AN INTRODUCTION TO NONTRANSITIVE GAMES

ZACH STRONG

Date: November 29, 2023.

Contents

1. Introduction

For many games, there exist "things" that are more "powerful" than all other "things". What exactly these "things" are will vary depending on the game. Take, for instance, the game of poker. The rarity of a poker hand dictates its power. The royal flush—an ace, king, queen, jack, and 10, all of the same suit—is the rarest of all named hands, and is thus the most powerful, beating all other hands. The second rarest is the straight flush—a sequence of five numerically consecutive cards of the same suit—and is thus the second most powerful, beating all hands except the royal flush. The third most powerful hand—the four-of-a-kind—beats all hands except the royal flush and straight flush. This pattern continues, with each named hand beating all hands more common than it, yet losing to all hands less common.

As another example, take the game *Scrabble*. Ignoring any score bonuses resulting from specific tile placements on the board, each valid *Scrabble* word is given a score based on the letters that compose it. If the power of a word is taken to be its score, then all valid *Scrabble* words can be ranked according to their power. Thus, there must be (at least) one word sitting at the top of this hypothetical hierarchy which is at least as powerful as every other word.[1]

The specifics regarding these most powerful "things" aren't important in these examples. What *is* important to recognise is that these most powerful "things" exist at all. Initially, this observation seems obvious. After all, in any situation where objects can be ranked in terms of power (whatever that definition of power may be), it seems reasonable, maybe even trivial, to assume that one object must be more powerful than the rest. There are endless examples of such situations. If we take our definition of power to be preference—that is, out of a set of objects, the one that people prefer to choose the most often is deemed the most powerful—then any plurality-based election system can be viewed as an example of such a situation. Clearly, elections always result in one candidate being ranked the highest, and thus attaining office. It seems pointless, therefore, to draw attention to the existence of a "most powerful" object.

However, there's a very good reason for considering these "most powerful" objects: they don't always exist. In fact, some of the most common situations where objects are ranked by some form of power don't always have a "most powerful" object... including the plurality-based election system mentioned above! This isn't merely the case of a tie, either. There are certain situations where, when pairing together any two different objects, there will always be one object which is more powerful,

[1] According to `scrabblewizard.com`, ref [2023] the most powerful *Scrabble* word is oxyphenbutazone, worth 41 points.

but when considering all the objects together, there isn't a clear answer for which object is the *most* powerful.

This project aims to introduce the concept of a "nontransitive" game, showing how the property of "nontransitivity" can give rise to counter-intuitive situations where, despite a clear notion of "power" existing, no one object can be considered the most powerful. Some specific games will be introduced and explained in sections 3-5 and 7, paying extra attention to aspects of the games that are nontransitive. Along the way, some specific results will be referenced or proved relating to optimal strategies and surprising properties that arise as a result of nontransitivity. Then, section 8 aims to detail some real-world consequences related to nontransitivity, such as voting paradoxes and cyclical food chains.

Ultimately, this project aims to demystify the apparent contradictions of nontransitivity, instead showing how it is a very natural property that arises in systems involving object comparisons. This project also aims to be as self-contained as possible, though numerous references are provided for the curious reader.

2. Transitivity

Before tackling the seemingly-strange situations involving nontransitivity, it helps to first refresh our memory on the concept of transitivity.

Say we're given a set of characters

$$C = \{☺, ☯, 📱, ★, ☻\}$$

where each character has a designated credit score:

- ☺ has a credit score of 44.
- ☯ has a credit score of 17.
- 📱 has a credit score of 0 (because cell phones don't have credit scores).
- ★ has a credit score of 64.
- ☻ has a credit score of 67.

If $a, b \in C$, then we say ”$a \geq b$” if a has an equal or higher credit score than b. Similarly, we say ”$a \leq b$” if a has an equal or lesser credit score than b. For instance, the following two statements are true:

$$📱 \leq ☺, \quad ☺ \leq ★.$$

Given our credit-score-comparing relation "$\leq$", we can ask whether it is transitive or not. Any relation "R" is said to be *transitive* if, given any three objects a, b, and

c, the following statement is always true:

$$\text{If } a\text{R}b \text{ and } b\text{R}c, \text{ then } a\text{R}c.$$

In other words, if a is related to b, and b is related to c, then it must be the case that a is related to c.

For our credit score example, we can see that the ”$\leq$” relation is indeed transitive. From above, we saw that

$$☎ \leq ☺$$

since the phone’s credit score of 0 is less than the smiley’s credit score of 44. Similarly, we also saw that

$$☺ \leq ★$$

since the smiley’s credit score is less than the star’s credit score of 64. Then, it’s clear to see that

$$☎ \leq ★$$

since the phone’s credit score was less than the smiley’s, which was a lesser value than the star’s, implying that the phone’s credit score must necessarily be lower than the star’s. It’s easy to verify that this logic holds for any characters in the set C. Therefore, if $a, b, c \in C$ and we have that

$$a \leq b \text{ and } b \leq c,$$

then it must be the case that

$$a \leq c.$$

This means our ”$\leq$” relationship is transitive.

Transitive relationships are fairly common. A simple example of a transitive relation is height comparison. If Alice is taller than Bob, and Bob is taller than Chris, then it must be the case that Alice is taller than Chris (see figure 1).

The reason for introducing transitivity via our convoluted character set C is to demonstrate how transitive relations can be defined on any objects we wish, be them numbers, words, political figures, etc. No matter how complicated the setup, as long as our defined relation obeys the criteria for being transitive, it’s transitive. Similarly, if the relation doesn’t obey the criteria for transitivity, it isn’t transitive, no matter how much it *seems* to be transitive.

Once we’ve confirmed that a relationship is transitive, we can use the relationship to impose a sort of ranking on our objects.[2] From the table of credit scores given

[2]If there are certain objects in our set that can’t be directly compared using the relevant transitive relation—as may be the case for a preference relationship where preferences for certain objects

FIGURE 1. From left to right: Alice, Bob, and Chris. Clearly, Alice is taller than Bob, and Bob is taller than Chris. Because height comparisons are transitive, we can conclude that Alice is taller than Chris.

for our characters, it isn't difficult to see that the following relationship holds:

$$📱 \leq ☯ \leq ☺ \leq ★ \leq ☻.$$

From this, we can conclude that ☻ has the highest credit score, and is therefore the "highest ranked" using our credit score relation. The next highest would be ★, who beats all characters except for ☻. This situation mirrors the case of poker hands described in section 1, where all objects can be put onto a hierarchy with the "more powerful" object resting at the top. Whenever we have a transitive relationship, there naturally arises such a hierarchy, and thus a most powerful object also arises.

However, it's important to note that there are plenty of relationships that are not transitive. Romantic relationships happen to be one such relationship. If Alice loves Bob, and Bob loves Chris, we *cannot* then conclude that Alice loves Chris. Section 8 of this project delves into some more non-game examples of nontransitivity.

We conclude this section by recognising that both transitive and nontransitive relationships are not only possible, but reasonable. While there may be reasons for believing a certain relationship to be transitive, unless it's consistent with the definition of transitivity above, we cannot assume transitivity. In the next section, we'll introduce a relationship between dice that'll help us analyse a particular game and determine which player in this game has a better chance of winning. Although this

are unknown or unspecified, for instance—creating a ranking among objects may become more complicated. The relationships considered in this project will never have this problem, though, so we can safely ignore this nuance.

new relationship we introduce may seem obviously transitive at first, we'll show that this assumption is false.

