ escape sequences are recognized, i.e. '`\n`' for newline, '`\t`' for tab etcetera. The settings of `allow_multi_line_sentences` and `ignore_abbreviation_periods` are ignored.

  Although this is a more expressive mechanism for indicating end-of-sentence markers, it is not more readable: for instance, the default settings of the normal mode correspond to the regular expression

  ```
  [.!;?]|\n[␣\r\n\t]*\n|\r[␣\r\n\t]*\r|^\\.|[^a-zA-Z]\\.|[^␣\r\n\t].\\.
  ```

Notes When using regular expressions to mark the end of sentences, newlines have to be explicitly included in the regular expression in order to be taken into account. Note that on non-Unix systems, the '`\n`' symbol refers to the end-of-line marker customary on that system. The carriage return symbol '`\r`' can be used if there is a need to explicitly take into account line endings of non-Unix files when working on a Unix system.
An end-of-sentence marker is considered to be part of the sentence it is ending.
A sentence is converted into a sequence of words before it is searched for context/expression pairs. A word is a nonempty sequence of alphanumeric characters, or a single non-alphanumeric, non-whitespace symbol. Whitespace characters (spaces, tabs, newlines and carriage returns) function as word separators where necessary and are otherwise ignored.
For reasons of efficiency, contexts and expressions are not directly used as elements of the matrix. Instead, the actual contexts and expressions are stored in a table, and references to the entries in the table are used in the matrix.

There is a compilation option to put Emile in 'Morpheme Analysis' mode. In this mode, each word (using whitespace as a delimiter) is treated as a separate element of S, and is split into single characters for analysis. The settings related to end-of-sentence marking are ignored.

A.2 Extracting the grammatical types from the matrix

Synopsis A type $T = (T_C, T_E)$ is considered to have sufficient *support* if it satisfies the following three conditions:

1. $\#(M \cap (T_C \times T_E)) \geq \#(T_C \times T_E) \cdot$ `total_support%`
2. $\forall c \in T_C : \#(M \cap (\{c\} \times T_E)) \geq \#T_E \cdot$ `context_support%`
3. $\forall e \in T_E : \#(M \cap (T_C \times \{e\})) \geq \#T_C \cdot$ `expression_support%`

The program maintains a set G of grammatical types with sufficient support, and with contexts and expressions of length at most `max_primary_context_length`, or `max_primary_expression_length`, respectively. All these types are of maximal size (under the constraint of having sufficient support), and all are at or above a certain minimum size (as indicated by the settings `min_contexts_per_type` and `min_expressions_per_type`).
An element $(c, e) \in M$ is considered *covered* by a type T if $(c, e) \in T_C \times T_E$. This routine updates and enlarges G so that every element $(c, e) \in M$ (that can be covered by a type of minimum size) is covered by a type in G.

Algorithm

```
sub expand_grammar(G)
    for each T ∈ G do
        call enlarge_grammatical_type(T);
        if ((#T_C < min_contexts_per_type)
                or (#T_E < min_expressions_per_type)) then
            remove T from G;
        end if
    end for
    for each (c, e) ∈ M do
        if (¬∃T ∈ G : (c, e) ∈ T_C × T_E) then
            set T := ({c}, {e});
            call enlarge_grammatical_type(T);
            if ((#T_C ≥ min_contexts_per_type
                    or (#T_E ≥ min_expressions_per_type)) then
```

```
            insert T into G;
        end if
    end if
  end for
end sub

sub enlarge_grammatical_type(T)
  repeat
    set X := {c' a context | (T_C ∪ {c'}, T_E) has sufficient support,
                      length(c) ≤ max_primary_context_length},
           ∪{e' an expression | (T_C, T_E ∪ {e'}) has sufficient support,
                      length(e) ≤ max_primary_expression_length};
    if (X ≠ ∅) then
      nondeterministically select x from X;
      if (x is a context) then
        insert x into T_C;
      else
        insert x into T_E;
        end if
      end if
    until (X = ∅);
  end sub
```

User definable settings

- `max_primary_context_length`, `max_primary_expression_length`: in order to ensure that the set of types found will converge if sufficiently many sentences are read, the search space can be limited to primary contexts and expressions of bounded size.

