

Characterising Score Distributions in Dice Games

Aaron Isaksen Tandon School of Engineering, New York University
Christoffer Holmgård Tandon School of Engineering, New York Univ.
Julian Togelius Tandon School of Engineering, New York University
Andy Nealen Tandon School of Engineering, New York University

We analyze 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, rerolling 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 like *Risk* [1], two players roll dice at the same time, and then compare their values, with the higher value eliminating the opponent’s unit. Others use a hit-based system, like in *Axis and Allies* [2], where a die roll of a target value or less is a successful hit, with stronger units simulated by larger target 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 quite frustrating when rolling poorly against an opponent: players often blame the dice, or themselves, for bad rolls, even though logic and reason indicates that everyone has the same skill at rolling dice. Game designers may want to avoid or reduce this kind of negative player experience in their games.

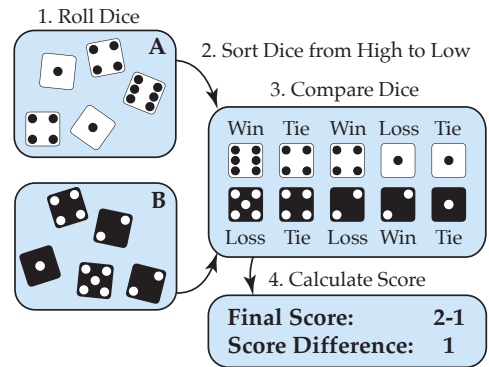


Figure 1. An example of a dice battle.

Although there are thousands of games based on dice (BoardGameGeek lists over 7,000 entries for dice¹, and hundreds of games are described in detail in [6, 7]), we examine games where players roll and compare the individual dice values, as in Figure 1. The dice are sorted in decreasing order and then paired up. Whichever player rolled a higher value on the pair wins a point. The points are summed, and whomever has more points wins the battle. We use the term *battle* to imply an event resolved 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.

We examine different variants and show how different factors affect the distribution of scores and other metrics which are helpful for evaluating a game. By adjusting the dice mechanics, a designer can influence the expected *closeness* of the outcomes of a battle, the *win bias* in favour of one of the players, and the *tie percentage*, the fraction of battles that end in a tie. The variants we examine include different numbers of dice,

¹<https://boardgamegeek.com/boardgamecategory/1017/dice>

various sided dice, different ways to sort the dice, and various ways to break ties.

Dice have come in various numbers of sides for millennia [8]: some of the oldest dice, dating back to at least 3500 B.C., were bones with 4 flat sides and 2 rounded ones. Eventually, 6-sided dice were created by polishing down the rounded sides. The dot patterns we see on today's 6-sided dice also come from antiquity. Ancient dice also come in the form of sticks with 4 long sides for Pachisi or 2 long sides for Senet [9]. The common dice in use for modern games are 4, 6, 8, 10, 12, and 20-sided, but other variants exist. In this paper, we use the notation ndk to mean a player rolls n dice which have k sides (e.g. 5d6 means rolling 5 dice, each with 6 sides).

By understanding how rules and randomness affect closeness, a designer can then choose the appropriate combination to try to achieve their desired game experience. A designer may prefer for their game to be highly unpredictable with large swings, intentionally increasing the risk for players to commit their limited resources. In addition, randomness can make a game appear to be more balanced because the weaker player can occasionally win against the stronger player [10]. Large swings may be more emotional and chaotic, and the "struggle to master uncertainty" can be considered "central to the appeal of games" [11]. Or, a designer may prefer for each battle to be close, to limit the feelings of one side dominating the other in what might be experienced as unfair or unbalanced, in a trait known as *inequity aversion* [12, 13]. Similarly, a designer may prefer to allow ties (simulating evenly matched battles), or wish to eliminate the opportunity for ties (forcing one side to win). Finally, a designer might want to vary the rules between each battle within a game, to represent changing strengths and weaknesses of the players and to provide aid to the losing player. A designer can adjust randomness to encourage situations appropriate for their game.

For most sections in this paper, we calculate the exact probabilities for each outcome by iterating over all possible rolls, tabulating the final score difference. Because each outcome is independent, we can parallelise the experiments across multiple processors to speed up the calculations (details about how many calculations are given in the Appendix). There are other methods one could use to computationally evaluate the odds, such a dice probability language like AnyDice [14] or Troll [15], or by using Monte-Carlo simulation (we use simulation when examining rerolls in Section 8). Writing the analytical probabilities becomes difficult for more complex games and we feel that presenting equations of this type has limited utility for most game designers.

2 Metrics for Dice Games

Quantitative metrics have been used to computationally analyze outcome uncertainty in games, typically for the purpose of generating novel games [16, 17]. Here we focus only on metrics that examine the final scores of the dice battle; we do not evaluate anything about how scores evolve during the battle itself (which we believe would be essential for more complicated games). But for simple dice battles, which are a component of a longer game, we can just focus on the end results. *Win bias* and *tie probability*, are similar to those used in previous work, but one of our metrics, *closeness*, is something we have not seen used before in game analysis.

We now define these metrics precisely. For the remainder of this paper, we use the terms "battle" and "game" interchangeably. Let s_A be the final score of a battle for Player A, and s_B be the final score for Player B. The *battle score* is then written as $s_A - s_B$. The *score difference*, d , is the numerical value of the battle score, so $d = s_A - s_B$. If we iterate over all the possible ways that the dice can be rolled, and count the number of times each score difference occurs, we can make a *score difference probability distribution*, $D(d)$. This describes the probability of achieving a score difference of d in the battle. We calculate $D(d)$ by first counting every resulting score difference in a histogram-like data structure, and then dividing each bin by the total sum of all the bins.

We now define the *win percentage* as the percent probability of Player A winning a battle. This can be calculated by summing the probabilities where the score difference is positive and is therefore a win for Player A. This is calculated as $100 \sum_{d>0} D(d)$ and will be between 0% and 100%. *Loss percentage* is the percent probability of Player A losing a battle, and is calculated as $100 \sum_{d<0} D(d)$, also between 0% and 100%.

We take the difference of the win and loss percentages to calculate **win bias**:

$$\text{win bias} = 100 \left(\sum_{d>0} D(d) - \sum_{d<0} D(d) \right) \quad (1)$$

This will vary between -100% and 100%. Games with a win bias of 0% are balanced, with no preference of Player A over Player B. If the win bias is $> 0\%$ then Player A is favoured; if $< 0\%$ then Player B is favoured. This metric is similar to the *Balance* metric in [16] but here we include the effect of ties and are concerned with the direction of the bias. A non-zero win bias is often desired, for example when simulating that one player is in a stronger situation than the other.