3. Efron Dice

3.1. A Simple Game of Dice. Let's say a friend approaches us with a set of dice and asks to play a game.[3] The game is deceptively simple. Both players choose a die. Once the dice are chosen, the players roll them at the same time. Whoever's die displays the highest number face-up wins. This friend is supposedly feeling quite generous today and allows us to pick our die before they do. Looking at the available dice, we notice that they happen to be quite unusual:

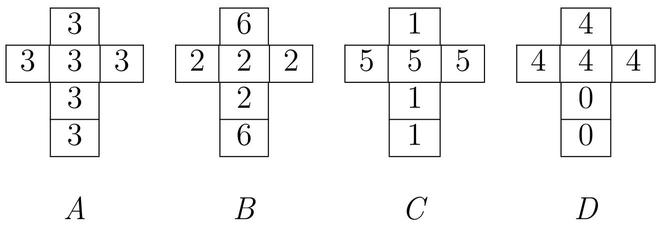

Figure 2. The nets for the four Efron Dice, named after the statistician Brad Efron who created them.

Looking at these dice, it's reasonable to assume that there must exist some "most powerful" die—that is, one that gives us the highest probability of rolling a number higher than our friend's die, whichever they happen to choose. However, as we've been thrown into this game so quickly, we have no time to calculate any of the probabilities.

Suppose we choose die A, maybe because it seemed more "balanced" than the others. Our friend then chooses die D, and we play a few rounds. After nine rolls of the dice, our friend has won six of the nine rolls—about 67%. While this percentage is somewhat close to a fair 50% (which we would expect if both dice had the same probability of winning), it's enough to suggest that maybe our friend's die has the advantage. So, we ask to use our friend's die. Surprisingly, they agree! Our friend then swaps to using die C, and we play a few more rounds.

Frustratingly, the same imbalance exists; our friend is still winning about 67% of the time. In vain, we try swapping to the other two dice, but it doesn't make a difference.

[3]This particular game is derived from Grime [2017].

Finally, we admit defeat. Our friend walks away with the glory, leaving us to brood.[4]

Thankfully, however, we happen to have a sharp memory and can remember exactly what the four dice looked like. We decide to do a few probability calculations to figure out exactly which die is the "more powerful" one. From our observations, it appeared that the dice, while having unusual numbering on their faces, weren't rigged in any way. Each face on each die was equally likely to land face-up. Each die had six faces, so we'd expect that each face appears face-up a sixth of the time. More formally, if D represents the outcome that a particular face on our die lands face-up, then this probability can be expressed as

$$\mathrm{P}(D) = \frac{1}{6}$$

for each possible outcome D of the die.[5]

As well, since the outcome of one die didn't affect the other, our friend's die was independent of our own. Therefore, if D_1 represents the outcome of our die roll, and D_2 represents the outcome of our friend's die roll, then

$$\mathrm{P}(D_1 \cap D_2) = \mathrm{P}(D_1)\mathrm{P}(D_2) = \frac{1}{6} \times \frac{1}{6} = \frac{1}{36},$$

and so the probability that each *pair* of faces lands face-up is as likely as any other pair of faces landing face-up.[6]

This gives us a relatively simple way to compare dice to each other: since each possible pair of outcomes between the two dice has an equal chance of occurring, if we list out all these pairs and count how many times one die rolls higher than the other (and vice versa), we can explicitly calculate which of the two dice rolls higher more often, i.e. which die is "more powerful".

---

[4]The author would like to apologise in advance for the dramatics.

[5]Here, we're implicitly using what's called the *principle of indifference* to assign probabilities to each outcome. The principle of indifference states that, when we have n possible outcomes and no reason to believe that one outcome should be more likely than any other, we assign each outcome a probability of $\frac{1}{n}$ so as to ensure that each outcome is equally likely, and to ensure that our defined probabilities are consistent with the probability axiom that $\mathrm{P}(S) = 1$ for our sample space S (see Weisstein [2023]). Here, the principle applies since each die is fair, meaning the likelihood of any outcome occurring *must* be as likely as any other outcome occurring. See pages 288-291 of Gardner [2001] for a more in-depth consideration of this principle.

[6]Intuitively, this can be interpreted as: for every outcome of the first die, there's a $\frac{1}{6}$th chance of it occurring. When that outcome does occur, there's yet another $\frac{1}{6}$th chance for the second die, since the first outcome occurring doesn't affect what outcome the second die will have. So, for $\frac{1}{6}$th of the time, we'd expect that $\frac{1}{6}$th of *that* time, both outcomes will occur, and $\frac{1}{6}$th of $\frac{1}{6}$th of the time is $\frac{1}{6} \times \frac{1}{6} = \frac{1}{36}$. In effect, independence of two events implies that the probability that both of them will occur is the product of their individual probabilities of occurring.

3.2. **Constructing a Hierarchy of Dice.** Equipped now with a thorough-enough understanding of the dice game, we begin creating tables of possible pairs of dice rolls to determine which die we should use the next time our friend challenges us. For convenience, we'll refer to these tables as "outcome tables".

Let's start by comparing die A and die B (see figure 2). If the tuple (D_1, D_2) represents the result of a pair of dice being rolled, where D_1 is the outcome of the first die and D_2 is the outcome of the second die, then for die A and B, the following dice roll outcomes are possible:

B \ A	3	3	3	3	3	3
2	(2, 3)	(2, 3)	(2, 3)	(2, 3)	(2, 3)	(2, 3)
2	(2, 3)	(2, 3)	(2, 3)	(2, 3)	(2, 3)	(2, 3)
2	(2, 3)	(2, 3)	(2, 3)	(2, 3)	(2, 3)	(2, 3)
2	(2, 3)	(2, 3)	(2, 3)	(2, 3)	(2, 3)	(2, 3)
6	(6, 3)	(6, 3)	(6, 3)	(6, 3)	(6, 3)	(6, 3)
6	(6, 3)	(6, 3)	(6, 3)	(6, 3)	(6, 3)	(6, 3)

Here, the first entry of each tuple represents the roll of die B, while the second entry represents the roll of die A (as denoted by the top-left cell of the outcome table). Simply by counting the number of tuples where die A's roll is higher than die B's, we see that die A beats die B in 24 out of the possible 36 outcomes. Because each outcome is equally likely, we can write that

$$\mathrm{P}(A > B) = \frac{24}{36},$$

where $\mathrm{P}(A > B)$ denotes the probability that die A's roll is greater than die B's roll. Similarly,

$$\mathrm{P}(B > A) = \frac{12}{36}.$$

It seems that die A is "more powerful" than die B, in the sense that die A is more likely to roll higher than die B. To represent this, we'll use the following notation:

$$A \succ B.$$

Now, let's compare the relative powers of die B and die C. Creating an outcome table for dice B and C, we get

C \ B	2	2	2	2	6	6
1	(1, 2)	(1, 2)	(1, 2)	(1, 2)	(1, 6)	(1, 6)
1	(1, 2)	(1, 2)	(1, 2)	(1, 2)	(1, 6)	(1, 6)
1	(1, 2)	(1, 2)	(1, 2)	(1, 2)	(1, 6)	(1, 6)
5	(5, 2)	(5, 2)	(5, 2)	(5, 2)	(5, 6)	(5, 6)
5	(5, 2)	(5, 2)	(5, 2)	(5, 2)	(5, 6)	(5, 6)
5	(5, 2)	(5, 2)	(5, 2)	(5, 2)	(5, 6)	(5, 6)