- `total_support%`, `context_support%`, `expression_support%`: these variables control the support required from the matrix for each type. The lower these values, the larger the size of the types found. Note that lowering one value while keeping the other values high will not have much effect.

- `min_contexts_per_type`, `min_expressions_per_type`: Types with fewer contexts or expressions than indicated by these settings are discarded.

Notes The selection of x from X is not nondeterministic, but based on the amount of support that would be added to the grammatical type.
For purposes of constructing a grammar, types of extremely small size usually are not very interesting. Types with less than `min_contexts_per_type` contexts or `min_expressions_per_type` expressions are discarded. This may cause some elements $(c, e) \in M$ to be uncoverable.

The type [0] is always set to the type of whole sentences, with '((),())' as a secondary context and all encountered sentences as secondary expressions.

Complexity The `enlarge_grammatical_type` subroutine maintains a significant amount of auxiliary data to avoid having to repeat calculations to collect the set X. Initialization of the auxiliary data has an execution time of order $O(\#\{(c,e) \in M \mid c \in T_C \vee e \in T_E\})$, while each iteration of the `repeat..until` loop has an average-case execution time of order $O(\#(T_C \cup T_E))$ and a worst-case execution time of order $O(\#\{(c,e) \in M \mid c \in T_C \vee e \in T_E\})$.

A.3 Eliminating superfluous types

Synopsis It is possible that the contribution of some type to the coverage of G is made (nearly) superfluous by types found later, i.e. most or all of the context/expression pairs that are covered by that type, are also covered by other types of G. This routine eliminates such types.

Algorithm

```
sub eliminate_superfluous_types(G)
   set cover(*) = 0;
   for each T ∈ G do
      for each (c,e) ∈ M ∩ (T_C × T_E) do
         increment cover(c,e);
      end for
   end for
   for each T ∈ G do
      if (#{(c,e) ∈ M^(T_C × T_E) | cover(c,e) = 1}
            < type_usefulness_required) then
         remove T from G;
         for each (c,e) ∈ M ∩ (T_C × T_E) do
            decrement cover(c,e);
         end for
      end if
   end for
end sub
```

User definable settings

- `type_usefulness_required`: this variable determines how useful a type has to be (in terms of contributions to the coverage of G) in order to not be discarded. Setting this to 0 will prevent types from being discarded, setting this to a high value will eliminate all but a few types of large size. The default value is 1, which eliminates only types which do not contribute anything.

Notes The types are checked in order of increasing absolute total support. This means that if the matrix can be covered by either a lot of small types or a single big one, probability favors the latter result.

A.4 Identifying characteristic and secondary contexts and expressions

Synopsis This routine finds, for each type T,

- the characteristic expressions of T, defined as those expressions which only appear with contexts of no type or of type T.
- the characteristic contexts of T, defined analogously.
- the characteristic* expressions and contexts of T, defined as the characteristic expressions and contexts of T if there are any, otherwise defaulting to the primary expressions and contexts of T.
- the secondary expressions of T, defined as its primary expressions and those expressions e satisfying

$$\#\{(c,e) \mid c \in T_C^*, (c,e) \in M\} \geq \#T_C^* \cdot \texttt{sec\_context\_support\%} \quad \text{(A.1)}$$

- the secondary contexts of T, defined as its primary contexts and those contexts c satisfying

$$\#\{(c,e) \mid e \in T_E^*, (c,e) \in M\} \geq \#T_E^* \cdot \texttt{sec\_context\_support\%} \quad \text{(A.2)}$$