Next, we have the **tie percentage** which tells us the percent probability of the battle ending in

a tie, which we define as:

$$\text{tie percentage} = 100 (D(0)) \quad (2)$$

Some designers may want a possibility of ties, while others may not. This metric is analogous to *drawishness* in [16].

Finally, we present **closeness**, our new metric which measures how much the final score values centre around a tied game. Game that often ends within 1 point should have higher closeness than games that often end with a score difference of 5 or -5. The related statistical term *precision* is defined as the inverse of variance about the mean. For closeness, we define this as the square root of the inverse of variance in the score difference distribution about the tie value $d = 0$:

$$\text{closeness} = \frac{1}{\sqrt{\sum_d d^2 D(d)}} \quad (3)$$

To explain this, we look at the denominator, which is similar to the *standard deviation* as the square root of *variance*. However, we do not want this to be centred about the mean as in the typical formulation. A game that always ends tied 0-0 would have a variance of 0, but so would a game that always ends in 5-0 because the outcome is always the same. Yet 5-0 is certainly not a close score. Thus, we centre the second moment around 0 since close games are those where the final score differences are almost 0. Finally, we take the inverse because we want the metric to increase as the scores become closer and decrease as the scores become further apart.

This formulation mirrors the well-known term “close game”, and the values of closeness have some intuitive meaning. Closeness approaching 0 means that the final score differences are very spread out. Closeness approaching ∞ means the scores are effectively always tied. A closeness of C means that a majority of the score differences will fall between $-1/C$ and $1/C$. If a game can only have a score difference of -1 or 1, its closeness will be exactly 1, no matter if it is biased or unbiased. If we also allow tie scores (score differences of -1, 0, or 1), we would expect the game to have more closeness – in fact for this case closeness will always be > 1 .

3 Rolling Sorted or Unsorted

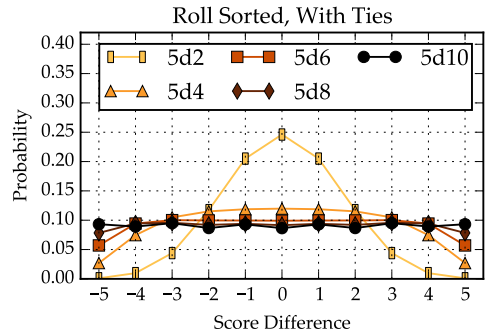
Many games ask the players to roll a handful of dice. A method to assign the dice into pairs is required. *Risk* sorts the dice in numerical order, from largest value rolled to smallest, which is the approach we will take here. We also consider

games where the dice are rolled one at a time (or perhaps one die is rolled several times) and left unsorted. We now show how these two methods of rolling dice significantly change the distribution of score differences.

3.1 Sorting Dice, With Ties

We first look at the case where each player rolls all n of their k -sided dice and then sorts them in decreasing order. The two sets of dice are then matched and compared. If a player rolls more than one copy of the same number, the relative order of those two dice doesn’t matter.

Figure 2 shows the distribution of score differences when each player rolls $n = 5$ dice and sorts them. We vary k , the number of sides. Ties are allowed, with neither player earning a point. We see the games all have a win bias of 0, as expected from the symmetry where the players have the same rules. Additionally, tie percentage decreases as we increase sides: the more possible numbers to roll, the less likely the players will roll the same values. Increasing sides also decreases closeness, making higher score differences more likely to occur. For the case of 5d8 and 5d10, it’s approximately equally likely to have every score difference: wide differences in scores are equally common to close scores.



Game	win bias	tie %	closeness
5d2	0.00	24.61	0.632
5d4	0.00	11.97	0.384
5d6	0.00	9.91	0.340
5d8	0.00	9.15	0.323
5d10	0.00	8.64	0.315

Figure 2. Rolling 5dk sorted, with ties.

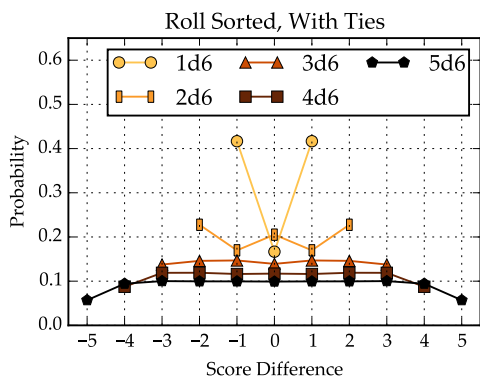
5d2 stands out as having a bell shaped curve with significantly higher closeness: close games are more likely, but ties are also more likely as well. Nonetheless, two-sided dice, which we know as coins, are not typically used in games

²<https://boardgamegeek.com/boardgame/146130/coin-age>

³<https://laboratory.vg/shift/>

partly because standard coins are difficult to toss and keep from rolling off the table (*Coin Age*² and *Shift*³ are notable counter-examples, and some countries use square coins). However, stick dice – elongated dice that only land on the two long sides – do not roll away, and might be something interesting for more game designers to investigate for future games.

In Figure 3, we see how changing the number of six-sided dice rolled affects the distribution of score differences. They remain symmetric with a win bias of 0, and after 2d6, adding more dice decreases the tie percentage. Closeness decreases as we add more dice, which makes sense as with more dice there is a higher probability of the score differences tending away from 0. 1d6 has a closeness greater than 1, because it allows 0-0 ties as well as games that end 1-0 or 0-1; without any ties, it would be exactly a closeness of 1.



Game	win bias	tie %	closeness
1d6	0.00	16.67	1.095
2d6	0.00	20.52	0.680
3d6	0.00	13.91	0.504
4d6	0.00	11.71	0.404
5d6	0.00	9.91	0.340

Figure 3. Rolling $nd6$ sorted, with ties.

3.2 Dice Unsorted, With Ties

We now examine the case where the dice are rolled and left unsorted. The dice could be rolled one at a time, possibly bringing out more drama as the battle is played out in single die rounds. Both players still roll n dice, but the order they were rolled in is used when comparing, as shown in Figure 4a. As before, the player with the higher value earns a point and if tied then neither player earns a point.

⁴An anonymous reviewer mentioned their preferred method for rolling unsorted $nd6$ is to throw dice against a sloped box lid: the dice line up in a random order as they slide against the lid wall. Occasionally one die might stop against another die instead of the wall; in that case, simply jiggle the lid slightly until they all slide against the wall.