Again, simply by counting how many times the roll of die B is higher than the roll of die C, we get that

$$\mathrm{P}(B > C) = \frac{24}{36}$$

and

$$\mathrm{P}(C > B) = \frac{12}{24},$$

so die B is more powerful than die C. Using our notation introduced above, we see that

$$A \succ B \succ C.$$

Now, creating another outcome table for dice C and D, we get the table

D \ C	1	1	1	5	5	5
0	(0, 1)	(0, 1)	(0, 1)	(0, 5)	(0, 5)	(0, 5)
0	(0, 1)	(0, 1)	(0, 1)	(0, 5)	(0, 5)	(0, 5)
4	(4, 1)	(4, 1)	(4, 1)	(4, 5)	(4, 5)	(4, 5)
4	(4, 1)	(4, 1)	(4, 1)	(4, 5)	(4, 5)	(4, 5)
4	(4, 1)	(4, 1)	(4, 1)	(4, 5)	(4, 5)	(4, 5)
4	(4, 1)	(4, 1)	(4, 1)	(4, 5)	(4, 5)	(4, 5)

with corresponding probabilities

$$\mathrm{P}(C > D) = \frac{24}{36}$$

and

$$\mathrm{P}(D > C) = \frac{12}{36}.$$

Therefore, die C is more powerful than die D, thus completing our hierarchy of dice:

$$A \succ B \succ C \succ D.$$

From this, it seems die A is the best possible choice. From our calculations, we see that, against die B, we should expect die A to win about 67% of the time, the exact win percentage our friend was achieving! This is good news; it seems we've completely figured the game out. Out of curiously, we may wonder how often we'd expect die A to win against die C. Maybe 70% of the time? 80%? Even higher? It's a simple enough thing to check. All we have to do is create another outcome table and count how many times die A beats die C:

C \ A	3	3	3	3	3	3
1	(1, 3)	(1, 3)	(1, 3)	(1, 3)	(1, 3)	(1, 3)
1	(1, 3)	(1, 3)	(1, 3)	(1, 3)	(1, 3)	(1, 3)
1	(1, 3)	(1, 3)	(1, 3)	(1, 3)	(1, 3)	(1, 3)
5	(5, 3)	(5, 3)	(5, 3)	(5, 3)	(5, 3)	(5, 3)
5	(5, 3)	(5, 3)	(5, 3)	(5, 3)	(5, 3)	(5, 3)
5	(5, 3)	(5, 3)	(5, 3)	(5, 3)	(5, 3)	(5, 3)

This time, however, there isn't a clear winner, as

$$\mathrm{P}(A > C) = \frac{18}{36}$$

and

$$\mathrm{P}(C > A) = \frac{18}{36}.$$

The probability that die A beats die C is *equal* to the probability that die C beats die A. This seems... wrong. If $A \succ B$, and $B \succ C$, then why isn't it the case that $A \succ C$? Maybe we should use a different relative power comparison, something like

$$A \succcurlyeq B \succcurlyeq C \succcurlyeq D.$$

Here, $A \succcurlyeq B$ means die A is *at least* as powerful as die B. Using this notation solves our apparent problem: $A \succcurlyeq B$, $B \succcurlyeq C$, and $A \succcurlyeq C$, as we expect.

After seeing that dice A and C are equally powerful against each other, we may wonder about the probability of die A beating die D. Making the outcome table and counting, we get the table

D \ A	3	3	3	3	3	3
0	(0, 3)	(0, 3)	(0, 3)	(0, 3)	(0, 3)	(0, 3)
0	(0, 3)	(0, 3)	(0, 3)	(0, 3)	(0, 3)	(0, 3)
4	(4, 3)	(4, 3)	(4, 3)	(4, 3)	(4, 3)	(4, 3)
4	(4, 3)	(4, 3)	(4, 3)	(4, 3)	(4, 3)	(4, 3)
4	(4, 3)	(4, 3)	(4, 3)	(4, 3)	(4, 3)	(4, 3)
4	(4, 3)	(4, 3)	(4, 3)	(4, 3)	(4, 3)	(4, 3)

and the corresponding probabilities

$$\mathrm{P}(A > D) = \frac{12}{36}$$

and

$$\mathrm{P}(D > A) = \frac{24}{36}.$$

This... also seems wrong. From the table, we can clearly see that die D is more powerful than die A. However, from our established hierarchy of dice, we see that

$$A \succcurlyeq B \succcurlyeq C \succcurlyeq D,$$

so die A should be more powerful than die D.

What's going on here? From direct calculations, we know that $A \succcurlyeq B$, $B \succcurlyeq C$, and $C \succcurlyeq D$, yet $D \succcurlyeq A$! No matter which die is chosen out of the bunch, there will always exist another die that's more powerful!

Suddenly, our friend's game makes a lot more sense. By allowing us to choose our die first, our friend was ensuring that they could choose the die that was more powerful than ours. Our friend wasn't being generous at all!

4. Nontransitivity

After the initial rage of being duped by our friend passes, we take a step back and reflect on the strange structure of the game itself. *How is it possible that no single die is better than all the rest?* What is it about our particular notion of "powerful dice" that disallows a hierarchy to be created? The computations we did were certainly correct. From the tables we created, it's obvious that die A is indeed more powerful than die B—it rolls higher more often—and die B is certainly more powerful than die C. However, from the same tables, we see that dice A and C are equally matched; both have an equal probability of beating the other.

In our original notation, we have that

$$A \succ B \text{ and } B \succ C,$$

but

$$A \not\succ C,$$

where "$\not\succ$" represents a die *not* being more powerful than another. From our definition of transitivity in section 2, then, our "powerful dice" relation isn't transitive. Because of this, the fact that

$$A \succ B \succ C \succ D,$$

but

$$A \not\succ D,$$

is perfectly valid. Without transitivity, there's no reason why die A must necessarily be more powerful than die D.

All of this may feel extremely counter-intuitive. It seems blatantly obvious that one die should have an inherent advantage over all others, yet our outcome tables show that each die will always be inferior to at least one other die.

However, after a bit of pondering, we realise that the "nontransitivity" of our friend's dice game isn't so strange after all. There are plenty of games where objects can't be nicely organised into a hierarchy. One extremely common game with this property is Rock, Paper, Scissors, where two players simultaneously choose between three objects—a rock, a sheet of paper, or a pair of scissors—to attack the other player with.

If rock is chosen, paper can cover it, so paper is more powerful than rock. If paper is chosen, scissors can cut it, so scissors are more powerful than paper. However, scissors aren't more powerful than rock, since rock can crush scissors. No matter which object one chooses, there always exists another object which is more powerful than the one chosen.[7]

The dice game, then, is just another example of such a nontransitive situation. In fact, if we draw a diagram similar to figure 3 for the set of dice used in our game, we'll see that the behaviour of the dice, in terms of which die is more likely to beat

[7]For a more in-depth analysis of Rock, Paper, Scissors and other explicitly-nontransitive games, see section 3.1 of Klimenko [2015].