Algorithm

```
sub identify_characteristic_and_secondary_aspects(G)
    for each T ∈ G do
        set T_C^ch := ∅
        for each context c ∈ T_C do
            if ∀e : [(c,e) ∈ M → (e ∈ T_E ∨ ¬∃U ∈ G : e ∈ U_E)] then
                insert c into T_C^ch;
            end if
        end for
        set T_E^ch := ∅
```

```
        for each expression e ∈ T_E do
            if ∀c : [(c,e) ∈ M → (u ∈ T_C ∨ ¬∃U ∈ G : c ∈ U_C)] then
                insert e into T_E^ch;
            end if
        end for
        if (T_C^ch ≠ ∅) then
            set T_C^* := T_C^ch;
        else
            set T_C^* := T_C;
        end if
        if (T_E^ch ≠ ∅) then
            set T_E^* := T_E^ch;
        else
            set T_E^* := T_E;
        end if
        set T_C^se := T_C;
        for each context c do
            if (#{e ∈ T_E^* | (c,e) ∈ M} ≥ #T_E^* · sec_context_support%) then
                insert c into T_C^se;
            end if
        end for
        set T_E^se := T_E;
        for each expression e do
            if (#{c ∈ T_C^* | (c,e) ∈ M} ≥ #T_C^* · sec_expression_support%) then
                insert e into T_E^se;
            end if
        end for
    end for
end sub
```

User definable settings

- `sec_context_support%`, `sec_expression_support%`: lower values for these settings will increase the number of secondary contexts or expressions found. Note that the effects of these settings are independent (unlike with the settings for primary contexts and expressions).

- `scsp_for_no_characteristics`, `sesp_for_no_characteristics%`: if a type has no characteristic expressions, the characteristic* expressions default to the primary expressions. If required support percentages are low, this can result in a single expression being taken as indicative of several types. To prevent this, EMILE provides the setting `scsp_for_no_characteristics`, to be used instead of `sec_context_support%` when a type has no characteristic expressions. Similar for `sesp_for_no_characteristics%`.

A.5 Deriving grammatical rules

Synopsis Emile uses the grammatical types it finds to infer grammatical derivation rules. A rule $r : [T] \Rightarrow s_0[T_1]s_1[T_2] \ldots s_k$ is considered to be *supported*, if $k = 0$ and $s_0 \in T_E^{se}$, or if $k \geq 1$ and for some $i \leq k$,

$$\#\{e \in T_E^* \mid r \text{ with } [T_i] \text{ replaced by } e \text{ has support}\} \geq \#T_E^* \cdot \texttt{rule\_support\%} \quad \text{(A.3)}$$

An *instantiation* of a rule $r : [T] \Rightarrow s_0[T_1]s_1 \ldots s_k$, $k \geq 0$, is an expression $e \in T_E^{se}$ which can be obtained by replacing the type references $[T_1], \ldots, [T_k]$ in r by expressions $e'_1 \in (T_1)_e^{se}, \ldots, e'_k \in (T_k)_e^{se}$. A rule r for some type $[T]$ is considered to be *covered* by other rules for $[T]$ if all of its instantiations are also instantiations of one or more of the other rules.
EMILE tries to find a set of supported rules which contains no rules covered by other rules, is capable of generating the original sample, and cannot easily be reduced in size. To do this, the program maintains a set of used types V_{used}, which initially contains only the whole-sentence type [0]. Whenever a type is added to V_{used}, the program gathers all supported rules for all types of V_{used} which only use types from V_{used} (note that for this purpose, rules for a type are considered to be using that type). Then rules which are covered by other rules are eliminated, until a set of rules R is obtained in which every rule has at least one instantiation which is not shared with any other rule. The program adds types to V_{used} as long as this will not result in a large increase in the size of the resulting rule-set.