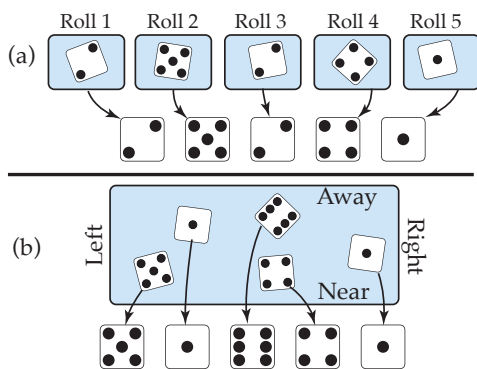



Figure 4. Two methods to roll unsorted dice.

Although we will think of the dice being rolled one at a time (and actually generate them in our simulations this way), it's also possible for players to roll a handful of dice to quickly create a sequence, as shown in Figure 4b. A player first rolls a handful of dice on the table. The dice are then put in order from left to right as they settled on the table. If two dice have the same horizontal position on the table (as the \square and \square do in the example), the die further away from the player will come before the die that is near.⁴



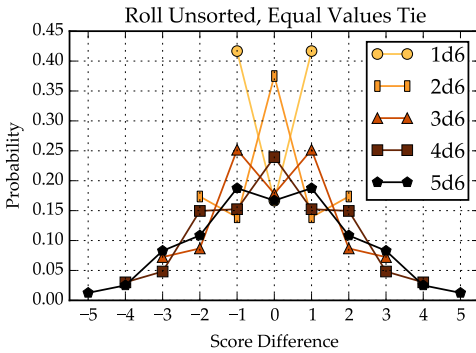
Game	win bias	tie %	closeness
5d2	0.00	24.61	0.632
5d4	0.00	19.32	0.516
5d6	0.00	16.69	0.490
5d8	0.00	14.49	0.478
5d10	0.00	12.71	0.471

Figure 5. Rolling $5dk$ unsorted, with ties.

In Figure 5, we examine how changing the number of sides of dice changes the distribution of ties and close games. We compare 2-sided dice (coins), 4-sided, 6-sided, 8-sided, and 10-sided dice. In all cases, the game is balanced, because the win bias is 0. We can see that more sides decreases the odds of the battle ending in a tie

score. We can also see that more sides decreases closeness and therefore increases the odds of a lopsided victory with more extreme score differences between the players.

In Figure 6, we examine how changing the number of dice rolled affects the score difference. All games are balanced, since the win bias remains 0 for these games no matter how many dice are rolled. Ties are much more common when rolling an even number of dice. When comparing with Figure 3 we see that rolling unsorted increases the percentage of ties. As for closeness, more sides decrease the closeness, as we've also seen when rolling sorted.



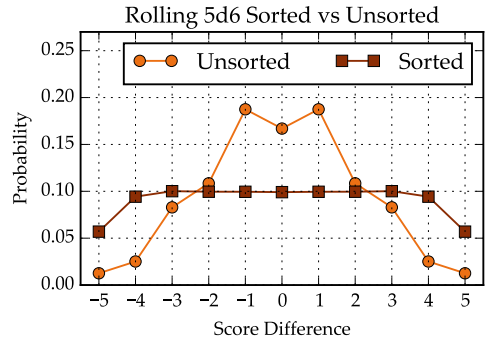
Game	win bias	tie %	closeness
1d6	0.00	16.67	1.095
2d6	0.00	37.50	0.775
3d6	0.00	17.82	0.632
4d6	0.00	23.95	0.548
5d6	0.00	16.69	0.490

Figure 6. Rolling n 6-sided unsorted, with ties.

3.3 Sorted Vs. Unsorted

In Figure 7 we review the effect of changing the way that dice are rolled, while keeping the same number of dice and number of sides. Rolling sorted has a flat distribution that leads to a higher likelihood of larger score differences, while rolling unsorted has a more normal-like distribution where closer games are more likely and closeness is higher. However, higher closeness increase tie percentage.

The game designer can choose the method they find more desirable for the particular game they are creating. In addition to choosing between rolling sorted or unsorted, the designer can change the number of dice and number of sides on the dice. Using fewer sides on the dice increases closeness, but also increases the tie percentage. Using fewer dice increases closeness, but again generally increases the tie percentage. We address ties in the next sections.



Game	win bias	tie %	closeness
Unsorted 5d6	0.00	16.69	0.490
Sorted 5d6	0.00	9.91	0.340

Figure 7. 5d6 rolled sorted vs unsorted.

4 Resolving Tied Battles

In the previous section, when dice were rolled with the same values, neither player received a point for that pair of dice. This leads to some situations where the players get a 0 score difference and tie the game (with as much as 24.6% for the 5d2 case). For games where n is even, a score difference of 0 can occur (becoming less likely as k increases).

A game designer might wish that tie games are not allowed. One simple way would be to have Player A automatically win whenever the battle ends with a score difference of 0 – however this would have a massive bias in favour of Player A. In the above example, this would add an additional 24.6% bias which is likely unacceptable when trying to make the games close. To eliminate the bias over repeated battles, Player A and B could take turns receiving the win (perhaps by using a two-sided disk to indicate who will next receive the tiebreak).

Another simple way that would not have bias would be for the players to flip a coin (or some other random 50% chance event) to decide who is the winner of the battle. Using dice, the players could roll $1dk$ and let the player with the higher value win the battle. If they tie again, they repeat the $1dk$ roll until there is not a tie – we analyze this type of rerolling in Section 8.

In the next few sections, we will examine other ways to change the rules of the game so that score differences of 0 will not occur for games when n is odd. When n is even, score differences of 0 can still occur, and one of the above final tiebreaker methods can be used.

5 Favouring One Player

We now investigate breaking tied dice by always having one player winning a point when two dice are equal. We examine the case where Player A will always win the point (as in *Risk* where defenders always win ties against attackers), but in general the same results apply if A and B are swapped. Favouring one player causes a bias, helping that player win more battles, so we also examine several ways to address this bias.