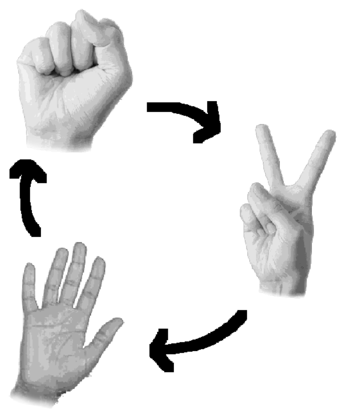

FIGURE 3. The three "objects" of Rock, Paper, Scissors. The arrows indicate which objects beat which. Starting from the top and working clockwise, rock beats scissors, scissors beats paper, and paper beats rock. As no one object is more powerful than all the others, the "power" relation here is nontransitive.

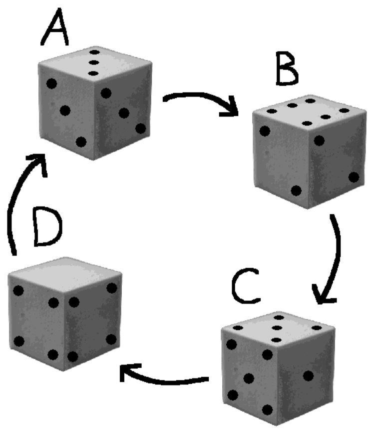

FIGURE 4. The four Efron Dice. The arrows indicate which die is "more powerful" than the other. For example, die A is more powerful than die B. As with Rock, Paper, Scissors, no one object is more powerful than all the rest, so the "power" relation here is nontransitive.

which, is nearly identical to the case of Rock, Paper, Scissors, just with four objects rather than three.

As soon as we acknowledge that such power relations (as we've defined for the dice) need not be transitive, all confusion disappears. The structure of figure 4 is fairly common (as Rock, Paper, Scissors demonstrates) and easy to understand. The counter-intuitive nature of this dice game arises solely from the assumption that we should be able to rank a die's power using some transitive relation on pairs of dice. Hopefully, we can start to see why assuming transitivity isn't always justified.

However, despite the assumption of transitivity in our dice game being false, it's a common-enough assumption to make. The assumption is so common, in fact, that in the 1990s, a US patent was filed by Robert Page for a particular dice game to be played in casinos which took advantage of this very assumption. This game made use of nontransitivity within sets of dice so that, if players entered the game with the assumption of transitivity, they would be unable to properly deduce optimal strategies and odds of winning. From column 3, line 9 of Page [1992], Page's dice game

> advantageously utilizes one or more sets of [non]transitive dice in order to ensure that a player may bet on any die or any of the available dice combinations and still leave a die or dice combination available for the "house" (i.e. casino) which is more likely to win than the player's selection.

This was the exact strategy employed by our friend in getting us to play their dice game.

5. Double Whammy Dice and Other Variants

5.1. **Won't Be Fooled Again.** Equipped with a better understanding of the dice game introduced in section 3, we confidently accept our friend's invitation to another game. This time, they've brought a new set of dice—three instead of four.

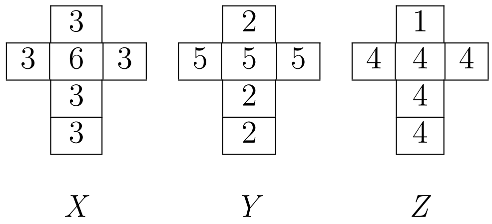

Figure 5. The nets for our friend's three new dice. This particular set of dice was taken from Grime [2017], in which no particular name was given to them. So, borrowing from the header under which they were introduced, we'll refer to these dice as the "Double Whammy Dice" for the remainder of the project.

However, this time we aren't going to be fooled. Putting our fist down, we *demand* that our friend picks their die first. Surprisingly, they agree! Immediately, we're suspicious of this, but we reluctantly continue. Perhaps our friend doesn't know that we're aware of the game's nontransitivity?

Our friend chooses die X (figure 5). Performing the same outcome table calculations as before, we see that $\mathrm{P}(X > Y) = \frac{21}{36}$, $\mathrm{P}(Y > Z) = \frac{21}{36}$, and $\mathrm{P}(Z > X) = \frac{25}{36}$. In our fancy notation, we can describe these probabilities as

$$X \succ Y \succ Z \succ X.$$

Therefore, we choose die Z, as according to the above probabilities, die Z rolls higher than die X in 25 of the 36 possible outcomes.

With the dice selected, our friend makes a suggestion: instead of each player rolling one of each die, each player rolls two of the same dice, just to make things more interesting. When each player's two dice are rolled, their face-up values are summed, and the player with the higher sum wins. Our friend presents the second set of dice and, after a close inspection, we decide that they are, in fact, identical to the first set. Thus, we agree to rolling two of the same dice. After all, how could it be any different than just rolling one of each die twice? We know die Z is more likely to roll higher than die X, so we should still have the advantage, no matter how many dice each player rolls... right?

We dismiss these extra rules as distractions. Surely our friend is simply trying to make the game *seem* more complicated so that we second-guess our strategy. However, our strategy is fool-proof, so we play.

Unfortunately, luck doesn't seem to be on our side. After many games, our friend seems to be winning *slightly* more than we are. Is this just the result of random chance? To be sure, we ask our friend to switch their dice. They decide to roll two of die Y. From our calculated probabilities, we choose to roll two of die X. However, even with this new selection, our friend is still beating the odds: they seem to be winning more often than us.

Once our friend is satisfied with their number of accumulated wins over us, they call it quits, leaving us to wonder, once again, how in the world they were able to gain an edge over us. The dice all seem to be fair (insomuch as each face is equally likely to roll face-up), and the outcome of one die is clearly independent of the other[8], so we should be able to model the game in exactly the same way as before, except now the outcomes we're comparing to each other are *sums* of dice rolls, rather than pure dice rolls.

It seems unlikely to us that rolling two of the same die instead of one should change anything. *However,* we tell ourselves, *that's just an assumption. Remember that transitivity was also an assumption we made, and that turned out to be wrong...*

The only way to know for sure is to look at the numbers.

5.2. **Random Variables and Other Formalities.** To calculate the probability of one *sum* of dice rolls beating another *sum* of dice rolls—say, comparing two rolls of die X to two rolls of die Y—we could create an outcome table as we did in section 3.2, listing all possible sums for two rolls of die X and all possible sums for two rolls of die Y, and seeing which sum is higher than the other more often.

There are six possible outcomes for each individual die, and each player's die rolls independently from the other. If D_1 represents the outcome of a player's first die, and D_2 represents the outcome of a player's second die, then

$$\mathrm{P}(D_1 \cap D_2) = \mathrm{P}(D_1)\mathrm{P}(D_2) = \frac{1}{6} \times \frac{1}{6} = \frac{1}{36}.$$

[8]This is assuming the dice aren't quantumly entangled with each other, a feat we're pretty sure our friend isn't able to accomplish...

Each possible sum of a player's two dice is equally likely with a probability of $\frac{1}{36}$. Because there are six possible outcomes for each individual die, and because the two die are independent, then there are a total of $6 \times 6 = 36$ possible sums for each player's pair of dice. So, if we were to create the outcome table for comparing two rolls of die X to two rolls of die Y, the table would have 36 rows for each possible sum for the first player, and 36 columns for each possible sum of the second player. This means our table would have $36 \times 36 = 1296$ cells we'd have to count in order to get our probabilities!

