

The Characterization of Chance and Skill in Games



Dissertation

zur Erlangung des akademischen Grades
doctor rerum naturalium (Dr. rer. nat.)

vorgelegt dem Rat der Fakultät für Mathematik und Informatik der
Friedrich-Schiller-Universität Jena

von Diplom-Informatiker Jakob Erdmann
geboren am 11.07.1980 in Königs Wusterhausen

Gutachter

1. Prof. Dr. Ingo Althöfer, FSU Jena
2. Prof. Dr. Matthias Löwe, WWU Münster

Tag der öffentlichen Verteidigung: 9. Juli 2010

Contents

Zusammenfassung	iv
Overview	vi
Acknowledgements	vii
I The Chanciness of Two-Player Zero-Sum Games	1
1 Introduction	2
2 Previous Work	4
3 Terms and Concepts	5
4 The Concept of Influence	7
5 The Relative Influence of Chance	13
6 Chanciness	16
7 Simulation Results	18
8 Relevance of <i>Chanciness</i> for Human Agents	22
9 Conclusions and Open Problems	25
II Chanciness and Controllability of N-Player General-Sum Games	29
10 Introduction	30
11 Previous Work	32
12 Terms and Concepts	36
13 The Concept of Influence	42

14 Chanciness	57
15 Controllability	65
16 Other Influences	72
17 Fundamental Results	76
18 Approximation of <i>Chanciness</i> and <i>Controllability</i>	82
19 Model Games	89
20 “EinStein würfelt nicht!” and Variants Thereof	95
21 Conclusion	108
 Bibliography	 113
 Appendix	 118
A Lottery Chess	118
B Lookup Sampling	119
C Bounding <i>chanciness</i> and <i>controllability</i> for a Small Game.	121
D Proof that $ \mathbb{E}(\text{IC}) \equiv 0$	123
E Exemplary Calculations of cha for a match of Game G_{example}	124
F Exemplary Calculations of CHA for Game G_{example}	126
G Proof that $\text{cha}_1(g) \equiv \text{cha}_2(g)$ for Two-Player Zero-Sum Games	128
H Estimation of the Approximation Error of CHA and CON	130
I Result Statistics for “EinStein würfelt nicht Tetra”	131

Zusammenfassung

Der Mensch spielt nur, wo er in voller Bedeutung des Wortes Mensch ist, und er ist nur da ganz Mensch, wo er spielt.

Friedrich Schiller

Die große Bedeutung des Spielens in der menschlichen Kultur steht sicherlich außer Frage. Allein die Geschichte der Brettspiele reicht tausende von Jahren in die Vergangenheit. Die frühesten Erwähnungen des Spiels *Go* sind wenigstens 2500 Jahre [Pot 1984] alt, und noch heute erfreut sich dieses Spiel großer Beliebtheit. Das bei Kindern beliebte *Leiterspiel* (auch *Schlangen und Leitern*) nimmt seinen Ursprung bei dem alt-indischen Spiel *Moksha Patamu* und ist wenigstens 700 Jahre alt [Par 1999, Nat 2005]. Das Brettspiel *Backgammon*, um noch ein weiteres Beispiel zu nennen, leitet sich möglicherweise von den Römischen Spiel *Duodecim Scripta* ab und wäre damit über 2000 Jahre alt [Aus 1934].

Die drei genannte Spiele illustrieren unterschiedliche Ausprägungen eines wichtigen Merkmals. Während das Spiel *Go* als „Strategiespiel“ bezeichnet wird, handelt es sich beim *Leiterspiel* um ein sogenanntes „Glücksspiel“. Diese Einordnung drückt aus, welchen Einflüssen die Ergebnisse einer Spielpartie ausgesetzt sind. Beim *Leiterspiel* wird der Spielausgang ausschließlich durch Würfeln bestimmt und ist damit Ausdruck von unbeeinflussbarem „Glück“ bzw. „Pech“. Bei dem Spiel *Go* hingegen wird kein Würfel gebraucht, und das Ergebnis hängt ausschließlich von den Entscheidungen und Strategien der Spieler ab. Das Spiel *Backgammon* nimmt eine interessante Zwischenposition bei dieser Einordnung ein, da der Partieausgang sowohl von Spielerentscheidungen als auch vom Würfelglück abhängt¹. Es stellt sich daher die Frage, ob sich die Einteilung in Strategiespiele einerseits und Glücksspiele andererseits sinnvoll verfeinern lässt. Wo ließe sich beispielsweise das beliebte *Mensch Ärgere Dich Nicht!* einordnen? Liegt es eher zwischen *Go* und *Backgammon* oder zwischen *Backgammon* und dem *Leiterspiel*?

Aus mathematischer Sicht werden Brettspiele zumeist unter dem Gesichtspunkt der Optimierung betrachtet. An erster Stelle steht die Frage: ob und mit welcher Vorgehensweise sich ein Spiel gewinnen lässt [BCG 1982]. Die vorliegende Arbeit weicht von diesem Muster ab und widmet sich stattdessen der Frage:

Wie lassen sich Brettspiele zwischen den Kategorien „Glücksspiel“ und „Strategiespiel“ einordnen?

¹Das ägyptische Spiel *Senet* liefert ein noch älteres Beispiel für die Kombination aus Strategie und Glück. Sein Alter wird auf 5000 Jahre geschätzt. Die exakten Regeln sind jedoch nicht bekannt [Pic 1980].

Zur Klärung dieser Frage widmen wir uns dem Konzept der Beeinflussung. Wir führen die Beeinflussung von Spielergebnissen auf den Einfluss einzelner Spielzüge zurück. Einige „Spielzüge“ gehen auf den Ausgang von Zufallsereignissen (z. B. Würfelergebnisse) zurück. Die anderen Spielzüge lassen sich den verschiedenen Spielern zuordnen. Diese Zuordnung ermöglicht uns schließlich den Vergleich des Zufallseinflusses mit dem Einfluss der verschiedenen Spieler.

Die vorliegende Arbeit gliedert sich in zwei Teile. Im ersten Teil betrachten wir die Klasse der Zwei-Personen-Nullsummenspiele mit vollständigen Informationen. Fußend auf unserem Vergleich der verschiedenen Einflüsse entwickeln wir die Maßzahl *chanciness* (Glücksbestimmtheit). Der *chanciness*-Wert eines Spieles ist nahe null falls das Spiel kaum von Würfelergebnissen und ähnlichen Zufallselementen beeinflusst wird. Spiele, deren Ausgang stark vom Zufall beeinflusst wird, haben hingegen einen *chanciness*-Wert nahe bei eins. Ein wichtiges Ergebnis unserer Betrachtungen ist die Abhängigkeit dieser Maßzahl vom Verhalten der Spieler. Am Beispiel von computergesteuerten Spielern berechnen wir näherungsweise die *chanciness* für ausgewählte Spiele. Dieser Teil der Arbeit basiert auf einer vorherigen Publikation des Autors [Erd 2009].

Im zweiten Teil der Arbeit verallgemeinern wir unsere Vorgehensweise auf die Klasse der Nicht-Nullsummenspiele mit beliebiger Spielerzahl und vollständigen Informationen. Dies führt uns zu der Berücksichtigung verschiedener Spielerperspektiven. Es sind Spiele konstruierbar, die aus der Sicht eines Teilnehmers den Charakter eines Glücksspiels haben, während sie aus der Perspektive anderer Teilnehmer als Strategiespiel erscheinen. Wir betrachten daher die *chanciness* eines Spiels aus Sicht der verschiedenen Spielerperspektiven (Startpositionen). Wir zeigen außerdem, dass eine einzelne Maßzahl nicht ausreicht, um alle denkbaren Spiele und Perspektiven einzuordnen. Daher führen wir die zusätzliche Maßzahl *controllability* (Steuerbarkeit) ein. Der *controllability*-Wert eines Spiels beschreibt den Einfluss eines Spielers auf sein privates Spielresultat. Der Wert liegt nahe bei null falls dieses Resultat hauptsächlich vom Einfluss der Mitspieler sowie vom Zufall abhängt. Ist andererseits das Ergebnis jedes Spielers hauptsächlich von dessen eigenem Handeln bestimmt ergibt sich ein *controllability*-Wert nahe eins. Die Summe von *chanciness* und *controllability* ist kleiner oder gleich eins. Bei Spielen, deren Ergebnis stark durch eine „ungerechte“ Ausgangssituation bestimmt ist, liegt die Summe aus *chanciness* und *controllability* nahe bei null.

Die Ergebnisse unserer Arbeit haben einen ausgeprägt theoretischen Charakter. Die starke Abhängigkeit unserer Maßzahlen von einem Verhaltensmodell der Spielteilnehmer erschwert die praktische Anwendung. Dennoch weisen wir am Beispiel von computergesteuerten Spielern nach, dass eine näherungsweise Berechnung möglich ist.

Overview

Games are an important part of mankind's cultural heritage. Board games such as go and backgammon are thousands of years old, yet they still have an enthusiastic following. The fact that more than a thousand new board games are published each year[Ald 2000] underlines the public interest in game playing.

One of many properties, used to classify and discuss games is the dependence of game results on chance and player skill. Typically, games are described as “game of chance”, “game of skill”, or they are said to lie somewhere in between those extremes. The quantitative assessment of the relationship between chance and skill has received little scholarly attention.

In this thesis we offer a novel approach to the characterization of chance and skill in games. Our approach is based on the analysis of individual moves. This allows us to analyze matches (single game rounds) as well as games.

The thesis is organized in two parts. In the first part, we present our analysis for the class of two-player zero-sum games with perfect information. We introduce the measure of *chanciness* to describe the relative influence of chance on games. This part is based on a previous publication by the author [Erd 2009].

In Part II we generalize our approach and apply it to the class of general-sum games of perfect information with an arbitrary number of players. We also supplement the measure of *chanciness* with the new measure *controllability*. The *controllability of a game* describes the relative influence of the players on their personal game results.

Both parts can be read independently of each other. However, familiarity with our basic concepts from Part I helps with understanding the generalizations in Part II.

Acknowledgements

The initial impulse for this work grew out of many pleasant discussions about games and their properties with Prof. Dr. Ingo Althöfer. Were it not for his guidance, this thesis would not have been written. I also wish to thank him for his patience and for helping me to see the big picture.

My friends and colleagues have supported this work with countless fruitful discussions, critical comments and with their help in proofreading. I am particular grateful to Rico Walter, Martin Dörnfelder, Wolfgang Schnerring, Andreas Schäfer, and Philipp Kratz.

The second person without whom this thesis could not have been written is my wife Johanna. Her constant encouragement gave me the strength to persevere. I also wish to thank my mother for her enormous support in the final phase of writing.

The author of this thesis was supported by the *3-Hirn-Grant* and by a stipend from the *Jena Graduate Academy*.

This thesis is dedicated to my son Anton Ijon.

Part I

The Chanciness of Two-Player Zero-Sum Games

1 Introduction

Games are conventionally classified as games of chance, games of skill, or games with mixed characteristics. The desire for such a classification stems from various sources:

- Many states have laws to regulate gambling, in other words, playing games of chance with money at stake. It is up to legislators and lawyers to decide which games are within the scope of such laws. In Germany, the gambling law “Glücksspielgesetz” defines games of chance thus: A game of chance is one in which outcomes are fully or mostly determined by chance [Sta 2007]. The term “mostly” is not elucidated. This distinction becomes even less clear when the accumulated outcomes of numerous game rounds are considered. When the game in question involves even a tiny bit of skill, the central limit theorem may be used to argue that the accumulated result of a long sequence of games is determined mostly by skill [Alo 2007].

The controversial nature of the skill / chance distinction is apparent from US-American court rulings regarding poker [BDS 2002]. Poker was ruled a *game of skill* in Pennsylvania [Cyp 2009] and a *game of chance* in North Carolina [Bur 2007]. This contradiction leaves open the task of quantifying the skill and chance inherent in a game.

- Naturally, the people who play games are interested in matters of chance and skill. Some players prefer games of chance over games of skill, while others have opposite tastes. In this context, a precise classification of games may not be important. However, the results of individual matches (instances) lends itself to analysis. Having won a game with a chance component, we may ask ourselves whether our victory was due to our skillful moves or simply caused by a string of lucky dice rolls [EL 2009]. The answer to this question is central to the feeling of accomplishment of the players and thus to the enjoyment of the game.
- A third group of people who concern themselves with the balance of chance and skill in games are game inventors. When designing a game for a particular audience, the inventor may wish to realize a particular level of chance and skill. Unfortunately, it is hard to measure the effects that subtle rule variations have on this balance.

We are interested in *computer-aided game inventing* as proposed by Althöfer [Alt 2003]. We wish to use computers in place of human players to test newly designed games. A prerequisite of such schemes is the existence of computer-readable rule descriptions. The field of *general game playing* [LHH⁺ 2008, Pel 1994, FB 2008] provides computer algorithms capable of using such rule descriptions to play games against each other.

Our eventual goal is an algorithm for automatically determining the extent of chance and skill in a game. In this paper we present the concept of *chanciness* C , our theoretical groundwork for such an undertaking. The *chanciness* of a game is high if game results are strongly influenced by chance, and it is low if the game allows for skillful play. We will show that our *chanciness* depends on the game players (agents). This leads to a measure of $C(\text{game}, X, Y)$ where X and Y are specific agents who compete in a two-player game. We narrow our focus to two-player zero-sum games of perfect information, sequential moves, finite length, and a finite number of moves in every game state.

In Chapter 2 we discuss previous works on formalizing the extent of chance and skill in games. We introduce terms and concepts for our discussion in Chapter 3. Our concept of *chanciness* relies on the analysis of individual game instances (matches). To what extent did player moves and random events influence the match result? We deal with the subject of *influence* in Chapter 4. Based on our notion of influence we define the *relative influence of chance* in Chapter 5. In Chapter 6 we generalize from matches to games and define *chanciness*. Chapter 7 contains exemplary measurements of *chanciness* for different games that are played by various computer agents. Chapter 8 relates our concept of *chanciness* to games played by humans. In Chapter 9 we summarize our findings and propose further lines of research.

2 Previous Work

The question whether games should be legally classified as games of skill or games of luck has been addressed by Borm and Dreef since 2001 [BvdG 2001, DBvdG 2004, Dre 2005]. Based on three types of playing agents they introduce the concept of *Relative Skill* (RS) and give values for the games of *Roulette*, *Black Jack*, *Poker* and other games. Games that allow for skillful play receive a higher rating of RS than games in which chance dominates the result. Thus, RS expresses the opposite idea of *chanciness* and is closely linked to our question.

RS is measured by comparing expected game rewards of beginners, optimal players and fictive players. “Fictive players” know in advance the result of any kind of random event.

$$RS = \frac{\text{gain}_{\text{optimal}} - \text{gain}_{\text{beginner}}}{\text{gain}_{\text{fictive}} - \text{gain}_{\text{beginner}}} \quad (2.1)$$

Dreef recommends that the playing style of the beginner is to be modelled in one of three ways depending on what is called “the structure” of the game.

- random move selection
- quantitative analysis of human beginners playing style
- using domain expert knowledge

Our objections to the measure of *relative skill* are the following. First of all, RS depends on the analysis of optimal play. Discovery of optimal play, however, requires resources beyond the scope of automatic analysis. For some games, even advanced human play will be far from optimal, which diminishes the relevance of the measure for analysis of human behavior. Secondly, the notion of “fictive players” allows the creation of games that receive ratings of RS which contradict common sense (see the game of *Lottery Chess* in Appendix A).

However, as our most important objection to *relative skill*, we wish to argue that an assessment of chance and skill for a given game depends strongly on the particular behavior of the playing agents. Therefore, a measure that only considers three agents gives an incomplete picture.

3 Terms and Concepts

In this section, we introduce the necessary terms and concepts for our discussion of *chanciness*. The starting point is the central term *game*. We wish to discuss non-deterministic¹ finite zero-sum games of two players, perfect information and sequential moves².




Figure 3.1: An exemplary game tree with chance nodes.

As our game model we use the extensive form with chance moves [SLB 2008]. A game G is represented by a finite tree as in Figure 3.1. Nodes of the tree correspond to game states. The leafs of the tree represent terminal game states and are labeled with the deterministic result r of the game. The game result r is a real number that gives the reward for player *Max* (his goal is to maximize r). As G is a zero-sum game, the other player, *Min*, receives $-r$ as reward. The inner nodes represent non-terminal states that require a move. They are labeled as either *Max* or *Min* nodes if that player is to move. If the game state requires a “chance move”, such as a roll of the dice or the drawing of cards from a randomized deck, the node is labeled as a “chance node”. The edges of the tree correspond to the available moves by either players (“player moves”) or chance (“chance moves”). Chance moves are labeled with their respective probabilities. For every chance node s we define $\mathcal{C}(s, s')$ as the probability that chance will move from s to the child node s' .

A match g of game G is given by a tuple $g = (X, Y, (s_0, s_1, \dots, s_k))$. X and Y are agents who take on the roles of players *Max* and *Min*. (s_0, \dots, s_k) is a path of nodes (game states) through the tree from the root to a leaf node. As we are interested in *computer-aided game*

¹We treat deterministic games as a special case of non-deterministic games.

²Games with sequential moves are also known as *Dynamic Games*.

3 Terms and Concepts

inventing, our agents are computer algorithms. The terms “player” and “agent” are almost synonymous. We use the term “player” to emphasize the roles in a game (*Max* and *Min*). The term “agent” emphasizes the algorithm that is used to fill the role of a player.

Since we are not interested in *how* our agents form their move decisions, we can use a very simple agent model. It suffices that the agents are somehow able to select their move from the finite set of available moves. To reflect our ignorance of the selection process, we use a stochastic model. For every game G let S_G denote the set of all possible states of G . Let $S_G^{Max} \subseteq S_G$ denote the set of all states in which player *Max* has to move. For every game state $s \in S_G$ let $\text{Children}(s) = \{c_1, c_2, \dots, c_n\}$ denote the set of states that can be reached by one legal move. An agent X , capable of playing G as *Max* is given by a function from $S_G^{Max} \times S_G$ to the real numbers. $X(s, s') = p$ is the probability that agent X will move from state s to the child node s' . The probabilities given by X must meet the following constraint:

$$\left[\sum_{c \in \text{Children}(s)} X(s, c) \right] = 1, \quad X(s, c) \geq 0 \quad (3.1)$$

Likewise, agents capable of playing G as *Min* are given by a function from $S_G^{Min} \times S_G$ to the real numbers.

Deterministic agents are handled as a special case of stochastic agents. For the sake of simplicity our agents only consider the current state s when selecting their move. This stochastic independence of moves is not a requirement for the calculation of *chanciness*. In principle, agents could consider previous states and opponent moves when making their move³

Using the terms defined above, we can describe a wide selection of board games such as chess, backgammon and *Pachisi* (reduced to two players). When viewed as a two-player game between player and bank, even *Roulette* can be modeled in these terms. Care must be taken when modelling some of these games as a tree of finite depth. While all games have finite length in practice, some games such as backgammon allow for arbitrarily long matches⁴. One solution lies in modifying the rules slightly, similar to the *50 move rule* [FID 2005] in tournament chess (without this rule, chess would also allow matches of arbitrary length). For the practical application of our game measures, it is sufficient that matches terminate eventually.

³This assumes that the game is represented by a real tree. (Two chess positions with the same board and different histories would be represented by two different nodes).

⁴However, it has been proven that backgammon ends with probability 1 [McC 1994].

4 The Concept of Influence

When describing a game G as a game of chance or a game of skill humans employ an intuitive understanding of influence. To qualify as a game of chance, G must involve some sort of randomizing device and that device must have an influence on the game result. Likewise, to deserve the label “game of skill”, there must be room for making “good” and “bad” move decisions by the players. In other words, players must be able to influence game results through skillful decisions. If the result of a game is subject to influence of both chance and skill, we can try to compare their relative importance. For a scientifically meaningful comparison we must develop a rigorous definition of *influence*.

In Section 4.1 we explain our approach for formalizing the concept of *influence*. This leads us to the need for evaluating game states. We discuss problems related to positional evaluation in Section 4.2 and introduce *relative positional values* in Section 4.3. We explain how *relative positional values* solve the aforementioned problems in Section 4.4 and finally present our definition of *influence* in Section 4.5.

4.1 Approach to Formalization

To understand influence we must distinguish between the game G and a match g of G that took place between specific players. When talking about the game G we actually discuss *possible* influence. When we say G is a game of chance, we are saying, a typical match g of G is strongly influenced by chance. However, not every match plays out the same. If one player gets very lucky, he may win without effort. In another match, the influence of dice rolls may favor both players equally and the result is then decided by the skill of the players. Therefore, we will define *influence* for matches g of game G . After we have understood how chance and skill affect matches, we generalize our findings to talk about the properties of G .

What exactly is influence? When we talk about the influence of moves (player moves or chance moves), we talk about their influence on the game result. A single move rarely decides the game. Instead, moves lead to game states in which a player has improved or reduced chances of winning. From the perspective of player *Max*, a move has a positive influence if it improves his position and hence increases his expected reward from the game. This interpretation is crucial because it allows us to quantify influence. We evaluate consecutive game states and define *influence* as the change in evaluation. This approach leads us to the problem of evaluating game states.

4 The Concept of Influence

4.2 Difficulties of Positional Evaluation

For deterministic 2-player zero-sum games such as chess, the *minimax* algorithm allows us to compute the game-theoretic value of every position by backward induction [Zer 1913, vN 1927]. These values represent the game result for agents who play optimally. For non-deterministic games, the *expectimax* algorithm (a generalization of minimax for trees with chance nodes) computes the expected values for optimal agents [SLB 2008]. The expectimax value $V(s)$ of a position s is given by a recursive formula:

$$V(s) := \begin{cases} r(s) & \text{if } s \text{ is terminal} \\ \max_{c \in \text{Children}(s)} V(c) & \text{if Max moves in } s \\ \min_{c \in \text{Children}(s)} V(c) & \text{if Min moves in } s \\ \sum_{c \in \text{Children}(s)} \mathcal{C}(s, c) * V(c) & \text{if Chance moves in } s \end{cases} \quad (4.1)$$

Here $r(s)$ is the game result for a terminal position s . By starting at the leafs and backing up towards the root, the expectimax value can be computed for every node. Figure 4.1 gives the expectimax values for all game states of our example game from Chapter 3.




Figure 4.1: A game tree with expectimax values.

The minimax and expectimax algorithm are widely used to analyze games. They also form the basis for many game playing agents. However, in light of our goals, the positional values computed by these algorithms have two important drawbacks.

- The *Interpretation Problem* arises when using expectimax values to reason about games between imperfect agents. Consider the game tree in Figure 4.1. In the root position a , an optimal agent playing as *Max* moves to position b which leads to the maximum reward $r = 1$. For that agent, moving to state c instead would be a “bad” move (adverse influence) because his expected reward would drop from 1 to -1. However, if an imperfect agent was playing as *Max*, that assessment might

change. Consider an imperfect agent X who does not play optimal in position d . Let X have the probabilities $(0.25, 0.25, 0.5)$ for selecting the left, middle and right child of d respectively (based on Figure 4.1). Thus, in position d , agent X would have an expected reward of -2 . If we assume further that X plays optimal in position f and Y plays optimal as well, we must conclude that the move to b is “bad” for agent X while the move to c is “good” (beneficial influence). This example demonstrates that positions must be evaluated with respect to the playing agents. As another example, imagine a chess master giving advice to a novice. A risky gambit that the chess master would play is not advisable for the novice who cannot exploit the situation. Thus, the gambit would be a “good” move for the master, but a “bad” move for the novice. When using expectimax values to calculate the influence of moves, the interpretation problem crucially effects our results.

- The practical use of expectimax values is further hampered by the *Resource Problem*: expectimax values cannot be accurately approximated with generic algorithms for large game trees. Since the expectimax algorithm requires the evaluation of almost¹ the complete game tree, it is unfeasible to compute the exact values for complex games. In practice, expectimax values are approximated by using heuristics. These heuristic methods can be very efficient but must often be tuned for every new game. General game playing algorithms avoid this tuning, but they give no guarantees as to the accuracy of their approximations.

In the next section we introduce a different approach to positional evaluation which avoids these problems.

4.3 Relative Positional Values

We define the *relative positional value* $\tilde{V}(s, X, Y)$ for a game state s and the agents X, Y . The term *relative* emphasizes that the positional value depends on the agents.

The result r of a match that is started in state s and then played out by the agents X (as *Max*) and Y (as *Min*) can be seen as the realization of a random variable R . When s is a terminal state the realization of R is a constant. Otherwise the value of r depends on the stochastic behavior of X, Y , and on chance. We define the *relative positional value* as the expected value of R .

Definition 4.2

$$\tilde{V}(s, X, Y) := \mathbb{E}(R)$$

Figure 4.2 shows an annotated game tree. The move probabilities for chance and for exemplary agents X, Y are given for every move in the tree. The *relative positional values*

¹A tree may allow safe pruning of some branches (See [HBS 2006]).

4 The Concept of Influence

for every node in the tree were computed by a recursive formula that looks similar to Formula 4.1 but replaces the *min* and *max* terms with expectation terms.

$$\tilde{V}(s, X, Y) := \begin{cases} r(s) & \text{if } s \text{ is terminal} \\ \sum_{c \in \text{Children}(s)} X(s, c) * \tilde{V}(c, X, Y) & \text{if } X(\text{Max}) \text{ moves in } s \\ \sum_{c \in \text{Children}(s)} Y(s, c) * \tilde{V}(c, X, Y) & \text{if } Y(\text{Min}) \text{ moves in } s \\ \sum_{c \in \text{Children}(s)} \mathcal{C}(s, c) * \tilde{V}(c, X, Y) & \text{if Chance moves in } s \end{cases} \quad (4.3)$$

Again, $r(s)$ is the game result for a terminal position s .




Figure 4.2: A game tree with move probabilities for specific agents X and Y . The *relative positional values* for these agents are given for every node.

4.4 Relative Positional Values Compared to Expectimax Values

Our new concept of *relative positional values* solves the two problems from Section 4.2. The *Interpretation Problem* is resolved by taking the behavior of imperfect agents into account. Figure 4.2 above shows the same game tree as Figure 4.1, this time annotated with *relative positional values* for the imperfect agents X and Y . Our intuitions of “good” and “bad” moves for X are reflected by increasing and decreasing positional values.

The move from root node a to node c is “good” because $\tilde{V}(c, X, Y) = -1 > -1.4 = \tilde{V}(a, X, Y)$. Likewise, the move from a to the other successor b is “bad” because $\tilde{V}(b, X, Y) = -2 < -1.4 = \tilde{V}(a, X, Y)$. It must be stressed that these values (as well as the judgements of “good” and “bad”) are specific to the agents X and Y .

The *Resource Problem* from Section 4.2 stems from the algorithmic complexity of computing minimax and expectimax values. The exact computation of *relative positional values* is just as expensive as the computation of expectimax values.

4.5 Influence Defined

However, it is possible to approximate *relative positional values* with Monte Carlo methods thus saving computing time. As in Section 4.3 let R be a random variable that gives the result of a match between X and Y with start in state s . Each simulation of such a match yields a realization R_i of R . The average over n such realizations R_i can be used to approximate the expected value of R and hence the *relative positional value* of s for X and Y .

$$\tilde{V}(s, X, Y) \approx \tilde{V}_{\text{approx}} = \frac{1}{n} \sum_{i=1}^n R_i \quad (4.4)$$

Since we intend to use computer algorithms as agents, simulation games can be repeated as often as desired (given enough computing resources).

In Section 4.2 we noted that approximations of expectimax values either rely on game specific knowledge or fail to give guarantees on accuracy.

In contrast, approximations of *relative positional values* via Monte Carlo methods are general and give such guarantees. They are general because we already assume computer agents that can be used for simulation games (these agents are generally imperfect and therefore unsuitable for approximating expectimax values).

Monte Carlo methods also give guarantees on the accuracy of approximation. The standard deviation of $\tilde{V}_{\text{approx}}$ (and thus the error in approximation) varies inversely with the square root of the number of observations: $\sigma(\tilde{V}_{\text{approx}}) = \sigma(R)/\sqrt{n}$ (central limit theorem).

In conclusion we deem *relative positional values* suitable for evaluating the influence of game moves.

4.5 Influence Defined

Our concept of *relative positional values* from Section 4.3 allow us to define the *influence* I of moves. Let $m = (s, s')$ be a move that transforms position s into s' . We define the influence of that move in a game between the agents X, Y as follows.

Definition 4.5

$$I(m, X, Y) := \tilde{V}(s', X, Y) - \tilde{V}(s, X, Y)$$

This definition does not distinguish between “player moves” and “chance moves”. Figure 4.3 is based on the same game tree and agents as Figure 4.2. Additionally, it shows the influence of all moves.

As a consequence of Definition 4.5 we observe:

Observation 4.6 “Deterministic moves” always have influence 0.

We call a move $m = (s, s')$ “deterministic move” for agent X if state s is always followed by the specific move s' . In other words $X(s, s') = 1$. As an example see the move (b, d) in Figure 4.3. Here, the agent Y playing as *Min* moves from b to d with probability 1. Observation 4.6 follows directly from Equation 4.3 and Definition 4.5.

4 The Concept of Influence

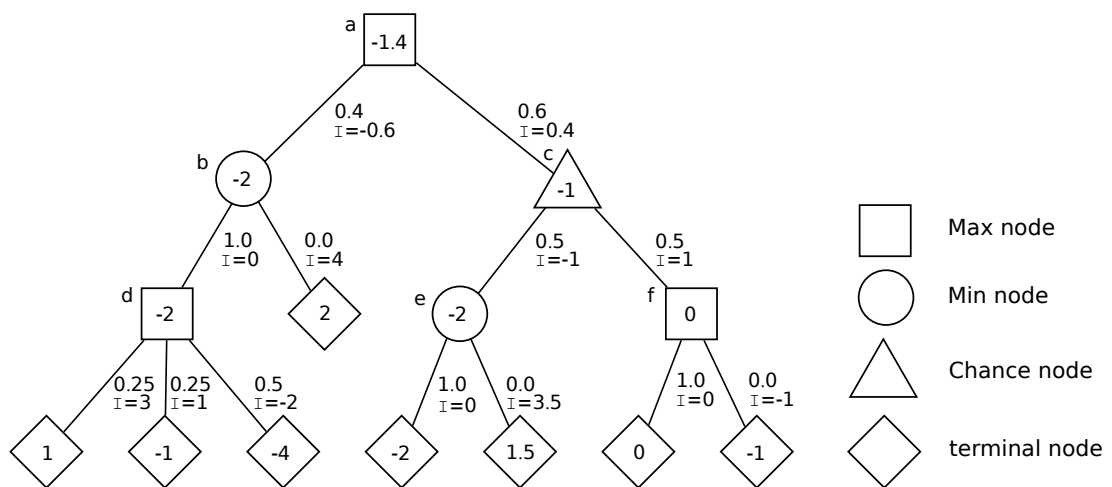


Figure 4.3: A game tree with move probabilities, *relative positional values* and the influence I for all moves.

5 The Relative Influence of Chance

In the previous section we have defined the *influence* $I(m, X, Y)$ of a move m in a game between the agents X and Y . Using this definition we can compare the influence of chance moves with the influence of player moves on the result of a match.