Algorithm

```
sub derive_rules(G, R)
    for each T ∈ G do
        set R_T^sup := {[T] ⇒ e | e ∈ T_E^se};
        set R_T^sup := R_T^sup ∪ {r | r is supported by R_T^sup. r uses [T] and only [T]};
        set R_T := R_T^sup;
        for each r ∈ R_T do
            if (∃r' ∈ R_T : r' ≠ r ∧ r' covers r) then
                remove r from R_T;
            end if
        end for
    end for
    set V_used := ∅;
    set R := ∅;
    set R^sup := ∅;
    set T_add := [0];
    repeat
        insert T_add in V_used;
        set R^sup := R^sup ∪ R_{T_add}^sup;
```

```
        set R := R_T;
        for each T ∈ G do
            if (T ∉ V_used) then
                set R_T^add := { r | r is supported by R^sup ∪ R_T^sup,
                                     r uses both [T] and [T_add],
                                     r uses no types outside V_used ∪ {T} };
                set R_T^sup := R_T^sup ∪ R_T^add;
                set R_T := R ∪ (R_T ∩ R_T^sup) ∪ R_T^add;
                for each r ∈ R_T do
                    if (∃r' ∈ R_T : r' ≠ r ∧ r' covers r) then
                        remove r from R_T;
                    end if
                end for
            end if
        end for
        if (∃T : #R_T < #R + ruleset_increase_disallowed) then
            select T_add from G such that T_add ∉ V_used and #R_T is minimal
        end if
    until (¬∃T : #R_T < #R + ruleset_increase_disallowed);
    for each r ∈ R do
        if (R−{r} covers r) then
            remove r from R;
        end if
    end for
end sub
```

User definable settings

- `rule_support%`: this variable controls the support required for a rule before it is considered for inclusion in the grammar. The lower this value, the more rules will be found.

- `rsp_for_no_characteristics`: if a type has no characteristic expressions, the characteristic* expressions default to the primary expressions. If required support percentages are low, this can result in a single expression being taken as indicative of several types. To prevent this, EMILE provides the setting `rsp_for_no_characteristics`, to be used instead of `rule_support%` when a type has no characteristic expressions.

- `ruleset_increase_disallowed%`: this variable determines how useful a type has to be (in terms of the resulting reduction in the number of rules) in order to be used. Setting this to 0 requires types to reduce the number of rules, setting this to 1 will include types as long as they don't actually increase the size of the rule-set. Higher values will allow more types to be included, at the expense of increasing the size of the rule-set.

Notes The sets R_T are actually stored as changes to be applied to R.
For reasons of efficiency this routine only checks whether rules are covered by single other rules, everywhere except at the very end. This decreases computation time, and also allows for some other optimizations.
Covered rules are eliminated in order of increasing complexity. This means that the final result will not contain any rule of which an abstraction exists that uses only types in V_{used} and is supported.
Rules of the form $[T] \Rightarrow [U]$ are only allowed if $\#T_E^{se} > \#U_E^{se}$, to prevent loops.
For types $T \notin V_{used}$, the rules in R_T for $[T]$ are retained for reference purposes.

A.6 Short-circuiting superfluous types

Synopsis If a type $T \in G$ has only a single rule in R, or if there is only one reference to $[T]$, then we can decrease the size of the rule-set, and remove T from the set of used types, by substituting the rules of $[T]$ for all references to $[T]$.

Algorithm

```
sub short-circuit_types(G, R)
    for each [T] ∈ V_used do
        if (R contains only one rule for [T]) then
            let r ∈ R be the unique rule for [T];
            for each rule r′ ∈ R referring to [T] do
                remove r′ from R;
                substitute r for [T] in r′;
                insert r′ into R;
            end for
            remove [T] from V_used;
        end if
        if (R contains only one reference to [T]) then
            let r′ ∈ R be the unique rule referring to [T];
            remove r′ from R;
            for each rule r ∈ R for [T] do
                substitute r for [T] in r′;
                insert r′ into R;
            end for
            remove [T] from V_used;
        end if
    end for
```

```
end sub
```

A.7 Parsing a sentence

Synopsis Given a set of rules, we will sometimes want to see whether a given sentence is parsable with those rules. Furthermore, we will want to see what types unknown words must be assigned in order to make the sentence parsable. EMILE can check for parsability while allowing up to a user-settable number of words to be assigned arbitrary types, using a recursive algorithm.