5.1 Rolling Sorted, Player A Wins Ties

In Figure 8 we see the score distributions that occur when tied dice give a point to Player A. First, we see these distributions are not symmetric, and are heavily skewed towards Player A, as reflected in the positive win bias. As one would expect, giving the ties to Player A causes that player to have an advantage over B. Increasing the number of sides on the die decreases the win bias – this is expected as with more sides on a die, it's less likely for the players to both roll the same number. When n is odd, we also see that even score differences are no longer possible, and most importantly a tied score difference of 0 is no longer possible so tie percentage is always 0%. For the first time, we see an example of closeness *increasing* as the number of sides increases, because the distributions are less skewed towards large 5-0 lopsided wins.

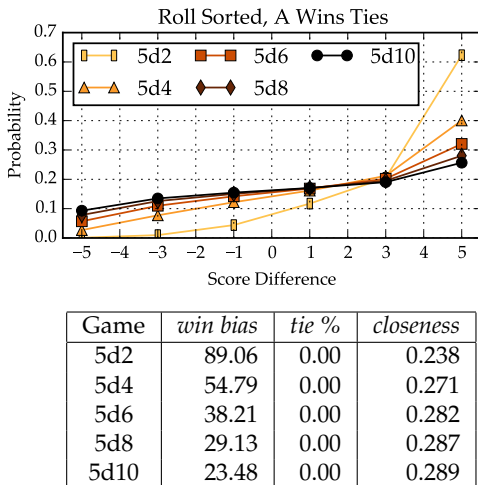
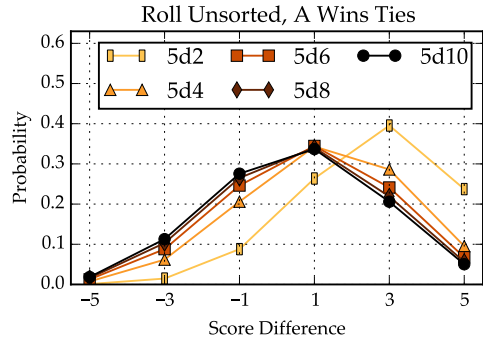


Figure 8. Rolling k -sided dice sorted, A wins ties.

5.2 Rolling Unsorted, A Wins Ties

By switching to rolling dice unsorted, the closeness is increased for all numbers of dice, and the distribution is more centred, but there is still a

significant bias towards Player A, as we can see from Figure 9. This is an improvement, but one might desire another way to eliminate the bias.



Game	win bias	tie %	closeness
5d2	79.30	0.00	0.316
5d4	44.96	0.00	0.400
5d6	30.68	0.00	0.424
5d8	23.19	0.00	0.434
5d10	18.63	0.00	0.439

Figure 9. k -sided dice unsorted, A wins ties.

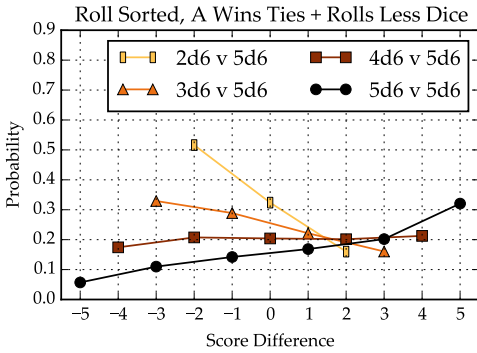
In conclusion, breaking ties in favour of one player eliminates ties, but creates a large win bias. However, this can be reduced with more sides on the dice. This bias occurs for both rolling sorted and unsorted, although rolling unsorted results in higher closeness and slightly lower win bias. We now examine ways to reduce this bias in various ways.

6 Reducing Bias With Fewer Dice

The bias introduced by having one player win ties can be undesirable for some designers and players, so we now look at a method of reducing this bias by having Player A roll fewer dice than Player B, to make up for the advantage they earn by winning ties. This is the strategy used in *Risk*: the winning-ties bias towards the Player A (defender) is reduced by allowing Player B (attacker) to roll an extra die when both sides are fighting with large armies. When rolling sorted, the dice are sorted in decreasing order, and the lowest-valued dice which are not matched are ignored. When rolled unsorted, if one player rolls fewer dice then there is no way to decide which dice should be ignored. We therefore only examine the case of rolling sorted.


We examine the effect of requiring Player A to roll fewer dice in Figure 10. Rolling two or three fewer dice significantly favours Player B, and rolling the same number of dice favours Player A. However, Player A rolling 4d6 against Player B rolling 5d6 has a relatively balanced distribution,

no longer significantly favouring one player over the other. Unfortunately, ties once again occur for 4d6 vs 5d6 – they occur for any battle where A rolls an even number of dice – with a significant likelihood of a final tie score.



Game	win bias	tie %	closeness
2d6 v 5d6	-35.61	32.37	0.608
3d6 v 5d6	-23.63	0.00	0.451
4d6 v 5d6	3.23	20.40	0.357
5d6 v 5d6	38.21	0.00	0.282

Figure 10. Player A rolls fewer dice to control bias.



Game	win bias	tie %	closeness
1d6 v 2d6	-15.74	0.00	1.000
2d6 v 3d6	-7.91	33.58	0.613
3d6 v 4d6	-2.80	0.00	0.450
4d6 v 5d6	3.23	20.40	0.357

Figure 11. Rolling 1 fewer die to control bias.

Since having one fewer die made Player A and Player B relatively balanced when B rolls 5 dice, we can look at more cases when Player B rolls n dice. In Figure 11, we have more cases where Player A has one fewer die than Player B. Most of these cases are relatively balanced, although 1d6 vs 2d6 still gives a significant advantage to Player B. Note that the cases of 1d6 v 2d6 and 2d6 v 3d6 are the ones that occur in *Risk*.


To reduce the win bias introduced by having Player A win all ties, we reduced this bias by having Player A roll fewer dice. As we have seen, rolling one fewer dice is the best choice that leads to the smallest win bias, and having both players roll more dice also reduces the win bias, but decreases the closeness. Instead of having the players rolling different numbers of dice, we now will examine having the players roll different number of sides for the dice.

7 Reducing Bias With Mixed Dice

Another way we can reduce the bias towards Player A when they always win ties is to give Player B some dice with more sides. For example, we could have Player A roll 5 6-sided dice and have Player B roll 3 6-sided dice and 2 8-sided dice, to give them a small advantage to help eliminate the advantage A receives for winning ties. Because bias does not occur when we allow ties, we will only examine using mixed dice for games where Player A wins ties.

7.1 Mixed Dice Sorted, A Wins Ties

In Figure 12, we show the distribution of score differences for different mixes of d6 and d8 for Player B, while Player A always rolls 5d6. We can see that adding more d8 adjusts the bias in favour of Player B, but adding too many then biases B's win rate too far.