Let $g = (X, Y, (s_0, s_1, \dots, s_n))$ be a match of game G between X and Y (X playing as *Max*, Y as *Min*). Let m_i denote the move (s_i, s_{i+1}) . The result of g is the relative value of its final position $\tilde{V}(s_n, X, Y)$. As the value of a terminal game state does not depend on the agents anymore, *relative positional value* and expectimax value are the same, hence $\tilde{V}(s_n, X, Y) = V(s_n)$. Using our definition of *influence* we can express the result of g by the influences of the moves s_i .

Lemma 5.1

$$V(s_n) = \tilde{V}(s_n, X, Y) = \tilde{V}(s_0, X, Y) + \sum_{i=0}^{n-1} I(m_i, X, Y)$$

Idea of the proof: Applying Definition 4.5 leads to a telescoping sum that leaves only the term $\tilde{V}(s_n, X, Y)$ □

Lemma 5.1 “explains” how the combined influence of all moves leads to the result of match g . The term $\tilde{V}(s_0, X, Y)$ is special because it is not the influence of a move but the expected value of the starting position. It reflects the “fairness” of the game G combined with the skill difference of the agents X and Y . We call $\tilde{V}(s_0, X, Y)$ the *match expectation* ME. We can group all the moves that were made in g according to their mover. This gives use the set of player moves M_P and the set of chance moves M_C .

$$\begin{aligned} M_P &:= \{(s_i, s_{i+1}) \mid \text{Max or Min to move in } s_i\} \\ M_C &:= \{(s_i, s_{i+1}) \mid \text{chance to move in } s_i\} \end{aligned}$$

The sets M_P and M_C allow us to sum the combined influences of player moves and chance moves as IP and IC.

$$\begin{aligned} \text{IP} &= \sum_{m \in M_P} I(m, X, Y) \\ \text{IC} &= \sum_{m \in M_C} I(m, X, Y) \end{aligned}$$

The sets M_P, M_C and the values ME, IP, IC all depend on a concrete match g .

5 The Relative Influence of Chance

We can now rephrase Lemma 5.1 as

$$V(s_n) = \text{ME} + \text{IP} + \text{IC} \quad (5.2)$$

Based on these terms we define $c(g)$, the *relative influence of chance* on the result of match g .

Definition 5.3

$$c(g) := \frac{|\text{IC}|}{|\text{ME}| + |\text{IP}| + |\text{IC}|} \text{ if } |\text{ME}| + |\text{IP}| + |\text{IC}| \neq 0$$

We intentionally leave the special case of $|\text{ME}| + |\text{IP}| + |\text{IC}| = 0$ undefined. Chapter 6 will show that the definition of this special case has no influence on our measure of *chanciness*.

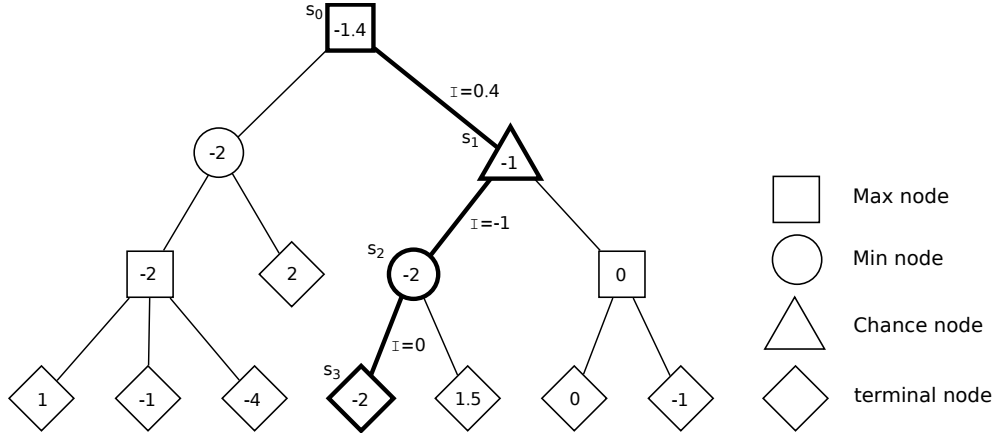


Figure 5.1: A game tree with and a single match (path) in that tree with move influence for the agents X and Y .

Figure 5.1 shows an exemplary match $g = (X, Y, (s_0 \dots s_3))$. We observe that Lemma 5.1 holds:

$$\begin{aligned} V(s_3) &= -2, \quad \text{ME} = -1.4, \quad \text{IP} = 0.4, \quad \text{IC} = -1 \\ -2 &= -1.4 - 1 + 0.4 \end{aligned}$$

We calculate the *relative influence of chance* for g .

$$c(g) = \frac{|\text{IC}|}{|\text{ME}| + |\text{IP}| + |\text{IC}|} = \frac{1}{1.4 + 1 + 0.4} \approx 0.36$$

5.1 Properties of $c(g)$

For a match $g = (X, Y, (s_0 \dots s_n))$, the *relative influence of chance* $c(g)$ is a real number from the interval $[0, 1]$. We can analyze how the components of Definition 5.3 are affected by match properties and how they affect the value $c(g)$.

5.1 Properties of $c(g)$

IC: If a game does not allow for chance moves, the influence of chance IC will be zero for any match g and hence $c(g)$ will be zero as well. For a game that does allow chance moves IC will generally vary from match to match. It may still be zero (or close to zero) for some matches if the influence of chance moves cancel out each other. For example, if a sequence of lucky dice rolls for player *Max* is followed by another sequence of dice rolls that favor player *Min*. Thus, even for games of chance, some matches may be completely decided by player actions.

IP: The term IP quantifies the combined influence of the players for a match g . Whereas games without chance moves are common, games without player moves are uncommon and may not be at all recognized as “games”. Even though player moves occur they may not be influential to the game result. For example, picking red or black in *Roulette* has no impact on the expected result. Thus, games that do not allow for “bad” and “good” moves will have low values of IP for any match g . According to Observation 4.6, a game where “good” moves are obvious to the agents will also allow for little player influence. This resonates strongly with a depreciation of games by humans who have fully mastered those games. Few people return to the game of *Tic-Tac-Toe* after mastering its “strategy”. High values of IP indicate that a game allows for influential moves and that the agents have not yet mastered the strategy. Consequently such matches will receive lower values of $c(g)$.

ME: The *match expectation* ME is defined as the expected result of a match (starting at the root position) between specific agents. Its value does not depend on the moves of a match because all matches for a game share the same root position. It is effectively a function of the game G and the agents X, Y . This function $ME(G, X, Y)$ combines two different aspects of a match: the skill difference between the agents and the asymmetry of game G . If two agents of differing skill play a game that allows for influential moves, the better player can expect a higher reward. This usually leads to a higher value of $|ME|$. If the asymmetry of the player positions confers an advantage to one of the players, this also tends to increase the value of $|ME|$. Either effect is independent of chance and tends to increase the predictability of game results. Consequently, $|ME|$ appears in the denominator of Definition 5.3.

As a remarkable consequence of combining skill difference and game asymmetry in the term ME, the playing order of the agents is relevant for the value of $c(g)$. In general $|ME(G, X, Y)| \neq |ME(G, Y, X)|$. To understand this, we must imagine an asymmetric game G (favoring player *Max*), a skilled agent X and a less skilled agent Y . If X plays as *Max*, ME will be high because the stronger agent plays in an advantageous position. If X plays as *Min*, ME will be closer to zero because skill differential and game asymmetry cancel each other out. In this case we have $|ME(G, X, Y)| > |ME(G, Y, X)|$.

6 Chanciness

In the previous section we have defined the *relative influence of chance* to analyze a specific match g of the game G . In this section we discuss how those results can be aggregated for the analysis of game G .

Before the agents X and Y play a match g of game G , the moves that will be taken and the result of the match are generally unknown. Moves and result depend on stochastic processes, thus IP, IC , and likewise the *relative influence of chance* $c(g)$ are random variables (ME is a constant).

It might seem intuitive to define *chanciness* as the expected value of $c(g)$. However, that approach would lead to counterintuitive results for some games. Consider a game G and two deterministic agents X, Y that lead to the following situation:

$$\begin{aligned} ME &= 1, \quad IP = 0 \quad \text{for all matches} \\ |IC| &= 0 \quad \text{for 90\% of the matches} \\ |IC| &= 10^{10} \quad \text{for 10\% of the matches} \\ c(g) &= \frac{0}{1+0+0} = 0 \quad \text{for 90\% of the matches} \\ c(g) &= \frac{10^{10}}{1+0+10^{10}} \approx 1 \quad \text{for 10\% of the matches} \end{aligned}$$

Most single matches are decided by the skill difference of the agents (X having an advantage over Y) and the expected value of $c(g)$ would be approximately 0.1. This naive approach however, fails to take the different magnitudes of influence into account. In 10% of the matches, chance has a dramatic influence on the result. If the agents were to play a series of 100 matches, it is very likely that the influence of chance would dominate the cumulative result. In this light a *chanciness* of 0.1 seems inappropriate.

The problem exemplified here is related to the *Fallacy of the Averages* [Wag 1969]. The discrepancy in interpretation boils down to the difference of applying the expectation operator \mathbb{E} before or after a nonlinear transformation (in this case division)¹. To promote the second interpretation we define the *chanciness* C of a game when played by X and Y as follows:

¹Another non-linear operation in the definition of $c(g)$ is the *absolute value* operation. Applying the expectation operator \mathbb{E} before taking the absolute value gives another candidate formula for *chanciness*. Since $\mathbb{E}(IC) = \mathbb{E}(IP) = 0$ this formula would lead to a constant *chanciness* of zero for any game G and agents X, Y (undefined for $ME = 0$). Obviously, this definition is not useful for discriminating games.

Definition 6.1

$$C(G, X, Y) := \frac{\mathbb{E}(|IC|)}{|ME| + \mathbb{E}(|IP|) + \mathbb{E}(|IC|)}$$

Applying this definition to the example above gives a *chanciness* of $\frac{10^9}{1+0+10^9} \approx 1$.

In Chapter 5 we left the *relative influence of chance* $c(g)$ undefined for a denominator of zero. This does not present a problem for Definition 6.1 as long as some matches are influenced by chance or skill. In theory, the denominator of Definition 6.1 might be zero as well. In our opinion, this is not a problem either, because games where $ME, \mathbb{E}(|IP|)$, and $\mathbb{E}(|IC|)$ all equal zero are rather uninteresting. $\mathbb{E}(|IC|) = 0$ implies that IC is always zero and thus chance moves are meaningless for the outcome of the game. Likewise $\mathbb{E}(|IP|) = 0$ implies that IP is always zero and hence each player move is either deterministic or meaningless for the outcome of the game. In fact, every state in the game tree that is ever visited would have the same *relative positional value* of zero.

Our *chanciness* $C(G, X, Y)$ measures the propensity of matches to be (relatively) influenced by chance. The *chanciness* for a specific game and specific agents can be computed exactly by calculating the influence and probability of every possible move in G . For the game and agents from Figure 4.3 we receive a *chanciness* of ≈ 0.197 .

For obvious reasons, this exhaustive method of computation quickly becomes unfeasible when considering larger game trees. Instead, we approximate *chanciness* via sampling. A number of matches g_i is played and the weighted sample mean of $c(g_i)$ is taken as approximation of *chanciness*.

$$C(G, X, Y) \approx \frac{\sum_{g_i} |IC|}{\sum_{g_i} |ME| + |IP| + |IC|} \quad (6.2)$$

In the next section we present the results of such measurements.

7 Simulation Results

In this Chapter we present measurements of *chanciness* for various combinations of games and agents. The combinations of games and agents are chosen so as to highlight important properties of *chanciness*. *Relative positional values* were calculated with the *Lookup Sampling* algorithm (see Appendix B), a more efficient version of the simple sampling algorithm described in Section 4.4.

7.1 PaCRaWa

We introduce the game PaCRaWa (*Partially Controlled Random Walks*). It serves as a very simple game model. More precisely $\text{PaCRaWa}[d, T, a]$ is a class of games, as the rules contain several parameters: the “power” of chance $d \in \mathbb{R}_0^+$, the number of turns $T \in \mathbb{N}$, and the beginner’s advantage $a \in \mathbb{R}$. The playing material consists of a counter variable $c \in \mathbb{R}$ and a fair coin. At the beginning of the game, c is set to zero. The two players *Max* and *Min* take turns. In every turn there are exactly two legal moves: increment c by one or decrement c by one. After a player move, the coin is tossed and according to its outcome, c is either incremented or decremented by d (the “power” parameter). Player *Max* takes the first move and the game ends after T turns ($2 \times T$ moves). At the end of the game, player *Max* is awarded $c + a$ as reward (*Min* receiving $-(c + a)$). For $d > 0$ this game contains both elements of skill and elements of chance as its result is influenced by player choices as well as random events. The trivial choice of moves in *PaCRaWa* serves to illustrate an important point. It is possible for agents to master the strategy of a game, whereupon the game may turn into a pure game of chance. Thus, the perceived nature of the game changes with the skill of the agents. This again justifies our belief that a game should be judged in reference to the agents.

We conducted sets of experiments with three different agents. Agent *X* plays perfectly, agent *Y* plays the correct move with probability 0.75 and the wrong move with probability 0.25. Agent *Z* plays the correct move with probability 0.5.

In a first set of experiments we tested the effect of different agent combinations and “chance powers” d on *chanciness*. Figure 7.1 generally shows a positive correlation between d and *chanciness*, which is consistent with our intuition regarding the influence of chance. The effect of different agents on *chanciness* is a bit more complicated. In the two-player matches, stronger agents experience more *chanciness*. This is to be expected, as the strong moves cancel each other out and gains can only be made with lucky flips of the coin. When a strong agent faces a weak opponent, the game has little *chanciness*, because the skill difference results in a definite score advantage for the stronger agent.

In a second set of experiments we tested the effect of different game lengths T on

$(Max, Min) \backslash d$	0.1	0.5	1.0
(X, X)	0.97	0.98	0.98
(Y, Y)	0.11	0.37	0.53
(Z, Z)	0.09	0.33	0.47
(X, Y)	0.06	0.24	0.37
(Y, Z)	0.05	0.21	0.34
(X, Z)	0.04	0.16	0.27

Figure 7.1: *chanciness* of PaCRaWa[$d, T = 10, a = 0$] for different chance powers d , and different agents (1000 matches per value, $\sigma < 0.01$).

chanciness. Figure 7.2 shows the *chanciness* for different game lengths and agents. We observe that games with an odd number of turns are asymmetric and give an advantage to player *Max*. If *Max* is played by a strong agent, that advantage translates into a high *match expectation* ME and thus reduced *chanciness*. ME also increases with game length, when the agents have differing skills.

$(Max, Min) \backslash T$	1	2	5	6	31	32
(X, X)	0.50	0.97	0.65	0.97	0.80	0.98
(Y, Y)	0.45	0.54	0.48	0.58	0.51	0.54
(Z, Z)	0.49	0.52	0.50	0.48	0.47	0.52
(X, Y)	0.50	0.45	0.37	0.41	0.29	0.30
(Y, Z)	0.45	0.38	0.36	0.36	0.27	0.27
(X, Z)	0.50	0.33	0.32	0.31	0.18	0.18

Figure 7.2: *chanciness* of PaCRaWa[$d = 1, T, a = 0$] for different game lengths T and various players \mathcal{P} (1000 matches per value, $\sigma < 0.01$).

The error in our approximations of *chanciness* is obvious in games between the agents (X, X) and even game lengths T . We would expect the player influence IP to be zero because *Max* and *Min* moves always cancel each other out. Likewise ME should be zero because both players have the same number of moves and the coin is fair. Consequently, *chanciness* should equal 1. However, in Figure 7.2 the highest value of *chanciness* is 0.98.

Even though PaCRaWa is an exceedingly simple game, it took 200 minutes to compute the data tables in this section (albeit with unoptimized code).

7.2 EinStein Würfelt Nicht

“EinStein würfelt nicht!” is a popular board game designed by Ingo Althöfer [Alt 2004a, Ald 2004]. The title can be translated either as “A single stone does not play dice” or

7 Simulation Results

“Einstein does not play dice”. This play on words alludes to the game’s rules and is also a reference to Albert Einstein’s famous quotation about God and dice. The game is played by two players who rely on chance and skill to fulfill their objectives. “EinStein würfelt nicht!” can be played online for free at [Mal 2002].




Figure 7.3: Win rates in EWN of Monte Carlo Agents with different strengths versus an opponent who moves at random.

We introduce a game with slightly different rules and examine the effects of that rule change on *chanciness*. We will denote the original game as EWN and our variant as $\overline{\text{EWN}}$.

In EWN, players alternate in moving their pieces which are numbered from 1 to 6. Before a move can be undertaken, a six-sided die is cast and the number on the die indicates the piece that must be moved. During the course of the game pieces are removed from play. If the active player has lost piece i and the die shows that number, the player may choose between the piece with the closest number to i , higher than i and the piece with the closest number lower than i (for detailed rules see [Alt 2004a]).

In our variant $\overline{\text{EWN}}$, this choice is abolished. The player always has to move the piece with the next available piece after i in the cyclic order $(1, 2, \dots, 6, 1, 2, \dots)$. Because of this rule change, some situations in $\overline{\text{EWN}}$ offer fewer move choices than they would in EWN.

For our experiments we have used agents that employ the *Monte Carlo* algorithm to select their moves [Abr 1990, BH 2003]. Game states are evaluated by the mean result of k random matches. The child move with the maximum value is then played. We denote these *Monte Carlo Agents* as MC_k , with k indicating the number of random matches. The playing strength of *Monte Carlo Agents* in game of EWN grows with k as shown in Figure 7.3.

Figure 7.4 shows the *chanciness* of different player match-ups for the games EWN and $\overline{\text{EWN}}$. It can be seen that in both games the *chanciness* is higher when the opponents

$(Max, Min) \setminus$ game	EWN	$\overline{\text{EWN}}$	σ
(MC ₁ , MC ₁)	0.51	0.57	0.028
(MC ₈ , MC ₈)	0.62	0.68	0.025
(MC ₆₄ , MC ₆₄)	0.71	0.76	0.022
(MC ₅₁₂ , MC ₅₁₂)	0.79	0.84	0.015
(MC ₁ , MC ₈)	0.44	0.53	0.023
(MC ₁ , MC ₆₄)	0.35	0.41	0.020
(MC ₁ , MC ₅₁₂)	0.30	0.40	0.017
(MC ₈ , MC ₁)	0.43	0.47	0.021
(MC ₆₄ , MC ₁)	0.33	0.35	0.018
(MC ₅₁₂ , MC ₁)	0.24	0.32	0.017
(MC ₅₁₂ , MC ₁₀₂₄)	0.82	0.85	0.013

Figure 7.4: *Chanciness* of EWN and $\overline{\text{EWN}}$ for different agents (100 matches per value). For each row the larger standard deviation is given.

are more skilled and higher when the players are closer to each other in strength. The *chanciness* of $\overline{\text{EWN}}$ is consistently higher than the *chanciness* of the original game EWN. We propose that the lack of movement choices in our variant lessens the influence of the players and thus makes for a more chance-dependent game. The least *chanciness* is found for matches between agents of greatly differing skill (MC₁, MC₅₁₂). This can be explained by a high value for the *match expectation* ME. The table also shows a consistently higher *chanciness* for matches in which the weaker agent takes the first move rather than the second move. This can be explained by the different values of ME as explained in Section 5.1.

Each data point in Figure 7.4 is the result of 100 simulation matches for the respective agents. The combined simulations required about 1000 CPU hours (1GHz Opteron). The high resource requirements when compared to our measurements for the game of PaCRaWa can be explained by the increased complexity of EWN and of the *Monte Carlo Agents* (especially those with high k).

8 Relevance of *Chanciness* for Human Agents

In the preceding sections we have introduced our concept of *chanciness* for games played by computer agents. What relevance does this concept have for matches between humans?

This difficult question has philosophical and practical aspects. Is the expected result of a game between human agents and thus the *relative positional value* (Section 4.3) well defined? The answer to this philosophical question ultimately depends on one's interpretation of probability [Haj 2009].

One important interpretation is the *frequentist interpretation*. It defines probability as the relative frequency of an event for a large number of trials. A game between humans cannot be repeated because the agents change as they learn from previous games. For that reason our definition of *chanciness* is meaningless for human agents under a *frequentist* interpretation of probability.

Another influential interpretation of probability is called *Bayesian Probability* (also *Bayesianism*). This interpretation defines probability as the “degree of belief” of rational agents and assigns probability to repeatable as well as to non-repeatable events. This interpretation gives meaning to *relative positional values* for human agents and hence to our definition of *chanciness* when applied to matches between human agents.

An equally pressing issue is the matter of computability of *chanciness* for human agents. Given a game G and a pair of human agents X, Y , can we discover the *chanciness* $C(G, X, Y)$ with an algorithm that uses the results of matches involving X and Y ?

As in the case for computer agents we should not expect an exact calculation of *chanciness* for most games. We therefore rephrase the question: given a game G and a pair of human agents X, Y , can we approximate the *chanciness* $C(G, X, Y)$ with an algorithm that uses the results of matches involving X and Y ?

As humans play much slower than computers, are less willing to repeat games and generally learn while they play, direct mass simulations (as with computer agents) can be ruled out. Instead, we suggest an indirect approach: if we had access to computer agents U, V that played very similar in style to our human agents X and Y , we could approximate *chanciness* for these computer agents. $C(G, U, V)$ would arguably be an approximation for $C(G, X, Y)$.

Of course, a general approximation of “style” is in itself a daunting task. If taken to the extreme this could be translated into passing the Turing Test [Tur 1950]. In fact, in an earlier version of his famous test, Turing explicitly talks about comparing the playing style of humans and machines for the game of chess [Tur 1948].

We conjecture that for some games it is sufficient to approximate the playing strength

of human agents to approximate their style. We would consider an approximation of style sufficient, if matches between the computer agents U, V covered similar regions of G 's game tree as matches between the humans X, Y . How this *similarity* should be defined is an open question. It should somehow involve the distribution of chance nodes and their potential influence. Games for which chance nodes are distributed homogeneously in the tree would then automatically fulfill the condition of match similarity.

When faced with such a homogeneous game G , we would measure the playing strengths of the humans X in G by conducting a limited number of matches between X and computer agents of different strengths. We then pick a computer agent U that matches the strength of X (and likewise a computer agent V that matches the strength of human Y). We would then take $C(G, U, V)$ as an approximation for $C(G, X, Y)$.

While we have no clear concept of homogeneity, we can construct games in which equal playing strength does not imply equal style (and equal *chanciness*). As an example, consider a hybrid of chess and backgammon: The first move that is taken is a decision whether chess or backgammon should be played. Obviously the tree of this game is not homogeneous as it has a chess subtree and a backgammon subtree which are quite different in their distribution of chance nodes.

The move of the first player decides whether the chess- or the backgammon-subtree is used. Given agents with the right combination of skills, this first move decision is independent of playing strength and only a matter of "style". Yet, the first decision determines whether the influence of chance for the match will be zero (chess) or non-zero (backgammon). Fortunately, common games appear to be far more homogeneous in this regard. Another way to increase the representativeness would be to rate actual human matches (e.g. from tournament records) with the help of computer agents.

Eventually we wish to use *chanciness* to automatically compare games that are played by larger groups of agents (e.g. buyers of board games). This forces us to generalize from particular agents to larger groups. We point to a list of games that were played by general¹ computer agents (based on Monte Carlo methods).

Würfelbingo [Wue 2007] Agents match the playing strength of average human players. Group trials by this author indicate indistinguishability from equally skilled humans.

EinStein würfelt nicht Agents match the playing strength of average human players. Monte Carlo agents played against humans on the game server www.inetplay.de from 2005 to 2008. User feedback indicates indistinguishability from human style.

Gelb Voran [Alt 2005] Agents match the playing strength of advanced human players. Personal experience by this author suggests indistinguishability in style from human play.

Seasons of the Sun [Alt 2007] Agents match the playing strength of advanced human players. Reports from various players indicate indistinguishability in style from human play.

¹Not tuned to the specific properties of a game.

8 Relevance of *Chanciness* for Human Agents

Hex [Hei 1942] Agents play considerably below top human levels. They displayed a tendency to play obvious losing moves that only prolong the match (“chains of resignation” [Gue 2008]). This assessment only refers to a plain UCT algorithm. Various heuristics have been successfully combined with UCT [Hay 2009].

Big Balls [AL 2007] In group trials by the author, agents played below average human levels. This allows distinguishing playing style for a large group of humans.

Based on these examples, approximating the playing style of casual game players with general game playing agents of sufficient strength appears plausible to us.

Therefore, we think that the concept of *chanciness* is useful in *computer-aided game inventing*. Althöfer lists several game features that are amenable to automatic evaluation: game length (average and maximum), drawing quota, balancedness, and deepening positivity (the impact of search depth on playing strength) [Alt 2003]. Browne adds further features such as rule complexity, convergence, and drama ² [Bro 2008]. We add *chanciness* to this list of features for automatic game evaluation.

²*Convergence* measures the trend in the number of (good) move choices as a match progresses. *Drama* measures how well agents can recover from self-perceived disadvantageous positions.

9 Conclusions and Open Problems

We have defined the *chanciness* of a non-deterministic ¹ perfect-information zero-sum game in extensive form that is played between two specific agents. This definition is based on the *relative influence of chance* for specific matches of that game by those same agents. We have approximated those *chanciness*-values for exemplary games played by computer agents.

The *chanciness* of a game G for some agents X, Y is a real number from the interval $[0, 1]$. It is close to 1 if the influence of chance on the outcome of the game is high. In this case, G played by X and Y is a game of chance. The *chanciness* is low or close to zero for games where chance has little influence on the game result. In that case, G played by X and Y is a game of skill. The *relative influence of chance* can be used to analyze specific matches of a game. If the *relative influence of chance* is low, the match result should be attributed to the skill difference of the agents or to the inherent unfairness of game G .

A necessary condition for the calculation of *chanciness* is the (frequent) repeatability of games between the agents. This does not exclude the use of learning agents as long as those agents can be reset to a previous internal state. As human agents can neither be reset nor do they tolerate many repeated games, *chanciness* can not be computed directly for human agents.

For computer-aided game development, the existence of repeatable automatic players can be assumed. The calculation of *chanciness* aids in the process of game inventing when adjusting the role of chance in a game. If, for example a game inventor wishes to reduce the influence of chance in one of his invented games, he can compute the *chanciness* for various rule variations and some set of agents. He can then discard those rule variations which do not lower the *chanciness* value.

Outside of computer-aided game development, we see *chanciness* and the *relative influence of chance* as a tool for game appreciation and as a stepping stone to mathematically founded game legislation.

An open problem of our method is the resource requirement for approximating *relative positional values*. The Algorithm described in Appendix B might be further improved through the use of hash tables. Measuring the variance during approximation would allow us to cut off some simulations while maintaining overall accuracy. It would be interesting to compare our results to the *RS*-values of Borm and Dreef for a large number of games. Since our objections to *RS* were mostly on theoretical grounds, the results could still agree in many cases.

¹Deterministic games are a special case of non-deterministic games [Now 1996]. Their *chanciness* is always zero.

9 Conclusions and Open Problems

Chanciness can be calculated for one-player games (puzzles) by modeling them as zero-sum games against a passive opponent. Additional research is undertaken to generalize the concept of *chanciness* to other game models:

Games with more than two players. These games require a vectorial representation of positional values and influence. Furthermore, these games may force an agent to select among moves which are of equal value for him but differ in their influence on his opponents. It may be appropriate to attribute such influence to chance.

Non-zero-sum games. These games can be mapped onto zero-sum games with an additional player and present all the challenges of games with more than two players.

Games in normal form (matrix games). Players typically employ randomized (mixed) strategies. This gives an additional source of chance which has to be taken into account.

Sequential games with imperfect information. Players have to make educated guesses regarding the game state they are in. The unpredictability of some game state aspects is from their perspective a form of randomness. This should be reflected in a generalization of *chanciness*.

As we continue to explore these questions, we expect that the concept of *relative positional values* will be useful as it is easily applied to a large class of games with any number of players.

Part II

Chanciness and Controllability of N-Player General-Sum Games

10 Introduction

Games are traditionally classified as games of chance, games of skill, or games with mixed characteristics [Par 2008]. As was already explained in Part I of this thesis, the desire for such a classification stems from various sources:

- Many states have laws to regulate gambling, in other words, playing games of chance with money at stake. It is then up to legislators and lawyers to decide which games are within the scope of such laws. In Germany, the gambling law “Glücksspielgesetz” defines games of chance [Sta 2007]. According to this law a game of chance is characterized by offering monetary (or equivalent) rewards and the extend of these rewards depends fully or mostly on chance. The term “mostly” is not elucidated. This distinction becomes even less clear when the accumulated outcomes of numerous game rounds (matches) are considered. When the outcome of a single match depends even a tiny bit on skill, the central limit theorem may be used to argue that the accumulated result of a sequence of matches is determined mostly by skill [Alo 2007]. The controversial nature of the skill / chance distinction is apparent from US-American court rulings regarding poker [BDS 2002]. Poker was ruled a *game of skill* in Pennsylvania [Cyp 2009] and a *game of chance* in North Carolina [Bur 2007]. This contradiction is arguably caused by the vague distinction between games of chance and skill.
- Naturally, the people who play games are interested in matters of chance and skill. Some players prefer games of chance over games of skill, while others have opposite tastes. In this context, a precise classification of games may not be important. However, the results of individual matches lends itself to analysis. Having won a game with a chance component, we may ask ourselves whether our victory was due to our skillful moves or simply caused by a string of lucky dice rolls [EL 2009]. The answer to this question is central to the feeling of accomplishment of the players and thus to the enjoyment of the game.
- A third group of people who concern themselves with the balance of chance and skill in games are game inventors. When designing a game for a particular audience, the inventor may wish to realize a particular level of chance and skill. Unfortunately, it is hard to measure the effects that subtle rule variations have on this balance. In the emerging field of *computer-aided game inventing* [Alt 2003, Bro 2008] computer simulations are used to measure game attributes. This reduces the need of a game inventor to find human play testers. Measuring the extent in which a game is influenced by chance or skill then requires a mathematical clarification of those concepts.

In Part I of this thesis we looked at the problem of rating two-player zero-sum games for their dependency on chance and skill. In this part, we expand our scope to include games with an arbitrary number of players and with generalized scoring rules. In the course of our analysis we will be concerned with the following questions:

- What assumptions underly such a analytic rating of chance and skill?
- How can the different influences on a game be quantified?
- How can those influences be compared and interpreted?
- What are the practical limits to such an analysis?

We introduce the concepts of *chanciness* and *controllability* to describe the relative impact of chance and player actions. By applying these concepts to matches (single rounds of play)¹ as well as to games we obtain measures on a $[0,1]$ scale.

In Chapter 11 we discuss previous works on formalizing the extent of chance and skill in games. We introduce terms and concepts for our discussion in Chapter 12. Our concepts of *chanciness* and *controllability* rest on the analysis of individual matches. To what extent did player moves and random events influence the match result? We deal with the subject of *influence* in Chapter 13. Based on our notion of influence we define the concept of *chanciness* in Chapter 14 and the concept of *controllability* in Chapter 15. These definitions refer to the participating agents. In Chapter 16 we deal with the interpretation of games that have at the same time low *chanciness* and low *controllability*. In Chapter 17 we justify the dependency of our measures on the participating agents. Restrictions on resource use and restrictions on knowledge of the agents necessitate the approximation of our proposed measures. Chapter 18 deals with these issues. In Chapter 19 we analyze *chanciness* and *controllability* for two parameterized game models. In Chapter 20 we present approximations of *chanciness* and *controllability* for the game “EinStein würfelt nicht!” [Alt 2004a] and for variants of this game. We summarize our results in Chapter 21.

