Characterising Score Distributions in Dice Games

Aaron Isaksen, Tandon School of Engineering, NYU

Christoffer Holmgård, Tandon School of Engineering, NYU

Julian Togelius, Tandon School of Engineering, NYU

Andy Nealen, Tandon School of Engineering, NYU

We analyse a variety of ways that comparing dice values can be used to simulate battles in games, measuring the 'win bias', 'tie percentage', and 'closeness' of each variant, to provide game designers with quantitative measurements of how small rule changes can significantly affect game balance. Closeness, a metric we introduce, is related to the inverse of the second moment, and measures how close the final scores are expected to be. We vary the number of dice, number of sides, rolling dice sorted or unsorted, biasing win rates by using mixed dice and different number of dice, allowing ties, re-rolling ties, and breaking ties in favour of one player.

1 Introduction

Dice are a popular source of randomness in games. We examine the use of dice to simulate combat and other contests. While some games have deterministic rules for exactly how a battle will resolve, many games add some randomness, so that it is uncertain exactly who will win a battle. In games such as Risk [1], two players roll dice at the same time, and then compare their values, with the higher value eliminating the opponent's unit. Other games use a hit-based system, such as Axis and Allies [2], where a die roll of a strength value or less is a successful hit, with stronger units simulated by larger strength values and larger armies rolling more dice. In both games, stronger forces are more likely to win the battle, but lucky or unlucky rolls can result in one player performing far better, making a wide difference in scores.

Given a very large number of games played between players, unlucky and lucky rolls will balance out such that the player who has better strategy will probably end up winning; however, people might not play the same game enough times for the probabilities to even out. Instead, they play a much smaller, finite number of rolls spread across one session, or perhaps a couple of play sessions. The *gambler's fallacy* is the common belief that dice act with *local representativeness*: even a small number of dice rolls should be very close to the expected probabilities [3, 4, 5]. Therefore, it can often be frustrating when rolling poorly against an opponent: players often blame the dice, or themselves, for bad rolls, even though logic and reason dictate that everyone has the same skill at randomly rolling dice. Game designers may want to avoid or reduce this kind of negative player experience in their games.

Although there are thousands of games based on dice (the BoardGameGeek online database lists over 7,000 entries for dice games,[1] and hundreds of games are described in detail in [6, 7]), we specifically examine games where players roll and compare the individual dice values, as in Figure 1. Each player's dice are sorted in decreasing order and then paired up. Whichever player rolled a higher value in a pair wins a point. The points are summed, and whoever has more points wins the battle.

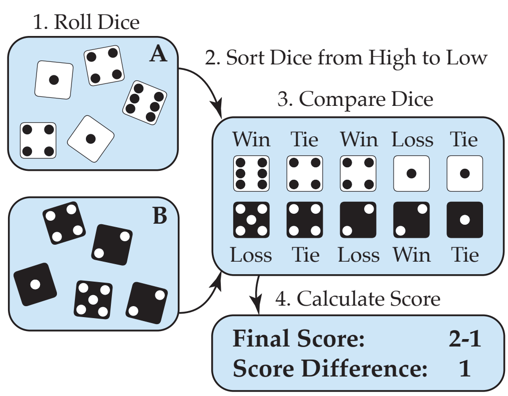

Figure 1. An example of a dice battle.

We use the term *battle* to denote any event resolved randomly within a larger game. The word is normally used to refer to combat, but our analysis can be used any time players compare dice outcomes in a contest. In this paper, we will use the terms *battle* and *game* interchangeably.

[1]https://boardgamegeek.com/boardgamecategory/1017/dice

Isaksen, A., Holmgård, C., Togelius, J. and Nealen, A., 'Characterising Score Distributions in Dice Games', *Game & Puzzle Design*, vol. 2, no. 1, 2016, pp. 24–37. © 2016