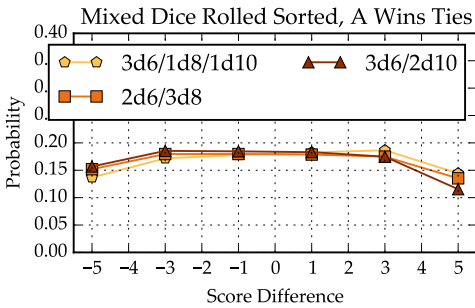
Game	win bias	tie %	closeness
5d6/0d8	38.21	0.00	0.282
4d6/1d8	24.36	0.00	0.294
3d6/2d8	10.80	0.00	0.302
2d6/3d8	-2.24	0.00	0.305
1d6/4d8	-14.56	0.00	0.305
0d6/5d8	-25.98	0.00	0.301

Figure 12. Mixed d6 and d8 to control bias.

The most balanced position is to have Player B roll 2d6 and 3d8 against Player A's 5d6 (this is

drawn as a solid line in the figure), with win bias of -2.24%.

We tried all possible mixes of 5 dice made of 6-sided, 8-sided, and 10-sided dice, and found that only 3 cases have a win minus loss bias under 10%; these cases are shown in Figure 13. The bias is still most balanced when Player B rolls 2d6 and 3d8 against Player A's 5d6. However, by rolling 3d6/1d8/1d10, we can get a slight bias towards Player A, if that is desired.



Game	win bias	tie %	closeness
3d6/1d8/1d10	2.67	0.00	0.307
2d6/3d8	-2.24	0.00	0.305
3d6/2d10	-5.36	0.00	0.310

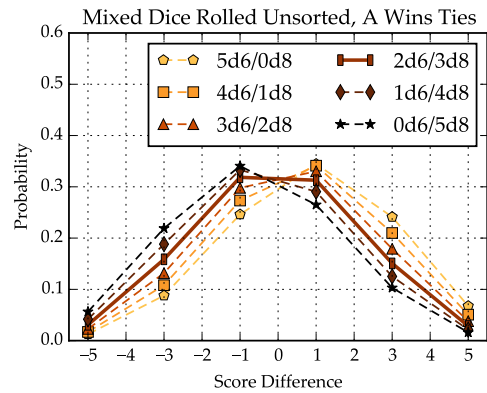
Figure 13. Least biased mixes of d6, d8, and d10.

7.2 Mixed Dice Unsorted, A Wins Ties

We can do the same type of experiment for all variations of Player B rolling unsorted a mix of 5 d6s and d8s against Player A's 5d6, getting the results as shown in Figure 14. By using 2d6 and 3d8, we can reduce the bias down to a small 1.61% in favour of Player B.

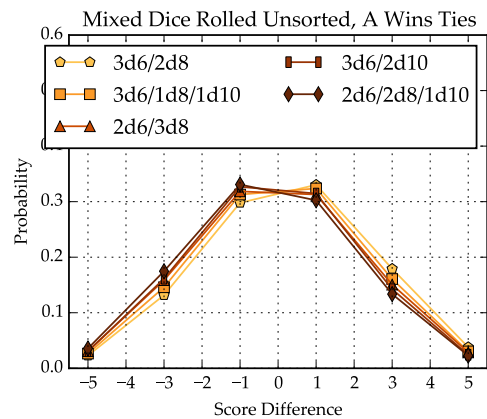
By trying all variations of 5 d6s, d8s, and d10s, we find that there are 5 cases where the bias is kept under 10%, which are shown in Figure 15. Rolling 2d6 and 3d8 is still the lowest overall bias; rolling 3d6/1d8/1d10 is the lowest bias that favours Player A.

In conclusion, we can reduce the win bias by having the unfavoured player roll different sided dice. Looking at all mixes of five dice composed of d6, d8, and d10, we found that rolling 5d6 against 2d6/3d8 produced the smallest win bias, for both rolling sorted and unsorted. In fact, there was no way to completely eliminate the win bias. Nonetheless, we will examine one final way to break ties that will lead to a zero win bias.



Game	win bias	tie %	closeness
5d6/0d8	30.68	0.00	0.424
4d6/1d8	20.34	0.00	0.440
3d6/2d8	9.48	0.00	0.450
2d6/3d8	-1.61	0.00	0.452
1d6/4d8	-12.60	0.00	0.446
0d6/5d8	-23.19	0.00	0.434

Figure 14. Mixed d6 and d8 to control bias.



Game	win bias	tie %	closeness
3d6/2d8	9.48	0.00	0.450
3d6/1d8/1d10	2.97	0.00	0.456
2d6/3d8	-1.61	0.00	0.452
3d6/2d10	-3.73	0.00	0.459
2d6/2d8/1d10	-8.26	0.00	0.453

Figure 15. Mixing d6, d8, and d10 to control bias.

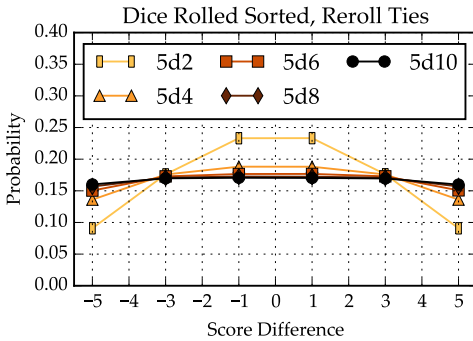
8 Rerolling Tied Dice

We now examine rerolling tied dice as a final way to break ties. For example, it is quite common to reroll 1d6 at the start of a game to decide who goes first. This can be generalized to ndk , but it is quite cumbersome and this section exists mainly

as an explanation on why we believe this is inadvisable in practice. Because rerolling can go on for many iterations, we use Monte-Carlo simulation to evaluate the odds empirically instead of exactly, since these games can theoretically continue indefinitely with increasingly unlikely probability. We used $N=6^{10}=60,466,176$ simulations per game, as this is the same number of cases that evaluated for the other sections (see Appendix for this calculation). When these are simulated and not exact values, we use the \approx symbol in the figures.

8.1 Rolling Sorted, Rerolling Tied Dice

We first examine the case where we roll a handful of dice and then sort them from highest to lowest. Any dice that are not tied are scored first. Then, any remaining dice that are tied are rerolled by both players at once in a sub-game.