¹Our use of the term *match* should not be confused with a tournament. The appropriate term in German would be “Partie”. The terms *match* and *game* are defined in Chapter 12.

11 Previous Work

The scholarly analysis of games has been mostly focused on solution concepts for the class of two-player zero-sum games. The quantitative analysis of the chance/skill distinction has received comparatively low attention. In this chapter we discuss a selection of previous works that are closely related to our topic

11.1 Parlett 2008

A good introduction to the many facets of chance and skill in games is given by David Parlett [Par 2008]. He wants to contribute to the classification of games, game inventing and game appreciation. Parlett frames the chance/skill distinction as a question of uncertainty (or unpredictability). He then discusses various sources of uncertainty and lists skills that help to control uncertainty¹.

As major sources of uncertainty he lists randomization (with the prime example of dice rolling), incomplete information (e.g. hidden cards), imperfect opponent modelling and difficulties in foreseeing the future even when the mechanisms are known (chaotic systems).

He also mentions the role of subjective perception in evaluating the uncertainty/controllability of a game. Parlett notes that games could be rated on a continuous scale that measures the uncertainty of their outcomes. However, he does not attempt a quantitative analysis of uncertainty.

11.2 Browne 2008

Cameron Browne is leading the field in quantifying aesthetic game properties [Bro 2008]. His work focuses on players' perceptions in 2-player zero-sum games without chance. He defines and calculates 57 game properties such as *Balance*, *Drawishness*, and *Depth*. Two notable properties of his are *Uncertainty* [Bro 2008, 113] and *Control* [Bro 2008, 123]. *Uncertainty* quantifies the difficulty of the players in foreseeing the result of the game. His property *Control* describes the ability of one player to reduce the number of his opponent's viable moves.

Browne uses the game properties developed by him to automatically rate the quality of games and even direct automated evolutionary search in a subspace of combinatorial

¹Parlett uses the term “controllability” as an opposite to uncertainty. His uses differs from our concept of *controllability* as defined in Chapter 15. In Parlett's work “controllability” is highest if the outcome of the game is most certain while for us *controllability* is high if a player has a strong influence on his result.

games. Using a sample population of human game players he demonstrates a correlation between his mathematical definitions and human perception.

11.3 Borm, Dreef and Genugten

The question whether games should be legally classified as games of skill or games of luck has been directly addressed by Borm, Dreef, and Genugten (hereafter, collectively BDG) [BvdG 2001, DBvdG 2004, Dre 2005]. They developed a measure called *Relative Skill (RS)* to place zero-sum games on a $[0, 1]$ scale. Using this measure they give ratings for the games of Roulette (0.0004), Black Jack (0.006), poker variants (0.4) and various coin games (e.g. 0.33). Games that allow for skillful play receive a higher rating of *RS* than games in which chance dominates the result. Thus, *RS* expresses the opposite idea of *chanciness* and is closely linked to our question.

The measure of BDG is based on three types of playing agents: the *Optimal Player*, the *Fictive Player*, and the *Beginner*. The *Optimal Player* player plays according to the expectimax algorithm [SLB 2008, 127]. The *Fictive Player* is assumed to know beforehand the outcome of all chance events. Furthermore, the *Fictive Player* is assumed to possess complete information about the game state. For this player type all games reduce to deterministic games of perfect information and he can thus use the simpler minimax algorithm [SLB 2008, 125]. The *Beginner* is intended to mimic a human player who has a very limited grasp of the game. Dreef [Dre 2005] recommends² that the playing style of the *Beginner* is to be modelled in one of three ways depending on what is called “the structure” of the game.

- random move selection
- quantitative analysis of human beginners playing style
- using domain expert knowledge to describe how a beginner would play

Games with more than two players are reduced to multiple two-player cases by assuming that one player faces a coalition of all the others.

The measure *RS* is obtained by comparing expected game rewards of the three Agent models.

$$RS = \frac{\text{reward}_{\text{Optimal}} - \text{reward}_{\text{Beginner}}}{\text{reward}_{\text{Fictive}} - \text{reward}_{\text{Beginner}}} \quad (11.1)$$

RS can be understood as the relative advantage of the *Fictive Player* over the *Optimal Player* when each is faced with a *Beginner*. If beforehand knowledge of chance events gives little advantage while the skill difference between *Beginner* and *Optimal Player* is large, then the game will receive a high value of *Relative Skill*.

²In [BvdG 2001] the definition of the *Beginner* remains rudimentary: “It is our experience that agreement can be obtained as far as it concerns the type beginner”.

11 Previous Work

11.4 Erdmann

The characterization of chance for the case of two-player zero-sum games is treated by Erdmann [Erd 2009]. This work is reproduced verbatim in Part I of this thesis. In this section we highlight the key differences between Part I and Part II. Differences between this thesis and the work of other authors will be discussed in the following section.

- The class of games under consideration is expanded from two-player zero-sum games to general-sum games with an arbitrary (but fixed) number of players³. One consequence of this is the need for multiple values to evaluate a game state (one value for every player). Consequently, the influence of a move must also be described with multiple values (influence on every player's value).
- The *chanciness* measure from Part I is supplemented with the new measure *controllability*. This additional measure reflects the observation that some games have low *chanciness*, yet should not be regarded as games of skill (for some or all players).
- New methods of approximating *chanciness* (and *controllability*) are discussed.
- The *chanciness* of a game is discussed in reference to a population of agents, while Part I only discusses pairs of agents.

11.5 Differentiation from Previous Work

Our main goal is the quantitative analysis of games and individual matches in regards to chance and skill. In contrast to the multiple sources of uncertainty given by Parlett, we focus on explicit randomization by dice and similar means.

Compared to Browne we analyze less properties, yet we consider a hugely expanded class of games. Whereas Browne only discusses two-player zero-sum games, we look at games with an arbitrary number of players and without restrictions on the scoring function. The work of Browne required speedy computation of his game properties to be useful within a genetic algorithm. In contrast, we are more interested in the theoretical background of our game properties and do not emphasize speed.

The goals of BDG are closely related to our own. We wish to improve on some of the limitations exhibited by their measure of *Relative Skill (RS)*.

First of all, our work leads us to the conclusion that an assessment of chance and skill for a given game depends strongly on the particular behavior of the playing agents. The measure *relative skill* is based on only three agents and thus can only give an incomplete picture.

Secondly the measure *RS* depends on a notion of optimality for their *Optimal Player*. Optimal behavior is well defined for two-player zero-sum games. The Minimax-Algorithm as well as the Expectimax-Algorithm [Zer 1913, vN 1927, SLB 2008] conform to the concept

³See Chapter 12.

11.5 Differentiation from Previous Work

of *subgame perfect equilibrium* [SLB 2008, 121]. However, in the context of general-sum games the optimality of this concept is disputed⁴.

Thirdly, the concept of the *Fictive Player* allows the creation of games, that receive ratings of *RS* which contradict core intuitions about games of chance and skill. For an example consider the game of *Lottery Chess* in Appendix A).

Also, the work of BDG does not cover the subject of analyzing individual matches. Our work on the other hand bases game measures upon the analysis of individual matches. While game measures are the ultimate goal, we consider match measures worthy of analysis in their own right.

Another important difference between this thesis and the work of BDG lies in a different concept of skill. BDG base their measure of *Relative Skill* on the skill difference among players. In contrast, our related measure of *controllability* is based on the variability of play by individual players. Our measure is able to describe the changing perception of a player as he masters a game's strategy.

We take the dependence of *RS* on modelling the *Beginner* as a valuable pointer. Assumptions about the game playing agents are intrinsic to the judgement about the relative impact of chance and skill on game outcomes. With our proposed measures of *chanciness* and *controllability* we make these assumptions more explicit. We point out methods for analysing games where optimal play is yet unknown. Furthermore, we avoid the contradictions that can be constructed with the help of the *Fictive Player*.

⁴The *Centipede Game* [MP 1992] is often used to demonstrate that players who violate equilibrium play obtain higher payoffs.

12 Terms and Concepts

12.1 Game Model

In this chapter, we introduce the necessary terms and concepts for our discussion of *chance* and *controllability*. The starting point is the central term *game*. For our game model we use the extensive form with chance moves [SLB 2008, 117]. An n -player¹ game G is represented by a finite rooted tree as in Figure 12.1. Nodes of the tree correspond to game states. We will use the terms *node*, *state* and *position* synonymously. Edges of the tree correspond to state transitions which we will call moves. Edges are directed from a parent node to its child node, thus moves are always directed away from the root and towards the leafs.




Figure 12.1: An exemplary game with three players 1,2,3.

The leafs of the tree represent terminal game states and every leaf t is labeled with an n -tuple of real numbers $\text{Result}(t) = (r_1, r_2, \dots, r_n)$. Each tuple represents the rewards for the players when a match ends in that state. For a match ending in in state t with the result $\text{Result}(t) = (r_1, r_2, \dots, r_n)$, player i receives a reward of r_i . We will use the terms *reward*, *result* and *payoff* synonymously.

The inner nodes represent non-terminal states that require a move. These nodes are either marked as player nodes or as chance nodes. A player node is further labeled with an index i from $\{1 \dots n\}$ to represent a node in which a decision by player i is required.

¹The term n -player game is sometimes used to refer to a specific game that may be played by different numbers of players. We use the term to denote a game that is be played with a fixed number of n players ($n \in \mathbb{N}$).

12.1 Game Model

Chance nodes on the other hand represent game states in which the successor state is selected by a random event such as the roll of a dice or the drawing of cards from a randomized deck. Correspondingly, the edges leading to child nodes are either called *player moves* or *chance moves*. Chance moves are further labeled with their respective probabilities. For every chance node s we define $\mathcal{C}(s, s')$ as the probability that chance will move from s to the child of s' . Without loss of generality we require that $\mathcal{C}(s, s')$ is positive.

We will repeatedly employ the game from Figure 12.1 as an example and refer to it as G_{example} .

Using this game representation, we discuss the class of games with the following properties:

Non-deterministic: The game may contain chance nodes. We treat deterministic games as a special case of non-deterministic games.

Finite: The game tree is finite in breadth as well as depth.

Arbitrary number of players : The game has a fixed number of n players with $n \in \mathbb{N}$.

Perfect information: All players know the history of past moves and the exact game state they are in [SLB 2008].

Complete information: All players know the game's rules including the results (payoffs) of all players for every terminal state [SLB 2008].

Sequential moves: Players move in sequence. Games with this property are also known as *dynamic games*.

General-sum: There are no constraints on the sum of the payoffs for each terminal node. In contrast to *constant-sum games* where the payoff sums for all terminal nodes must sum to the same constant C . For $C = 0$ these games are called *zero-sum games*. General-sum games which do not meet the zero-sum constraint are also referred to as *non-zero-sum games*.

Our model of game play is based on the terms *agent*, *population*, *lineup*, and *match*. Matches for a game are played by a lineup of agents taken from a finite population Π . We will explain all these terms in the following sections.

12 Terms and Concepts

12.2 Agent Model

Our model of an agent is meant to represent humans as well as computer algorithms (bots). Since we are not interested in *how* our agents form their move decisions, we can use a very simple agent model. It is sufficient that the agents are somehow able to select their move from the finite set of available moves. To reflect our ignorance, we use a stochastic model. For every game G let S_G denote the set of all possible states of G . Let $S_G^i \subseteq S_G$ denote the set of all states in which player i has to move. For every game state $s \in S_G$ let $\text{Children}(s) = \{c_1, c_2, \dots, c_m\}$ denote the set of states that can be reached by one legal move. An agent A capable of playing G as player i is a function from $S_G^i \times S_G$ to the real numbers. $A(s, s') = p$ is the probability that agent A will move from state s to the child node s' . The probabilities given by A must meet the following constraint:

$$\left[\sum_{c \in \text{Children}(s)} A(s, c) \right] = 1, \quad A(s, c) \geq 0 \quad (12.1)$$

Deterministic agents are a special case of stochastic agents. For simplicity's sake, our agents only consider the current state s when selecting their move. Since our model is a proper tree, every node has a unique history and therefore knowledge about past decisions can be included in our agent model².

Figure 12.2 shows the move probabilities for three agents A_1, A_2, A_3 . In our example each agent A_i plays the role of player i . Our agents are defined as follows:

$$\begin{aligned} A_1 &:= \{((s_0, s_1), 0.25), ((s_0, s_2), 0), ((s_0, s_3), 0.75)\} \\ A_2 &:= \{((s_1, s_4), 0.8), ((s_1, s_5), 0.2), ((s_3, s_6), 0.8), ((s_3, s_7), 0.2)\} \\ A_3 &:= \{((s_7, s_{10}), 0.5), ((s_7, s_{11}), 0.5)\} \end{aligned}$$

Our agent model only supports a limited representation of opponent modelling. By our assumption of perfect information, agents know the history of moves during a match. This knowledge could be used to build a model of opponent behavior. However, the value of $A(s, s')$ is independent of the lineup in which A plays (more sophisticated opponent modelling could retain knowledge about other agents across matches).

We furthermore assume that our agents are interested in maximizing the expectation of their payoff. They are indifferent regarding the payoffs awarded to other players inasmuch as these are decoupled from their own rewards. This assumption only refers to the *intentions* of the agents and not to their capabilities!

Actually, human players sometimes appear to be interested in the results of other players. Humans may be jealous of other player's rewards and try to diminish them even to the point of hurting their own payoff. In contrast, humans may also seek cooperation in ways that are detrimental to their own payoffs. Such attitudes can be expressed in our

²Two chess positions with the same board and different histories would be represented by two different nodes.




Figure 12.2: The tree of Game G_{example} with move probabilities for three exemplary agents A_1, A_2, A_3 . Agent A_i performs the role of player i . Each node is labeled for reference.

model by changing the payoff structure of the game. Malevolence among the agents corresponds to a zero-sum payoff structure while the intent on cooperation can be expressed with non-zero-sum payoffs.

12.3 Population, Lineup and Match

We presume that there exists a set of agents Π . We call this set the *population*. From this set Π , n -tuples of agents are taken. We call such a tuple $L = (A_1, A_2, \dots, A_n)$ a *lineup*.

A match g of game G is given by a pair $g = (L, P)$. $L = (A_1, A_2, \dots, A_n)$ is a lineup and the agent A_i takes on the roles of player i . An agent who plays as player i is responsible for selecting the successor node for any nodes marked with index i . We use the term “player” to describe an agent who has been assigned a player role (an index from $\{1, 2, \dots, n\}$). $P = (s_0, s_1, \dots, s_k)$ is a sequence of nodes. s_0 is the root node of game G , s_k is a leaf node and for every i , node s_{i+1} is a child of node s_i . The sequence P thus forms a path of nodes (game states) through the tree that leads from the root to a leaf node.

A match g of Game G always starts in the root node of the game tree. However, according to our game model, a subtree of G rooted in state s is also a game (a subgame of G). Accordingly, we can discuss matches that are started in any state s of G (matches in the subgame of G). We call such matches *partial matches*.

We further assume that there exists for every n -player game G a sampling process that yields lineups from the population Π . This unspecified process yields a probability mass function $\Psi : \Pi^n \rightarrow [0, 1]$ over the set of n -player lineups. $\Psi(L)$ then gives us the probability

12 Terms and Concepts

that a match with the lineup L will be played.

Using the terms defined above, we can describe matches for a wide selection of board games such as *Chess*, *Backgammon*, *Pachisi*, *Ludo* and *Roulette*. Our confinement to games of perfect information requires that all players may know the exact history of moves leading to a game state and can thus uniquely identify the game state they are in. This demand excludes many card games such as *Poker* or *Black Jack*. Some card games, such as *Concentration* (*Memory*, *Pairs*) can be modeled as games of perfect information when using the appropriate probability distributions to represent the uncovering of never-before-seen cards.

Care must be taken when modelling some of these games as a tree of finite depth. While all games have finite length in practice, some games such as backgammon allow for arbitrarily long matches³. One solution lies in modifying the rules slightly, similar to the *50 move rule* [FID 2005] in tournament *chess* (without this rule, *chess* would also allow matches of arbitrary length). For the practical application of our game measures, it is sufficient that matches terminate eventually (see Chapter 18).

12.4 Confined Perspective

The analysis of *influence* in Chapter 4 was simplified by the fact that game positions could be rated with a single number. For general-sum games with n players, game positions must be rated with n -tuples. Our method for reducing the complexity of the general case lies in temporarily confining our perspective to a single player index (player role).

We call this player i the *affected player*. The remaining $n - 1$ players will be called *other players*. Figure 12.3 shows our example game G_{example} from Figure 12.1 reduced to the perspective of player $i = 2$. When using the confined perspective, all agents are assumed to behave as in the real game. However, each positional value is reduced to the i th position of the n -tuple.

12.5 Probabilistic modelling of Human Agents

Our agent model assumes that agents can be described in terms of move probabilities. One might question the validity of this approach to describe human agents. The answer to this question is ultimately a matter of philosophy. It hinges on the chosen interpretation of probability.

One important interpretation is the *frequentist interpretation* [Haj 2009]. It defines probability as the relative frequency of an event for a large number of trials. This definition is difficult to apply to the act of deciding on a move. Such a decision depends on many elusive factors, some of which will shift over time as the human player is learning and adapting.

³However, it has been proven that backgammon ends with probability 1 [McC 1994].

12.5 Probabilistic modelling of Human Agents




Figure 12.3: The game G_{example} reduced to the perspective of one player 2.

Another influential interpretation of probability is called *Bayesian Probability* (also *Bayesianism*) [Haj 2009]. This interpretation defines probability as a “degree of belief” of rational agents and assigns probability to repeatable as well as to non-repeatable events. This interpretation of probability can be used to justify the application of our model to human agents.

13 The Concept of Influence

13.1 Introduction

The fact that state laws distinguish between games of chance and games of skill suggests that people have at least some criteria for accomplishing such a classification. We propose that this intuitive understanding is based on a concept of *influence* (or synonymously *importance*, *impact*). For a game (with perfect information) to qualify as a “game of chance”, it must involve some sort of randomizing device and that device must have an influence on the game results. Likewise to be called a “game of skill” a game must allow for players to make move decisions and those decisions must have an influence on game results. The problem of correctly classifying a game can then be restated as the problem of comparing the relative influences of chance and skill. For a scientifically meaningful comparison we must develop a rigorous definition of *influence*.

In Section 13.2 we look at the assumptions underlying the concept of *influence*. At the end of that section we declare the fundamental axiom crucial to this thesis. In Section 13.3 we give a quantitative definition of the influence of a move. We use the influence of individual moves in Section 13.4 to discuss the combined influences of moves on the result of a match.

13.2 Fundamental Assumptions

To understand influence we must distinguish between the game G and a match g of G that took place between specific players. When talking about the game G we actually discuss *possible* (or potential) influence. When we say G is a “game of chance”, we are saying, a typical match g of G is strongly influenced by chance. Likewise, we would call G a “game of skill” if it is influenced by the players. However, not every match plays out the same. If one player gets very lucky, he may win without effort. In another match, the influence of dice rolls may favor both players equally and the result is then decided by the skill of the players. Therefore, we will define *influence* for matches g of a game G . After we have understood how chance and skill affect matches, we generalize our findings to talk about the properties of G . The assumption behind this approach can thus be explicitly stated:

Assumption 1 *The properties of a game are determined by the properties of its matches.*

In the preceding paragraph we have used the phrase “influence of dice rolls”. We turn this manner of speaking into another assumption:

13.2 Fundamental Assumptions

Assumption 2 *The influence of chance on a match is determined by the influence of chance moves taken during that match.*

And accordingly:

Assumption 3 *The influence of the players on a match is determined by the influence of player moves taken during that match.*

In Section 15.4 we will discuss subtle distinctions between “player influence” and “skill”. In the conclusion of Part I (Chapter 9) we suggested that chance moves might not be the only source of chance influence. In Section 16.1 we will reexamine this suggestion and justify Assumption 2.

13.2.1 Influence Depends on the Game Tree

Based on our three assumptions from the preceding section we can already draw some conclusions regarding the classification of games. Moves by chance or skill can only be taken during a match if they appear in the game tree.

Figure 13.1 shows 5 small game trees. Game 13.1 a) does not contain any chance moves. By Assumption 2, chance cannot influence match results and therefore Game a) is not a game of chance. Using the same argument we arrive at the obvious conclusion that *Chess* and *Go* are not games of chance. Similarly, by Assumption 3 game 13.1b) cannot be a game of skill, since it contains no player moves.

What about the games 13.1c),d) and e)? These “games” contain neither chance nor player moves. If we accept that these trees constitute games, then by our assumptions they are neither games of chance nor of skill.




Figure 13.1: Small games. Games a)-d) are single-player games and game e) is a game of two players.

The common understanding of chance and skill as suggested by Parlett [Par 2008], puts games on a one-dimensional scale with complete chance (uncertainty) on one end and complete skill (controllability) on the other. By our game model and by our previous

13 The Concept of Influence

assumptions, this scale has no place for games 13.1c),d) and e). We must therefore abandon the idea of expression the nature of chance and skill with a single one-dimensional property. In Chapter 17 we will find that this issue is not confined to the (pathological) case of games without moves.

13.2.2 Influence Depends on the Players

The classification of games that allow only one type of moves (chance moves or player moves) is hardly interesting. If we wish to further our understanding of the nature chance and skill we must look at games that allow for both types of moves. However, there is quite a difference between games that allow a certain type of move and games which enforce the occurrence of these moves. As one important consequence of Assumptions 2 and 3, only moves that actually occur during a match can be considered influential.

The game in Figure 13.2 is meant as a thought experiment to illustrate this issue. We call this game *Chess-Or-Coin* (CC). The tree of CC is composed of two very different subtrees (subgames). In a match of CC, player 1 must decide whether play will continue according to the rules of *Chess* or whether the result will be decided by the flip of a coin.




Figure 13.2: The game *Chess-Or-Coin* (CC). The left subtree is meant to represent the complete tree of the game *Chess*.

For a player who prefers the left subtree, chance moves will never occur and CC will be a game of skill. For a different player who prefers the right subtree the game will be a game of chance. The nature of CC therefore depends on the behavior of the players.

While the game CC may appear rather contrived, nothing prevents the creation of games where different parts of the game tree are influenced to different degrees by chance and skill. Therefore, the issue of player behavior must generally be considered when discussing the nature of a game.

13.2.3 Influence Depends on the Values of Game States

By Assumptions 2 and 3, individual moves are considered influential. When we talk about the influence of moves (player moves or chance moves), we talk about their influence on

the match result. However, a single move rarely decides the game. Instead, moves lead to game states in which a player has improved or reduced chances of “winning” (increased or decreased expectations for their result). Following our player model, a move has a desirable influence for a given player if it increases the expected result for that player. This leads us to our *Axiom of Influence*:

Axiom 1 Fundamental Axiom of Influence

The Influence of a move from state s to state s' is the difference between the expected value of s and the expected value of s' .

The **Fundamental Axiom of Influence** forms the basis of this thesis¹. In the next section we will use Axiom 1 to quantify the influence of moves.

13.3 The Influence of a Move

By the **Fundamental Axiom of Influence** (Axiom 1) the influence of a move from state s to state s' depends on the expected values of both states. However, Axiom 1 does not specify how this expected value is to be calculated.

13.3.1 The Prescriptive Approach to the Evaluation of Game States

The evaluation of game states is a central problem of game research [BCG 1982]. The motivation for this task usually arises from the question of *how* to maximize one’s result against opponents who attempt the same. For the class of perfect-information games with sequential moves, the concept of *subgame perfect equilibrium* [SLB 2008, 121] (SPE) is the most prominent answer on *how* to play a game optimally. This concept allows the evaluation of game states with an algorithm called *backward induction* [SLB 2008, 124]². The positional values assigned by this algorithm are known as *minimax values* (for deterministic games) and *expectimax values* (for nondeterministic games).

We call this approach to the evaluation of game states the *prescriptive approach*, because it prescribes how we *should* play if we want to maximize our results. Under the *prescriptive approach*, the value of a game state is defined as the expected result for players who all play perfectly as prescribed.

13.3.2 The Descriptive Approach to the Evaluation of Game States

In contrast to the *prescriptive approach*, this thesis defines the *descriptive approach* to the evaluation of game states. In the *descriptive approach*, we presume there exists a description of how the players behave. The expected value of a state is then based on this behavioral description. The behavior of optimal agents (according to the SPE) is just one special case of agent behavior. For this reason, the *descriptive approach* is more general than the *prescriptive approach*.

¹ Axiom 1 was implicitly used in Part I of this thesis. .

²This algorithm is defined for two-player games in Section 4.2.

13 The Concept of Influence

In the real world, a multitude of agents with different behaviors can be observed. These real world agents rarely play optimal. Even though agents seek to maximize their result, their behavior often differs from the optimal behavior for at least two reasons:

- Agents lack the resources to determine optimal choices. Human agents as well as computer agents are generally unable to analyze the complete tree of a game. Instead they rely on heuristic shortcuts of different sophistication (counting the remaining pieces of a chess position would be a simple example).

The statement that real-world agents fail to play optimally should need no further justification. However, we wish to mention the amusing example of humans playing against chicken in the game of *Tic-Tac-Toe*. In this game optimal play guarantees a draw. Nevertheless, the trained chickens are said to have lost nine matches out of 270,000 matches played, according to Kinney [Kin 2003].

- Agents know that other agents sometimes do not play optimally. This gives them a reason to deviate from optimal play in hopes of exploiting their opponent's weakness. An often mentioned example is the *Centipede game*. Figure 13.3 shows the tree of the *Four Stage Centipede Game*. In this game, the expected result would be (1,0) (for player 1 and player 2 respectively) according to the *prescriptive approach*. However, according to studies by McKelvey [MP 1992] humans reach even higher expected results by deviating from SPE-behavior.




Figure 13.3: Game tree of the *Four Stage Centipede Game*.

Since the *descriptive approach* gives a more accurate picture of the actual expected rewards of real world agents, we use this approach to define the values of game states and the influence of moves ³. As we have pointed out in Section 13.2.2, the influence of chance and skill on a game depends on the behavior of the players anyway. Therefore, the reliance on a behavioral description does not increase the demands upon our agent model.

³Section 18.2 discusses the consequences of using another approach.

13.3.3 Relative Positional Values (RPV)

In this Section we define the *relative positional value* (RPV) of a game state according to the *descriptive approach*. The term *relative* emphasizes that the positional value depends on the behavior of the agents.

For the n -player game G , the values of all terminal nodes (the leafs of the game tree) are well defined. According to our game model (Chapter 12) the value $\text{Result}(t)$ of such a terminal game state t is given by an n -tuple of numbers. Position i in the tuple denotes the result of player i , should a match g terminate in state t . Accordingly, we must differentiate the values of any game state s for the perspectives of the different players.

The result of player i for a match or a *partial match*⁴ that is started in state s and played by the lineup L can be seen as the realization of a random variable R . When s is a terminal state the realization of R is the constant value of $\text{Result}(s)_i$. Otherwise the value of R depends on the stochastic behavior of the agents in L and on the stochastic behavior of chance. We define the RPV as the expected value of R . The RPV $\tilde{V}_i(s, L)$ is defined for a game state s from the perspective of player i in a match played by the lineup of agents L :

Definition 13.1

$$\tilde{V}_i(s, L) := \mathbb{E}(R)$$

Figure 13.4 shows our example game G_{example} from the perspective of *affected player* 2. (This tree was already used in Chapter 12). The lineup $L = (A_1, A_2, A_3)$ consists of the example agents defined in Section 12.2. Using this lineup the RPV $\tilde{V}_2(s, L)$ are given for every node s in the tree.

The value state of $\tilde{V}_i(s, L)$ can be computed based on the values of the children of s . Since the values of terminal nodes (leafs) are well defined, we can compute the value of any state s recursively by starting at the leafs and backing up towards the root. This gives us the following recursive definition of $\tilde{V}_i(s, L)$:

$$\tilde{V}_i(s, L) := \begin{cases} \text{Result}(s)_i & \text{if } s \text{ is a leaf} \\ \sum_{c \in \text{Children}(s)} A_k(s, c) \cdot \tilde{V}_i(c, L) & \text{if player } k \text{ moves in } s \\ \sum_{c \in \text{Children}(s)} \mathcal{C}(s, c) \cdot \tilde{V}_i(c, L) & \text{if chance moves in } s \end{cases} \quad (13.2)$$

Here $\text{Result}(s)_i$ is the match result for player i in the terminal state s .

The RPV defined in Part I (Definition 4.2) are a special case of the RPV defined in this Section. For the special case of two-player zero-sum games, treated in Part I, there is no need to differentiate among the player perspectives. Due to the zero-sum restriction on match results, the equivalence $\tilde{V}_1(s, L) = -\tilde{V}_2(s, L)$ holds for any game state s .

⁴See Section 12.3.

13 The Concept of Influence




Figure 13.4: The tree of Game G_{example} reduced to the perspective of player 2 with move probabilities for a specific lineup L . Each node s is annotated with the RPV $\tilde{V}_2(s, L)$.

13.3.4 Definition of the Influence of a Move

In the previous Section we introduced the concept of RPV to evaluate game states. Based on Definition 13.1 we will define the *influence* of a move $m = (s, s')$ from state s to state s' .

Definition 13.3 The influence I of move $m = (s, s')$ from the perspective of player i in a match with the lineup L is defined as:

$$I(m, L, i) := \tilde{V}_i(s', L) - \tilde{V}_i(s, L) \quad (13.4)$$

Note that there is no distinction between player moves and chance moves.

This definition of influence is very similar to Definition 4.5 in Part I. Instead of only two agents we now have the lineup L , an n -tuple of agents. We must also specify a player perspective i . In contrast to the two-player zero-sum situation in Part I, the positional values for different players are now generally independent.

One important consequences of using RPV in our definition of influence is this:

Lemma 13.5 “Deterministic moves” always have influence 0.

We call a move $m = (s, s')$ “deterministic move” for agent A if state s is always followed by the specific state s' . In other words $A(s, s') = 1$.

Proof: Let s be a game state with player k to move and $d \in \text{Children}(s)$. Let $L = (A_1, A_2, \dots, A_n)$ be a lineup such that agent $A_k(s, d) = 1$ (move (s, d) is a deterministic move for agent A_k).

13.4 Influences in a Match

Using Definition 13.2, $\tilde{V}_i(s, L)$ (for an arbitrary perspective i) is as a weighted sum of the values of $\text{Children}(s)$.

$$\begin{aligned}\tilde{V}_i(s, L) &= \sum_{c \in \text{Children}(s)} A_k(s, c) \cdot \tilde{V}_i(c, L) \\ &= 1 \cdot \tilde{V}_i(d, L), \text{ hence} \\ I((s, d), L, i) &= \tilde{V}_i(d, L) - \tilde{V}_i(s, L) = 0\end{aligned}$$

□

More generally speaking, the absolute influence of a move $m = (s, d)$ decreases when the probability of being taken p increases. This can be shown by a slight variation of the above proof:

Let $A_k(s, d) < 1$ be the probability that agent A_k takes the move $m = (s, d)$. Let A'_k be a second agent who differs from A_k only in the move probabilities for s , such that $A'_k(s, d) = q, q > p$. We consider the lineup L that contains agent A_k and compare it to the lineup L' which contains the agent A'_k instead of agent A_k .