Algorithm

```
function parse_phrase(R,s,[T])
   if (([T] ⇒ s) ∈ R) then
      return 0;
   else if (length(s) = 1) then
      return 1;
   else
      set n := ∞;
      for each rule ([T] ⇒ s_0[T_0]s_1 ... s_k) ∈ R do
         for each sequence (s'_1, ..., s'_k) with s_0^s'_1^s_1^ ... ^s'_k^s_k = s do
            set n = min(n, Σ_{i=1}^{k} parse_phrase(R, s'_i, T_i));
         end for
      end for
      return n;
   end if
end function

function parse_sentence(R,s)
   if (parse_phrase(R, s, [0]) ≤ parser_tolerance) then
      return true;
   else
      return false;
   end if
end function
```

User definable settings

- parser_tolerance: This setting is the maximum number of words to which Emile will assign or reassign a type in order to make parsing of a sentence possible.

Notes Types are assigned 'on the fly' only to single words, not to larger expressions.
Emile keeps track of the parser-tolerances used in the search: if a particular search cannot possibly result in a better parsing than the best one found up to now, the searching doesn't take place.

Bibliography

Blackwell Games

[1] **V. Allis**, *A Knowledge-based Approach of Connect-Four,* Report IR-163, Faculty of Mathematics and Computer Science at the Vrije Universiteit Amsterdam, The Netherlands, 1988.

[2] **D. Blackwell**, *Infinite G_δ games with imperfect information*, Zastosowania Matematyki Applicationes Mathematicae, Hugo Steinhaus Jubilee Volume X, p. 99-101, 1969.

[3] **D. Blackwell**, *Operator Solution of Infinite G_δ Games of Imperfect Information*, in: T.W. Anderson, K.B. Athreya and D.L. Iglehart (editors), *Probability, Statistics and Mathematics. Papers in Honor of S. Karlin*, p. 83-87, Academic Press, New York, 1989.

[4] **D. Blackwell** and **M.A. Girshick**, *Theory of Games and Statistical Decisions*, Dover Publications Inc., New York, 1954.

[5] **M. Davis**, *Infinite games of perfect information*, in: M. Dresher, L.S. Shapley and A.W. Tucker (editors), *Advances in game theory*, Annals of Mathematics Studies 52, p.85-101, Princeton University Press, Princeton, N.J., 1964.

[6] **D. Gale** and **F.M. Stewart**, *Infinite Games with Perfect Information*, in: *Contributions to the theory of Games*, Annals of Mathematics Studies 28, p.245-266, Princeton University Press, Princeton, N.J., 1953.

[7] **A.J.C. Hurkens**, *Borel determinacy without the axiom of choice*, Universiteitsdrukkerij Nijmegen, Nijmegen, The Netherlands, 1993.

[8] **A. Kanamori**, *The Higher Infinite. Large cardinals in set theory from their beginnings*, Perspectives in Mathematical Logic Series, Springer-Verlag, Berlin, 1994.

[9] **A.S. Kechris**, *Classical Descriptive Set Theory*, Graduate Texts in Mathematics 156, Springer-Verlag, New York, 1994.

[10] **E.M. Kleinberg**, *Infinitary Combinatorics and the Axiom of Determinateness*, Lecture Notes in Mathematics 612, Springer-Verlag, Berlin, 1977.

[11] **B. Leuwe**, *The Projective Theory of the Reals under the Axiom of Blackwell Determinacy*, in preparation.

[12] **A. Maitra**, **R. Purves** and **W. Sudderth**, *Approximation theorems for gambling problems and stochastic games*, in: B. Dutta, D. Mookherjee, T. Parthasarathy, T.E.S. Raghavan, D. Ray and S. Tijs (Editors), *Game Theory and Economic Applications (New Delhi, 1990)*, p. 114-132, Lecture Notes in Economics and Mathematical Systems 389, Springer-Verlag, Berlin, 1991.

[13] **D.A. Martin**, *Borel Determinacy*, Annals of Mathematics 102, p. 363-371, 1975.