Game	win bias	tie %	closeness
5d2	≈ 0	0.00	≈ 0.350
5d4	≈ 0	0.00	≈ 0.311
5d6	≈ 0	0.00	≈ 0.302
5d8	≈ 0	0.00	≈ 0.298
5d10	≈ 0	0.00	≈ 0.296

Figure 16. Rerolling ties with 5dk sorted.

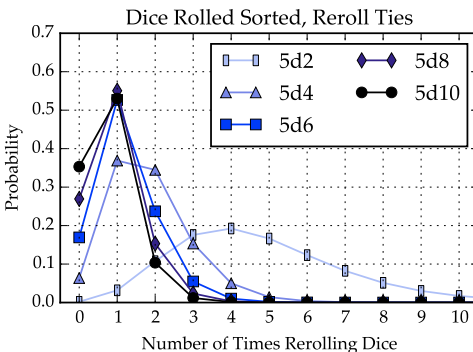
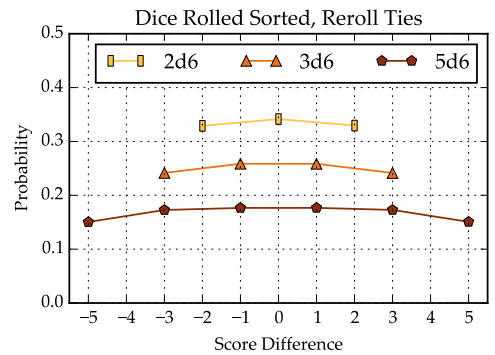


Figure 17. Reroll probability for Figure 16.

This process is repeated for any remaining tied dice in the sub-game, until there are no more ties. All the scores from the first game and all sub-games are summed together for the final score.

The resulting score difference distributions are shown in Figure 16. The battles are all unbiased and without ties. For 5d4, 5d6, 5d8, and 5d10 the distributions are effectively flat with low closeness and have approximately the same shape as when rolling sorted with ties (as in Figure 2) but now do not permit tie games. Compared to 5d4 and higher, 5d2 has a higher closeness. However, this closeness comes at a significant cost of requiring many rerolls, as demonstrated in Figure 17. This shows that more sides decreases the probability of a reroll, and with 5d2 or 5d4 there are significant chances at rolling 2 or more rerolls for a single battle, which could be cumbersome for the players in practice. Higher sided dice are less likely to tie, so the probability of rerolling decreases quickly when using six or more sides.



Game	win bias	tie %	closeness
2d6	≈ 0	≈ 34.16	≈ 0.616
3d6	≈ 0	0.00	≈ 0.453
5d6	≈ 0	0.00	≈ 0.302

Figure 18. Rerolling ties with $nd6$ sorted.

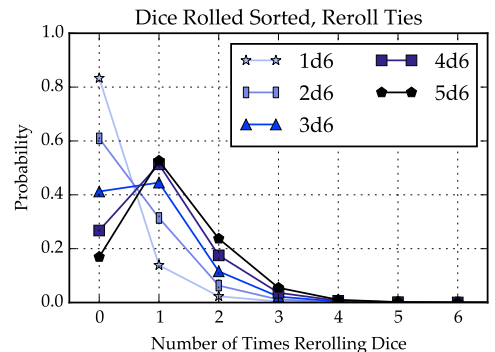


Figure 19. Reroll probability for Figure 18.

We also examine the effect of changing the number of dice while holding the number of sides fixed in Figure 18. The distributions are all flat, but closeness can be increased by using fewer dice, as we've seen in previous sections. The probability of rerolls is also affected by the number of dice, as shown in Figure 19. For 1d6 and 2d6, the most common outcome is no rerolls. Increasing the number of dice makes rerolls more likely, but the probabilities of having additional rerolls decreases rapidly.

8.2 Rolling Unsorted, Rerolling Ties

Finally, we examine the case of rolling n k -sided dice unsorted when rerolling ties. The dice are rolled one at a time, and any time there is a tie, the two dice must be rerolled until they are no longer tied. This occurs for each of the n dice. In practice, this is unlikely to be much fun for the players, but we present the analysis here for completeness.




Figure 20. Rerolling ties for 5dk unsorted.

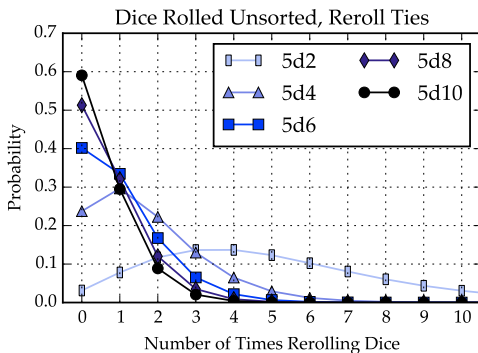


Figure 21. Reroll probability for Figure 20.

Because neither player is favoured, the metrics can be analytically calculated from the binomial distribution $\binom{n}{w} p^w (1-p)^{n-w}$ with w being the number of wins for Player A in the battle, n dice rolled, and probability $p = .5$ (no matter the value of k) of Player A winning each point. Given a score difference d , we can calculate $w = (n + d)/2$.

In Figure 20 we see that the score difference distribution is identical for all dice, no matter how many sides. The battle is unbiased, with no ties, and has a closeness of .447. However, they do not have the same number of rerolls, as shown in Figure 21, generated with Monte Carlo simulation. To reduce rerolls, the game designer can use higher sided dice.

The closeness can be increased by reducing the number of dice rolled, as shown in Figure 22. Rolling fewer dice also reduces the probability of rerolls, as shown in Figure 23.

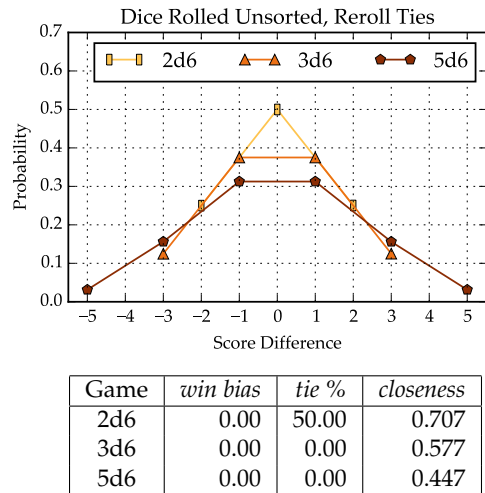


Figure 22. Rerolling ties for $nd6$ unsorted.