We note that $\tilde{V}_i(d, L) = \tilde{V}_i(d, L')$ since the agents differ only in their probabilities for s . The RPV of s , $\tilde{V}_i(s, L)$ is a convex sum of the value of its children. An increased probability for move (s, d) means an increased weight for the child d . Thus, the value of $\tilde{V}_i(s, L')$ lies between $\tilde{V}_i(d, L)$ and $\tilde{V}_i(s, L)$. Influence is the change in value between s and d thus, the influence for the move by A'_k is lower. □

A player who understands a game position well knows which move he should take in that position. Consequently, such a move will have a high probability of being taken and the influence of that move will be low. Thus, the influence of a player move is a good measure for the interestingness (difficulty) of deciding on that move.

Important: By tying our evaluation function \tilde{V} to the participating agents, we allow that a move may have a different influence, depending on the agent who took that move. While this may seem counter-intuitive, it later allows us to explain why games change their character once we master them. Once all move decisions in a game are obvious, we might as well just roll a die instead of playing the game. If the game did not contain random elements we would probably abandon it at this point.

13.4 Influences in a Match

By Assumption 1 we will derive the properties of a game G from the properties of matches. In this section we lay the groundwork for the analysis of a match.

Let $g = (L, P)$ be a match of an the n -player game G . It was played by the lineup $L = (A_1, A_2, \dots, A_n)$. During the match g the game states $P = (s_0, s_1, \dots, s_k)$ were visited in

13 The Concept of Influence

sequence. We will analyze g from the perspective of player i .

By our model assumptions from Chapter 12, player i only cares about his own result. This allows us to focus our discussion on the influence of moves from the perspective of player i . We call this player the *affected player*.

13.4.1 The Combined Influence of Moves

We have previously defined the influence $I(m, L, i)$ for a single move m from the perspective of player i in a match with the lineup L . Since a match usually consists of more than one move, we must be able to reason about the combined influence of those moves.

In this section we introduce the concept of a *move-sequence* and show that the influence I is additive when applied to *move-sequences*.

Definition 13.6 For a match $g = (L, P)$ with the sequence of states $P = (s_0, s_1, \dots, s_k)$ For $0 \leq x < y \leq k$ we call the pair (s_x, s_y) the *move-sequence* from state s_x to state s_y .

The move sequence (s_x, s_y) represents the combination of the moves $((s_x, s_{x+1}), (s_{x+1}, s_{x+2}), \dots, (s_{y-1}, s_y))$ by sequential execution.

Definition 13.7 For two move-sequences of the following form, we define the addition $(+)$

$$(s_x, s_y) + (s_y, s_z) := (s_x, s_z) \quad (13.8)$$

Definition 13.9 We define the influence of a move-sequence from the perspective of player i in a match with the lineup L :

$$I((s_x, s_y), L, i) := \tilde{V}_i(s_y, L) - \tilde{V}_i(s_x, L) \quad (13.10)$$

This definition of influence directly mirrors Definition 13.3.

Lemma 13.11 The influence I applied to move-sequences is additive.

Proof:

$$\begin{aligned} I((s_x, s_y) + (s_y, s_z), L, i) &= I((s_x, s_z), L, i) \\ &= \tilde{V}_i((s_z), L) - \tilde{V}_i((s_x), L) \\ &= \tilde{V}_i((s_z), L) - \tilde{V}_i((s_y), L) + \tilde{V}_i((s_y), L) - \tilde{V}_i((s_x), L) \\ &= I((s_x, s_y), L, i) + I((s_y, s_z), L, i) \end{aligned}$$

□

In the following section we use the additivity of *move-sequences* to interpret the result of match g .

13.4.2 The Match Equation

Let match $g = (L, P)$, for the n -player game G be given by a lineup of agents $L = (A_1, A_2, \dots, A_n)$ and a sequence of game states $P = (s_0, s_1, \dots, s_k)$. Let $m_j = (s_j, s_{j+1})$ denote the j th move in P .

The additivity of influence for *move-sequences* allows us to express the result of match g as the following sum:

Lemma 13.12

$$\tilde{V}_i(s_k, L) = \tilde{V}_i(s_0, L) + \sum_{j=0}^{k-1} I(m_j, L, i) \quad (13.13)$$

We call this equation the *match equation*.

Proof: Every move m_j is also *move-sequence*. Thus we can apply Lemma 13.11.

$$\begin{aligned} & \tilde{V}_i(s_0, L) + \sum_{j=0}^{k-1} I(m_j, L, i) \\ = & \tilde{V}_i(s_0, L) + I\left(\left[\sum_{j=0}^{k-1} m_j\right], L, i\right) \\ = & \tilde{V}_i(s_0, L) + I((s_0, s_k), L, i) \\ = & \tilde{V}_i(s_0, L) + \tilde{V}_i(s_k, L, i) - \tilde{V}_i(s_0, L, i) \\ = & \tilde{V}_i(s_k, L, i) \end{aligned}$$

□

This *match equation* “explains” how the combined influence of all moves leads to the result of match g .

The term $\tilde{V}_i(s_0, L)$ is special because it is not the influence of a move but rather the value of the starting position. We call $\tilde{V}_i(s_0, L)$ the *match expectation* for the perspective of player i (ME_i). Since all matches of a game start in the same state s_0 , the value of ME_i depends only on the lineup L .

The *match expectation* can well be the dominant term of the *match equation*. We will therefore include it in our comparison of the different influences in the following chapter. Later, in chapter 16, we will discuss the interpretation of ME_i in detail.

13.4.3 Sources of Influence (SI)

In the previous section we have defined the influence $I(m, L, i)$ of a move m that was taken during a match with lineup L (Definition 13.3). This influence relates to the perspective of player i as it is based on the RPV from the perspective of player i . We will use this definition to quantify the influence of chance and the influences of the players.

We intend to arrive at measures that describe the influence of chance in comparison to other sources of influence (SI). To conduct such a comparison we must identify sources of

13 The Concept of Influence

influence and assign each element of the *match equation* 13.13 (Section 13.4.2) to one of these sources.

For a game with n players we propose to the following $n + 2$ distinct sources of influence:

- chance
- player 1
- player 2
- ...
- player n
- *match expectation*

The source *chance* is rather obvious. To this source we assign the combined influence all moves that are labeled as chance moves. (We have yet to define what we mean by combined influence). In line with Assumption 2 we consider chance moves to be the only source of chance influence. This implies that player moves are *not* influenced by chance. We will discuss this assumption later in Section 16.1 and again in Section 21.2.4.

To every player we assign the influence of all moves taken by that player (in accordance with Assumption 3). The separation of player influences into n distinct *SI* differs from the approach taken in Part I. This separation is a necessary prerequisite for the concept of *controllability*, discussed in Chapter 15.

Since the *match expectation* describes a situation in which the moves of a match have not yet taken place, it is independent of the move influences. Therefore it does not contribute to the influence of chance nor to the influence of the players. That we have named the *match expectation* (ME_i) among the sources of influence may appear confusing. After all, we have defined *influence* as a property of moves. Our choice can be justified by interpreting ME_i as the influence of “meta-moves”. Namely, selecting a game G and sampling a lineup L . This concept is elaborated in Chapter 16.

Besides, the *match expectation* is simply one of the contributing terms in the *match equation* and must thus be accounted for. “Influence” in this case, can be understood in the colloquial sense of word.

13.4.4 Combined Influence of an *SI*

In the previous section we have defined distinct sources of influence (*SI*). We must now discuss the combined influence that was exercised by each *SI* over the course of match g . Doing this requires additional assumptions.

In Section 13.4.1 we have shown the additivity of the influence I for *move-sequences*. If we wish to talk about the combined influence of all the chance (or player) moves in a game we must extend our definition of influences to arbitrary sets of moves. We propose that the following assumption is compatible with an intuitive understanding of influence:

Assumption 4 *The combined influence of any set of moves is the sum of the individual move influences.*

This assumption allows moves with positive influence to cancel out with moves that have negative influence even when not executed directly after each other. This behavior assures consistency if we consider that a new game could be created by chaining several other games together. Whether we play these games individually or as a whole, we would be interested in the bottom-line influence of the *SI*.

Under Assumption 4 we can quantify the combined influence for each *SI* defined in the previous section. These numbers are specific to a match $g = (L, P)$ with $P = (s_0, s_1, \dots, s_k)$

Let M_i denote the set of all moves by player i and M_C denote the set of chance moves in match g .

$$\begin{aligned} M_C(g) &:= \{(s_j, s_{j+1}) \mid \text{chance is to move in state } s_j\} \\ M_i(g) &:= \{(s_j, s_{j+1}) \mid \text{player } i \text{ is to move in state } s_j\} \end{aligned}$$

Using the sets of moves defined above, we define the influence for the *SI*:

$$IC_i(g) := \sum_{m \in M_C(g)} I(m, L, i) \quad (13.14)$$

$$IP_{j,i}(g) := \sum_{m \in M_j(g)} I(m, L, i) \quad (13.15)$$

$$ME_i(g) := \tilde{V}_i(s_0, L) \quad (13.16)$$

$IC_i(g)$ denotes the influence of chance on the *affected player* i , $IP_{j,i}(g)$ denotes the influence of player j on the *affected player* i and $ME_i(g)$ denotes the *match expectation* of the *affected player* i .

We will use the terms IC, IP_j, ME as shorthands for $IC_i(g), IP_{j,i}(g), ME_i(g)$ whenever the match g and the *affected player* i are obvious from the context (e.g. in single-player games).

We can now rephrase the *match equation* (Equation 13.13 from Section 13.4.2) with the terms from above.

$$\tilde{V}_i(s_n, L) = ME + IC + IP_1 + IP_2 + \dots + IP_n \quad (13.17)$$

Whereas Equation 13.13 “explained” how the influence of individual moves leads to the match result, we can now give an explanation in terms of *SI*.

Figure 13.5 shows an exemplary match $h = (L, (s_0, s_1, s_2, s_3))$ for the game G_{example} from the perspective of the *affected player* 2. Game G_{example} and the lineup L were already used

13 The Concept of Influence




Figure 13.5: An exemplary match h . The path through the game tree is highlighted and every node is marked with the RPV for the lineup L from the perspective of player 2.

in Figure 13.4. The influences of the SI are as follows:

$$M_C(h) = \{(s_2, s_3)\}$$

$$M_1(h) = \{(s_0, s_1)\}$$

$$M_2(h) = \{(s_1, s_2)\}$$

$$M_3(h) = \{\}$$

$$ME_2(h) = 1.1$$

$$IC_2(h) = -0.5$$

$$IP_{1,2}(h) = 0.1$$

$$IP_{2,2}(h) = 0.3$$

$$IP_{3,2}(h) = 0$$

We observe that Equation 13.17 holds:

$$\tilde{V}_1(s_3, L) = 1 = 1.1 - 0.5 + 0.1 + 0.3 + 0 \quad (13.18)$$

Given the influence value for each SI we can already discuss some interesting match properties: The leading sign of IC shows whether a player was advantaged or disadvantaged by chance moves. If IC is a positive number, than the *affected player* cannot claim that his result is due to “bad luck”.

The leading sign of ME points to the existence of “unfairness” or skill disparity among the agents or both (the difficulty of distinguishing between these cases is discussed in Section 16.2). The term IP_i (short for $IP_{i,i}$) denotes the influence of the *affected player* i upon his own result. It reveals whether the *affected player* played below or above his usual strength. To reach further conclusions we must proceed to compare these values.

13.4.5 Comparing different SI

The influence of an *SI* can be a negative or a positive number. However, when comparing influence values of different *SI*, the leading sign is irrelevant. Consider match h from Figure 13.5. The influence of player 1 is positive ($IP_1 = 0.1$) and the influence of chance is negative ($IC = -0.5$). A direct comparison of influence values is misleading. The fact that $IP_1 > IC$ does not imply that chance had less “influence” on the match result than player 1 (“influence” in this case is a synonym for importance). On the contrary, chance produced a much greater change in value than player 1 did.

We turn this observation into an assumption:

Assumption 5 *SI must be compared according to their absolute influence.*

We will therefore assemble our measures from the absolute values $|ME|$, $|IC|$ and $|IP_j|$.

We wish to note that the last two assumptions in this section may not be psychologically accurate. Regarding Assumption 4 it may be that people do not weigh all moves equally when judging the combined influence of moves. One might argue that moves with higher influence are remembered more and therefore dominate the perception of the combined value. As a thought experiment, consider a match where IC arose as the following sum: $IC = 1 + 1 + 1 + 1 + 1 + 1 - 5 = 1$. It is easy to image the frustration of the *affected player* over his misfortune and the end of the match. Yet, when looking at the complete match, this player must admit that he was lucky after all.

Regarding Assumption 5 it may also be that people judge the influence (or rather importance) of an *SI* differently depending on the sign. One could argue for a bias to understate the importance of favorable chance and to overstate the importance of favorable moves taken by oneself. Such a perceptual bias is understandable as it makes the *affected player* appear more skilled. However, the goal of this thesis is a measure that could be agreed on by rational people after a discussion. Such discussion should uncover perceptual biases and lead to assumptions as taken above.

13 The Concept of Influence

13.5 Summary

In this chapter we have discussed the general concept of *influence*. Based on three fundamental assumptions we have shown that the influence of chance and skill on a game depends on the game tree, the behavior of the players, and a concept for evaluating game states.

We have stated our **Fundamental Axiom of Influence** and used it to define the influence I of a move. The influence I depends on the match in which the move occurred. More specifically, it depends on the agents who played in that match. For a game with multiple players, the influence I also depends on the perspective of the player who is affected by the move. Based on the influence of a move, we have analyzed the different influences on a match result and introduced the concept of a source of influence (SI). In the next chapters we will define the measures of *chanciness* and *controllability* based on a comparison of these SI .

14 Chanciness

14.1 Introduction

In the previous Chapter we have quantified the influence of different sources of influence (*SI*). In this chapter we define the measure of *chanciness* based on comparison of these influences. We distinguish between the *chanciness*-values of different player perspectives. We also distinguish between the *chanciness* of a single match and the *chanciness* of a game.

14.2 The *Chanciness* of a Match from the Perspective of Player i

In the previous chapter we have laid the groundwork for comparing the various influences on a match result. We have defined the sources of influences (*SI*) and assigned an influence value to each one of them. In accordance with our Assumption 5, we will now compare the absolute influence of the *SI*.

To shorten our forthcoming definitions we first introduce the term *total absolute influence*. $TI_i(g)$.

Definition 14.1

$$TI_i(g) := |ME_i(g)| + |IC_i(g)| + |IP_{1,i}(g)| + |IP_{2,i}(g)| + \dots + |IP_{n,i}(g)|$$

Our first measure relates the absolute influence of chance *IC* to the influence of all other sources. We thus define the *chanciness of a match*, $cha_i(g)$ for a match g of an n -player game G evaluated from the perspective of player i :

Definition 14.2

$$cha_i(g) := \frac{|IC_i(g)|}{TI_i(g)}$$

The measure cha reaches its minimum of zero iff the absolute influence of chance $|IC|$ reaches zero. It reaches its maximum of 1 iff *IC* is nonzero and all other *SI* have influence zero. The *chanciness of a match* is undefined for a *total absolute influence* TI of zero. For this to happen, all *SI* must have had influence zero. In this case, a comparison is arguably meaningless.

14 Chanciness

For the example match h from the previous chapter (Figure 13.5) we calculate the *chanciness of a match* for the *affected player 2*:

$$\begin{aligned}
 \text{ME}_2(h) &= 1.1 \\
 \text{IC}_2(h) &= -0.5 \\
 \text{IP}_{1,2}(h) &= 0.1 \\
 \text{IP}_{2,2}(h) &= 0.3 \\
 \text{IP}_{3,2}(h) &= 0 \\
 \text{TI}_2(h) &= 1.1 + 0.5 + 0.1 + 0.3 + 0 = 2 \\
 \text{cha}_2(h) &= 0.25
 \end{aligned}$$

For the perspectives of the other players we would generally receive different values. For the example match h however, all perspectives have the same *chanciness* of 0.25 (calculations in Appendix E).

Our Definition of cha is very similar to the definition of the *relative influence of chance* (Definition 5.3) from Part I. The only difference is that all players are considered as separate sources of influence in this part. For an n -player game G we will generally have n different values of cha .

However, for the special case of two-player zero-sum games it can be shown that $\text{cha}_1(g) \equiv \text{cha}_2(g)$ for any match g (The proof is given in Appendix G). This justifies the approach of Part I where a single value of *chanciness* was used to describe such games.

14.3 The *Chanciness* of a Game from the Perspective of Player i

In this section we define the measure *chanciness of a game*. For this measure we rely on Assumption 1: The properties of a game derive from the properties of matches. We must therefore define which matches are to be considered and furthermore how the measures of these matches are to be combined.

14.3.1 Sampling Matches

A game tree that contains chance moves may allow for matches that completely avoid these chance moves. The game *Chess-Or-Coin* was introduced in Section 13.2.2 to illustrate this. Whether matches with chance moves or matches without chance moves occur, depends on the behaviour of the players. Thus, the *chanciness* for such games also depends on the (stochastic) behaviour of the players. We make this dependency explicit by basing our sample of matches on a population of agents Π and a probability distribution Ψ over the lineups that can be sampled from Π . Both are part of our model given in Chapter 12.

As stated in Section 12.3 we assume that for the game G under discussion there exists a process that samples lineups according to some probability mass function Ψ . Since our

agents are also described stochastically, we can regard a match $g = (L, P)$ as the result of a stochastic process. Consequently the influences of the *SI* and the *chanciness of a match* $\text{cha}_i(g)$ are random variables.

14.3.2 Combining the Properties of Possible Matches

The probability for a given match $g = (L, P)$ to be played is well defined as the probability of the path P to be taken under the condition that the lineup L was sampled.

An obvious way to account for the many different matches that could be played among the population Π is to aggregate the values of cha with the expectation operator \mathbb{E} . Since cha is a nonlinear function of random variables it makes an important difference whether we apply \mathbb{E} at the outside of cha or at the inside of cha at its arguments. The failure to consider these differences goes by the name *Fallacy of the Averages* [Wag 1969].

The *chanciness of a match* $\text{cha}_i(g)$, itself a random variable, is a function of the random variables $\text{ME}, \text{IC}, \text{IP}_1, \text{IP}_2, \dots, \text{IP}_n$. To these variables, the non-linear operations of division and of taking the absolute value are applied. This gives us three candidate formulas for the *chanciness of a game*.

$$\text{Candidate}_A(G, \Pi, \Psi) := \frac{|\mathbb{E}(\text{IC})|}{\left[|\mathbb{E}(\text{IC})| + |\mathbb{E}(\text{IP}_1)| + |\mathbb{E}(\text{IP}_2)| + \dots + |\mathbb{E}(\text{IP}_n)| + |\mathbb{E}(\text{ME})| \right]} \quad (14.3)$$

$$\text{Candidate}_B(G, \Pi, \Psi) := \frac{\mathbb{E}(|\text{IC}|)}{\mathbb{E}(\text{TI})} \quad (14.4)$$

$$\text{Candidate}_C(G, \Pi, \Psi) := \mathbb{E} \left[\frac{|\text{IC}|}{\text{TI}} \right] \quad (14.5)$$

In the definition of Candidate_A , the expectation operator is applied directly to the influence values of the *SI*. This candidate is of little use, since the numerator term $|\mathbb{E}(\text{IC})|$ is zero for any game G regardless of the population (Proof in Appendix D). Thus Candidate_A is also constantly zero. Obviously, a measure that is constantly zero cannot be used to differentiate games.

The problem with Candidate_C already appeared in Part I (Chapter 6). We demonstrate the undesirable behavior of this measure with the exemplary game H shown in Figure 14.1.

This game leads to the following situation (regardless of the population Π and sampling

14 Chanciness




Figure 14.1: A single-player game in which some matches are strongly influenced by chance. In this game, the RPV of all nodes are independent of the agent.

probabilities Ψ since there are no player nodes):

$$\begin{aligned}
 \text{ME} &= 1, \quad \text{IP} = 0 \quad \text{for all matches} \\
 |\text{IC}| &= 0 \quad \text{for 90\% of the matches} \\
 |\text{IC}| &= 100 \quad \text{for 10\% of the matches} \\
 \text{cha}_1(g) &= \frac{0}{1+0+0} = 0 \quad \text{for 90\% of the matches} \\
 \text{cha}_1(g) &= \frac{100}{1+0+100} \approx 1 \quad \text{for 10\% of the matches} \\
 \text{Candidate}_C(H, \Pi, \Psi) &= 0.9 \cdot 0 + 0.1 \frac{100}{101} \approx 0.1
 \end{aligned}$$

Nine out of ten matches are not influenced by chance at all (the first chance move has influence zero). The result of these matches is only determined by the *match expectation*. Every tenth match however, is very strongly influenced by chance. The definition of Candidate_C suppresses the different magnitudes of these influences.

The definition of Candidate_B on the other hand, would assign to this game a value much closer to 1 (regardless of Π, Ψ):

$$\text{Candidate}_B(H, \Pi, \Psi) = \frac{10}{10+0+1} \approx 0.909 \quad (14.6)$$

Why is Candidate_B preferable to Candidate_C ? To answer this question we must look at series of matches.

14.3.3 Series of Matches

So far, our model of a match and our definitions of influence were all centered on the analysis of *single* matches (played by a random lineup of agents). If we know that more than one match will be played by the same lineup and we are interested in the expected influence for such a series of matches, we should expect different results.

As Noga Alon argues (based on the central limit theorem), the expected influence of chance diminishes when the number of consecutive matches grows [Alo 2007]. If we are

14.3 The Chanciness of a Game

interested in the remaining influence of chance on the added result of a fixed number of k matches, we can still use the measures defined in this chapter. We can model k matches of game G as a single game G' by assembling a new tree from copies of the tree of G . To model $k + 1$ matches every leaf node for the tree of G' must be replaced with a copy of the tree of G as new subtree. Furthermore, the result of the leaf that is replaced must be added to the results of all leaves in the new subtree.

Figure 14.2 shows the Game H' which results from modelling a series of two matches of Game H (Figure 14.1).




Figure 14.2: Two consecutive matches of Game H (Figure 14.1) modeled as a single match.

Using this approach we should expect that the *chanciness* of a series of matches will be lower than for a single match. Instead, the *match expectation* will come to dominate the combined result for such series of matches.

Based on this trend we argue that the game of *Roulette*, which is commonly classified as a game of pure chance, ceases to lose this character when the number of matches grows. For a sufficiently large number of matches it is practically certain that the combined result for the bank will be a large positive number (which explains the continued existence of this venture). On the other hand, the reduced influence of chance does *not* imply that roulette is a game of skill. For *roulette* it just underlines that the game is “unfair”. The idea that low *chanciness* does not imply a high level of skill is a key motivation for our concept of *controllability*, defined in the next chapter.

We now calculate the values of both candidates Candidate_B and Candidate_C in regard to the game H' (again, all values are independent of Π and Ψ). For every terminal node in the tree of H' there is a unique match which ends in that node. The probability for each match is the product of the probabilities of all moves in that match. By summing over all

14 Chanciness

possible matches we obtain the following values (Pr denotes the probability):

$$\begin{aligned}
\text{ME} &= 2, \quad \text{IP} = 0 \quad \text{for all matches} \\
\Pr(|\text{IC}| = 0) &= 0.9^2 + 2(0.1^2 \cdot 0.5^2) = 0.815 \\
\Pr(|\text{IC}| = 100) &= 0.9 \cdot 0.1 + 2(0.1 \cdot 0.9 \cdot 0.5) = 0.18 \\
\Pr(|\text{IC}| = 200) &= 2(0.1^2 \cdot 0.5^2) = 0.005 \\
\text{Candidate}_B(H', \Pi, \Psi) &= \frac{0.18 \cdot 100 + 0.005 \cdot 200}{(0.18 \cdot 100 + 0.005 \cdot 200) + 0 + 2} \approx 0.905 \\
\text{Candidate}_C(H', \Pi, \Psi) &= 0.18 \frac{100}{102} + 0.005 \frac{200}{202} \approx 0.18
\end{aligned}$$

We can now observe the problem of the Candidate_C measure. Instead of decreasing, as we would expect for a longer series of matches (two instead of one), the value of Candidate_C has increased! The measure Candidate_B on the other hand, behaves as expected. For this reason we abandon Candidate_C and use Candidate_B for our definition of the *chanciness of a game*.

14.3.4 The Definition of CHA

We define the *chanciness of the game* G from the perspective of player i as the expectation over random matches $g = (L, P)$ sampled according to Ψ from a population Π and played according to the rules of G and the agents in L :

Definition 14.7

$$\text{CHA}_i(G, \Pi, \Psi) := \frac{\mathbb{E}(|\text{IC}_i(g)|)}{\mathbb{E}(\text{TI}_i(g))}$$

In Section 14.2 we left the *chanciness of a match* $\text{cha}_i(g)$ undefined for a denominator of zero. This does not present a problem for Definition 14.7 as long as some matches are influenced by chance or by the players.

In theory, the denominator of Definition 14.7 might be zero as well. In our opinion, this is not a problem either, because games where $\mathbb{E}(|\text{IC}|)$, ME , $\mathbb{E}(|\text{IP}_j|)$ all equal zero are rather pathological. $\mathbb{E}(|\text{IC}|) = 0$ implies that IC is always zero and thus chance moves are meaningless for the outcome of the game. Likewise $\mathbb{E}(|\text{IP}_j|) = 0$ implies that none of the agents have influenced the result. In this case, player moves must either be deterministic or meaningless for the outcome of the game. In fact, every state in the game tree that is ever visited would have the same *relative positional value* (RPV) of zero. Arguably, such a game need not even be analyzed.

The *chanciness of a game* measures the propensity of matches to be (relatively) influenced by chance. The value of $\text{CHA}_i(G, \Pi, \Psi)$ can be computed exactly by first calculating the probability for every possible match g and then calculating the expected absolute influence for all SI .

14.4 Condensing Player Perspectives

We will now calculate the value of $\text{CHA}_2(G_{\text{example}}, \Pi, \Psi)$ for our example game G_{example} . We keep our population Π and lineup probability Ψ simple by assuming that the single lineup $L = (A_1, A_2, A_3)$ occurs with probability 1. For the agents in L we again use our example agents from Section 12.2. Thus we have $\Pi = \{A_1, A_2, A_3\}$ and $\Psi = \{(L, 1)\}$.

This special case of having only a single lineup is very similar to the situation discussed in Part I (Chapter 6).

The forthcoming values are calculated for the perspective of *affected player 2*. Each terminal node in the tree of G_{example} corresponds to a possible match. The probability that one of these matches occurs is the product of the move probabilities along the sequence of nodes for that match. For reference, Figure 14.3 shows these probabilities. The expectation operator \mathbb{E} applies to the random match g .

$$\begin{aligned}
 \mathbb{E}(|\text{ME}_2(g)|) &= 1.1 \\
 \mathbb{E}(|\text{IC}_2(g)|) &= 0.25 \cdot 0 + 0.75 \cdot 0.8 \cdot 0.5 = 0.3 \\
 \mathbb{E}(|\text{IP}_{1,2}(g)|) &= 0.25 \cdot 0.3 + 0.75 \cdot 0.1 = 0.15 \\
 \mathbb{E}(|\text{IP}_{2,2}(g)|) &= 0.25 \cdot 0.8 \cdot 0.2 + 0.25 \cdot 0.2 \cdot 0.8 \\
 &\quad + 0.75 \cdot 0.8 \cdot 0.3 + 0.75 \cdot 0.2 \cdot 1.2 = 0.44 \\
 \mathbb{E}(|\text{IP}_{3,2}(g)|) &= 0.75 \cdot 0.2 \cdot 2 = 0.3 \\
 \mathbb{E}(\text{TI}_2(g)) &= 1.1 + 0.3 + 0.15 + 0.44 + 0.3 = 2.29 \\
 \text{CHA}_2(G_{\text{example}}, \Pi, \Psi) &= \frac{\mathbb{E}(|\text{IC}_2(g)|)}{\mathbb{E}(\text{TI}_2(g))} \approx 0.13
 \end{aligned}$$

For the perspectives of player 1 and player 3 we get the following values (calculations in Appendix F):

$$\begin{aligned}
 \text{CHA}_1(G_{\text{example}}, \Pi, \Psi) &= 0.12 \\
 \text{CHA}_3(G_{\text{example}}, \Pi, \Psi) &= 0.16
 \end{aligned}$$

The exact computation of the *chanciness of a game* is unfeasible for large game trees (trees with a large number of nodes). For this reason we must resort to approximations of our measure. This topic is discussed in Chapter 18.

14.4 Condensing Player Perspectives

The *chanciness of a match* and the *chanciness of a game* are thus far defined for the perspective of the *affected player i*. For an n -player game G , we must therefore deal with a n -tuples of values. This tuple is interesting because it enables us to compare the different player perspectives and uncover potential asymmetries¹.

¹The idea that different player perspectives may perceive the impact of chance and skill differently is also discussed by Dreef [Dre 2005, 42].

14 Chanciness




Figure 14.3: The tree of Game G_{example} reduced to the perspective of player 2 with move probabilities for a specific lineup L . Each node s is annotated with the RPV, $V_2(s, L)$.

For human players who wish to argue about a match g and for game inventors seeking to analyze their creations, this n -tuple is the appropriate measure. If we seek to classify games from the perspective of a lawmaker, it may be preferable to condense this tuple into a single number.

An obvious choice for describing n values with a single number is the mean value. This approach could be justified if matches are typically repeated with rotated lineups such that every agent experiences the different perspectives (in this case it may also be interesting to model the series of matches as a new game as explained in Section 14.3.3). Otherwise one could take the conservative approach of bounding the *chanciness of a game* with the maximum or minimum value for the different perspectives. Comparing the minimum value of *chanciness of a game* (across all player perspectives) with a given threshold is conservative in declaring a game as “game of chance”. In contrast, comparing the maximum value *chanciness of a game* is conservative in declaring a game as “not a game of chance”.

14.5 Summary

In this chapter we have defined the *chanciness of a game* to describe a game from the perspective of a specific player. This measure is based on the *chanciness of a match* which describes a match, again from the perspective of a specific player.

Ultimately, these measures depend on a population of agents and detailed models of individual agent behavior. These dependencies may appear inconvenient but as we will see in Chapter 17, they cannot be neglected.

15 Controllability

15.1 Introduction

In the previous chapter we have defined the *chanciness of a match* and the *chanciness of a game* to describe the influence of chance on a match g and a game G . Both measures are specific to a player perspective i (the *affected player*).

If we know that the *chanciness* from the perspective of player i is high, we also know that the influence of the players is low. On the other hand, interesting questions remain when the *chanciness* is low. In this case we cannot generally deduce that a specific player had much influence. It may well be that some players did nothing of relevance.

As an example, consider match h from the perspective of player 3 in Figure 15.1 (match h was already shown from the perspective of player 2 in Figure 13.5). In match h , player 3 did not even get to move. If we examine the game tree closer, we notice that the only choice player 3 could have had has no influence either! Both alternative moves have the same terminal value of 0. Nevertheless, the *chanciness* of match h for player 3 has the low value of 0.25 (calculation in Appendix E).




Figure 15.1: An exemplary match h . The path through the game tree is highlighted and every node is marked with the *relative positional values* for the specific lineup from the perspective of player 3.

15 Controllability

Thus for matches (or games) with low *chanciness*, the question remains: How much influence did the players have? To answer this question we supplement our concept of *chanciness* with the new concept of *controllability*.