[14] **D.A. Martin**, *A purely inductive proof of Borel determinacy*, in: A. Nerode and R.A. Shore (Editors), *Recursion Theory*, Proceedings of Symposia in Pure Mathematics 42, p. 303-308, American Mathematical Society, Providence, R.I., 1985.

[15] **D.A. Martin** and **J.R. Steel**, *A proof of projective determinacy*, Journal of the American Mathematical Society 2, p. 71-125, 1989.

[16] **D.A. Martin**, *The determinacy of Blackwell games*, Journal of Symbolic Logic 63, no. 4, p. 1565-1581, 1998.

[17] **J. Mycielski** and **H. Steinhaus**, *A mathematical axiom contradicting the axiom of choice*, Bulletin de l'Académie Polonaise des Sciences, Série des sciences mathématiques, astronomiques et physiques 10, p. 1-3, 1962.

[18] **J. von Neumann**, *Zur Theorie der Gesellschaftsspiele*, Mathematischen Annalen, Volume 100, p. 295-320, 1928.

[19] **M. Vervoort** *An Elementary Construction of an Ultrafilter on $\aleph_1$, using the Axiom of Determinateness*, ILLC Research Report ML-1994-13, Institute for Logic, Language and Computation, Amsterdam, 1994.

[20] **M. Vervoort**, *Blackwell Games*, in: T.S. Ferguson, L.S. Shapley and J.B. MacQueen (editors), *Statistics, probability and Game Theory. Papers in Honor of David Blackwell*, p. 369-390, IMS Lecture Notes - Monograph Series 30, Institute of Mathematical Statistics, Hayward, California, 1996.

Random Walks

[21] **L.O. Arkeryd**, **N.J. Cutland** and **C.W. Henson (editors)**, *Nonstandard Analysis: Theory and Applications*, NATO Advanced Science Institutes Series C: Mathematical and Physical Sciences, Volume 493, Kluwer Acad. Publ, Dordrecht, the Netherlands, 1997.

[22] **D. Coppersmith** and **P. Diaconis**, *Random Walk with Reinforcement*, unpublished, 1987.

[23] **B. Davis**, *Reinforced Random Walk*, Probability Theory and Related Fields 84, page 202-229, 1990.

[24] **J.L. Doob**, *Stochastic Processes*, John Wiley and Sons, New York, 1953.

[25] **P. Doyle** and **L. Snell**, *Random Walks and Electric Networks*, Carus Mathematical Monographs, Mathematical Association of America, Washington, D.C., 1984.

[26] **S. Kakutani**, *Markov Processes and the Dirichlet Problem*, Proceedings of the Japan Academy Series A: Mathematical Sciences, Volume 21, page 227-233, 1945.

[27] **E. Nelson**, *Internal Set Theory: A New Approach to Nonstandard Analysis*, Bulletin of the American Mathematical Society 83, page 1165-1198, 1977.

[28] **E. Nelson**, *The Syntax of Nonstandard Analysis*, Annals of Pure and Applied Logic 38, page 123-134, 1988.

[29] **R. Pemantle**, *Phase Transition in Reinforced Random Walk and RWRE on Trees*, Annals of Probability 16, page 1229-1241, 1988.

[30] **P. Pólya** and **F. Eggenberger**, *Über die Statistik verketter Vorgänge*, Zeitschrift für Angewandte Mathematik und Mechanik 3, page 279-289, 1923.

[31] **T. Sellke**, *Recurrence of Reinforced Random Walk on a Ladder*, Technical Report #93-10, Department of Statistics, Purdue University, West Lafayette, Indiana, 1993.

[32] **F. Spitzer**, *Principles of Random Walks*, The University Series in Higher Mathematics, D. Van Nostrand Co., Inc., Princeton, N.J.-Toronto-London, 1964.

EMILE

[33] **P. Adriaans**, *Language Learning from a Categorial Perspective*, PhD thesis, University of Amsterdam, 1992.