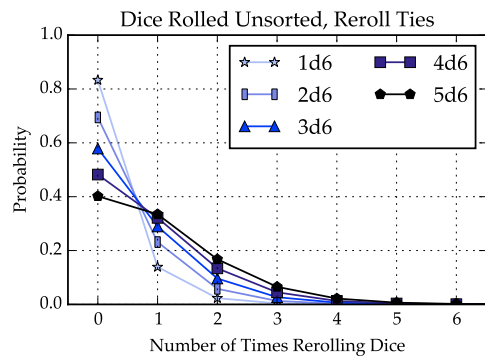


Figure 23. Rerolling ties when rolling unsorted.

8.3 Sorted, A Wins Ties, Rerolls Highest

We can make a hybrid case, where A wins all ties but must reroll when A rolls the die's highest value (e.g. a 6 on a 6-sided die). This effectively means that Player A is rolling a $k - 1$ sided die while Player B is rolling a k sided die. This gives an advantage to Player B to make up for the advantage that Player A has when breaking ties.

Interestingly, this has the same effect as in the previous reroll sections, for both rolling sorted or unsorted. Therefore, the plots are the same as in Figures 16, 18, 20, and 22. However, only Player A has to reroll dice, and Player B can keep the dice untouched no matter what they roll. Therefore, there are many less rerolls in total.

We can show analytically why this is unbiased for the simple case of one k -sided die. Player A will reroll when rolling a k , which is the same as rolling a $k - 1$ sided die. If Player A rolls a value of i with probability $1/(k - 1)$, then they win when Player B rolls a value $\leq i$ with probability i/k , since A wins ties. Calculating the expected number of wins for Player A, over all values of i from 1 to $k - 1$ we have:

$$\sum_{i=1}^{k-1} \frac{1}{k-1} \frac{i}{k} = \frac{1}{k(k-1)} \frac{(k-1)(k)}{2} = \frac{1}{2} \quad (4)$$

which is independent of k , and always $1/2$.

When breaking ties by rerolling, we get unbiased results, but at the cost of requiring the players to reroll, which can take longer. However, by using higher sided dice or fewer dice, the designer can mitigate the expected number of rerolls. Because other tie-breaks presented in this paper do not require extra rolls, we suggest following another approach to breaking ties.

9 Risk & Risk 2210 A.D.

We can use the results of this paper to examine how the original Risk compares with the popular variant Risk 2210 A.D. [19]. In both games, the players roll sorted dice and the defender wins tied dice, which we showed in Section 5.1 gives a strong advantage to the defender when rolling the same number of dice. A game with the defender having an advantage can lead to a static game where neither player wants to attack.

To counteract this, both games allow the attacker to roll an extra die (3d6 v 2d6). We show in Section 6 this flips the advantage towards the attacker. This advantage encourages players to play more aggressively, as it's better to be the attacker than the defender.

In Risk 2210 A.D. special units called commanders and space stations will swap in one or more d8 instead of d6 when engaging in battles.

As we showed in Section 7, using mixed dice biases the win rate towards the player rolling higher valued dice, which can be either used by attackers to have a stronger advantage (less closeness but more predictability) or by defenders to even out the bias inherent in letting the attacker roll more dice.

As we've shown in this paper, the rules in dice games require careful balancing as the exact number of dice and number of sides can often have a large impact on the statistical outcome of the battles. Risk and Risk 2210 A.D. are no exception and they appear to have carefully tuned dice mechanics to have reasonable win bias and closeness values.

10 Conclusion

We have demonstrated the use of *win bias*, *tie percentage*, and *closeness* to analyze a collection of dice battle variants for use as a component in a larger game. We introduce closeness, which is related to the precision statistic about 0, and matches the intuitive concept of a game being close. We have not seen this statistic used before to analyze games.

By examining the results of the previous sections, we can make some general statements about this category of dice battles where the number values are compared.

In Section 3, we showed that when allowing ties, rolling dice unsorted results in higher closeness, and therefore a lower chance of games with large point differences; however, this comes at the cost of increasing the tie percentage. Using fewer sides on the dice increases closeness, but also increases the tie percentage. Using fewer dice increases closeness, but again generally increases the tie percentage.

Battles that end tied with a score difference of 0 can be broken with a coin flip or other 50/50 random event, as discussed in Section 4. However, we also wanted to explore rule changes that would cause odd-numbers of dice to never end in a tie. Breaking ties in favour of one player, as shown in Section 5, eliminates ties but creates a large win bias, although this can be reduced with more sides on the dice. This bias occurs for both rolling sorted and unsorted, although rolling unsorted results in higher closeness and slightly lower win bias.

To reduce this win bias, in Section 6 we have the favoured player roll fewer dice. Rolling one fewer die is the best choice that leads to the smallest win bias, and having both players roll more dice also reduces the win bias (but decreases the closeness).

In Section 7, we reduced the win bias by hav-

ing the unfavoured player roll different sided dice. Looking at all mixes of five dice composed of d6, d8, and d10, rolling 5d6 against 2d6/3d8 produced the smallest win bias, for both rolling sorted and unsorted. However, there was no way to completely eliminate the win bias.

Finally, in Section 8, we examined a method of breaking ties by rerolling. This gives unbiased results, but at the cost of requiring the players to reroll, which can take longer. By using higher sided dice or fewer dice, the designer can reduce the expected number of rerolls that will occur, although we recommend other tie breaking methods that are less cumbersome for the players.

One surprising outcome of this study is that $nd2$ sorted with ties may be an under-used dice mechanic for games. This has high closeness, and can be easily done by throwing 2-sided coins or stick dice, subtracting the number of heads from the number of tails. Stick dice do not have the practical problems that round coins do, as dice sticks with flat edges don't roll off the table easily.

The analysis in this paper focuses on comparing dice values, but we are also doing a similar study of hit-based dice games including analyzing the effect of critical hits, following the same framework presented here.

Additionally, for finer grained control over the experience, a game can instead use a bag of dice tokens (e.g. small cardboard chits with a dice face printed on them) or a deck of dice cards to enforce that certain distributions are obeyed with local representation – this is choosing without replacement instead of the typical choosing with replacement that occurs with dice. We are currently experimenting with examining similar games that use bags of dice tokens, using an exhaustive analysis similar to that done here.

In summary, there is no perfect solution to the dice battle mechanic, and a designer must make a series of tradeoffs. We hope that this paper can provide some quantitative guidance to a designer looking for a specific type of game feel when using dice. For a designer that wishes to use rules that we did not discuss in this paper, we hope it would not be difficult to use the same technique to evaluate how the players might experience the distribution of score differences by measuring win biases, ties, and closeness.