In Section 15.2 we will define the *controllability of a match* and in Section 15.3 we will define the *controllability of a game*. Both measures are structurally similar to *chanciness* measures in the previous chapter. In Section 15.4 we will discuss the reasons for naming our measure “controllability” and contrast it with the concept of “skill”.

15.2 The Controllability of a Match from the Perspective of Player i

In Section 13.4 we have discussed how to group the influences on a match according to sources of influence (*SI*). In this chapter we will use the same terminology.

Let $g = (L, P)$ be a match of an n -players game G . The result of match g is then subject to the influence of $n + 2$ distinct *SI*.

$$\text{ME}_i(g), \text{IC}_i(g), \text{IP}_{1,i}(g), \text{IP}_{2,i}(g), \dots, \text{IP}_{n,i}(g) \quad (15.1)$$

All *SI* are specific to the perspective of the *affected player* i .

Our measure *controllability* is intended to describe the relative importance of player influence in comparison to all sources of influence. Specifically, the influence of the *affected player* i . There are n terms of the form $\text{IP}_{i,j}$ that describe the influence of player i on the results of the players 1 to n .

Based on our agent model (Section 12.2), we know that every agent only cares about his own result¹. Therefore, the influence of player i on the result of the *other players* is irrelevant. The only important term for the influence of player i is $\text{IP}_{i,i}$, his own influence on his *own* result. This idea is central to our concept of *controllability*.

Analogously to our measure of *chanciness*, we relate the absolute influence of the *affected player* i to the influence of all other *SI*. We thus define the *controllability of a match*, $\text{con}_i(g)$ for a match g of an n -player game G evaluated from the perspective of player i :

Definition 15.2

$$\text{con}_i(g) := \frac{|\text{IP}_{i,i}(g)|}{\text{TI}_i(g)}$$

The measure con reaches its minimum of zero iff the absolute influence of the *affected player* $|\text{IP}_{i,i}(g)|$ is zero. It reaches its maximum of 1 iff $\text{IP}_{i,i}(g)$ is nonzero and all other *SI* have influence zero. For the measure con to yield values from the range of $[0, 1]$ it was necessary for influences of the players to be separated. If the influences of the players had been combined into a single *SI*, the numerator might have exceeded the denominator².

¹ Jealousy or malevolence among the agents can be expressed through a zero-sum reward structure.

² As an example, consider the case of $\text{ME}_i = \text{IC}_i = 0, \text{IP}_{1,i} = 2, \text{IP}_{2,i} = -1$. For the *affected player* $i = 1$ we would have a numerator of 2 and a denominator $\text{TI}_i = |2 - 1| = 1$

15.3 The Controllability of a Game

Just like the measure cha , the *controllability of a match* is undefined for a *total absolute influence* TI of zero.

Since $\text{cha}_i(g)$ and $\text{con}_i(g)$ use the same denominator $\text{TI}_i(g)$ it makes sense to compare them directly. This tells us the relation of chance influence to player control.

For the example match h in Figure 15.1 we calculate the *controllability* from the perspective of the *affected player* 3:

$$\begin{aligned}
 \text{M}_C(h) &= \{(s_2, s_3)\} \\
 \text{M}_1(h) &= \{(s_0, s_1)\} \\
 \text{M}_2(h) &= \{(s_1, s_2)\} \\
 \text{M}_3(h) &= \{\} \\
 \text{ME}_3(h) &= 1.15 \\
 \text{IC}_3(h) &= -0.5 \\
 \text{IP}_{1,3}(h) &= 0.05 \\
 \text{IP}_{2,3}(h) &= 0.3 \\
 \text{IP}_{3,3}(h) &= 0 \\
 \text{con}_3(h) &= \frac{0}{1.15 + 0.5 + 0.05 + 0.3 + 0} = 0
 \end{aligned}$$

This was to be expected as player 3 did not even get to move.

For the perspectives of the other players we receive different values (calculations omitted):

$$\begin{aligned}
 \text{con}_1(h) &= 0.4 \\
 \text{con}_2(h) &= 0.15
 \end{aligned}$$

The *chanciness of a match* defined in the previous chapter is only interesting for games with chance moves (cha is trivially zero for deterministic games). By contrast, the *controllability of a match* can be used to compare asymmetries among the different player perspectives even for deterministic games.

15.3 The Controllability of a Game from the Perspective of Player i

In Section 14.3 we have defined the *chanciness of a game* based on the properties of matches. Our reasoning in regard to the sampling of matches and the combining of their properties applies unchanged.

Thus, we define the *controllability* of game G from the perspective of player i as the expectation over random matches $g = (L, P)$ sampled according to Ψ from a population Π and played according to the rules of G and the agents in L :

15 Controllability

Definition 15.3

$$\text{CON}_i(G, \Pi, \Psi) := \frac{\mathbb{E}(|\text{IP}_{i,i}(g)|)}{\mathbb{E}(\text{TI}_i(g))}$$

The *controllability of a game* measures the propensity of players to influence their own match results (relatively). The value of $\text{CON}_i(G, \Pi, \Psi)$ can be computed exactly by first calculating the probability for every possible match g and then calculating the expected absolute influence for all SI .

As already explained for the match measures cha and con , the game measures of CHA and CON can be compared directly, since both definitions use the same denominator $\mathbb{E}(\text{TI}_i(g))$. This comparison helps us to understand the relation of chance influence to player control for a game G .

In Section 14.3 we have calculated $\text{CHA}_2(G_{\text{example}}, \Pi, \Psi)$ for exemplary values of Π and Ψ . In Appendix F we have also calculated the values for the perspectives of players 1 and 3. The intermediate values used in those calculations can be used directly to calculate $\text{CON}_i(G_{\text{example}}, \Pi, \Psi)$. Figure 15.2 gives the expected absolute influences of the different SI from all three player perspectives³.

$SI \setminus i$	1	2	3
$\mathbb{E}(\text{ME}_i(g))$	0.4	1.1	1.15
$\mathbb{E}(\text{IC}_i(g))$	0.3	0.3	0.3
$\mathbb{E}(\text{IP}_{1,i}(g))$	1.2	0.15	0.075
$\mathbb{E}(\text{IP}_{2,i}(g))$	0.36	0.44	0.36
$\mathbb{E}(\text{IP}_{3,i}(g))$	0.3	0.3	0.0

Figure 15.2: Expected values of the absolute influences of all SI for random matches g of the Game G_{example} . The lineup $L = (A_1, A_2, A_3)$ was assumed to occur with probability 1.

Using the influence values from the above table we compute the *controllability* of G_{example} for all three player perspectives.

$$\begin{aligned} \text{CON}_1(G_{\text{example}}, \Pi, \Psi) &= \frac{1.2}{0.4 + 0.3 + 1.2 + 0.36 + 0.3} \approx 0.47 \\ \text{CON}_2(G_{\text{example}}, \Pi, \Psi) &= \frac{0.44}{1.1 + 0.3 + 0.15 + 0.44 + 0.3} \approx 0.19 \\ \text{CON}_3(G_{\text{example}}, \Pi, \Psi) &= \frac{0}{1.15 + 0.3 + 0.075 + 0.36 + 0} = 0 \end{aligned}$$

In Section 15.1 we pointed out that the moves of player 3 in the game G_{example} are irrelevant to his own result. Consequently, the *controllability* from the perspective of player 3 is zero. While the *chanciness* of G_{example} is very similar for the different players, our

³One way to think about *chanciness* and *controllability* is as tools for summarizing this table.

15.4 Controllability and the Concept of “Skill”

measure of *controllability* reveals striking asymmetries among them. It can be shown that the *controllability* of player 3 is zero in all cases. However, the *controllability* for players 1 and 2 would be quite different for other populations of agents. This dependency on the population (and the sampling of lineups) is discussed in Chapter 17.

Just as for the *chanciness* of a game, the exact computation of the *controllability* of a game is unfeasible for large game trees (trees with large number of nodes). Another consideration that applies to *chanciness* as well as *controllability* is the need for dealing with n -tuples of values. The ideas discussed in Section 14.4 apply to both measures.

15.4 Controllability and the Concept of “Skill”

The title of this thesis mentions the concepts of “chance” and “skill”. Our measure of *chanciness* was so named because it is intended to measure the influence of chance on a game. Why then have we not yet spoken of the “skilliness” of a game? The reason for naming our measure of player influence “controllability” is the wish to avoid confusing two different concepts.

On the one hand there is the concept of “skill”. Skill is usually understood as the attribute of an agent. The skill of an agent is his aptness for achieving high results in a game. In the context of two-player zero-sum games such as *Chess*, the skill of an agent is understood as his ability to defeat other agents. This ability is measured by analyzing the results of many matches between different agents. It can be expressed in a skill rating such as the *Elo rating* [Elo 1986]. In Section 16.2.2 we explain how this agent “skill” contributes to the *match expectation*.

On the other hand we have the concept of “influence”. This describes a relationship between actions (moves) and outcomes (match results). By applying the concept of influence to chance moves we obtain the measure of *chanciness* as defined in Chapter 14. Applying the concept of influence to player moves however, yields something distinct from “skill”!

As an example, consider a player who is able to play the game chess (or any two-player zero-sum game⁴) perfectly. Such a player would certainly be considered “skilled”. However, the moves of this perfect players would have zero influence. Since the perfect player always selects the move with the optimum reward he must either play deterministically or select at random among moves that are equally optimal. In either case the *relative positional value* of his position does not change. Consequently, the influence of a player cannot be the same concept as his “skill”.

We have chosen to name our measure “controllability” to emphasize the intention of the players. The alternative term “influenceability” is not specific enough since it does not distinguish between the influence of chance and the influence of the players. There are some similarities to the term “controllability” as used in control theory [KFA 1969]. A control system is controllable if the inputs allow the system to be moved through its entire

⁴As mentioned before, perfect play is not easily defined for more general classes of games [MP 1992].

15 Controllability

state space. Likewise, a game can only be controllable if there are player moves that allow different game results to be achieved.

The question remains, why the concept of “skill” appears in the title of this thesis, even though we discuss the distinct concept of *controllability*. This is in reference to another use of the word “skill”. We use this word to describe games, saying “game G is a game of skill”, “game G is not a game of skill”, and “game G is partly a game of skill”. There are basically two approaches of quantifying this usage of the word “skill”

The approach of BDG (Section 11.3): In this approach the “skill of a game” is quantified by measuring how agents with different skill ratings affect the expected result of matches (the *match expectation* ME, using our terminology). The “skill of a game” is thus linked to the “skill” of the agents.

Our concept of controllability: In our approach, a game is a “game of skill” to the degree that it has high *controllability*.

We can expect both approaches to arrive at similar classifications. If a game allows for agents of different skill ratings, it will also allow for populations that have high *controllability*.

The approach of this thesis has the advantage of applying to matches as well as games. Using our approach we can, for example, recognize that a game G is a “game of chance”, yet some matches g of G are decided by the influences of the players rather than by chance.

We propose that our measure of *controllability* captures the essence of that which humans players seek in a “game of skill”. For one, humans who seek to apply their skill avoid games that they find trivial. People rarely go back to the game of *Tic-Tac-Toe* once they have mastered its strategy. This state of having mastered a game is reflected in being certain of one’s moves and being able to act almost deterministically. Hence, the true master of a game has zero *controllability* over its result (unlike his novice opponent who may or may not discover a viable move)⁵. Thus our concept of *controllability* resonates with the ideas of Lasker who said that chess would be “dead” for perfect agents. He suggested that a game is “alive”, only if agents sometimes make erroneous moves [Las 1925, 66,171].

One might object that such a master player could deliberately play worse and thus negatively influence his result. This kind of control should not be considered meaningful. Otherwise we could add a “lose the game instantly”-move to any game of chance and thus transform it into a controllable game.

The character of a game changes with the proficiency of its players. Once we understand a game so thoroughly that all move decisions turn into trivialities, we can no longer improve our results by our conscious effort. Thus we lose the sense of control we once had when the game was still new. If all participating players should reach such an advanced level of play they may as well just roll a die to discover their results. Unless they

⁵ The master may still control the course of the game, yet his victory (and hence his result) is certain.

discover that the game, once mastered, gives an unfair advantage to one of the players, their results will be dominated by the influence of chance⁶.

The *controllability of a game* can be used to express this changing character of a game as the population gains in proficiency. One simply has to compare the *controllability* value for a population of beginners to the value for a population of advanced agents. The latter population will most likely have reduced controllability⁷.

By assuming that players seek games with a high *controllability* we are able to justify the common practice of giving the stronger player a handicap. As an example consider the game of Go. An experienced player would tend to win most games against a player of lesser rank. He would therefore have a high *match expectation*. By playing with a handicap, he reduces the *match expectation* (preferably to zero). Since the *match expectation* appears in the denominator of Definition 15.2, the *controllability* is increased and the game becomes more interesting. Finally, we prefer our concept of *controllability* because it elegantly treats the influence of chance and the influence of the players in the same way.

15.5 Limitations

Our concept of *controllability* does not include the aspect of learning by individual agents. A human agent may willfully modify his playing behavior through practicing a particular game. If successful, this will be reflected in an increased value for the *match expectation* ME_i whenever this agent plays in the player role i . By influencing the *match expectation* this player also influences his result. One could argue that the decision to practice a game is also an instance of controlling game results. Since our model only considers static agents, our measure cannot account for this kind of control.

15.6 Summary

In this chapter we have defined the *controllability of a game* to describe a game from the perspective of a specific player, and the *controllability of a match* which describes a match from the perspective of a specific player. Both measures are based on the comparison of influence and follow the ideas developed in Chapter 14.

The concept of *controllability* complements the concept of *chanciness* to describe the influence of chance and skill on a game. Since the *controllability* may differ among the player perspectives this measure can be a useful tool for describing deterministic games as well as non-deterministic games.

⁶A deterministic game would become meaningless once all agents are masters.

⁷There may also be games where the beginners fail to find influential moves. In this case advanced agents may have increased *controllability*. However, as soon as the agents approach optimal play, *controllability* will go down.

16 Other Influences

As we have argued in the previous chapters, the result of a match is affected by different sources of influences (*SI*). Two of those *SI* have been given special treatment:

The relative influence of chance IC_i has been used as the main term in our definition of *chanciness*. If the *chanciness* is high for a match (or a game), we know that chance is important.

The influence of the affected player on his own result $IP_{i,i}$ has been used as the main term in our definition of *controllability*. If the *controllability* is high for a match (or a game), we know that skillful decisions are important.

However, there may be matches (and games) where both the *chanciness* and the *controllability* are low. In this case the other *SI* must be dominant. In this chapter we discuss the interpretation of the other sources of influence.

16.1 Influence of the Moves of *Other Players*

What does it mean if the result of a match (or a game) is dominated by the influence of the *other players*? In the concluding remarks of Part I we suggested that such influence may be related to chance.

In a match g of a two-player zero-sum game G the influence of the *other player* on the result of the *affected player* is mirrored by the influence of the *other player* upon his own result (for every move m , $I_1(m) = -I_2(m)$).

As a consequence every move decision by the *other player* which influences the *affected player* is a relevant decision for the *other player*. This makes it easy to argue that the influence of the *other player* is unrelated to chance.

In general-sum games or games with more than two players a different type of situation may occur. There may be moves that influence the result of the *affected player* but do not influence the *other player* who is responsible for that move. An example of this is known as the *king maker situation*. It can occur in games with more than two players: The player who moves has effectively lost the game but with his move, he has the power to decide which one of the other players will win. In this situation the rules of the game give no incentive for him to prefer either choice. From the perspective of an *affected player* the choice may thus appear random.

However, the fact that the *affected player* may have no way to predict the behavior of the *other player* based on the games rules does not allow the inference that the choice actually

is random. Without detailed knowledge about the thought process (or algorithm) of the *other player*, it is impossible to know if randomization is involved. Indeed one might argue that true randomness is impossible for humans and every action is based on a (potentially subconscious) reason.

If on the other hand we admit that players may occasionally randomize their choices, they could very well do so even in a game of chess¹. In that case even games without chance nodes could be influenced by chance. Thus, the well accepted concept of a “pure game of skill” would lose its meaning.

The concepts developed in this thesis acknowledge that the influence of *other players* can appear similar to the influence of chance: Our measure of *controllability* is equally diminished when either the influence by the *other players* or the influence by chance is high.

16.2 The Match Expectation

All matches of game G start in the same root position s_0 . Accordingly, the value of *match expectation*, $ME_i(g) = \tilde{V}_i(s_0, L)$ depends on two things: the lineup L and the game G itself.

16.2.1 The Influence of the Game

For some games, the value of the starting position s_0 is unaffected by the lineup L . As an obvious example, consider games without player nodes. In these cases, the *match expectation* appears to be a fundamental property of the game G . We call this property the *influence of the game G* .

The *influence of a game* is closely related to the concept of “unfairness”. In a zero-sum game without player nodes, a player with a negative *match expectation* is disadvantaged through no fault of his own (there must be other players who have a positive expectation). In a non-zero-sum game all players may have a negative (or positive) *match expectation* at the same time. For these games, bad (or good) “fate” may be a better description of the situation.

If a game has low *chanciness* and low *controllability*, it is helpful to know that results are influenced by “unfairness” or “fate” rather than by chance or by player decisions. It would therefore be useful if the *influence of a game* could be generalized to games where the lineup L does affect the *match expectation*.

One possible way to generalize the *influence of a game* comes from our population Π and the sampling probability for lineups Ψ . In Chapter 12 we assume that there exists a procedure to sample lineups. Since this sampling is also a sort of “random move” which precedes the state s_0 , we can posit the “meta-state” s_{-1} . Namely, the state in which the

¹One could argue, that agents would *only* randomized their choices when perceiving situations similar to the *king maker situation*. Following this assumption would require a more elaborate agent model.

16 Other Influences

sampling of the lineup has yet to take place. The expected value of state s_{-1} is well defined through the sampling probability Ψ and the different values of $\tilde{V}_i(s_0, L)$. Since this value of s_{-1} is independent of any particular lineup it can be considered a fundamental game property (in respect to Π and Ψ , of course).

Definition 16.1 *The population value of state s , $V^{pop_i}(s, \Pi, \Psi)$ is the expected value of $\tilde{V}_i(s, L)$ when the lineup L is sampled with probability $\Psi^n(L)$ from the population Π .*

Definition 16.2 *The influence of the game I_G for game G with root node s_0 is defined as:*

$$I_G(\Pi, \Psi) := V^{pop_i}(s_0, \Pi, \Psi) \quad (16.3)$$

The *influence of the game G* is exactly the value of the “meta-state” s_{-1} in which the lineup L has yet to be sampled.

That *influence of a game*, or the “fairness” of a game is population dependent can be seen when considering games such as *Connect-Four* [All 1988] and *Hex*. If these games were played by a population of agents that play optimal according to the game-theoretic solution, all matches for these games would end in a win for the player who moves first².

On this level of playing strength the games are obviously unfair. If we assume that games which are obviously unfair are avoided by most players, we can conclude from the enduring interest in *Hex* and *Connect-Four* tournaments that the games are quite fair for the population of tournament participants.

16.2.2 The Influence of the Lineup

We have already described the sampling of a lineup as a “meta-move” from state s_{-1} to state s_0 . This suggests another definition.

Definition 16.4 *The influence of the lineup L on game G with root node s_0 , I_L*

$$I_L(\Pi, \Psi) := \tilde{V}_i(s_0, L) - I_G(\Pi, \Psi)\Psi \quad (16.5)$$

This definition follows Axiom 1: the influence of the “meta-move” (s_{-1}, s_0) is defined as the change in value between s_0 and s_{-1} .

The *influence of the lineup I_L* can be seen as an indicator of skill (for one-person-games) or as skill disparity for games with multiple players. BDG have proposed that such skill disparity within a population is a necessary property for games of skill (see Section 11.3). While we prefer the measure of *controllability* for describing games of skill, it is certainly interesting to know how strongly match results are influenced by a particular lineup.

²Such agents exist for *Connect-Four*, while a non-constructive proof exists for *Hex*.

16.2.3 Concluding Remarks

Using the above definitions it is plain to see that for any match $g = (L, P)$ of the game G the following equality holds:

$$\text{ME}_i(g) = \text{I}_G(\Pi, \Psi) + \text{I}_L(\Pi, \Psi) \quad (16.6)$$

Similar to our approach towards the *match equation*, we can “explain” the value of the *match expectation* with the combined influence of the *influence of the game* and the *influence of the lineup*.

We have yet to justify our choice of using the word “influence” for the *influence of a game*. Can the *influence of a game* be interpreted as the influence of a “meta-move”, just as the *influence of the lineup* is the influence of the “meta-move” (s_{-1}, s_0) ? What would be the state s_{-2} and how could we define its value?

These questions are well outside the usual scope of game analysis. They border on the metaphysical. One suggestion that presents itself is to consider the move (s_{-2}, s_{-1}) as the act of selecting a game among the set of all possible games. It is not clear to us who should perform this move, since there is not even a lineup of agents yet. However, for reasons of universal symmetry and beauty we are convinced that the value of the state s_{-2} should be 0. In this sense the *influence of the game* G could indeed be considered as the change in value caused by the selection of game G .

17 Fundamental Results

In the previous chapters we have defined the measures *chanciness*, CHA and *controllability* CON in reference to a population of agents. This dependency complicates practical use, since it is difficult to create a detailed population model for a population of human agents. One could easily argue that it is impossible to obtain all the necessary move probabilities to model even one human agent accurately. For this reason we should obviously ask: what conclusions can be drawn about a game G in ignorance of a population?

17.1 Games Without Player Nodes

Our whole framework of analyzing influence relies on the evaluation of game states with *relative positional values* (RPV). Fortunately there exist some game states (nodes) s for which $\tilde{V}_i(s, L)$ can be evaluated without knowledge of the lineup L (and hence without knowledge of agents or population). By Definition 13.1 we can trivially evaluate terminal nodes. We can also evaluate a chance node s if we know the RPV of all children of s . Taken together, these two facts allow us to completely analyze games that have no player nodes (and hence no player moves).




Figure 17.1: Small games without player nodes. Games a)-c) are single-player games and game d) is a game of two players.

Figure 17.1 shows four small games which can be fully analyzed independently of a population. For all four games, the *controllability* is obviously zero, since there are no player moves (and therefore no player influence). However, all games receive different values of *chanciness*. Game 17.1a) is a game for one player. It has a *match expectation* of

0 and an expected influence of chance (on player 1) of 1. Accordingly, the *chanciness* of that game is 1. By contrast, game 17.1b) has a *match expectation* of 99. This value is quite large in comparison to the expected influence of chance, $\mathbb{E}(|IC_1|) = 1$. This discrepancy in importance is expressed in the low *chanciness* value for this game of 0.01.

Game 17.1c) demonstrates that the existence of chance moves is not a sufficient criteria for chance influence. Since both terminal nodes have the same value, the influence of chance is always zero. Hence, the *chanciness* of this game is 0 as well. The two-player game 17.1d) shows another aspect of *chanciness*. While the *chanciness* of player 1 is 1 (as analyzed in game a), the game has a *chanciness* of 0 from the perspective of player 2. A simple “yes/no” answer to the question: “Is this a game of chance?”, could not adequately describe this game.

17.2 Bounded Influence

By Definition 13.1 The RPV of a game state s is a convex combination of the values of $\text{Children}(s)$. This gives us upper and lower bounds for the influence of a move $m = (s, s')$ from the perspective of player i .

$$I_i(m) \leq \max \tilde{V}_i(\text{Children}(s)) - \min \tilde{V}_i(\text{Children}(s)) \quad (17.1)$$

$$I_i(m) \geq \min \tilde{V}_i(\text{Children}(s)) - \max \tilde{V}_i(\text{Children}(s)) \quad (17.2)$$

The influence of a move $m = (s, s')$ approaches this bound if three conditions are met:

- state s' has the highest (lowest) value among $\text{Children}(s)$,
- move m has a probability close to zero,
- move (s, s'') has a probability close to one, where s'' is the child of s with the lowest (highest) value.

Since move m must have a non-zero probability, the extremal value is never achieved. As an obvious consequence, a move $m = (s, s')$ has influence zero if all children of s have the same value. This allows us to fully analyze some games even though they contain player moves.

Figure 17.2 shows two small games which can be fully analyzed independently of a population. For both games, the *chanciness* is obviously zero, since there are no chance moves. For the single-player game 17.2a) *controllability* is zero as well due to the bounding of influence.

Game 17.2b) is a game for two players. Even though player 1 may influence the result of player 2, neither player can influence his *own* result. Therefore, both players have *controllability* zero, independent of the population.

17.3 Bounded Expectations

The influence of a single move that starts in state s can be bounded by considering the spread in the RPV of the children of s . The contribution of that move to the expected

17 Fundamental Results




Figure 17.2: Small games without chance nodes. Games a) is a single-player game and game b) is a game for two players.

importance of chance or to the expected importance of a player is also affected by the likelihood that s will occur during a match.




Figure 17.3: A small single-player game with chance and player nodes.

The single-player game G in Figure 17.3 cannot be fully analyzed without knowledge about the population. However, we can compute tight bounds on the *chanciness* and the *controllability* of this game due to the extreme bias of one chance node.

The *controllability* of this game can be as low as zero if all agents in our population move deterministically. This trivial lower bound of *controllability* holds for any game. Computing the other bounds is more difficult.

An agent X_p who plays this game can be fully characterized by the probability p of taking the left child node. When p varies, the values of ME , $\mathbb{E}(|\text{IC}|)$ and $\mathbb{E}(|\text{IP}_1|)$ are all affected. The *controllability* as well as the *chanciness* are non-linear functions of these non-linear terms.

Figure 17.4 shows the *chanciness* as a function of p . $\text{CHA}(G, \{X_p\}, \{(\{X_p\}, 1)\})$ is computed for a population that consists only of the agent X_p (the lineup (X_p) occurs with probability 1). The *controllability* of Game G has its peak of 0.01 for the uniform random agent $X_{0.5}$. The *chanciness* has two maxima of ≈ 0.99098 at $p \in 0, 1$ and two minima of ≈ 0.98759 at $p \in 0.25, 0.75$. The uniform random agent with $p = 0.5$ only achieves a local *chanciness* maximum of 0.99. The derivation is given in Appendix C. These tight bounds




Figure 17.4: The *chanciness* of Game *G* as a function of the agent parameter *p*.

justify the conclusion that Game *G* from Figure 17.3 is mostly a game of chance and hardly a game of skill.

As a general rule, bounding *chanciness* and *controllability* necessitates solving highly dimensional non-linear optimization problems. This quickly becomes unfeasible because the dimensionality of the solution space is at least equal to the number of player nodes in the game tree. Each player move probability is an input to the optimization problem. Assuming a binary tree, one variable per player node is sufficient to express the probabilities for both available moves. For trees with a higher branching factor, even more variables are needed.

17.4 Sensitivity of Boundaries

Computing exact bounds for the *chanciness* and *controllability* of a game is not just difficult. Even worse, the relevance of exact bounds is questionable. When reasoning about *chanciness* and *controllability* without reference to a population, all possible agents must be considered. However, not all possible agents are plausible. If “implausible” agents are responsible for the extremum values then those values are misleading.

To demonstrate the high sensitivity of the bounds towards “implausible” agent behavior we will look at the term $\mathbb{E}(|IP_{i,i}|)$ (the expected influence of a player upon his own result).

As a first disappointment, the lower bound for $\mathbb{E}(|IP_{i,i}|)$ (and thus for *controllability*) is trivially zero. This value is achieved with a population of deterministic agents (all agents act deterministically at all player nodes). It is highly unlikely that a real life game would be played if all people always knew (or presumed to know) exactly how to act. Therefore,

17 Fundamental Results

this lower bound for *controllability* is most likely useless.

The upper bound for *controllability* is dubious as well. This can be seen by considering a situation we call “instant-loss-button” (ILB). An ILB is any move which causes the acting agent to end the game with an extremely low payoff. Figure 17.4 gives an example of this situation.

In a real game, most players would quickly learn to avoid an ILB. Thus, a real population would contain few agents who use ILBs. In a game that contains an ILB, the upper bound of $\mathbb{E}(|IP_{i,i}|)$, will most likely be caused by a population of “silly” agents who use the ILB move with probability 0.5 (An agent who always uses the ILB has less influence). The term $\mathbb{E}(|IP_{i,i}|)$ also appears in the denominator of the definition of *chanciness*. Thus, the upper and lower bounds of *chanciness* are dubious as well.




Figure 17.5: A small game with an “instant-loss-button” (ILB). The move towards the right child constitutes the ILB.

17.5 Summary

The computation of upper and lower bounds for *chanciness* and *controllability* avoids the need to specify a population. For small games these bounds are easily computed and for some games the bounds are sufficiently tight.

Unfortunately, the computation is unfeasible for larger games and the boundaries are very sensitive to “improbable” agents (extremely deterministic or extremely incompetent agents). This underlines the need of determining the *chanciness* and the *controllability* of a game in reference to a population of agents.

Necessary conditions for the influence of chance on game results are

- existence of chance nodes in the game tree,
- chance moves must be relevant (the spread in the RPV of the children of a chance node must not be too low),
- chance nodes must actually appear during matches.

The necessary conditions for the influence of skill (player decisions) are very similar:

- existence of player nodes in the game tree,
- player moves must be relevant (the spread in the RPV of the children of a player node must not be too low),
- player nodes must actually appear during matches,
- player moves must not be trivial (agents must have plausible alternative moves).

Of these conditions, only the existence-condition of nodes can be tested without reference to a population. All other conditions depend on the behavior of actual agents.

18 Approximation of *Chanciness* and *Controllability*

In the previous chapters we have defined the measures of *chanciness* and *controllability* for matches as well as games. These measures depend on a detailed behavioral agent model. The results in Chapter 17 have shown that it is not generally useful to reason about *chanciness* and *controllability* without taking the actual behavior of an agent population into account.

This presents us with a practical problem. For real-world populations of human agents and computer agents (bots) we generally do not possess such a detailed behavioral model¹. In this chapter we present different approaches for approximating *chanciness* and *controllability* that do not require an agent model. Due to the similarity in definition of *chanciness* and *controllability*, all topics in this chapter apply to both measures equally. For the sake of brevity, we will only speak of the *chanciness* measures *cha* and *CHA*.

In Section 18.1 we describe two sampling algorithms for the approximation of *chanciness*. In Section 18.2 we introduce the concept of *subjective chanciness* as an alternative approach to the approximation of *chanciness*. We also discuss how both approaches relate to different types of agents. Section 18.3 pertains to the special case of all agents being bots. In Section 18.4 we discuss the case where human agents are involved.

18.1 Approximation by Sampling

Section 18.1.1 describes a sampling algorithm for RPV. This is the main requirement for the approximation of the *chanciness of a match*. Additional considerations are necessary for approximating the *chanciness of a game*. These are explained in Section 18.1.2.

18.1.1 Approximating RPV

The *chanciness of a match* depends on the influence of moves which is defined in terms of *relative positional values* (RPV). These RPV are defined as the expected value of a position for specific agent behavior (this is the *descriptive approach* as explained in Section 13.3.2).

We can approximate the RPV, $\tilde{V}_i(s, L)$ for state s from the perspective of player i and a lineup of agents L in the following way.

¹ The lack of an agent model is sufficient to frustrate the exact calculation of our measures. However, this is not the only problem. The calculation of RPV, using Definition 13.2, requires a recursive evaluation of the complete game tree, which is unfeasible for large trees.

18.1 Approximation by Sampling

1. We start a *partial match*² in state s and record the result for player i .
2. We repeat the previous step a number of m times.
3. We average the results.