Perhaps because we have better things to do, or perhaps because we have a lack of patience[9], we decide that there must be a better way of calculating these probabilities. Our friend was clearly able to find the probabilities in some fashion—or else they wouldn't have suggested their variation on the game so confidently—so we should be able to find them, too!

After staring blankly into space for a bit, perhaps we recall the concept of a random variable. Instead of directly handling the sums of pairs of dice, could we not reframe our thinking in terms of some bivariate distribution of two random variables? Each die we roll is essentially already a random variable since we interpret each outcome as a *number* rather than some arbitrary outcome in some abstract outcome space. Adding the outcomes of two die rolls, then, is essentially equivalent to adding two random variables together.

If we're to use the language of random variables to analyse this new variant of our friend's game, it would probably be beneficial to describe our dice probabilities in terms of random variables. Let X be a random variable that outputs the numerical outcome of rolling die X in the set of Double Whammy Dice. Similarly, let Y and Z be random variables that output the numerical outcome of rolling die Y and die Z, respectfully, of the Double Whammy Dice.[10] Since there are a finite number of outcomes for each die, our random variables happen to be *discrete* random variables. Thus, to describe the probability distribution of each die, it's sufficient to simply list out all the possible outputs of the random variables, along with their respective probabilities for being output.

[9]Or, perhaps, because the author doesn't want to type out a table with 1296 cells in LaTeX. Maybe they could've made a script of some sort to generate it, but that would be a lot of work for not a lot of pay-off. Besides, the method used in place of the pure outcome table is much better, anyway.

[10]Throughout the rest of this project, we'll treat a die's random variable as if it *is* the die. The die and its random variable essentially represent the same object, just viewed under a different lens. After all, what is a die if not but a physical vessel through which to sample some probability distribution?

Let $f_A(D)$ represent the probability of some random variable A outputting a value of D. Then, using the fact that each of the three Double Whammy Dice are fair dice, we see that

$$f_X(x) = \begin{cases} \frac{5}{6} & \text{if } x = 3, \\ \frac{1}{6} & \text{if } x = 6, \\ 0 & \text{otherwise.} \end{cases}$$

$$f_Y(y) = \begin{cases} \frac{1}{2} & \text{if } y = 2, \\ \frac{1}{2} & \text{if } y = 5, \\ 0 & \text{otherwise.} \end{cases}$$

and

$$f_Z(z) = \begin{cases} \frac{1}{6} & \text{if } z = 1, \\ \frac{5}{6} & \text{if } z = 4, \\ 0 & \text{otherwise.} \end{cases}$$

Our dice are independent, so the probability that $X = x$ and $Y = y$ (i.e., die X rolls x and die Y rolls y) for some values x and y should be the probability of each one happening separately, but multiplied together (just like in section 3.1):

$$\mathrm{P}(X = x \wedge Y = y) = f_X(x) f_Y(y).$$

Since probabilities of the above form will be quite important to us (since our ultimate goal is to compare probabilities of *pairs* of dice rolled together), it makes sense to give probabilities of this form their own special notation, just as we did for the probabilities of our Double Whammy Dice. The notation

$$\mathrm{P}(X = x \wedge Y = y) = f_{\langle X,Y\rangle}(x, y) = f_X(x) f_Y(y)$$

seems fitting, so we'll go with that.

As an example, the function $f_{\langle X,Y\rangle}(x, y)$ is described by

$$f_{\langle X,Y\rangle}(x, y) = \begin{cases} \frac{5}{12} & \text{if } x = 3 \wedge y = 2, \\ \frac{5}{12} & \text{if } x = 3 \wedge y = 5, \\ \frac{1}{12} & \text{if } x = 6 \wedge y = 2, \\ \frac{1}{12} & \text{if } x = 6 \wedge y = 5, \\ 0 & \text{otherwise.} \end{cases}$$

Formally, what we've defined are the *probability mass functions* (pmfs) of our random variables. The functions that describe probabilities for single random variables are called *marginal probability mass functions*, while the functions that describe probabilities for multiple random variables are called *joint probability mass functions*. With these pmfs, the probability notation we were using in earlier sections now has a more concrete meaning. For instance, when we write $\mathrm{P}(X > Y)$, we mean the probability that the random variable X has a higher value than the random variable Y. Instead of computing this using those large outcome tables, we can now make use of our newly-defined functions:

$$\mathrm{P}(X > Y) = \sum_{x>y} f_{\langle X,Y\rangle}(x,y) = f_{\langle X,Y\rangle}(3,2) + f_{\langle X,Y\rangle}(6,2) + f_{\langle X,Y\rangle}(6,5) = \frac{7}{12}.$$

As the possible outputs of our pmfs represent probabilities of *disjoint* events, i.e. a set of events where only one can occur for any given outcome, we can simply add the probabilities of each event where $X > Y$ to get the probability that $X > Y$.[11] Using our joint pmf for X and Y, we see that $\mathrm{P}(X > Y) = \frac{7}{12}$. Note that this answer agrees with the answer we obtained in section 5.1 by using our outcome tables.

Now, the game our friend forced us to play involved adding the result of two identical dice. Therefore, rather than only concern ourselves with joint pmfs of the form $f_{\langle A,B\rangle}$ for two different dice A and B, it may be helpful to consider joint pmfs of the form $f_{\langle A,A\rangle}$. Can we construct such a function?

Let's try and construct $f_{\langle X,X\rangle}$. Even though X and X are clearly the same random variable, we should keep in mind that, in our friend's game, we're still rolling two *different* dice. The faces on the dice are the same, and the chances of rolling each face are the same, though the actual values we get when rolling the dice may differ between them. This means that our joint pmf for X and X should take in two values: one for the first die and one for the second. Other than this little caveat, calculating the values for $f_{\langle X,X\rangle}$ is exactly the same process as for the other joint pmfs. Since our two dice are still independent of one another, we see that $f_{\langle X,X\rangle}(x_1, x_2) = f_X(x_1)f_X(x_2)$. Hence,

$$f_{\langle X,X\rangle}(x_1,x_2) = \begin{cases} \frac{25}{36} & \text{if } x_1 = 3 \wedge x_2 = 3, \\ \frac{5}{36} & \text{if } x_1 = 3 \wedge x_2 = 6, \\ \frac{5}{36} & \text{if } x_1 = 6 \wedge x_2 = 3, \\ \frac{1}{36} & \text{if } x_1 = 6 \wedge x_2 = 6, \\ 0 & \text{otherwise.} \end{cases}$$

This expression for $f_{\langle X,X\rangle}$ seems about right: 3 is by far the most common face on die X, so it makes sense that the probability that *both* die roll a 3 is the highest

[11]Our ability to add probabilities like this is simply a result of the axioms of a probability space, mainly the axiom which states that the probability measure must be countably additive for pairwise disjoint events. See Weisstein [2023] for a complete list of all probability axioms.