Acknowledgements

We wish to thank Steven J. Brams and Mehmet Ismail for comments on an early draft of the paper, and to the editor-in-chief Cameron Browne and anonymous reviewers for their insightful reviews and constructive suggestions.

References

- [1] *Risk!*, Parker Brothers, 1963.
- [2] *Axis and Allies*, Milton Bradley, 1984.
- [3] Tversky, A. and Kahneman, D., 'Belief in the Law of Small Numbers', *Psychological Bulletin*, vol.76, no. 2, 1971, pp. 105–110.
- [4] Bar-Hillel, M. and Wagenaar, W.A., 'The Perception of Randomness', *Advances in Applied Mathematics*, vol. 12, no. 4, 1991, pp. 428–454.
- [5] Nickerson, R.S., 'The Production and Perception of Randomness', *Psychological Review*, vol. 109, no. 2, 2002, pp. 330–357.
- [6] Bell, R.C., *Board and Table Games from Many Civilizations, Vol. 1 & 2*, Courier Corp., 1979.
- [7] Knizia, R., *Dice Games Properly Explained*, Blue Terrier Press, 2010.
- [8] Jay, R. "The Story of Dice: Gambling and Death from Ancient Egypt to Los Angeles", *New Yorker*, December 11, 2000, pp. 90-95.
- [9] Falkener, E., *Games Ancient and Oriental and How to Play Them*, 1892, pp. 85-86.
- [10] Elias, G. S., Garfield, R., Gutschera, K. R. and Whitley, P., *Characteristics of Games*, Massachusetts, MIT Press, 2012.
- [11] Costikyan, G. *Uncertainty in Games*, MIT Press, 2013.
- [12] Fehr, E. and Schmidt, K. M., 'A Theory of Fairness, Competition, and Cooperation', *Quarterly Journal of Economics*, vol. 114, no. 3, 1999, pp. 817–868.
- [13] Isaksen, A., Ismail, M., Brams, S.J., and Nealen, A. 'Catch-Up: A Game in Which the Lead Alternates', *Game and Puzzle Design*, vol. 1, no. 2, 2015, pp. 38–49.
- [14] Flick, J., 'AnyDice Dice Probability Calculator', <http://anydice.com>.
- [15] Mogensen, T. Æ. 'Troll, a Language for Specifying Dice-rolls.' *ACM Symposium on Applied Computing*, 2009, pp. 1910-1915.
- [16] Browne, C., *Automatic Generation and Evaluation of Recombination Games*, PhD Dissertation, Queensland University of Technology, 2008.
- [17] Althöfer, I., 'Computer-Aided Game Inventing', Friedrich Schiller Universität, Jena, Germany, Tech. Rep., 2003.
- [18] Brualdi, R. A., *Introductory Combinatorics*, 5th edition, 2009, pp 52–53.
- [19] *Risk 2210 A.D.*, Avalon Hill, 2001.

Aaron Isaksen is a PhD Candidate in Computer Science at the NYU Tandon School of Engineering, researching computer-aided game design, and a professional independent video game developer. **Address:** NYU Tandon School of Engineering, 6 MetroTech Center, Brooklyn, NY 11201, USA. **Email:** aisaksen@nyu.edu

Christoffer Holmgård is a Postdoctoral Associate at the NYU Tandon School of Engineering. He researches player modeling, human-like agents, procedural content generation, and applied games and does video game development. **Address:** NYU Tandon School of Engineering, 6 MetroTech Center, Brooklyn, NY 11201, USA. **Email:** holmgard@nyu.edu

Julian Togelius is Associate Professor of Computer Science at the NYU Tandon School of Engineering and co-director of the Game Innovation Lab. **Address:** NYU Tandon School of Engineering, 6 MetroTech Center, Brooklyn, NY 11201, USA. **Email:** julian.togelius@nyu.edu

Andy Nealen is Assistant Professor of Computer Science at the NYU Tandon School of Engineering, co-director of the Game Innovation Lab, and co-creator of the video game Osmos. His research interests are in game design and engineering, computer graphics and perceptual science. **Address:** NYU Tandon School of Engineering, 6 MetroTech Center, Brooklyn, NY 11201, USA. **Email:** nealen@nyu.edu

Appendix

In this appendix, we give analytical results for the probabilities and number of possible outcomes for many of the games studied in this paper. A more complete coverage of these probabilities and combinatorics can be found in [18].

A fair k -sided die has equal probability of rolling each of its k sides, so the probability of rolling any particular number is $1/k$. Therefore, the total probability of rolling a value v or less is $\sum_{i=1}^v 1/k = v/k$.

If we roll n dice unsorted, there are k^n dif-

ferent ways to roll the dice. Each way of rolling the dice, since the order matters, has an equal k^{-n} chance. For example, if we roll 5 6-sided dice unsorted, there are $6^5 = 7,776$ possible outcomes each with $1/7,776$ probability. If Player A is rolling a dice, and player B is rolling b dice, then there are $k^a k^b = k^{a+b}$ possible outcomes. So, if each side rolls 5 6-sided dice unsorted, there are $6^{10} = 60,466,176$ possible games that can occur, each equally likely. Rolling 5d10 against 5d10 has 10,000,000,000 different possible outcomes.

If we roll the n dice sorted, then we can describe the probabilities using the *multinomial distribution*, a generalization of the binomial distribution when there are k possible outcomes for each trial. If one knows the outcome of a sorted roll had x_i copies of i (i.e. x_1 1's, x_2 2's, etc.), such that $x_1 + x_2 + \dots + x_k = n$, the number of ways that particular outcome could have been rolled is:

$$\frac{n!}{x_1!x_2!\dots x_k!} \quad (5)$$

The probability of rolling that outcome is:

$$\frac{n!}{x_1!x_2!\dots x_k!} k^{-x_1x_2\dots x_k} \quad (6)$$

For rolling n k -sided dice sorted, the number of different possible results a player can roll is:

$$\binom{n+k-1}{k-1} \quad (7)$$

For example, for 5d6, there are $\binom{5+6-1}{6-1} = \binom{10}{5} = 252$ unique ways to roll the dice, although these are not of equal probability. For two players, there are $252^2 = 63,504$ ways to evaluate the game. This means that the rolling sorted calculations can be made much faster by only calculating each unique outcome once, but then multiplying the results by Equation 5, the number of ways each result can occur.