This gives us the sample mean of the match results, which is an approximation of the RPV of s . For match results with a sample standard deviation of d , the standard error³ of our approximation is given by d/\sqrt{m} .

To approximate the *chanciness* $\text{cha}_i(g)$ for match $g = (L, P)$ with $P = (s_0, s_1, \dots, s_k)$, we need to approximate the RPV for $k + 1$ game states. Each *partial match* gives the result for all player perspectives i , therefore the values of \tilde{V}_i and cha_i are obtained for all players at the same time.

Each *partial match* requires in the order of k moves to complete. Thus, the number of individual moves which must be simulated is proportional to $m \cdot k^2$. This becomes problematic for games with long matches, especially if the agents are slow movers. In Appendix B we discuss a slight improvement to the above algorithm.

18.1.2 Approximating the Distribution of Matches

The *chanciness of a game* is defined using the expectation operator, applied to influence values of random matches. These random matches depend on a sampling distribution of lineups Ψ as well as on the behavior of the agents. For a given population of agents Π and a given distribution of lineups Ψ we can approximate the value of CHA by sampling matches.

Our approximation algorithm consists of the following steps

1. We sample a lineup L from the population Π .
2. We Sample a match $g = (L, P)$ with the previously chosen lineup L and record the influence for each SI (see Section 13.4.3).
3. We repeat the above two steps m times.
4. We average the influence for each SI and substitute those averages for their expectations in the definition of CHA (Definition 14.7).

The sampling of lineups can be simplified for distributions where each lineup L_j has a probability p_j such that $p_j \cdot m \in \mathbb{N}$. In this case we can simply use the lineup $p_j \cdot m$ times to eliminate any sampling error from step 1. In Appendix H we describe how the error of approximation can be estimated for this special case.

Computing the influences of all SI for a match g from the perspective of player i , has the same algorithmic complexity as computing the *chanciness*, $\text{cha}_i(g)$ (The latter requires one additional division operation). Thus, the above algorithm requires m times as many agent moves to be simulated as the approximation of cha .

²See Section 12.3.

³We refer to the estimate of the true standard deviation.

18.2 Subjective *Chanciness* and Subjective *Controllability*

As mentioned in Section 13.3.1, the question of how to win has been a major focus of game research. This research has led to a wealth of algorithms for evaluating game states. We will refer to these algorithms as *evaluation functions* (EF). Ultimately, an *evaluation function* approximates the expected value of a game state under specific assumptions about agent behavior. If the agents in a lineup L were to behave in accordance with those assumptions, such an EF would yield values identical to the RPV.

That an EF yields a sufficient approximation of the RPV for particular agents, is generally hard to prove. What happens if we use a particular EF anyway?

We look back at the **Fundamental Axiom of Influence** and recall that it defines influence in terms of “expected” values. If we interpret the word “expected” as the subjective belief on an observer (who need not be rational in any sense), we arrive at a subjective interpretation of influence. This subjective interpretation allows all the analysis of match influence (Chapter 13) and *chanciness* (Chapter 14) to be carried out.

Previously the *chanciness* of a match g was defined, based on an objective description of agent behavior. We can define the *subjective chanciness* of g based on any EF. Just as cha_i was specific to a player perspective, different EF may be used to define the *subjective chanciness* for each perspective.

What might be the use of these subjective measures? We note that the *match equation* holds in the subjective view. Thus, the separation into different sources of influence is still reasonable. If the observer, whose beliefs are expressed through the EF, can be convinced of the reasoning behind this thesis, he should accept the subjective values of cha and con as a characterization of chance and skill. If this observer was somehow able to analyze a significant number of matches, he should also accept the characterization provided by the subjective versions of CHA and CON . At the same time, we can then interpret the subjective measures as a description of this observer’s perceptions.

While the subjective values of *chanciness* and *controllability* may be radically different from their objective counterparts, they have one important advantage: speed. As we will see in the forthcoming Section 18.3.2, the subjective approach can result in an enormous speed up over the sampling approach for the case of computer agents.

18.3 Computer Agents

In this section we look beyond the agent model of Chapter 12. In our discussion so far, agents were only understood as having a certain behavior which can be described stochastically. This was convenient for our theory of influence and the definition of our measures. Now we must look at more practical properties of agents.

18.3.1 The Structure of a Computer Agent (Bot)

The fundamental task of a bot is the selection of moves. For most computer agents this selection follows a similar two-step design pattern. Let s denote the game position that demands a decision.

1. Assign a value to every successor state of s . These values are given by an *evaluation function* (EF). We call this the *internal evaluation function* (IEF) of the bot.
2. Select the move, leading to the child of s which is rated highest⁴ by the IEF. Often, the selection will be randomized among the highest ranking options if the optimum is not unique.

The IEF of an agent may be either deterministic or non-deterministic. This allows for non-deterministic behavior by the agent. By our initial assumption that every agent cares about maximizing its *own* reward, the IEF evaluates game state from the perspective i in which the agent is playing.

In the state s , it is assured that a bot will move to a child with the optimum IEF value. Thus, when in this state, the agent can use the optimum IEF value as an estimate for the value of s .

We define the *subjectively perceived value* (SPV) of a bot A :

Definition 18.1

$$SPV_A(s) := \max_{c \in \text{Children}(s)} \{IEF_A(c)\} \quad (18.2)$$

To select a follow up move in the state s , the computer agent A must calculate the value of $SPV_A(s)$ once.

18.3.2 The Sampling Approach applied to Bots

For the approximation of our measures, it makes an enormous difference whether human agents are involved or whether we only deal with computer agents (bots). The distinguishing feature of bots is their capacity for playing a large number of matches. Human players will tire and, even worse, change their behavior due to learning. Bots on the other hand, are only subject to limited computer resources and will not change their behavior (assuming, that potential learning mechanisms can always be suppressed).

The sampling algorithms in Section 18.1 rely heavily on the feasibility of conducting great numbers of matches. The limiting factors are the average length of matches and the time, that agents require for their move decisions.

As shown by the results in Chapter 20, the sampling algorithms *can* be carried out for some short games (average length 25 moves) and some fast agents in practice.

⁴In two-player zero-sum games, one player usually uses a negated value function and thus selects the lowest child.

18 Approximation of *Chanciness* and *Controllability*

18.3.3 The Subjective Approach Applied to Bots

The *subjective chanciness* (and the *subjective controllability*) can principally be computed for an arbitrary observer and his EF. However, it can be safely assumed that the agents who play a match are particularly interested (and interesting) observers. It is therefore natural that we consider the SPV of our agents as an approximation for the RPV. The computation of the *subjective chanciness of a match* is much faster than the approximation of the objective *chanciness of a match* by sampling.

Approximating the RPV for a single state s requires frequent computations of the SPV_{A_j} . As explained in Section 18.1.1 this would entail the simulation of m *partial matches*, each with a number of moves in the order of k (the average length of a match). The subjective value of s from the perspective of player i is instead obtained by a single computation of $SPV_{A_i}(s)$. If we further assume that all SPV_{A_j} are equally fast, we obtain a speed up in the order of $m \cdot k$.

Consequently, the computation of the *subjective chanciness of a match* is faster than the approximation of the original *chanciness* in the order of $m \cdot k$. To compute the *chanciness of a game* (subjective or objective) we must still carry out a number of simulation matches. Nevertheless, the speedup of the subjective approximation over the sampling approximation applies here as well.

The idea of using the *subjective* values of game states to describe the perceptions of the playing agents is also used by Browne [Bro 2008]. He defines aesthetic properties of a game based on the IEF of the agents who play the game. In Browne's work, speed is of paramount importance since he feeds those aesthetic properties into the fitness function of a genetic algorithm.

18.4 Human Agents

The sampling algorithms in the previous Section cannot be usefully applied to human agents. The number of simulation matches required for reasonable accuracy is far above the endurance of humans. As another difference, only a few humans can express the value of a game state similar to the SPV of bots⁵. We therefore cannot use the two approaches of Section 18.1 and 18.2 directly.

18.4.1 Substituting Computer Agents

The problem of approximating *chanciness* for human agents was already discussed in Part I of this thesis. The main idea presented there was the substitution of human agents by bots with similar behavior. Such substitution would allow us to apply all the algorithms from Section 18.1.

The substitution approach creates the two problems of defining “similar behavior” and finding agents who satisfy the definition. How then might “similar behavior” be de-

⁵A *Chess* master may evaluate positions using fractions of pawns and a *Go* master will tally his territory. This kind of quantitative evaluation is uncommon for less advanced players.

defined? The ultimate answer to this question lies in a variation of the well known *Turing Test* [Tur 1950]. Two game players X and Y behave in a similar way if they cannot be distinguished by someone (or something) observing them play.

One criteria by which agents may be readily distinguished is their skill. If agent X performs consistently stronger than Y they cannot be the same. This criteria does not even require us to hide the identities of the agents from the observer (as in the Turing Test). We can simply let them play against each other and compare their results. We could also compare their performance against a specific lineup of opponents.

How does the criteria of skill relate to the feasibility of the substitution approach? For any game G human players encompass a certain range of skill from the worst players up to the master players⁶. At the same time, all the bots which are capable of playing G encompass another range of skill. Since bots can easily be made to have very low skill (e.g. by changing the leading sign of their EF), the most interesting question is the high end of their skill range. The set of human players which cannot be distinguished from bots by skill is indicated by the intersection of both skill ranges.

For simple games such as *Tic-Tac-Toe* we can certainly match every human with a bot of similar skill. The same is true for more elaborate games such as *Connect Four* and *Backgammon* [Tes 1995] and it is approximately true for the game of *Chess*. For other games such as *Go* and *Havannah*, only a subset of all human players can be matched.

All in all, it appears to be possible to match a significant subset of all human players to a bot with the same skill for a significant subset of all perfect information games.

Unfortunately, equal skill is not sufficient to declare two agents undistinguishable. There is also the matter of “style”⁷. One example for the effect of playing style comes from the game *Chess-Or-Coin* (CC), defined in Section 13.2.2. Depending on the “style” (or preferences) of player 1, a match of CC can have a *chanciness* of zero or a *chanciness* of one (for players who strictly prefer one of the move choices). For two agents with equal skill at chess, the critical move choice is completely independent of their skill.

Another example for agents distinguished by “style” is the “laziness” of *Monte Carlo Agents* in the game of *Go* [Alt 2008]. These agents can be readily distinguished from human players by playing unaggressive whenever they have an advantage.

We conjecture that finding a bot of similar “style” for a given human is much more difficult than matching skill. In fact, by modeling a natural conversation as a game of perfect information (each possible reply of at most C characters is a child move), we can translate the common version of the *Turing Test* into a comparison of style.

Nevertheless, we also conjecture that it will be possible to match humans and bots by “style” for commonly played games⁸. Thus, we think, the substitution approach will eventually be feasible.

⁶The notion of a “range” of skills is somewhat inexact since the relation “more skilled” need not be transitive. Nevertheless, it is common to measure skill on a linear scale (e.g. the *Elo Scale* [Elo 1986]).

⁷Other criteria for distinguishing agents could be provided by our measures of *chanciness* and *controllability*. However, this obviously does not help us to approximate these measures for human agents.

⁸We trust in the progress of computer science.

18 Approximation of *Chanciness* and *Controllability*

18.4.2 Avoiding Substitution

The substitution approach from the previous section is complicated because we need to match playing “styles”. We cannot approximate RPV by sampling and we cannot use SPV for most humans. However, we can still use the *subjective* approach and evaluate game states with an arbitrary EF. If the human players of a match can agree on the relevance of that EF, they can also agree on the *subjective chanciness* of their match, based on that EF.

Another issue is the *chanciness of a game*. For popular games there tend to exist tournament records of matches. Online gaming platforms such as *Little Golem* [Mal 2002] could also keep records of all matches played. Such samples of matches can be used, along with a *subjective* evaluation function, to compute the *subjective chanciness of a game*. Unfortunately, this does not help in the field of *computer-aided game inventing*, because newly invented games by definition do not have extensive human match records.

18.4.3 Partial Substitution

The substitution approach and the use of human match records can be combined. The issue of “style” makes the use of substitution difficult, because “style” affects the kind of matches that will be played. However, by using human match records the problem of “style” is eliminated. We can then substitute the human agents with bots of similar skill.

For bots of equal skill, the RPV will be identical to the RPV of the human lineup (since skill is defined by expected results). However, there is still an assumption involved: We usually measure the skill of agents by their expected results from the root position of a game. Thus, the accuracy of the substitution depends on the assumption, that equal playing strength at the root translates to equal playing strength for all other nodes.

We could of course select our substitute bots by comparing skill for a larger sample of nodes. However, since we cannot test all nodes, this assumption remains to a degree.

For a population of human players, for which we have match records, *partial substitution* allows the approximation of our measures *cha* and *CHA*.

18.5 Summary

The exact computation of *chanciness* and *controllability* depends on a detailed behavioral model for all agents in a population. Such models are usually not available. For populations of computer agents, one alternative lies in the use of sampling algorithms. With these sampling algorithms it is possible to approximate the *chanciness* of matches and the *chanciness of a game*.

The approximation of our measures for human players is more difficult. We propose that approximation may be accomplished by substituting appropriate computer agents.

Another possibility lies in the replacement of RPV with other evaluation functions. This leads to the concepts of *subjective chanciness* and *subjective controllability*. Those measure are applicable to human as well as computer agents.

19 Model Games

In this chapter we demonstrate our measures of *chanciness* and *controllability* by applying them to parameterized model games.

19.1 Sparrow and Pigeon

To illustrate how our measures *chanciness* and *controllability* can be used, we introduce the game class “Sparrow and Pigeon” ($SP[a, b]$). Figure 19.1 shows the parameterized game tree of this game class. We restrict the parameter b to the interval of the non-negative real numbers ($a \in \mathbb{R}, b \in \mathbb{R}_0^+$). The name of the game derives from the German proverb “A sparrow in the hand is worth more than a pigeon on the roof” (The analogue English proverb goes “A bird in hand is worth two in the bush.”). This name reflects the move choices in this single-player game. The player can either take a risk by selecting the left child node (hoping to catch the pigeon) or play it safe (take the sparrow) by selecting the right node.




Figure 19.1: The parameterized game tree of the $SP[a, b]$.

An agent X_p who plays the game SP can be characterized by its probability p to play the left node. We use the concept of *influence* and our measures of *chanciness* and *controllability* to analyze the game $SP[a, b]$ for different parameters of a and b when played by different agents (different parameters of p). As a simplifying assumption we consider the measures CHA, CON for a population $\Pi = \{X_p\}$ that consist of a single agent X_p , characterized by parameter p . The sampling distribution Ψ then by necessity gives probability 1 to the lineup (X_p) .

19 Model Games

Using Definition 13.2 we derive the values of ME , $\mathbb{E}(|IC|)$ and $\mathbb{E}(|IP_1|)$.

$$\begin{aligned} ME &= pa + (1-p)(a+1) = a+1-p \\ \mathbb{E}(|IC|) &= p(0.5|-b| + 0.5|b|) = pb \\ \mathbb{E}(|IP_1|) &= p(a-ME) + (1-p)(a+1-ME) = 2p(p-1) \\ \mathbb{E}(TI) &= pb + 2p(p-1) + |a+1-p| \end{aligned}$$

Using those terms we can express the *chanciness* and the *controllability* of SP when played by Π .

$$\begin{aligned} \text{CHA}(\text{SP}[a,b], \{X_p\}, \Psi) &= \frac{\mathbb{E}(|IC|)}{\mathbb{E}(TI)} = \frac{pb}{pb + 2p(p-1) + |a+1-p|} \\ \text{CON}(\text{SP}[a,b], \{X_p\}, \Psi) &= \frac{\mathbb{E}(|I_1|)}{\mathbb{E}(TI)} = \frac{2p(p-1)}{pb + 2p(p-1) + |a+1-p|} \end{aligned}$$

The measures *chanciness* and *controllability* are undefined for two regions in the parameter space: $\{(a,b,p) | a=0, b=0, p=1\}$ and $\{(a,b,p) | a=-1, b \in \mathbb{R}_0^+, p=0\}$. Both regions correspond to a deterministic agent who plays a deterministic game with reward zero. Outside these undefined regions we can find games with different characteristics:

- The *chanciness* of $\text{SP}[a,b]$ reaches its maximum of 1 in this regions: $\{(a,b,p) | a \in \mathbb{R}, b > 0, p=1\}$. This corresponds to a player who always selects the risky move with the risk being nonzero. In this case, SP is purely a game of chance.
- The *chanciness* is exactly zero in two regions: $\{(a,b,p) | (a \in \mathbb{R}, b=0, p \in [0,1]) \setminus \{(a,b,p) | (a=0, b=0, p=1)\} \text{ and } \{(a,b,p) | a \neq -1, b \in \mathbb{R}_0^+, p=0\}$. The first region corresponds to the case where the chance move has no influence, while the second region corresponds to the case where the player never takes a risk.
- The *controllability* of the game $\text{CON}(\text{SP}[a,b], \{X_p\}, \Psi)$ reaches a minimum of zero whenever the parameter p takes on one of the extreme values 0 or 1. This directly reflects Lemma 13.5. Agents X_0 and X_1 are completely deterministic and thus have influence zero.
- The *controllability* reaches its maximum at the point $(a = -0.5, b = 0, p = 0.5)$. At this point, the chance move has no influence and agent X_p exhibits uniform random behavior. While the minimum of *controllability* reflects a general rule, the observation for the maximum of *controllability* does not hold for all games. This can be seen with a slight modification of the SP game tree. The game tree in Figure 19.2 is derived from the game tree of SP by duplicating the right branch n times. An agent who plays uniformly random on this modified game experiences matches like the agent $X_{1/n}$ playing the regular game of SP. Thus, for increasing values of n the influence of the uniform random agent approaches zero and the *controllability* tends to zero as well.




Figure 19.2: A modified version of the SP game tree.

- In those parts of the parameter space where $|a| \gg b, |a| \gg 1$ we find games that have low values of *chanciness* as well as low values of *controllability*. These games are strongly influenced by the *match expectation* ME. Since the individual match results for these games are very predictable, this type of game is unsuited to recreational play. Why play at all if the outcome is (almost) certain? On the other hand, this type of game may represent the preferred scenario in the context of *stochastic control* when the goal is to maximize predictability.

19.2 PaCRaWa

The game PaCRaWa ($\text{PaCRaWa}[d, T, a]$) is a zero-sum game for two players, created by the author of this thesis. It was introduced in Part I to illustrate the measure of *chanciness*. The rules of PaCRaWa were defined in Section 7.1. In that section we also presented approximations of the *chanciness*, $C(\text{PCRW}[d, T, a], A_1, A_2)$ for different parts of the parameter space and for different lineups (A_1, A_2) . These approximations were obtained by the sampling algorithms from Section 18.1.

In the following Section 19.2.1 we present measurements of CHA and CON, applied to the same simulation data. In Section 19.2.2 we compare old and new measurements. In Section 19.2.3 we compare both sets of approximated results to analytic values.

19.2.1 Approximated CHA and CON for PaCRaWa

In the game of PaCRaWa, agents must repeatedly decide between one “good” move and one “bad” move. In Part I we used three agents who were each characterized by their probability of taking the “good” move. Agent *X* always selects the good move and represents deterministic optimal play. Agent *Y* plays the “good” move with probability 0.75 and agent *Z* selects both moves with equal probability 0.5. Agents *Y* and *Z* represent medium level and novice level of play, respectively.

19 Model Games

For our measurements of CHA and CON we will use the population $\Pi = (X, Y, Z)$. The lineups used in Part I will be represented by a special type of sampling probability Ψ . We define P_{XY} as the sampling probability which returns the lineup (X, Y) with probability one.

Figure 19.3 shows measurements of cha and CHA for PCRW[$d, T = 10, a = 0$]. Two different values of the “chance power” parameter d were used. Since PCRW is a two-player zero-sum game, both players experience the same amount of chance influence. Since in this case, $\text{CHA}_1 = \text{CHA}_2$ we give only one value (see Lemma G.7).

Just like in our previous PaCRaWa measurements, we see a positive correlation between *chanciness* and chance power. This is consistent with the intuition that an increased impact of coin flips on the payoff function increases the influence of chance. As in Part I, we also see that stronger agents experience more *chanciness*. This is understandable, when we recall that stronger agents also act more determined (see end of Section 13.3.4).

In contrast to Part I we can now appreciate the different perspectives of the players. In all instances of our experiment, the “master” agent X has zero *controllability*. His behavior is completely determined, and therefore his result is only influenced by chance and by his opponent. Weaker agents generally have more influential moves and thus receive higher *controllability* values. A notable exception can be found by comparing P_{XY} and P_{XZ} . In these lineups, the stronger agent Y has higher *controllability* than the random agent Z . This is caused by the different *match expectation* of both lineups. The weaker agent Z is sure to receive a low result against the master X . This in turn diminishes the *controllability* from the perspective of Z compared to Y . The advantage of X over Z also results in the low *chanciness* values for that lineup.

d	0.1			1.0		
Ψ	CHA	CON ₁	CON ₂	CHA	CON ₁	CON ₂
P_{XX}	0.98	0.00	0.00	0.98	0.00	0.00
P_{YY}	0.08	0.46	0.44	0.43	0.28	0.28
P_{ZZ}	0.07	0.45	0.47	0.39	0.30	0.31
P_{XY}	0.06	0.00	0.37	0.37	0.00	0.24
P_{YZ}	0.05	0.26	0.30	0.29	0.20	0.21
P_{XZ}	0.04	0.00	0.26	0.27	0.00	0.20

Figure 19.3: *chanciness* and *controllability* of PCRW[$d, T = 10, a = 0$] for different chance powers d and different agents (1000 matches sampled, approximation error ≈ 0.01).

19.2.2 Comparison of C and CHA

The definition of the *chanciness* C in Chapter 6 (Definition 6.1) is conceptually similar to the definition of the *chanciness* CHA in Chapter 14 (Definition 14.7). CHA is more general because it recognizes different player perspectives and is defined for a population rather than for a single lineup. Another difference between both definitions lies in the arrangement of the *sources of influence* (SI). The definition of C puts all player influences into a single SI. The definition of CHA, on the other hand, recognizes each player as a separate SI. Since SI are compared according to their absolute influence values, this results in a different denominator between both definitions.

Figure 19.4 shows a side by side comparisons of approximated values for C and CHA in the game of PCRW[$d, T = 10, a = 0$]. Approximation errors aside, the value of C is generally lower than the corresponding value of CHA. This is to be expected, due to the *triangle inequality* $|\text{IP}_1 + \text{IP}_2| \leq |\text{IP}_1| + |\text{IP}_2|$. The lower denominator for C results in a higher value of C compared to CHA. Since the master agent X has zero influence, the values of C and CHA are equal wherever he participates.

d	0.1		0.5		1.0	
Ψ	C	CHA	C	CHA	C	CHA
P_{XX}	0.97	0.98	0.98	0.98	0.98	0.98
P_{YY}	0.11	0.08	0.37	0.29	0.53	0.43
P_{ZZ}	0.09	0.07	0.33	0.24	0.47	0.39
P_{XY}	0.06	0.06	0.24	0.24	0.37	0.37
P_{YZ}	0.05	0.05	0.21	0.18	0.34	0.29
P_{XZ}	0.04	0.04	0.16	0.16	0.27	0.27

Figure 19.4: A comparison of *chanciness* measures C and CHA for the game of PCRW[$d, T = 10, a = 0$]. Both measures were computed from the same simulation data. (1000 matches sampled, approximation error ≈ 0.01).

Although the difference between the *chanciness* measures of Part I and Part II is notable, the relations among the different data points were preserved.

19.2.3 Analytic results for the two *chanciness* measures C and CHA

The *chanciness* of PaCRaWa according to both definitions was analyzed by Riedel [Rie 2010]. Figure 19.5 shows a side by side comparison of our approximative results for C and the analytic results by Riedel. Riedel's results confirm the accuracy of our approximated values.

Figure 19.6 shows a side by side comparison of our approximative results for the new measure CHA and the analytic results by Riedel. Again, the accuracy of our approxima-

19 Model Games

d	0.1		0.5		1.0	
Ψ	A	R	A	R	A	R
P_{XX}	0.97	1.00	0.98	1.00	0.98	1.00
P_{YY}	0.11	0.11	0.37	0.37	0.53	0.54
P_{ZZ}	0.09	0.09	0.33	0.33	0.47	0.50
P_{XY}	0.06	0.06	0.24	0.23	0.37	0.38
P_{YZ}	0.05	0.05	0.21	0.20	0.34	0.33
P_{XZ}	0.04	0.03	0.16	0.15	0.27	0.26

Figure 19.5: A comparison of approximated and analytic results for the *chanciness* measure C. Columns marked with “A” contain our approximations, obtained by sampling. Columns marked with “R” contain the analytic results of Riedel.

tions is confirmed by Riedel’s analysis.

d	0.1		0.5		1.0	
Ψ	A	R	A	R	A	R
P_{XX}	0.98	1.00	0.98	1.00	0.98	1.00
P_{YY}	0.08	0.07	0.29	0.28	0.43	0.44
P_{ZZ}	0.07	0.06	0.24	0.25	0.39	0.40
P_{XY}	0.06	0.06	0.24	0.23	0.37	0.38
P_{YZ}	0.05	0.04	0.18	0.17	0.29	0.29
P_{XZ}	0.04	0.03	0.16	0.15	0.27	0.26

Figure 19.6: A comparison of approximated and analytic results for the *chanciness* measure CHA. Columns marked with “A” contain our approximations, obtained by sampling. Columns marked with “R” contain the analytic results of Riedel.

20 “EinStein würfelt nicht!” and Variants Thereof

“EinStein würfelt nicht!” is a board game designed by Ingo Althöfer [Alt 2004a]. The game can be played at the online platform *Little Golem* [Mal 2002] where it is quite popular. The German title of the game is a play on words. It can be translated either as “A single stone does not play dice” or “Einstein does not play dice”. This play on words alludes to the game’s rules and is also a reference to Albert Einstein’s famous quotation about God and dice. In this chapter we present approximate measurements of *chanciness* and *controllability* for the original game and for several games variants.

20.1 Experimental Setup

20.1.1 Agents used in the Simulation

The results presented in this chapter were obtained by the sampling approach discussed in Section 18.1. This was possible because we used populations which consist entirely of computer agents.

For our experiments we have used agents that employ the *Monte Carlo* algorithm to select their moves [Abr 1990, BH 2003]. Game states are evaluated by the mean result of k random matches. The child move with the maximum value is then played. We denote these *Monte Carlo Agents* as MC_k , with k indicating the number of random matches. We have also used the random agent, RND who selects among all possible moves with equal probability. The playing strength of *Monte Carlo Agents* in the game of EWN grows with k as shown in Figure 20.1.

One notable feature of *Monte Carlo Agents* is their independence from a particular game’s rules. Such agents, if suitably programmed, can play a game based only on a description of the game’s rules and without knowledge about strategies. This makes them particularly suited in the field of *computer-aided game inventing* [Alt 2003] where such knowledge is not available for newly invented games.

20.1.2 Approximation of RPV

As explained in Chapter 18.1.1 we approximate the *relative positional values* (RPV) by sampling *partial matches*. This sampling is sped up by using the *Lookup Sampling*-algorithm (Appendix B). For every *relative positional value* needed, we conducted 1,000 *partial matches*¹.

¹Some of the *partial matches* were stopped early by the *Lookup Sampling*-algorithm.

20 “EinStein würfelt nicht!” and Variants Thereof




Figure 20.1: Win rates in EWN of Monte Carlo Agents with different strengths versus the random agent RND.

20.1.3 Sampling of Matches

For a game of n players and a population Π , the set of all possible lineups is given by the set Π^n with cardinality $|\Pi|^n$. When not indicated otherwise, we have assumed an equal sampling probability for all possible lineups.

For the population $\Pi = \{\text{RND}, \text{MC}_1, \text{MC}_8, \text{MC}_{64}\}$ and a game of 2 players, we have 16 possible lineups. In this case our sampling probability Ψ would therefore be the constant function with value $\frac{1}{16}$.

In order to approximate the expected absolute influence of the different *sources of influences*, matches must be sampled from the set of all possible matches. To this end we have simulated matches repeatedly for every possible lineup. To reflect our assumption of equal probability, every lineup was repeated the same number of times. The added number of simulations matches for the individual lineups yields the total number of matches (*#Matches*).

In order to conserve computing resources, the number of simulation matches had to be limited. Whenever we present our simulation results, we will give the total number of matches used (*#Matches*) and the estimated error (*err*) of our approximated measures. The definition of our error estimate *err* for the *chanciness of a game* and the *controllability of a game* is given in Appendix H. Wherever we only give one error value for CHA and CON, it is the maximum of both values.

20.2 “EinStein würfelt nicht!” (EWN)

20.2.1 Game Rules

The original game EWN is a two-player zero-sum game with perfect information. It is played on rectangular board with 5 rows and 5 columns. The two players “White” (player 1) and “Black” (player 2) seek to move their playing pieces towards a goal square or to eliminate their opponent.

Each player has six playing pieces, numbered one to six. Figure 20.2 shows a starting position for EWN. The 6 grid positions (squares) in which the pieces of each player are initially placed are fixed, while the order of the pieces is not fixed. One possibility is to leave the order of the pieces to the players discretion. Another possibility is to randomize the order of the pieces. This randomized approach is taken at the online platform *Little Golem* and was also used for our simulation matches.




Figure 20.2: A starting position for EWN.

The players take turn rolling a six-sided die and then moving the piece with the same number as the die result. Player White, who starts in the top-left corner, may move his piece to the adjacent right square, to the adjacent lower square or to the adjacent diagonal lower right square. Only moves that stay within the 5x5 grid are permitted. If the target square of a move is occupied by a piece of either player, that piece is removed from play. Player White wins the game if he moves a piece into the bottom-right corner square or if all black pieces are removed from play.

Player Black, who starts in the bottom-right corner, has mirrored move options. His pieces may move up, left, or up-left to an adjacent square. Black wins by moving a piece into the top-left corner square or if all white pieces are removed from play.

If the result of the die indicates a piece that was already removed from play, the player must choose another piece to make his move. If the die result was the number x , he may choose either the piece with the minimum number higher than x or the piece with the maximum number lower than x . If only one of those pieces remains, there is no choice. Figure 20.3 shows a situation in which Black can choose among two different pieces for his moves. Since the black pieces numbered 2 and 3 were removed, Black may move either with piece 1 or piece 4 on a die result of 3.

20 “EinStein würfelt nicht!” and Variants Thereof




Figure 20.3: Move options of Black for the die result 3. Every arrow corresponds to a legal move.

If a player has only a single piece left, that piece is moved every turn. It is then unnecessary to roll the die, hence the name of the game “A single stone does not play dice”. The winning player of EWN is rewarded with a score of +1 while the losing player receives a score of −1. The game EWN does not allow for a draw.

20.2.2 Analysis of an Example Match

The following match $g = (L, P)$ was played by a lineup of *Monte Carlo Agents*, $L = (\text{MC}_{64}, \text{MC}_8)$. Agent MC_{64} played as White (player 1) and agent MC_8 played as Black (player 2). During this match the states in the sequence $P = (s_0, s_1, \dots, s_{15})$ were visited.

Figure 20.4 shows the sequence of states in g from left to right and from top to bottom. The first move is used to determine the initial placement of the pieces. Afterwards, chance moves and player moves alternate. The game ends after 8 chance moves and 7 player moves with a win for White.