among the possible outcomes. As well, 6 is the most rare face, so it makes sense that both die rolling a 6 ends up being the least likely of the possible events. Creating joint pmfs of this kind for Y and Z, we get that

$$f_{\langle Y,Y\rangle}(y_1, y_2) = \begin{cases} \frac{1}{4} & \text{if } y_1 = 2 \wedge y_2 = 2, \\ \frac{1}{4} & \text{if } y_1 = 2 \wedge y_2 = 5, \\ \frac{1}{4} & \text{if } y_1 = 5 \wedge y_2 = 2, \\ \frac{1}{4} & \text{if } y_1 = 5 \wedge y_2 = 5, \\ 0 & \text{otherwise,} \end{cases}$$

and

$$f_{\langle Z,Z\rangle}(z_1, z_2) = \begin{cases} \frac{1}{36} & \text{if } z_1 = 1 \wedge z_2 = 1, \\ \frac{5}{36} & \text{if } z_1 = 1 \wedge z_2 = 4, \\ \frac{5}{36} & \text{if } z_1 = 4 \wedge z_2 = 1, \\ \frac{25}{36} & \text{if } z_1 = 4 \wedge z_2 = 4, \\ 0 & \text{otherwise.} \end{cases}$$

In theory, if we needed probability information regarding sets of three dice, or four, or five hundred, we could simply construct the joint pmfs for those particular sets in the exact same manner that we created the joint pmfs above. However, for our purposes, joint pmfs describing probabilities for sets of two dice are all we need.

5.3. **Joint Distributions and Added Dice.** Returning to our original task, we remember that we want to find the probability that the *sum* of two dice outcomes is higher than the sum of two other dice outcomes. For example, if we had chosen die X for our friend's game, and they chose die Y, we want to find $\mathrm{P}(X + X > Y + Y)$.[12] To make calculating this probability easier, it may help to first understand what values $X + X$ can have, and how likely each one is.

The random variable X can take on two different values: 3 and 6. Therefore, we could have that $X+X = 3+3 = 6$, $X+X = 3+6 = 6+3 = 9$, or $X+X = 6+6 = 12$. Each possible pair of dice outcomes is a disjoint event—we can't roll more than one possible outcome at a time—and so the probability that $X + X$ equals any of the sums listed above is just the sum of the probabilities of the outcomes that give that

[12]While it's tempting to write $X + X$ as $2X$, this would actually be incorrect. Remember that, although X and X represent the same random variable, their actual outcomes in the game can differ, so it's not necessarily the case that the outcome of two X dice added together equals the outcome of one X die multiplied by 2.

specific sum. So,

$$\mathrm{P}(X+X=6) = f_{\langle X,X\rangle}(3,3) = \frac{25}{36},$$
$$\mathrm{P}(X+X=9) = f_{\langle X,X\rangle}(3,6) + f_{\langle X,X\rangle}(6,3) = \frac{5}{18},$$
$$\mathrm{P}(X+X=12) = f_{\langle X,X\rangle}(6,6) = \frac{1}{36}.$$

As with our other probability expressions, it makes sense to give $\mathrm{P}(X+X=x)$ its own function. Following our previous notation, our probability function could be written as something like:

$$f_{X+X}(x) = \begin{cases} \frac{25}{36} & \text{if } x = 6, \\ \frac{5}{18} & \text{if } x = 9, \\ \frac{1}{36} & \text{if } x = 12, \\ 0 & \text{otherwise.} \end{cases}$$

Just like X, we see that $X+X$ can take on various values, where each value has a specific probability of occurring. It makes sense, then, to treat $X+X$ as a new random variable in its own right, just like the three dice X, Y, and Z. We can do the same thing for the other two Double Whammy Dice, creating new random variables $Y+Y$ and $Z+Z$. Using the same sort of calculations as above, we can write the pmfs for $Y+Y$ and $Z+Z$, too:

$$f_{Y+Y}(y) = \begin{cases} \frac{1}{4} & \text{if } y = 4, \\ \frac{1}{2} & \text{if } y = 7, \\ \frac{1}{4} & \text{if } y = 10, \\ 0 & \text{otherwise,} \end{cases}$$

and

$$f_{Z+Z}(z) = \begin{cases} \frac{1}{36} & \text{if } z = 2, \\ \frac{5}{18} & \text{if } z = 5, \\ \frac{25}{36} & \text{if } z = 8, \\ 0 & \text{otherwise.} \end{cases}$$

Notice that, despite each possible outcome of Y being equally likely, the possible outcomes of $Y+Y$ are *not* equally likely. It seems that the sum of random variables can have different properties than the individual random variables that make it. With this revelation, we start to understand why our friend was so eager to propose this new variant of their game: we were so distracted with the nontransitive nature of the dice that we ignored the idea that adding dice could change the underlying probabilities!

However, with our newly-created pmfs, we should be able to calculate the *true* probabilities for this new variant, without any implicit biases getting in the way. When we calculated $\mathrm{P}(X > Y)$, we added up the probabilities of all the disjoint events where the random variable X was greater than the random variable Y. So, to calculate $\mathrm{P}(X+X > Y+Y)$, we should do the same thing: add up the probabilities of all the disjoint events where the random variable $X+X$ is greater than the random variable $Y+Y$. Since we have the pmfs for both $X+X$ and $Y+Y$, this should be fairly straightforward:[13]

$$\begin{aligned}
&\mathrm{P}(X+X > Y+Y)\\
&= f_{X+X}(6)f_{Y+Y}(4) + f_{X+X}(9)[f_{Y+Y}(4) + f_{Y+Y}(7)]\\
&+ f_{X+X}(12)[f_{Y+Y}(4) + f_{Y+Y}(7) + f_{Y+Y}(10)]\\
&= \frac{25}{36}\left(\frac{1}{4}\right) + \frac{5}{18}\left(\frac{1}{4}+\frac{1}{2}\right) + \frac{1}{36}\left(\frac{1}{4}+\frac{1}{2}+\frac{1}{4}\right) = \frac{59}{144}.
\end{aligned}$$

So, the probability that the sum of two X dice rolls higher than the sum of two Y dice is $\frac{59}{144}$, which is less than 50%. Despite X being the superior die choice when only rolling one of each die, it seems choosing X against Y when rolling *two* dice doesn't give the greatest odds of winning.

Out of curiosity, we wonder about the chances of $Y+Y$ being higher than $X+X$...

$$\begin{aligned}
&\mathrm{P}(Y+Y > X+X)\\
&= f_{Y+Y}(7)f_{X+X}(6) + f_{Y+Y}(10)[f_{X+X}(6) + f_{X+X}(9)]\\
&= \frac{1}{2}\left(\frac{25}{36}\right) + \frac{1}{4}\left(\frac{25}{36}+\frac{5}{18}\right)\\
&= \frac{85}{144}.
\end{aligned}$$

The chance of $Y+Y$ beating $X+X$ is $\frac{85}{144}$, which is greater than 50%. When rolling two dice each, Y becomes the better die, even though X was the better die to choose when rolling only one die each. In our original "power" relation notation, we'd say $X \succ Y$, but $X+X \prec Y+Y$.

[13]Note that the expression $f_{X+X}(x)f_{Y+Y}(y)$ represents the probability that $X+X = x$ and $Y+Y = y$ (since the dice are independent). Another valid way to calculate the probability $\mathrm{P}(X+X > Y+Y)$ is to first find an expression for $f_{\langle X+X,Y+Y\rangle}$, then add up all terms $f_{\langle X+X,Y+Y\rangle}(x,y)$ where $x > y$.