[34] **P. Adriaans**, *Learning Shallow Context-Free Languages under Simple Distributions*, ILLC Research Report PP-1999-13, Institute for Logic, Language and Computation, Amsterdam, 1999.

[35] **P. Adriaans, M. Trautwein, M. Vervoort**, *The Emile Approach to Grammar Induction put into Practice*, in preparation.

[36] **E. Dörnenburg**, *Extension of the EMILE algorithm for inductive learning of context-free grammars for natural languages*, Master's Thesis, University of Dortmund, 1997.

Samenvatting

Deze dissertatie bestaat uit drie delen.
Het eerste deel, getiteld 'Blackwell Games', gaat over het probleem van gedetermineerdheid van Blackwell spellen, een klasse van oneindige spellen met imperfekte informatie, waarbij twee spelers tegelijkertijd zetten doen, de zetten gekozen worden uit een eindige verzameling, het spel oneindig veel rondes duurt, en de uitkomst van het spel wordt gegeven door een Borel-meetbare functie f op de verzameling mogelijke series zetten. Elementaire bewijzen van gedetermineerdheid worden gegeven voor Blackwell spellen wiens uitkomst wordt gegeven door de karakteristieke functie van een Borel verzameling met complexiteit tot $G_{\delta\sigma}$. Voor algemene Borel functies geven we een reductie, afkomstig van D.A. Martin[16], tot het al bekende resultaat van gedetermineerdheid van spellen met perfekte informatie. We beschouwen ook Blackwell spellen wiens uitkomst wordt gegeven door een functie die niet Borel meetbaar is, en formuleren een versie van het Axioma van Gedetermineerdheid voor deze spellen. Tenslotte vergelijken we sommige van de consequenties van dit 'Axioma voor Blackwell Gedetermineerdheid' met de consequenties van het oorspronkelijke Axioma van Gedetermineerdheid.
In het tweede deel, getiteld 'Random Walks', beschouwen we recurrentie in zelfversterkende willekeurige reizen, waarbij de kanten in een graaf worden bereisd met een waarschijnlijkheid die kan veranderen voor een tweede of latere doortocht. We concentreren ons op het geval dat de waarschijnlijkheid voor een kant maar één keer verandert, na de eerste keer dat de kant wordt bereisd. Als een speciaal geval laten we zien dat de een-keer-versterkte willekeurige reis bijna zeker recurrent is als de versterking voldoende klein is (een uitbreiding van een resultaat van T. Sellke[31]), en ook als de versterking voldoende groot is. Voor het laatste resultaat gebruiken we een toepassing van nonstandaard analyse op grafentheorie.
Het derde deel, getiteld 'The EMILE Grammar Inducer', gaat over het EMILE programma, een programma dat een tekst inleest, en zonder enige voorkennis probeert om de grammatikale structuur van de taal te bepalen. De basisconcepten en algoritmes die ten grondslag liggen aan het programma worden be-

handeld, zowel als de resultaten van deze benadering, zowel in theorie als in de praktijk. Het wordt betoogd dat natuurlijke talen voldoen aan een conditie genaamd *shallowness*, ondiepheid, en dat dit betekent dat het EMILE programma goed zal werken voor natuurlijke talen. In een aparte appendix wordt expliciete pseudo-code gegeven voor de sub-algoritmes gebruikt in EMILE.