Figure 20.5 gives the approximated RPV for the 16 game states from the perspectives of either player. Since EWN is a zero-sum game, the values for player 1 and 2 differ only in their leading sign. The standard error of these values is ≈ 0.01 (sample standard deviation of ≈ 0.03 , sample size of 1,000).

At first glance the table of RPV reveals that player White has quite an advantage right from the beginning. The positive influence of the first chance move (+0.02) suggests that White may have been further advantaged by getting an above-average placement of his pieces.

For further analysis of g we compute the influence of the *SI* (sources of influence, Page 51):

$$\begin{aligned} \text{ME}_1(g) &= 0.38, & \text{ME}_2(g) &= -0.38 \\ \text{IC}_1(g) &= 0.36, & \text{IC}_2(g) &= -0.36 \\ \text{IP}_{1,1}(g) &= 0.06, & \text{IP}_{1,2}(g) &= -0.06 \\ \text{IP}_{2,1}(g) &= 0.20, & \text{IP}_{2,2}(g) &= -0.20 \end{aligned}$$

20.2 “EinStein würfelt nicht!” (EWN)




Figure 20.4: All game states of match g from left to right and from top to bottom. States that require a chance move are marked with the symbol “?”.

20 “EinStein würfelt nicht!” and Variants Thereof

node-type	index i	$\tilde{V}_1(s_i, L)$	$\tilde{V}_2(s_i, L)$
chance	00	0.38	- 0.38
chance	01	0.40	- 0.40
player1	02	0.37	- 0.37
chance	03	0.40	- 0.40
player2	04	0.46	- 0.46
chance	05	0.55	- 0.55
player1	06	0.49	- 0.49
chance	07	0.50	- 0.50
player2	08	0.59	- 0.59
chance	09	0.60	- 0.60
player1	10	0.56	- 0.56
chance	11	0.58	- 0.58
player2	12	0.64	- 0.64
chance	13	0.74	- 0.74
player1	14	1.00	- 1.00
terminal	15	1.00	- 1.00

Figure 20.5: Approximated RPV for the match $g = (L, P), L = (\text{MC}_{64}, \text{MC}_8)$.

We see that the *match expectation* along with the influence of chance were the dominating influences on this match. The victory of White was also aided by tactical errors of Black. This assessment is reflected in the values of *chanciness* and *controllability*:

$$\begin{aligned}
 \text{cha}_1(g) = \text{cha}_2(g) &= 0.36 \\
 \text{con}_1(g) &= 0.06 \\
 \text{con}_2(g) &= 0.20
 \end{aligned}$$

From the perspective of White the match appears rather boring (if we put ourselves in the shoes of the computer agents). The low *controllability* indicates that White’s overall influence was low. That means he neither performed far below or above his usual skill.

On the other hand, the moves of player Black were more consequential, as shown by the higher *controllability* value from his perspective. In retrospect, this player should find the match interesting because some of his moves were worse than usual.

Either way, the influence of chance was more influential than the moves of either player. This can be seen by comparing the *chanciness* value to either *controllability* value.

20.2.3 Analysis of the Game

In this section we present our approximations of *chanciness* and *controllability* for the game EWN, using different populations. Figure 20.6 shows the values of CHA , CON_1 and CON_2 .

20.2 “EinStein würfelt nicht!” (EWN)

Since EWN is a two-player zero-sum game, the *chanciness* of both player perspectives is equal. For every population Π with cardinality $|\Pi| = k$ we have used the constant sampling probability $\Psi(L) \equiv k^{-2}$.

Π	CHA	CON ₁	CON ₂	#Matches	err
{RND}	0.43	0.26	0.26	128	0.022
{MC ₁ }	0.48	0.22	0.26	128	0.021
{MC ₈ }	0.54	0.20	0.21	128	0.019
{MC ₆₄ }	0.66	0.15	0.12	128	0.013
{MC ₅₁₂ }	0.77	0.08	0.08	100	0.016
{MC ₈ , MC ₆₄ }	0.54	0.16	0.16	512	0.009
{MC ₁ , MC ₈ , MC ₆₄ }	0.45	0.17	0.17	1152	0.006
{RND, MC ₁ , MC ₈ , MC ₆₄ }	0.39	0.17	0.17	2048	0.005
{MC ₁ , MC ₅₁₂ }	0.27	0.12	0.10	200	0.014

Figure 20.6: Measurements of CHA and CON for the game EWN using different populations Π . All lineups from $\Pi \times \Pi$ were sampled the same number of times. Since for this game CHA₁ equals CHA₂, only one value is given.

Our results reveal that both players roles have very similar *controllability*. While the equal *chanciness* is a direct result of the zero-sum payoff structure, equal *controllability* for both perspectives cannot generally be expected. In this case, it tells us that both player perspectives are balanced in respect to their self-influence². From the standpoint of a game inventor this is a sign of quality.

We also note the trend of increasing *chanciness* for populations of higher skill (represented by populations of a single agent). Just as the *chanciness* grows, the *controllability* of both perspectives is equally diminished. In the lower half of Figure 20.6 we see the trend of decreasing *chanciness* along with the increasing spread in population skill (spread of agent skill within the population). This decrease is not accompanied by a raise in *controllability*. We can explain this trend by the increased skill range within the different populations. This results in matches where the result is mainly influenced by the *match expectation*. While such unequal matches are not more controllable, they are much less subject to the influence of chance.

An important question not answered by our measurements relates to the interpretation of our results for human agents. Are any of the tested populations representative for a given human population? We suspect that the $\Pi = \{\text{RND}, \text{MC}_1, \text{MC}_8, \text{MC}_{64}\}$ could at most represent a population of novice players, while the population $\Pi = \{\text{MC}_{512}\}$ better represents average human skill. Due to resource constraints we did not test *Monte Carlo Agents* of higher playing strength. The required simulation matches for Figure 20.6 took

²And since EWN is a zero-sum game, the absolute influence on one’s own results is equal to the influence on the opponent’s result.

20 “EinStein würfelt nicht!” and Variants Thereof

approximately 400 CPU hours on 1GHz Opteron processors.

20.3 The Variant $\overline{\text{EWN}}$

20.3.1 Rules of $\overline{\text{EWN}}$

The game $\overline{\text{EWN}}$ is a close variant of the original game EWN. It was already introduced in Chapter 7 where it was used to measure the effect of reduced move choices on *chanciness*. The only rule change pertains to the selection of pieces whenever the die indicates a piece, that has been removed from play. In $\overline{\text{EWN}}$ the player always has to move the piece with the minimum number higher than the die result. If no such piece exists, the piece with the minimum number must be used. Because of this rule change, some situations in $\overline{\text{EWN}}$ offer fewer move choices than they would in EWN.

20.3.2 $\overline{\text{EWN}}$ Compared to EWN

Π	CHA		CON_1		CON_2	
	EWN	$\overline{\text{EWN}}$	EWN	$\overline{\text{EWN}}$	EWN	$\overline{\text{EWN}}$
{RND}	0.43	0.44	0.26	0.24	0.26	0.27
{MC ₁ }	0.48	0.52	0.22	0.23	0.26	0.22
{MC ₈ }	0.54	0.57	0.20	0.19	0.21	0.19
{MC ₆₄ }	0.66	0.70	0.15	0.12	0.12	0.12
{MC ₅₁₂ }	0.77	0.80	0.08	0.07	0.08	0.06
{MC ₈ , MC ₆₄ }	0.54	0.58	0.16	0.15	0.16	0.14
{MC ₁ , MC ₈ , MC ₆₄ }	0.45	0.49	0.17	0.16	0.17	0.16
{RND, MC ₁ , MC ₈ , MC ₆₄ }	0.39	0.45	0.17	0.16	0.17	0.16
{MC ₁ , MC ₅₁₂ }	0.27	0.35	0.12	0.10	0.10	0.10

Figure 20.7: Comparison between the games EWN and $\overline{\text{EWN}}$ for the measures CHA, CON_1 and CON_2 . The same lineups and number of matches were used as for the measurements from Figure 20.6. The estimated error was roughly the same as the estimate from Figure 20.6.

To compare the game variant $\overline{\text{EWN}}$ to the original game EWN we have repeated the measurements from Section 20.2.3 for $\overline{\text{EWN}}$. We used the same lineups and number of matches as reported in Figure 20.6. The error estimates for these new simulations were roughly the same as those reported before. Figure 20.7 shows a side-by-side comparison of our results.

The most obvious feature of Figure 20.7 are the consistently higher *chanciness* values for $\overline{\text{EWN}}$ compared to EWN. At the same time each player perspective generally experiences lower *controllability* (the exceptional data points can be attributed to the approximation error). The change in *controllability* between EWN and $\overline{\text{EWN}}$ appears less pronounced

20.4 “EinStein würfelt nicht Solo” (EWNS)

that the change in *chanciness*. One explanation for this phenomenon comes from the fact that the added influence-loss of *both* players is expressed in the higher *chanciness* value.

All in all, the results from Figure 20.7 confirm our findings from Part I of this thesis: Reducing the choice of moves in the game of EWN results in a game which is more dependent on chance.

20.4 “EinStein würfelt nicht Solo” (EWNS)

Ingo Althöfer, the author of EWN, has suggested a single-player version of his game. We present this game variant under the name “EinStein würfelt nicht Solo” (EWNS).

20.4.1 Rules of EWNS_k

In EWNS a single player attempts to move one of his pieces to the goal position in the far corner of the board as fast as possible. If the player succeeds in k or less moves he wins, otherwise the player loses. In regard to the parameter k , EWNS can be considered a set of games with the elements EWNS_k , $k \in \mathbb{N}$. If the player wins he is awarded with a score of $+1$, otherwise he receives a score of -1 . Figure 20.8 shows the initial board position that was used during all simulation matches (the order of the pieces was not randomized).

5	1	3		
2	6			
4				

Figure 20.8: The starting position for EWNS. All matches used the same layout of the pieces.

20.4.2 Analysis

In this section we present our approximations of *chanciness* and *controllability* for the game “EinStein würfelt nicht Solo” (EWNS_k), using different populations. We have also tested different values for the rule parameter k . Figure 20.9 shows the results of our measurements. A column-wise comparison reveals the strong dependence of *chanciness* and *controllability* on the parameter k . This should come as no surprise, since the game is extraordinarily hard if only 5 moves are allowed. On the other hand, the game becomes very easy if 15 moves are allowed to complete the objective.

20 “EinStein würfelt nicht!” and Variants Thereof

Matches which are almost invariably lost allow for little influence on the match result. Thus we see low values of CHA and CON in the first two columns of Figure 20.9. Nevertheless, we see that there is some influence, based on the skill of the playing agent. At the same time that the *controllability* rises for stronger agents, the *chanciness* rises as well. This is understandable, because chance can only affect the likelihood of winning if a win is not entirely impossible.

The situation is somewhat reversed in the rightmost columns. Here, the matches are mostly won, which leads to low *chanciness* and low *controllability*. Only the more incompetent agents actually manage to lose matches with parameter $k = 15$. Thus we see higher *controllability* only for agents RND and MC₁. If matches can actually be lost, then the influence of chance is more relevant as well.

k	5		7		10		12		15	
Π	CHA	CON	CHA	CON	CHA	CON	CHA	CON	CHA	CON
{RND}	0.04	0.02	0.12	0.10	0.36	0.25	0.50	0.38	0.36	0.30
{MC ₁ }	0.06	0.05	0.22	0.17	0.59	0.28	0.51	0.26	0.29	0.16
{MC ₈ }	0.10	0.04	0.38	0.18	0.52	0.23	0.27	0.16	0.07	0.04
{MC ₆₄ }	0.16	0.08	0.64	0.13	0.43	0.10	0.12	0.05	0.01	0.01

Figure 20.9: *Chanciness* and *Controllability* for the game EWNS _{k} , each lineup was sampled 100 times. The estimated error is ≈ 0.02 .

Most interesting are the middle columns of Figure 20.9. Here we notice peaks in *chanciness* and *controllability*. Within the table, the highest values of CHA and CON for agent MC₆₄ are recorded for $k = 7$. For weaker agents the peaks occur for increasingly higher values of k . Away from the peak value of k , the game is highly determined by the *match expectation*. When comparing the proportion of CHA and CON near their peaks among the different agents, we notice that agent MC₆₄ is on the verge of mastering the game. The *controllability* for that agent drops remarkably in comparison to the peak value for MC₈.

Based on our limited selection of parameters, the game EWNS _{k} appears most “exciting” for the agent RND given a value of $k = 12$. This is a point where the *match expectation* is close to zero³. Yet, even for this agent the *chanciness* is still higher than the *controllability* of EWNS₁₂. For agents of increasingly higher skill, the parameter k , where losing and winning are equally likely, will shift downwards. By necessity, a point will be reached where agents have mastered the game (zero *controllability*) and EWNS _{k} is completely dominated by chance.

³Since match results are only influenced by the SI, ME, IC, IP₁, we can calculate the relative expected influence of the *match expectation* as $1 - 0.50 - 0.38 = 0.12$.

20.5 “EinStein würfelt nicht Tetra” (EWNT)

“EinStein würfelt nicht Tetra” is a four player variant of EWN, also suggested by the original author of EWN, Ingo Althöfer.

20.5.1 Rules of EWNT

EWNT is a general-sum game played by four players. The initial position for this game is shown in Figure 20.10. The movement of the pieces for the top-left player (White) and the bottom-right player (Black) is the same as in EWN. The players that start in the top-right corner (Gray) and in the bottom-left corner (Ring) have the move options (down, left, down-left) and (up, right, up-right) respectively. The players take turn in clockwise order beginning with White (White, Gray, Black, Ring). As in the game of EWN, chance moves and player moves alternate.




Figure 20.10: The starting position for EWNT. All matches used the same layout of the pieces.

Each player seeks to be the first to move a piece towards the corner square opposite of his starting position. The first player to succeed ends the game. Should the pieces of all but one player be eliminated, the game ends as well.

A player who manages to reach the opposite corner receives a reward of +1. A player who has lost all pieces receives a result of -1. A player who keeps at least one piece but does not reach his goal receives a result of 0. A player who manages to eliminating all opponent pieces also receives a reward of +1. The number of players eliminated in any match ranges from zero to three. The sum of all player rewards is therefore in the range of $[-2, 1]$. Accordingly, the game EWNT is a non-zero-sum game.

EWNT bears strong similarities to the four player variant “EinStein würfelt nicht Quattro” [Alt 2004b] (EWNQ), also by Ingo Althöfer. However, EWNQ puts the players in two teams of two players each. Due to this team setup, EWNQ could actually be analyzed as a two-player zero-sum game.

20 “EinStein würfelt nicht!” and Variants Thereof

20.5.2 Analysis

In this section we present our approximations of *chanciness* and *controllability* for the game “EinStein würfelt nicht Tetra” (EWNT), using different populations. Since EWNT is a game of four players, we distinguish four values of CHA_i and four values of CON_i .

Π	CHA_1	CHA_2	CHA_3	CHA_4	#Matches	err
{RND}	0.20	0.21	0.19	0.23	100	0.018
{MC ₁ }	0.25	0.22	0.23	0.26	100	0.020
{MC ₈ }	0.31	0.31	0.33	0.34	100	0.023
{MC ₆₄ }	0.44	0.43	0.40	0.43	100	0.025
{MC ₅₁₂ }	0.52	0.52	0.50	0.45	100	0.026
{MC ₈ , MC ₆₄ }	0.35	0.35	0.35	0.36	256	0.015
{MC ₁ , MC ₈ , MC ₆₄ }	0.30	0.30	0.30	0.31	1296	0.006
{RND, MC ₁ , MC ₈ , MC ₆₄ }	0.27	0.27	0.27	0.27	4096	0.003
{MC ₁ , MC ₅₁₂ }	0.31	0.33	0.32	0.33	512	0.011

Figure 20.11: Measurements of CHA_1 for the game EWNT using different populations Π . All lineups from $\Pi \times \Pi$ were sampled the same number of times.

Figure 20.11 shows our measurements for *chanciness* together with the number of simulation matches and our estimated error. The player roles 1 to 4 correspond to the previous designation of (White, Gray, Black, Ring). The simulation matches, on which these values are based, required approximately 1600 CPU hours of computation time on 1GHz Opteron processors. In Appendix I we present additional result statistics for these simulations.

A column-wise comparison shows about equal *chanciness* for all four player perspectives. This is particularly obvious for the two bottom rows, where the estimated error is low. Similar to our observations for the game of EWN, we notice increasing *chanciness* as the skill level of the population grows (rows 1 to 4). However, the overall level of *chanciness* is lower by about 20% compared ⁴ to the same populations playing EWN (see Figure 20.6). This can be understood by considering that the influence of *other players* is much greater in EWNT (there are three *other players* instead of just one).

In the lower half of Figure 20.11 we see the familiar trend of decreasing *chanciness* along with the increasing spread in population skill. Again, we can explain this trend by the increased number of matches between agents of highly unequal skill. This leads to matches where the result is more strongly influenced by the *match expectation*. However, for the population $\Pi = \{MC_1, MC_{512}\}$ the trend is broken. This might be explained by the relatively “deterministic” behaviour of agent MC_{512} which lowers the overall level of player influence and thus increases the importance of chance. That this agent should act far more determined than the other agents can be inferred from the low *chanciness* values in the following Figure 20.12.

⁴Since *chanciness* is a relative measure on a fixed scale of $[0, 1]$ it makes sense to compare values in terms of percentages.

Π	CON ₁	CON ₂	CON ₃	CON ₄	#Matches	err
{RND}	0.20	0.18	0.21	0.17	100	0.018
{MC ₁ }	0.16	0.17	0.22	0.19	100	0.020
{MC ₈ }	0.17	0.15	0.15	0.16	100	0.015
{MC ₆₄ }	0.12	0.14	0.14	0.12	100	0.014
{MC ₅₁₂ }	0.12	0.11	0.11	0.11	100	0.015
{MC ₈ , MC ₆₄ }	0.15	0.16	0.14	0.14	256	0.010
{MC ₁ , MC ₈ , MC ₆₄ }	0.16	0.16	0.16	0.15	1296	0.004
{RND, MC ₁ , MC ₈ , MC ₆₄ }	0.16	0.16	0.16	0.16	4096	0.002
{MC ₁ , MC ₅₁₂ }	0.14	0.14	0.14	0.14	512	0.007

Figure 20.12: Measurements of CON₁ for the game EWNT using different populations Π . All lineups from $\Pi \times \Pi$ were sampled the same number of times.

Figure 20.12 shows our measurements for the *controllability* of EWNT, based on the same simulation matches as Figure 20.11. Again, a column-wise comparison shows about equal *controllability* for all four player perspectives. For populations of increasing skill level we also notice a slight trend for decreasing *controllability*. This is less pronounced than for the game of EWN. Presumably, because the *controllability* levels of EWNT are lower overall. For populations with different spread in agent skills we see only slight changes in *controllability*. We take this as an indication that the skill differences within our populations are not very important compared to the other influences⁵. Overall we notice that the game EWNT has rather lower *controllability*. While the *chanciness* of this game is not high either, prospective players should be aware that chance has significantly more influence on their results than their own actions.

20.6 Summary

In this chapter we have presented approximate measurements of *chanciness* and *controllability* for games from the “EinStein würfelt nicht!”-family. All these games require the frequent use of dice rolls during their matches. Our measurements have shown that the influence of these dice rolls makes a major contribution to match results. However, the relative importance of chance also depends on the populations who play these games. The results in this chapter also demonstrate that our measures can be effectively approximated using populations of computer agents.

⁵See the analysis of the *match expectation* in Section 16.2.

21 Conclusion

21.1 Results

This thesis aimed to characterize the influence of chance and skill on games. We were able to provide such a characterization based on two major assumptions:

- Influence is a property of game moves.
- Influence is the change in value between game states.

Using these assumptions we were able to quantify and compare different influences on game results. We have defined the measures of *chanciness* and *controllability* for general-sum games of perfect information with an arbitrary number of players. Both measures are defined in reference to a population of agents. A high value of *chanciness* indicates a game of chance while a high value of *controllability* indicates a game of skill. Our measures reveal the following important aspects of chance and skill in games:

- The character of a game depends on the features of the game tree.
- The character of a game depends on the behavior of the players.
- The character of an n -player game may differ among the n parties.
- There exist games that neither qualify as game of chance nor as game of skill.

The exact computation of *chanciness* and *controllability* is typically constrained by incomplete knowledge of agent behavior. The computation is unfeasible for games with deep trees because depth contributes exponentially to the cost of computation. However, we have shown that our measures can be approximated with reasonable accuracy using a polynomial time algorithm¹. This algorithm is based on sampling matches between computer agents. Additional effort is required to select agents with behavior representative of a human population.

21.2 Directions for Further Research

21.2.1 Empirical Studies

The measures developed in this thesis allow to make predictions about the relative characteristics of different games such as: game G is more strongly influenced by chance than

¹Computational cost is polynomial in the depth of the game tree.

game H . These predictions should be tested against comparisons undertaken by human game testers.

It seems to us that the concepts of *subjective chanciness* and *subjective controllability* (see Section 18.2) are a particularly promising area for empirical studies. These measures have the potential of being very fast approximations for *chanciness* and *controllability*.

21.2.2 Diplomatic Interactions

A crucial limitation of our game model can be found in its ignorance of communication and diplomacy. In games that allow for cooperation and alliance-building, communication is a major tool for influencing match results. When there is only one player, the effect of communication can obviously be ignored. Likewise, two-player zero-sum games do not allow for cooperation. Hence, it is justified to model such games without communication. However, as soon as one enters the realm of two-player general-sum games and multi-player games, communication can hardly be ignored.

We note that the effectiveness of communication depends on psychological factors, which are hard to grasp, as well as the tactical options of a game. If a game allows for obviously punishing or rewarding moves, they can be used to reinforce the effectiveness of verbal threats and promises.

Psychological agent models are probably out of reach in the foreseeable future. To fully account for the influence of communication, one would have to define a set of communication primitives. These basic elements of communication could be treated as a new kind of “move” available to each player. It appears even more difficult to us to account for the “internal state” of the participating agents which would be affected by those communication-moves.

We propose that properties of the bare game tree can be used to compute limits on the effectiveness of diplomacy. Diplomacy is concerned with the give and take of in-game advantages. Consequently, if a player i wishes to make demands towards another player j , he must be able to influence the results of player j . At the same time he can only gain from diplomacy if that other player is able to influence the results from the perspective of player i .

Such influence values may be obtained using the definition of influence proposed in this paper. A table such as the one in Figure 15.2 might then be the preferred tool for describing a game. This table holds the influence of every SI (*match expectation*, *chance* and *each player*) on every player.

21.2.3 Games of Imperfect Information

Our measures of *chanciness* and *controllability* cannot be applied to games with imperfect information. These games confront players with a novel type of unpredictability. Players do not always know the exact game state they are in. Instead they have to plan their moves based on a probability distribution over a set of states they might be in (this is called an information set [SLB 2008, 164]). During a match, information that is deduced

21 Conclusion

or directly uncovered may lead to a change in that probability distribution. From the perspective of a player such a change in probability distribution is very similar to the effect of a chance move. To describe the *chanciness* of a game of imperfect information would therefore require to account for the influence of such information changes. We see a major problem in modeling the way agents deduce information based on the behavior of their opponents. Since the opponent's actions may also be based on opponent models this potentially leads to an infinite regression of models about models about models.

The work of Borm, Dreef and Genugten (BDG) avoids this problem by focussing solely on the final payoffs of their three agent types. Dreef concludes that *Poker* is a game of skill: "Comparing the skill level, 0.3997, to the highest bound between games of chance and games of skill that was advised by Van der Genugten (1997), 0.15, leads undisputedly to the classification of the poker game with raising as a game of skill." [Dre 2005]. It would be desirable to extend our concepts to games of imperfect information and to compare results with those of BDG in that area.

21.2.4 Internal Chance

Our agent model does not discuss the internal mechanism by which agents select their moves. It simply assumes that their behavior can be described stochastically. This does not allow us to distinguish between fluctuations of agent skill and deliberate randomization. Deliberate randomization is called *internal chance* by BDG. *Internal chance* is encountered in mixed strategies employed to play games of imperfect information. However, it may also occur in multi-player games of perfect information because an agent is unable to discriminate among his move options. Randomization would then be a way to avoid systematic errors. Yet another reason to randomize moves could be found in diplomacy. Faced with the option of willfully harming one player or another, the active player may instead randomize his choice to absolve himself of diplomatic responsibility and avoid making an enemy. One could argue that such intentional randomization, (if openly acknowledged) contributes to the influence of chance and should therefore increase the *chanciness*.

Nevertheless, within our framework such randomization would lead to an increased influence of *other players* upon the *affected player* and thus decrease the *controllability* of the *affected player*.

21.2.5 Lawmaking

If the *chanciness* of a game (given the relevant population) is high, it should be called a game of chance. Deciding on a threshold to separate games in two classes is difficult, as it will lead to almost-similar games being treated completely different. One approach suggested by BDG [BvdG 2001] is to use games that are classified without or with little dispute as markers and classify other games in relation to these markers. This approach applies to our measures as well: If *Backgammon* is allowed as a game of skill and some other game *G* exhibits less *chanciness* than *Backgammon*, then *G* should not be classified as a game of chance.

The question of *controllability* is another matter: A game may have at the same time low *chanciness* and low *controllability*. One class of examples would be unfair games: In the German State Lottery, *Süddeutsche Klassenlotterie*, the payout ratio is about 50% [Nie 2008]. A player who is committed to bet a fixed amount faces an expected loss of half his stake. The player's choices for filling in his lottery ticket are independent of the process which selects lottery winners. Thus, we can assume a *controllability* of zero (the *controllability* may be slightly above zero when taking into account strategies which avoid shared winnings). Together with the high expectation of a loss, this results in a *chanciness* of ≈ 0.5 . Even though the lottery is less influenced by chance than a single pass at the *Roulette* table, we do not suggest, that this makes it more desirable from a societal standpoint. We would rather call into question the general concept of the gambling law.

Another example of games with low *chanciness* and low *controllability* comes from deterministic games with many players. The game of *Chinese Checkers* has zero *chanciness*, yet it also lacks *controllability* when three or even five other players constantly change the board. We find it difficult to select one of both attributes as more relevant: The game is clearly not a game of chance, yet it has similar *controllability* to a game of chance. If *Chinese Checkers* was played for money (a typical requisite for the application of gambling law), we could offer no help on classifying it.

21.3 Summary

In this thesis we have developed measures for the quantification of chance and skill in games. Our work was build on the assumption that *influence* is a property of game moves. This allowed us to describe the influence of chance and skill through the study of individual moves, taken in individual matches.

We have defined the measures of *chanciness* and *controllability* for general-sum games of perfect information with an arbitrary number of players. Our measures rate the influence of chance and the influence of players on their own results, on the interval $[0, 1]$. A high value of *chanciness* indicates a game of chance while a high value of *controllability* indicates a game of skill. We have shown that games can be low on both, chance and skill, and that a single one-dimensional measure is insufficient to describe all games.

Our analysis has further revealed that it does not suffice to look at the structure of a game (its game tree). It is also necessary to consider the behavior of the agents who play the game. For this reason, our measures are defined in reference to a population of agents. The dependency on agent behavior complicates the effective computation of our measures. It can also be understood as the reason why results on the relationship between chance and skill are scarce. Asking about this relationship without giving assumption about agents leads to an under-determined question. However, we have shown that the *chanciness* and *controllability* can be effectively approximated for example games when taking the agents into account.

The measures of *chanciness* and *controllability* contribute to the theoretical characterization of games. When used in conjunction with computer agents they also provide a tool in the arsenal of game inventors.

Bibliography

- [Abr 1990] Bruce Abramson. Expected-outcome: A general model of static evaluation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(2):182 – 193, 1990.
- [AL 2007] Ingo Althöfer and Jakob Lübke. Board game "Big Balls". 3-Hirn-Verlag, 2007. www.3-hirn-verlag.de/games.html.
- [Ald 2000] Scott Alden. 2000. BoardGameGeek, Gaming Unplugged Since 2000. www.boardgamegeek.com. Search interface was used to obtain a lower bound on the number of published games.
- [Ald 2004] Scott Alden. 2004. "EinStein würfelt nicht!" on boardgamegeek.com. www.boardgamegeek.com/boardgame/18699. accessed 2009/07/21.
- [All 1988] Victor Allis. *A Knowledge-based Approach of Connect-Four - The Game is Solved: White Wins*. PhD thesis, Vrije Universiteit Amsterdam, 1988. www.connectfour.net/Files/connect4.pdf.
- [Alo 2007] Noga Alon. 2007. Poker, chance and skill. www.cs.tau.ac.il/~nogaa/PDFS/publications.html. accessed 2009/01/19.
- [Alt 2003] Ingo Althöfer. 2003. Computer-aided game inventing. www.minet.uni-jena.de/preprints/althoefer_03/. accessed 2009/01/19.
- [Alt 2004a] Ingo Althöfer. Board game "EinStein würfelt nicht". 3-Hirn-Verlag, 2004. www.en.wikipedia.org/wiki/EinStein_würfelt_nicht. accessed 2010/02/26.
- [Alt 2004b] Ingo Althöfer. Board game "EinStein würfelt nicht Quattro". 3-Hirn-Verlag, 2004. www.3-hirn-verlag.de/games.html. Online play at www.devilplay.de.
- [Alt 2005] Ingo Althöfer. Board game "Gelb Voran". 3-Hirn-Verlag, 2005. www.ludorium.at/SPIELE/0_english_version/3_hirn/gelb_voran_e.htm. accessed 2010/02/26.
- [Alt 2007] Ingo Althöfer. Board game "Seasons of the Sun". 3-Hirn-Verlag, 2007. Playable online against a computer at www.devilplay.de/sots/sots.html.
- [Alt 2008] Ingo Althöfer. On the Laziness of Monte-Carlo Game Tree Search in Non-tight Situations, 2008. www.althofer.de/mc-laziness.pdf. accessed 2010/02/08.

Bibliography

- [Aus 1934] Roland G Austin. *Roman Board Games*, volume I. Cambridge University Press, 1934.
- [BCG 1982] Elwyn R. Berlekamp, John H. Conway, and Richard K. Guy. *Winning Ways for your Mathematical Plays*. Academic Press, ISBN 0-12-091152-3, 1982.
- [BDS 2002] Darse Billings, Aaron Davidson, and Duane Szafron. The challenge of poker. *Artificial Intelligence*, 134:201–240, 2002.
- [BH 2003] Bruno Bouzy and Bernard Helmstetter. Monte Carlo Go developments. In H. Jaap van den Herik, Hiroyuki Iida, and Ernst A. Heinz, editors, *Advances in Computer Games conference (ACG-10)*, pages 159–174. Kluwer, 2003. www.math-info.univ-paris5.fr/~bouzy/publications/bouzy-helmstetter.pdf.
- [Blu 2001] Dennis Blumenfeld. *Operations Research Calculations Handbook*. CRC Press, 2001.
- [Bro 2008] Cameron Browne. *Automatic Generation and Evaluation of Recombination Games*. PhD thesis, Queensland University of Technology, 2008. eprints.qut.edu.au/17025/.
- [Bur 2007] Earl Burton. 2007. North carolina appeals court says poker 'game of chance'. www.pokernews.com/news/2007/5/north-carolina-appeals-game-chance.htm. accessed 2009/01/19.
- [BvdG 2001] Peter Borm and Ben van der Genugten. On a relative measure of skill for games with chance elements. *Journal of the Spanish Society of Statistics and Operations Research*, 9(1):91–114, 2001.
- [Cyp 2009] Dan Cypra. 2009. Pennsylvania Judge Rules, Poker is a Game of Skill. www.pokernewsdaily.com/pennsylvania-judge-rules-poker-is-a-game-of-skill-941. accessed 2009/01/19.
- [DBvdG 2004] Marcel Dreef, Peter Borm, and Ben van der Genugten. Measuring skill in games: several approaches discussed. *Mathematical Methods of Operations Research*, 59(3):375–391, 2004.
- [Dre 2005] Marcel Dreef. *Skill and Strategy in Games*. PhD thesis, Tilburg University, Faculty of Econometrics and Operations Research, 2005.
- [EL 2009] Olivier Egger and Johannes Levermann. 2009. Analysis of Backgammon matches by the Snowie engine. www.bgsnowie.com/backgammon/matches.dhtml. accessed 2009/02/23.
- [Elo 1986] Arpad E. Elo. *The Rating of Chessplayers, Past and Present*. Arco Pub., New York, 1986.