In fact, calculating the probabilities for a few more dice pair-ups, we see that

$$\begin{aligned}
&\mathrm{P}(Z + Z > Y + Y)\\
&= f_{Z+Z}(5)f_{Y+Y}(4) + f_{Z+Z}(8)[f_{Y+Y}(4) + f_{Y+Y}(7)]\\
&= \frac{5}{18}\left(\frac{1}{4}\right) + \frac{25}{36}\left(\frac{1}{4} + \frac{1}{2}\right)\\
&= \frac{85}{144},
\end{aligned}$$

and

$$\begin{aligned}
&\mathrm{P}(X + X > Z + Z)\\
&= f_{X+X}(6)[f_{Z+Z}(2) + f_{Z+Z}(5)]\\
&+ [f_{X+X}(9) + f_{X+X}(12)][f_{Z+Z}(2) + f_{Z+Z}(5) + f_{Z+Z}(8)]\\
&= \frac{25}{36}\left(\frac{1}{36} + \frac{5}{18}\right) + \left(\frac{5}{18} + \frac{1}{36}\right)\left(\frac{1}{36} + \frac{5}{18} + \frac{25}{36}\right)\\
&= \frac{671}{1296} \approx 0.5177.
\end{aligned}$$

So, using our power notation, we have that

$$X \succ Y \succ Z \succ X,$$

but

$$X + X \prec Y + Y \prec Z + Z \prec X + X.$$

Rolling two dice per player instead of one reverses the power ordering for the dice! Now that we understand how joint pmfs work, it makes sense that the probabilities for the Double Whammy Dice would change once we add more dice to each roll, but it's still surprising that the "ordering" of the dice completely flips. In this way, if we chose a die against our friend that was more powerful in the single die game, it *ensures* that our friend has the more powerful die in the double dice game!

Perhaps this property seems surprising. It may even be more counter-intuitive than the original nontransitivity of the dice. It's likely that others found this "reversing of the chain" property just as counter-intuitive, as the US patent for the casino dice game mentioned in section 4 makes use of this very property to further mislead potential gamblers. As Page explains it on column 7, line 67 of Page [1992],

> If one were to try and predict the favored outcome ... based on the performance record of each individual die involved in [a] contest, it would be noted that the true odds favor SW to beat SR a majority of the time and that BG is favored to beat BW a majority of the time [...] Based on this, logic would seem to indicate that a pairing

> of SW–BG (the two favorites in individual die comparisons) would be favored to beat a pairing of SR–BW (the two underdogs in individual die comparisons). However, in actuality, the pairing of SR–BW is favored to beat the pairing of SW–BG a majority of the time

In the patent, SW, SR, BG, and BW are the names given to different dice used in the game. This quotation only further exemplifies the importance of carefully analysing probabilistic games... such as the ones forced upon us by our friend. Otherwise, we may fall victim to the (apparently numerous) foibles of our intuition.

6. Theoretical Considerations

6.1. **Maximum Guaranteed Winning Probabilities.** It's inevitable that our friend will try and trick us again. They've already duped us twice... what's stopping them from duping us a third time? To prepare for this inevitability, we figure that it's best to try and learn as much as we can regarding these nontransitive dice games. Knowledge is the only tool we have to defend ourselves against the cunning intellect of our friend.[14]

Maybe as a first curiosity, we wonder to what extent our friend could rig a game in their favour. As we've seen previously, for each die (or set of dice) that's "more powerful" than another die (or set of dice), it's not necessarily the case that the probability of rolling higher is the same for each of them. For the Double Whammy Dice, we calculated that

$$\begin{aligned} \mathrm{P}(X > Y) &= \frac{21}{36}, \\ \mathrm{P}(Y > Z) &= \frac{21}{36}, \text{ and} \\ \mathrm{P}(Z > X) &= \frac{25}{36}. \end{aligned}$$

Evidently, these probabilities aren't all the same. A natural question might be to ask what the "highest minimum guaranteed probability of winning" is for our friend. That is, assuming our friend chooses their die/dice optimally during play, what's the highest possible probability p their game can attain so that, no matter what die/dice we choose, their probability of rolling higher than us is always at least p?

This exact question has been tackled in multiple papers. An upper bound of this probability was derived in Usiskin [1964]. Their proof goes as follows:

[14] At this rate, they may soon become our ex-friend!

Theorem 6.1. *Say we're given n random variables $\langle X_1, X_2, \cdots, X_n \rangle$ representing the outcomes of die rolls. Then*

$$\max_{X_i} \min(\{\mathrm{P}(X_1 > X_2), \mathrm{P}(X_2 > X_3), \cdots, \mathrm{P}(X_{n-1} > X_n), \mathrm{P}(X_n > X_1)\}) = \frac{n-1}{n}.$$

Proof. Consider the sum

$$\sum_{i=1}^{n} \mathrm{P}(X_i \leq X_{(i \bmod n)+1})$$
$$= \mathrm{P}(X_1 \leq X_2) + \mathrm{P}(X_2 \leq X_3) + \cdots + \mathrm{P}(X_{n-1} \leq X_n) + \mathrm{P}(X_n \leq X_1).$$

By properties of probability spaces, we can say that

$$\sum_{i=1}^{n} \mathrm{P}(X_i \leq X_{(i \bmod n)+1}) \geq \mathrm{P}\left(\bigcup_{i=1}^{n} X_i \leq X_{(i \bmod n)+1}\right).$$

It turns out that

$$\mathrm{P}\left(\bigcup_{i=1}^{n} X_i \leq X_{(i \bmod n)+1}\right) = 1$$

since at least one of the events $X_i \leq X_{(i \bmod n)+1}$ in the union must always happen, no matter the choice of random variables X_i. To see why, imagine we take a sample $\langle x_1, x_2, \cdots, x_n \rangle$ from our joint distribution $\langle X_1, X_2, \cdots, X_n \rangle$. If the first $n-1$ events in the union haven't occurred, this implies that $x_1 > x_2 > x_3 > \cdots > x_{n-1} > x_n$. Then it must be the case that $x_n < x_1$ by the transitivity of inequalities, which means the n-th event in the union occurred. Therefore, at least one of the events in the union must occur, and thus the probability is 1 since the union of a set of events where one is guaranteed to occur is a certain event.

This means that

$$\sum_{i=1}^{n} \mathrm{P}(X_i \leq X_{(i \bmod n)+1}) \geq 1,$$

which implies that at least one $\mathrm{P}(X_i \leq X_{(i \bmod n)+1})$ has a value of at least $\frac{1}{n}$. Otherwise, there's no way the sum could be greater than or equal to 1 (since a sum of n values, all of which are less than $\frac{1}{n}$, must necessarily be less than 1).

By properties of probability spaces, there's at least one $\mathrm{P}(X_i > X_{(i \bmod n)+1})$ where

$$\mathrm{P}(X_i > X_{(i \bmod n)+1}) = 1 - \mathrm{P}(X_i \leq X_{(i \bmod n)+1}) \leq 1 - \frac{1}{n} = \frac{n-1}{n},$$

since $\mathrm{P}(X_i \leq X_{(i \bmod n)+1})$ and $\mathrm{P}(X_i > X_{(i \bmod n)+1})$ are complementary events.