Abstract

This dissertation consists of three disjunct parts.
The first part, titled 'Blackwell Games', is about the problem of determinacy of Blackwell games, a class of infinite games of imperfect information, where both players simultaneously select moves from a finite set, infinitely many rounds are played, and payoff is determined by a Borel measurable function f on the set of possible resulting sequences of moves. We give elementary proofs of determinacy for Blackwell games whose payoff function is an indicator function of a Borel set up to complexity $G_{\delta\sigma}$. For general Borel payoff functions, we give a reduction, found by D.A. Martin[16], to the known result of determinacy of Borel perfect information games. We also consider Blackwell games whose payoff function is not Borel measurable, and formulate an analogue of the Axiom of Determinacy for these games, Finally, we compare some of the consequences of this 'Axiom of Blackwell Determinacy' with those of the original Axiom of Determinacy.
In the second part, titled 'Random Walks', we consider recurrence in reinforced random walks, where edges in a graph are traversed with probabilities that may be different (reinforced) at second, third etc. traversals. We focus on the case where the probability for any edge only changes once, after its first traversal. As a special case, we show that the once-reinforced random walk on the infinite ladder is almost surely recurrent if reinforcement is small (extending a result by T. Sellke[31]), as well as when reinforcement is sufficiently large. For the last result, we use an application of nonstandard analysis to graph theory.
The third part, titled 'The EMILE Grammar Inducer', is about the EMILE program, a program that reads in a text, and without prior knowledge attempts to determine the grammatical structure of the language. The basic concepts and algorithms underlying the program are discussed, as well as the results of this approach, both in theory and in practice. It is argued that natural languages satisfy the condition of *shallowness*, and that this implies that the EMILE program will work well for natural languages. In a separate appendix, explicit pseudo-code for each of the sub-algorithms of EMILE is given.

Titles in the ILLC Dissertation Series:

ILLC DS-1996-01: **Lex Hendriks**
Computations in Propositional Logic

ILLC DS-1996-02: **Angelo Montanari**
Metric and Layered Temporal Logic for Time Granularity

ILLC DS-1996-03: **Martin H. van den Berg**
Some Aspects of the Internal Structure of Discourse: the Dynamics of Nominal Anaphora

ILLC DS-1996-04: **Jeroen Bruggeman**
Formalizing Organizational Ecology

ILLC DS-1997-01: **Ronald Cramer**
Modular Design of Secure yet Practical Cryptographic Protocols

ILLC DS-1997-02: **Nataša Rakić**
Common Sense Time and Special Relativity

ILLC DS-1997-03: **Arthur Nieuwendijk**
On Logic. Inquiries into the Justification of Deduction

ILLC DS-1997-04: **Atocha Aliseda-LLera**
Seeking Explanations: Abduction in Logic, Philosophy of Science and Artificial Intelligence

ILLC DS-1997-05: **Harry Stein**
The Fiber and the Fabric: An Inquiry into Wittgenstein's Views on Rule-Following and Linguistic Normativity

ILLC DS-1997-06: **Leonie Bosveld - de Smet**
On Mass and Plural Quantification. The Case of French 'des'/'du'-NP's.

ILLC DS-1998-01: **Sebastiaan A. Terwijn**
Computability and Measure

ILLC DS-1998-02: **Sjoerd D. Zwart**
Approach to the Truth: Verisimilitude and Truthlikeness

ILLC DS-1998-03: **Peter Grunwald**
The Minimum Description Length Principle and Reasoning under Uncertainty

ILLC DS-1998-04: **Giovanna d'Agostino**
Modal Logic and Non-Well-Founded Set Theory: Translation, Bisimulation, Interpolation

ILLC DS-1998-05: **Mehdi Dastani**
Languages of Perception

ILLC DS-1999-01: **Jelle Gerbrandy**
Bisimulations on Planet Kripke

ILLC DS-1999-02: **Khalil Sima'an**
Learning efficient disambiguation

ILLC DS-1999-03: **Jaap Maat**
Philosophical Languages in the Seventeenth Century: Dalgarno, Wilkins, Leibniz

ILLC DS-1999-04: **Barbara Terhal**
Quantum Algorithms and Quantum Entanglement

ILLC DS-2000-01: **Renata Wasserman**
Resource Bounded Belief Revision

ILLC DS-2000-02: **Jaap Kamps**
A Logical Approach to Computational Theory Building (with applications to sociology)

ILLC DS-2000-03: **Marco Vervoort**
Games, Walks and Grammars: Problems I've Worked On

ILLC DS-2000-04: **Paul van Ulsen**
E.W. Beth als logicus

ILLC DS-2000-05: **Carlos Areces**
Logic Engineering. The Case of Description and Hybrid Logics