- [Erd 2009] Jakob Erdmann. Towards a Characterization of Chance in Games. The Case of Two-Player Zero-Sum Games with Perfect Information, 2009. www.minet.uni-jena.de/Math-Net/reports/sources/2009/09-05report.pdf. Submitted to the International Computer Games Association (ICGA) Journal. Review process ongoing at Feb. 2010.
- [FB 2008] Hilmar Finnsson and Yngvi Björnsson. Simulation-based approach to general game playing. In *AAAI*, volume 23, pages 259–264, 2008.
- [FID 2005] FIDE. 2005. Fide handbook, laws of chess. www.fide.com. accessed 2010/02/26.
- [Gue 2008] Jörg Guenther. Evaluation of monte-carlo agents for the board game "hex". Unpublished private correspondence, 2008.
- [Haj 2009] Alan Hajek. Interpretations of probability. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. The Metaphysics Research Lab, Stanford University, spring 2009 edition, 2009.
- [Hay 2009] Ryan Hayward. 2009. Mohex wins at the 14th computer olympiad. www.cs.ualberta.ca/~hayward/hex/. accessed 2009/07/16.
- [HBS 2006] Thomas Hauk, Michael Buro, and Jonathan Schaeffer. Rediscovering *-minimax search. In *Lecture Notes in Computer Science*, volume 3846, pages 35–50. Springer Berlin, Heidelberg, 2006.
- [Hei 1942] Piet Hein. Board game "Hex", Reinvented by John Nash in 1947 under the name "Nash", 1942. [en.wikipedia.org/wiki/Hex_\(board_game\)](http://en.wikipedia.org/wiki/Hex_(board_game)). accessed 2010/02/26.
- [KFA 1969] Rudolf E. Kalman, Peter L. Falb, and Michael A. Arbib. *Topics in Mathematical System Theory*. McGraw-Hill, New York, 1969.
- [Kin 2003] Monica Yant Kinney. *A casine game full of fowl play*. The Philadelphia Inquirer, March 20, 2003. www.netdotstuff.com/tictactoe/talkoftown4.html. accessed 2010/02/07.
- [Las 1925] Emanuel Lasker. *Gesunder Menschenverstand im Schach*. Berlin Wertbuchhandel, 1925.
- [LHH⁺ 2008] Nathaniel Love, Timothy Hinrichs, David Haley, Eric Schkufza, and Michael Genesereth. General Game Playing: Game Description Language Specification. Technical report, Stanford University, 2008.
- [Mal 2002] Richard Malaschitz. 2002. Little golem online board games. www.littlegolem.net.
- [McC 1994] Curtis McCullen. 1994. Backgammon ends. www.math.harvard.edu/~ctm/expositions/bgends.html.

Bibliography

- [MP 1992] Richard McKelvey and Thomas Palfrey. An experimental study of the centipede game. *Econometrica*, 60(4):803–36, July 1992.
- [Nat 2005] National Portal Content Management Team. 2005. Interesting Facts about India. india.gov.in/myindia/facts.php. accessed 2010/02/12.
- [Nie 2008] Claus Niederalte. Lagebericht und Jahresabschluss Süddeutsche Klassenlotterie, 2008. www.thueringen.de/imperia/md/content/tfm/berichte/beteiligungsbericht2007/sueddeutsche_klassenlotterie_ja_2006_2007.pdf. accessed 2010/02/12.
- [Now 1996] Richard J. Nowakowski. *Games of No Chance*. Cambridge University Press, 1996.
- [Par 1999] David S. Parlett. *The Oxford history of board games*. Oxford University Press, June 1999.
- [Par 2008] David S. Parlett. Elements of chance and skill in games. In *Proceedings of the 11th Board Game Studies Colloquium*, pages 73–87, 2008. www.davidparlett.co.uk/gamester/chanceskill.html. accessed 2010/02/26.
- [Pel 1994] Barney Pell. A strategic metagame player for general chess-like games. In *AAAI'94: Proceedings of the twelfth national conference on Artificial intelligence (vol. 2)*, pages 1378–1385, Menlo Park, CA, USA, 1994. American Association for Artificial Intelligence.
- [Pic 1980] Peter Piccione. In Search of the Meaning of Senet. *Archaeology Magazine*, 33, July-August, 1980. http://www.cofc.edu/~piccione/piccione_senet.pdf. accessed 2010/02/12.
- [Pot 1984] Donald Potter. Go in the classics. *Go World*, 37, 1984. www.kiseido.com/classics.htm.
- [Rie 2010] Hagen Riedel. Partially Controlled Random Walks. (in German language). Diplomarbeit, tentative draft, to be completed in summer 2010, Friedrich-Schiller-Universität Jena, Fakultä für Mathematik und Informatik, Lehrstuhl für Mathematische Optimierung, 2010.
- [SLB 2008] Yoav Shoham and Kevin Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, Cambridge, 2008. www.masfoundations.org.
- [Sta 2007] Staatsvertrag. Staatsvertrag zum Glücksspielwesen in Deutschland (GlüStV). Hamburgisches Gesetz- und Verordnungsblatt, 2007. [by.juris.de/by/GlueStVtr_BY_P29.htm](http://juris.de/by/GlueStVtr_BY_P29.htm).
- [Tes 1995] Gerald Tesauro. Temporal difference learning and td-gammon. *Commun. ACM*, 38(3):58–68, 1995.

Bibliography

- [Tur 1948] Alan Turing. *Intelligent Machinery*. Oxford University Press, 1948.
- [Tur 1950] Alan Turing. Computing machinery and intelligence. *Mind*, 59:433–460, 1950.
- [vN 1927] J. von Neumann. Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100:295–320, 1927.
- [Wag 1969] Harvey M. Wagner. *Principles of operations research, with applications to managerial decisions [by] Harvey M. Wagner*. Prentice-Hall Englewood Cliffs, N.J., 1969.
- [Wue 2007] Heinz Wueppen. Board game "Würfel Bingo". Ravensburger, item number: 27131, 2007. www.boardgamegeek.com/boardgame/29355/high-score.
- [Zer 1913] Ernst Zermelo. Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels. *Proceedings of the Fifth International Congress of Mathematicians*, 2:501–504, 1913. www.math.harvard.edu/~elkies/FS23j.03/zermelo.pdf. English translation, accessed 2010/02/26.

A Lottery Chess

Lottery Chess is an artificially designed class of games to reveal an inconsistency in the notion of *Relative Skill* [Dre 2005]. Lottery chess has two parameters $n, k \in \mathbb{N}$. There exist realizations of the parameters n, k for which the resulting games receive a rating of *Relative Skill* which in our opinion contradicts common sense on the notions of skill and chance.

Lottery chess is almost identical to the classical game of chess. As the only difference, each player secretly selects an integer from $[0, 1, \dots, n]$ before the start of the chess match. The winner of the match receives a score of 1. Additionally a lottery takes part. A random integer from $[0, 1, \dots, n]$ is drawn. Should a player have guessed that number he receives an additional score of k .

The influence of chance on the game outcome can be made arbitrarily small by increasing n . As n grows, the likelihood that one of the players correctly guesses the lottery number decreases to zero. At the same time the *Relative skill* (in the sense of Dreef) can be made arbitrarily small by increasing k . This follows because the fictive player *always* wins the lottery no matter its odds. As k grows, this advantage of the fictive player overshadows any advantage of being good at chess.

By choosing large values for n and k (such as $n = 10^3, k = 10^6$) we can create a game that has a *Relative Skill* rating near zero. Equation 2.1 gives

$$\begin{aligned} RS &= \frac{\text{gain}_{\text{optimal}} - \text{gain}_{\text{beginner}}}{\text{gain}_{\text{fictive}} - \text{gain}_{\text{beginner}}} \\ &= \frac{1.001 - 0.001}{1000001 - 0.001} < 10^{-6} \end{aligned}$$

However, the result of that game when played by any human or algorithm will be mostly determined by skill at chess and will therefore be a game of skill. As a game cannot be a game of skill and a game of no skill at the same time, the definition of *Relative Skill* is contradictory. In the defense of the *RS*, it must be noted that this extreme behavior is unlikely to show up in a typical game.

B Lookup Sampling

B.1 Lookup Sampling for Two-Player Zero-Sum Games

In Part I (Section 4.3) we introduced *relative positional values* (RPV) and in Section 4.4 we explained how to approximate them with Monte Carlo Methods. The efficiency of that sampling algorithm can be improved when calculating RPV for consecutive game states of a match. We call our improved algorithm *Lookup Sampling*.

Let $g = (X, Y, (s_0, \dots, s_n))$ be a match of game G between the agents X and Y . To calculate the *relative influence of chance* of g , we need the values $\tilde{V}(s_i, X, Y)$ for all s_i in g .

According to Equation 4.3 we can calculate the $\tilde{V}(s_i, X, Y)$, given the RPV for all child nodes of s_i and the probabilities of all child nodes to be selected in a match between X and Y . *Lookup Sampling* exploits this Equation to speed up our sampling algorithm.

We approximate the RPV in the backward sequence for $(s_n, s_{n-1}, \dots, s_0)$. When approximating $\tilde{V}(s_i, X, Y)$ we start a simulation match and sample a child node s'_i for s_i (by querying an agent or simulating a chance move according to the rules of G). If that child node s'_i is identical to the node s_{i+1} from g we immediately return $\tilde{V}(s_{i+1}, X, Y)$ as the value of that simulation match. Otherwise we continue the simulation match until a leaf node is reached. The average result of all simulation matches is then used as our approximation of $\tilde{V}(s_i, X, Y)$.

Because the comparison of s'_i and s_{i+1} can be accomplished in constant time (via Hashing) we can save time by stopping our simulation early. This total gain in computing speed depends on the probability of s_{i+1} to be sampled. If (s_i, s_{i+1}) happens to be a “deterministic move” (Section 4.5) the sampling for $\tilde{V}(s_i, X, Y)$ always stops after the first move.

For the game of PaCRaWa played by two agents of type Z (fully random), the running time is roughly halved compared to the original algorithm. This is to be expected as every simulation game has a 50% chance of being replaced by a table lookup after the first move. Moreover, *Lookup Sampling* improves the accuracy of influence calculations. Let $m = (s_i, s_{i+1})$ be a “deterministic move” by agent X . According to Observation 4.6, the influence of m is zero. With regular Monte Carlo sampling, the influence of m is the difference of two approximated RPV. It is very unlikely that both approximations return the same value and thus yield an influence of zero. With *Lookup Sampling*, $\tilde{V}(s_i, X, Y)$ equals $\tilde{V}(s_{i+1}, X, Y)$ exactly, because every simulation match for s_i returns $\tilde{V}(s_{i+1}, X, Y)$ as its result. Even though $\tilde{V}(s_{i+1}, X, Y)$ is only an approximation, the influence of m is calculated as zero.

The game PaCRaWa[$d = 1, T = 10, a = 0$] (Section 7.1) played by two optimal (and deterministic) agents X is fair (ME = 0). Because chance generally has an influence on the match result and the deterministic agents have no influence at all, the *chanciness*

B Lookup Sampling

$C(\text{PaCRaWa}[d = 1, T = 10, a = 0], X, X)$ is 1. Given 1000 samples per game state and 100 matches, regular Monte Carlo sampling approximates the *chanciness* as 0.88 while *Lookup Sampling* achieves an approximation of 0.98.

B.2 Lookup Sampling for General-Sum Games with n Players

In Part II of this thesis we generalize the concept of RPV for the class of general-sum games with n players. This generalization can be applied to the *Lookup Sampling*-algorithm in a straightforward manner.

Every occurrence of the term $\tilde{V}(s, X, Y)$ in the preceding section must simply be replaced by the generalized term $\tilde{V}_i(s, L)$.

C Bounding *chanciness* and *controllability* for a Small Game.

In Section 17.3 we have introduced a game and asked about the bounds of *chanciness* and *controllability*. The game, which we denote as G is shown again in Figure C.




Figure C.1: A small single-player game with chance and player nodes.

Since game G is a single-player game, it is sufficient to consider all populations with a single agent. Furthermore, G only contains a single player node with two children. This allows us to characterize our agent with a single parameter p , the probability of selecting the left child node.

To express *chanciness* and *controllability* as a function of p we look at the input terms $ME, \mathbb{E}(|IC|)$, and $\mathbb{E}(|IP_1|)$ as functions of p (see Chapters 14, 15). We deduce the terms as follows:

C Bounding *chanciness* and *controllability* for a Small Game.

$$\begin{aligned}
Q &:= 0.01 \quad \text{prob. to enter the left subtree} \\
V(p) &= p \cdot (-1) + (1-p) \cdot (+1) \\
&= 1 - 2p \quad \text{expected value of the player node} \\
\mathbb{E}(|IP_1|) &= A(p) = Q[p(-1 - V(p)) + (1-p)(1 - V(p))] \\
&= Q \cdot 4p(1-p) \quad \text{expected abs. influence of the player} \\
|ME| &= M(p) = |Q \cdot V(p)| \quad \text{abs. match expectation} \\
\mathbb{E}(|IC|) &= C(p) \quad \text{expected abs. influence of chance} \\
&= Q|V(p) - M(p)| \quad \text{left subtree} \\
&\quad + 0.5(1-Q)|0 - M(p) - 1| \quad \text{right subtree, leaf node -1} \\
&\quad + 0.5(1-Q)|0 - M(p) + 1| \quad \text{right subtree, leaf node 1} \\
CHA(p) &= \frac{C(p)}{C(p) + A(p) + M(p)} \\
CON(p) &= \frac{A(p)}{C(p) + A(p) + M(p)}
\end{aligned}$$

Figure C shows the *chanciness* as a function of p . The *chanciness* $CON(G, \{X_p\}, \{(\{X_p\}, 1)\})$ is computed for a population that consists only of the agent X_p (the lineup (X_p) occurs with probability 1).




Figure C.2: The *controllability* of G as a function of the agent parameter p .

The locations of the extremum values given in Section 17.3 were obtained by visual inspection of the rather obvious plots.

D Proof that $|\mathbb{E}(\text{IC})| \equiv 0$

The term IC (short for $\text{IC}_i(g)$, defined in Chapter 14) is a sum of the influence of all chance moves in g from the perspective of player i . The expectation operator applies to the random match $g = (L, P)$ which is the result of a randomly selected lineup L and a random sequence of game states $P = (s_0, s_1, \dots, s_k)$. The parameter i is fixed.

Due to the linearity of the expectation operator it is sufficient to show that every summand of IC has an expected influence of zero.

We will first show that the j th summand of $\text{IC}_i(g)$ is zero under some conditions. The j th summand has the form $\text{I}(m, L, i)$ with move $m = (s, s')$. Let L and s both be fixed but arbitrary. Let S denote $\text{Children}(s)$ and let $p(s')$ denote the probability that chance will move from the fixed node s to an arbitrary child s' .

$$\begin{aligned}
 \mathbb{E}(\text{I}(m, L, i)) &= \sum_{s' \in S} p(s') \cdot (\tilde{V}_i(s', L) - \tilde{V}_i(s, L)) \\
 &= \sum_{s' \in S} p(s') \tilde{V}_i(s', L) - \sum_{s' \in S} p(s') \tilde{V}_i(s, L) \\
 &= \tilde{V}_i(s, L) - \sum_{s' \in S} p(s') \tilde{V}_i(s, L) \quad \text{by Definition 13.1} \\
 &= \tilde{V}_i(s, L) - \tilde{V}_i(s, L) \sum_{s' \in S} p(s') = 0
 \end{aligned}$$

The lineup L and the node s used above were fixed but arbitrary. Thus, the *law of total expectation* applies and the expectation of the j th summand is zero in the general case \square

Corollary: The expected absolute influence of any player on the result of any other player is zero.

$$|\mathbb{E}(\text{IP}_{j,i})| \equiv 0 \quad \text{for any player roles } j, i \quad (\text{D.1})$$

This follows from the above proof when the probabilities for chance moves are replaced by the probabilities for players moves. These probabilities are well defined as the node s is then a player node for the specific player j . The move probabilities of agent A_j are well defined by the lineup $L = (A_1, A_2, \dots, A_n)$. Both s and L are conditionally fixed. \square

E Exemplary Calculations of cha for a match of Game G_{example}

In Section 14.2 we have calculated the *chanciness* of match h from the perspective of player 2. Match $h = (L, P)$ is a match of game G_{example} . The lineup $L = (A_1, A_2, A_3)$ consists of the agents defined in Section 12.2.

Here we repeat this calculation for the perspectives of player 1 and player 3. Figures E.1 and E.2 show the move probabilities and the RPV from the perspective of players 1 and 3 respectively. The sequence of states $P = (s_0, s_1, s_2, s_3)$ is highlighted.

$$\begin{aligned} M_C(h) &= \{(s_2, s_3)\} \\ M_1(h) &= \{(s_0, s_1)\} \\ M_2(h) &= \{(s_1, s_2)\} \\ M_3(h) &= \{\} \end{aligned}$$

$$\begin{aligned} ME_1(h) &= 0.4 \\ IC_1(h) &= -0.5 \\ IP_{1,1}(h) &= 0.8 \\ IP_{2,1}(h) &= 0.3 \\ IP_{3,1}(h) &= 0 \end{aligned}$$

$$\text{cha}_1(h) = \frac{0.5}{0.4 + 0.5 + 0.8 + 0.3 + 0} = 0.25$$

$$\begin{aligned} ME_3(h) &= 1.15 \\ IC_3(h) &= -0.5 \\ IP_{1,3}(h) &= 0.05 \\ IP_{2,3}(h) &= 0.3 \\ IP_{3,3}(h) &= 0 \end{aligned}$$

$$\text{cha}_3(h) = \frac{0.5}{1.15 + 0.5 + 0.05 + 0.3 + 0} = 0.25$$




Figure E.1: An exemplary match h . The path through the game tree is highlighted and every node is marked with the RPV for the lineup L from the perspective of player 1.




Figure E.2: An exemplary match h . The path through the game tree is highlighted and every node is marked with the RPV for the lineup L from the perspective of player 3.

F Exemplary Calculations of CHA for Game G_{example}

In Section 14.3.4 we have calculated the *chanciness* of game G_{example} from the perspective of player 2. We used the population $\Pi = A_1, A_2, A_3$ with the agents defined in Section 12.2, and probability 1 for the lineup $L = (A_1, A_2, A_3)$ (the same Ψ as in Section 14.3.4).

Here we repeat this calculation for the perspectives of player 1 and player 3. Figures F.1 and F.2 show the move probabilities and the RPV from the perspective of players 1 and 3 respectively.




Figure F.1: The tree of Game G_{example} reduced to the perspective of player 1 with move probabilities for a specific lineup L . Each node s is annotated with the RPV $\tilde{V}_1(s, L)$.




Figure F.2: The tree of Game G_{example} reduced to the perspective of player 3 with move probabilities for a specific lineup L . Each node s is annotated with the RPV $\tilde{V}_3(s, L)$.

$$\begin{aligned}
\mathbb{E}(|\text{ME}_1(g)|) &= 0.4 \\
\mathbb{E}(|\text{IC}_1(g)|) &= 0.25 \cdot 0 + 0.75 \cdot 0.8 \cdot 0.5 = 0.3 \\
\mathbb{E}(|\text{IP}_{1,1}(g)|) &= 0.25 \cdot 2.4 + 0.75 \cdot 0.8 = 1.2 \\
\mathbb{E}(|\text{IP}_{2,1}(g)|) &= 0.25 \cdot 0 + 0.75 \cdot 0.8 \cdot 0.3 + 0.75 \cdot 0.2 \cdot 1.2 = 0.36 \\
\mathbb{E}(|\text{IP}_{3,1}(g)|) &= 0.75 \cdot 0.2 \cdot 2 = 0.3 \\
\mathbb{E}(\text{TI}_1(g)) &= 0.4 + 0.3 + 1.2 + 0.36 + 0.3 = 2.56 \\
\text{CHA}_1(G_{\text{example}}, \Pi, \Psi) &= \frac{\mathbb{E}(|\text{IC}_1(g)|)}{\mathbb{E}(\text{TI}_1(g))} \approx 0.12 \\
\\
\mathbb{E}(|\text{ME}_3(g)|) &= 1.15 \\
\mathbb{E}(|\text{IC}_3(g)|) &= 0.25 \cdot 0 + 0.75 \cdot 0.8 \cdot 0.5 = 0.3 \\
\mathbb{E}(|\text{IP}_{1,3}(g)|) &= 0.25 \cdot 0.15 + 0.75 \cdot 0.05 = 0.075 \\
\mathbb{E}(|\text{IP}_{2,3}(g)|) &= 0.25 \cdot 0 + 0.75 \cdot 0.8 \cdot 0.3 + 0.75 \cdot 0.2 \cdot 1.2 = 0.36 \\
\mathbb{E}(|\text{IP}_{3,3}(g)|) &= 0 \\
\mathbb{E}(\text{TI}_3(g)) &= 1.15 + 0.3 + 0.075 + 0.36 + 0 = 1.885 \\
\text{CHA}_3(G_{\text{example}}, \Pi, \Psi) &= \frac{\mathbb{E}(|\text{IC}_3(g)|)}{\mathbb{E}(\text{TI}_3(g))} \approx 0.16
\end{aligned}$$

G Proof that $\text{cha}_1(g) \equiv \text{cha}_2(g)$ for Two-Player Zero-Sum Games

For a match g of an n -player game G there will generally be n different values of $\text{cha}_i(g)$ for the different player perspectives i .

However, for the special case of two-player zero-sum we can show the following:

Lemma G.1 *For any match g of the two-player zero-sum game G the chanciness of both player perspectives is equal.*

$$\text{cha}_1(g) \equiv \text{cha}_2(g) \quad (\text{G.2})$$

Proof: Let $g = (L, P)$ be a match of the two-player zero-sum game G . Because of the zero-sum property, the equality $\tilde{V}_1(t, L) = -\tilde{V}_2(t, L)$ holds for any terminal node t of G .

By the definition of \tilde{V}_i (Definition 13.1), the value of $\tilde{V}_i(s, L)$ is a linear combination of the values for the children of s . Therefore, the equality $\tilde{V}_1(s, L) = -\tilde{V}_2(s, L)$ extends to all nodes s of G .

Using this equality we get $I(m, L, 1) = -I(m, L, 2)$ (Definition 13.3). Since the combined influence of an SI is a sum of move influences (Definition 13.14), this equality extends to the SI as well:

$$\text{ME}_1(g) = -\text{ME}_2(g) \quad (\text{G.3})$$

$$\text{IC}_1(g) = -\text{IC}_2(g) \quad (\text{G.4})$$

$$\text{IP}_{1,1}(g) = -\text{IP}_{1,2}(g) \quad (\text{G.5})$$

$$\text{IP}_{2,1}(g) = -\text{IP}_{2,2}(g) \quad (\text{G.6})$$

Since the definition of cha (Definition 14.2) uses the absolute influences of the SI , we have the equality of $\text{cha}_1(g)$ and $\text{cha}_2(g)$ \square

The above proof can easily be modified to show the invariance of cha and con under certain “rule changes” applied to G . If the results for player i are all scaled with the factor α , the influences of the SI on player i will scale accordingly. Since α then appears in the numerator as well as the denominator of cha (or con), α is canceled out.

Since the equality of two random variables implies the equality of their expectations, the following Lemma follows readily from the proof of Lemma G.1:

Lemma G.7 *For any two-player zero-sum game G , any population Π and any sampling distribution Ψ , the chanciness of both player perspectives is equal.*

$$\text{CHA}_1(G, \Pi, \Psi) \equiv \text{CHA}_2(G, \Pi, \Psi) \quad (\text{G.8})$$

H Estimation of the Approximation Error of CHA and CON

The measures CHA is define as quotient of random variables $\frac{X}{Y}$.

To approximate the variance of the quotient of these non-independent random variables we have used the formula by Blumenfeld [Blu 2001, 2.29]:

$$\text{Var}\left(\frac{X}{Y}\right) \approx \left(\frac{\mathbb{E}(X)}{\mathbb{E}(Y)}\right)^2 \left(\frac{\text{Var}(X)}{(\mathbb{E}(X))^2} + \frac{\text{Var}(Y)}{(\mathbb{E}(Y))^2} - \frac{2\text{Cov}(X,Y)}{\mathbb{E}(X)\mathbb{E}(Y)} \right) \quad (\text{H.1})$$

Since in our case the random variables X and Y are themselves obtained as the sample mean of other random variables, we have used the sample mean, the sample variance of the mean, and the sample covariance of both means in place of the operators \mathbb{E} , Var and Cov .

We have then used the root of this approximation of an approximation as our error estimate.

I Result Statistics for “EinStein würfelt nicht Tetra”

In Section 20.5 we described the game “EinStein würfelt nicht Tetra” (EWNT) and gave approximated measurements for *chanciness* and *controllability*. For whom it may concern, we present additional result statistics from our simulation matches.

result \ player	1	2	3	4
+1	1161	1097	968	869
−1	863	768	715	636
0	2071	2230	2412	2590

Figure I.1: Absolute frequency of the different payoff values for each player role. Each possible lineup from the population $\Pi = \{\text{RND}, \text{MC}_1, \text{MC}_8, \text{MC}_{64}\}$ was sampled 16 times.

result \ player	1	2	3	4
+1	147	122	121	122
−1	105	101	88	57
0	260	289	303	333

Figure I.2: Absolute frequency of the different payoff values for each player role. Each possible lineup from the population $\Pi = \{\text{MC}_1, \text{MC}_{512}\}$ was sampled 32 times.

I Result Statistics for “EinStein würfelt nicht Tetra”

Lineup	ME ₁	ME ₂	ME ₃	ME ₄
(MC ₁ MC ₁ MC ₁ MC ₁)	0.06	0.05	0.04	0.04
(MC ₁ MC ₁ MC ₁ MC ₅₁₂)	-0.07	-0.06	-0.06	0.39
(MC ₁ MC ₁ MC ₅₁₂ MC ₁)	-0.06	-0.06	0.40	-0.04
(MC ₁ MC ₁ MC ₅₁₂ MC ₅₁₂)	-0.16	-0.13	0.28	0.26
(MC ₁ MC ₅₁₂ MC ₁ MC ₁)	-0.07	0.43	-0.05	-0.05
(MC ₁ MC ₅₁₂ MC ₁ MC ₅₁₂)	-0.16	0.28	-0.11	0.26
(MC ₁ MC ₅₁₂ MC ₅₁₂ MC ₁)	-0.14	0.29	0.25	-0.09
(MC ₁ MC ₅₁₂ MC ₅₁₂ MC ₅₁₂)	-0.21	0.19	0.18	0.17
(MC ₅₁₂ MC ₁ MC ₁ MC ₁)	0.45	-0.05	-0.05	-0.04
(MC ₅₁₂ MC ₁ MC ₁ MC ₅₁₂)	0.29	-0.12	-0.11	0.26
(MC ₅₁₂ MC ₁ MC ₅₁₂ MC ₁)	0.29	-0.12	0.26	-0.08
(MC ₅₁₂ MC ₁ MC ₅₁₂ MC ₅₁₂)	0.17	-0.18	0.18	0.19
(MC ₅₁₂ MC ₅₁₂ MC ₁ MC ₁)	0.30	0.26	-0.10	-0.09
(MC ₅₁₂ MC ₅₁₂ MC ₁ MC ₅₁₂)	0.19	0.17	-0.15	0.18
(MC ₅₁₂ MC ₅₁₂ MC ₅₁₂ MC ₁)	0.18	0.19	0.17	-0.12
(MC ₅₁₂ MC ₅₁₂ MC ₅₁₂ MC ₅₁₂)	0.09	0.11	0.11	0.12

Figure I.3: *match expectation* for various lineups in the game EWNT (estimated error ≈ 0.01).

Ehrenwörtliche Erklärung

Hiermit erkläre ich:

- dass mir die Promotionsordnung der Fakultät für Mathematik und Informatik der Friedrich-Schiller-Universität Jena bekannt ist,
- dass ich die Dissertation selbst angefertigt habe, keine Textabschnitte oder Ergebnisse eines Dritten oder eigene Prüfungsarbeiten ohne Kennzeichnung übernommen habe und alle von mir benutzten Hilfsmittel, persönliche Mitteilungen und Quellen in meiner Arbeit angegeben sind,
- dass alle Personen die mich bei der Auswahl und Auswertung des Materials sowie bei der Herstellung des Manuskripts unterstützt haben im Abschnitt *Acknowledgements* (Seite vii) aufgeführt sind,
- dass die Hilfe eines Promotionsberaters nicht in Anspruch genommen wurde und dass Dritte weder unmittelbar noch mittelbar geldwerte Leistungen von mir für Arbeiten erhalten haben, die im Zusammenhang mit dem Inhalt der vorgelegten Dissertation stehen,
- dass ich die Dissertation noch nicht als Prüfungsarbeit für eine staatliche oder andere wissenschaftliche Prüfung eingereicht habe und dass ich weder die gleiche, eine in wesentlichen Teilen ähnliche noch eine andere Abhandlung bei einer anderen Hochschule als Dissertation eingereicht habe.

Schulzendorf, 14.07.2010

Tabellarischer Lebenslauf

Persönliche Daten

Name:	Jakob Erdmann
Geburtsdaten:	11.07.1980 in Königs Wusterhausen
Anschrift:	Dohlenstieg 48 15732 Schulzendorf
E-Mail:	jakob.erdmann@gmail.com
Familienstand:	verheiratet
Kinder:	Anton Ijon Erdmann Dora Elena Erdmann

Schulbildung

1979-1993	Humboldt-Grundschule Eichwalde
1993-2000	Humboldt-Gymnasium Eichwalde, Leistungskurse: Mathematik und Informatik, Schüleraustauschjahr 1997/98 in den USA, Abschluss: Abitur

Zivildienst

11/2000-10/2001	Zivildienst im Achenbach Krankenhaus Königs Wusterhausen
-----------------	---

Studium

10/2001-2007	Studium der Informatik an der Friedrich-Schiller-Universität Jena Vertiefungsrichtung: Künstliche Intelligenz und Mustererkennung
12/2006-05/2007	Lehrstuhl Rechnerarchitektur, FSU Jena Diplomarbeit: <i>Strategien des dynamischen Lastausgleichs für Multi-Domain-Cluster am Beispiel einer parallelen Dendriten-Simulation</i> Abschluss: Diplom, mit Auszeichnung
10/2007-07/2010	Promotion an der FSU Jena