So, for the set of probabilities $\{\mathrm{P}(X_1 > X_2), \mathrm{P}(X_2 > X_3), \cdots, \mathrm{P}(X_n > X_1)\}$, at least one is bounded above by $\frac{n-1}{n}$, which means that

$$\max_{X_i} \min(\{\mathrm{P}(X_1 > X_2), \mathrm{P}(X_2 > X_3), \cdots, \mathrm{P}(X_{n-1} > X_n), \mathrm{P}(X_n > X_1)\}) \leq \frac{n-1}{n}.$$

Now, consider the following joint distribution, with each outcome having a probability of $\frac{1}{n}$:

$$\begin{aligned}
\langle Y_1, Y_2, \cdots, Y_n \rangle \in \; & \{\langle n, n-1, n-2, \cdots, 3, 2, 1\rangle\} \\
& \cup \{\langle n-1, n-2, n-3, \cdots, 2, 1, n\rangle\} \\
& \cup \{\langle n-2, n-3, n-4, \cdots, 1, n, n-1\rangle\} \\
& \cup \; \cdots \\
& \cup \{\langle 2, 1, n, \cdots, 5, 4, 3\rangle\} \\
& \cup \{\langle 1, n, n-1, \cdots, 4, 3, 2\rangle\}.
\end{aligned}$$

This joint distribution is constructed so that each $\mathrm{P}(Y_i > Y_{(i \bmod n)+1}) = \frac{n-1}{n}$. Therefore, the upper bound for our maximum is attainable. Therefore,

$$\max_{X_i} \min(\{\mathrm{P}(X_1 > X_2), \mathrm{P}(X_2 > X_3), \cdots, \mathrm{P}(X_{n-1} > X_n), \mathrm{P}(X_n > X_1)\}) = \frac{n-1}{n}.$$

QED

What theorem 6.1 tells us is that, given any set of n nontransitive dice, the highest possible winning probability our friend can guarantee for themselves, no matter what die we choose, is bounded by $\frac{n-1}{n}$.

However, theorem 6.1 makes no assumptions about the underlying probability distributions of the given random variables; the theorem applies to *all* possible joint probability distributions, not just those representing the rolls of fair dice. In fact, the true, "achievable" upper bound on our friend's maximum guaranteed winning probability is quite different from the general result listed here. An exact formula exists for this "achievable" upper bound, given in Vuksanovic and Hildebrand [2021] as

$$p_n = 1 - \frac{1}{4\cos^2(\frac{\pi}{n+2})},$$

where n is the number of dice in the game.

As an example, the value p_4 can be calculated as

$$p_4 = 1 - \frac{1}{4\cos^2(\frac{\pi}{4+2})} = 1 - \frac{1}{4(\frac{\sqrt{3}}{2})^2} = 1 - \frac{1}{3} = \frac{2}{3}.$$

If we refer back to the probabilities for the Efron dice listed in section 3.2, we notice that

$$\mathrm{P}(A > B) = \mathrm{P}(B > C) = \mathrm{P}(C > D) = \mathrm{P}(D > A) = \frac{2}{3},$$

meaning the Efron dice give our friend the highest possible guaranteed winning probability against us! If, instead, we calculate p_3, we get that

$$p_3 = 1 - \frac{1}{4\cos^2(\frac{\pi}{3+2})} = 1 - \frac{1}{4(\frac{1+\sqrt{5}}{4})^2} = 1 - \frac{4}{6+2\sqrt{5}} = \frac{\sqrt{5}-1}{2} \approx 0.6180.$$

Looking at the probabilities we calculated for the Double Whammy Dice, we see that

$$\mathrm{P}(X > Y) = \frac{21}{36} = 0.58\overline{3},$$
$$\mathrm{P}(Y > Z) = \frac{21}{36} = 0.58\overline{3},$$
$$\mathrm{P}(Z > X) = \frac{25}{36} = 0.69\overline{4},$$

so the guaranteed winning probability for these dice is $\frac{21}{36} < \frac{\sqrt{5}-1}{2}$, meaning the Double Whammy Dice don't give our friend the theoretical highest possible guaranteed winning probability.

Unfortunately, there isn't much more we can do with these guaranteed winning probabilities without getting a fair bit more abstract. The papers Richard P. Savage [1994], Usiskin [1964], Trybula [1965], and Funkenbusch and Saari [1983] go much more in-depth regarding the theoretical limitations of what sorts of probabilities are and aren't possible with joint distributions of random variables (such as our collections of dice). The purpose of this project is only to give an overview of the landscape of nontransitive games, and focusing *too* much on the nitty-gritty[15] takes away from the reader's general understanding of the subject.[16]

6.2. **The Extent of Nontransitivity.** Another question we may ask ourselves is whether there's any sort of limit to the "types of behaviour" that nontransitive dice can exhibit. For example, with the Double Whammy Dice, we had that

$$X \succ Y \succ Z \succ X,$$

but

$$X + X \prec Y + Y \prec Z + Z \prec X + X.$$

[15]In the author's humble opinion...

[16]Due to the relative simplicity of theorem 6.1, it was included for completeness' sake. The theorem gives the reader a taste of the sorts of considerations that go into proving properties of nontransitive games without demanding an unnecessary amount of thought. As well, should the reader have no interest in the proof, it can easily be skipped without sacrificing any understanding of the concepts discussed herein.

Rolling two dice switched our "power" ordering. However, is it possible to have our power ordering swap in *any* possible way? For example, could it be possible to construct a set of three dice S, T, and U such that

$$S \succ U, \quad T \succ U$$

but

$$S + S \prec U + U, \quad T + T \prec U + U?$$

Note that such a relation between S, T, and U isn't necessarily nontransitive in the way we've been talking about, but it illustrates the idea: what restrictions, if any, exist on the types of power relations we can have between dice?

This exact question is answered by the theorem on page 48 of Saari [1995]. When dealing specifically with one or more triplets of dice (such as the Double Whammy Dice), it's possible to choose a set of dice so that, by adding these various dice together, *any* sort of power relation ordering is possible, while also preserving the nontransitivity between triplets (or groups of dice added together)!

This seems a very unlikely result. The fact that *any* power ordering is possible by adding some dice together in some way seems "too loose", somehow. If any sort of power ordering is possible with our dice games, is any power ordering possible with other ranked objects? Could we apply this same logic to elections, perhaps? We'll address these concerns a little later in the project.[17]

6.3. **Various Nontransitive Dice.** One final question plagues our mind.[18] Up to now, we've only been considering cubic nontransitive dice. Nothing in our investigations have suggested that *only* cubic nontransitive dice exist. Could we potentially design nontransitive dice for other polyhedra?

Say we wanted to find a set of three nontransitive tetrahedral dice. Tetrahedral dice have fewer possible outcomes compared to cubic dice—four as opposed to six—so maybe finding a nontransitive set isn't so bad? Thankfully, this assumption turns out to be true. By fiddling with the Double Whammy Dice, it turns out that their face values can be adapted into nontransitive tetrahedral dice (see figure 6). Furthermore, if we were to take each tetrahedral die and copy-paste their face values onto both halves of an octahedral die, we'd have another valid set of nontransitive dice, since the probability of landing on each face would be identical for the corresponding tetrahedrons and octahedrons. This same technique can be used to create a set of icosahedral dice as well. However, it's also possible to create unique sets of

[17]See section 8.

[18]In actuality, there are many, many more questions that could still be plaguing our mind. To keep the narrative on-track, our hypothetical questions have been limited to a select few.

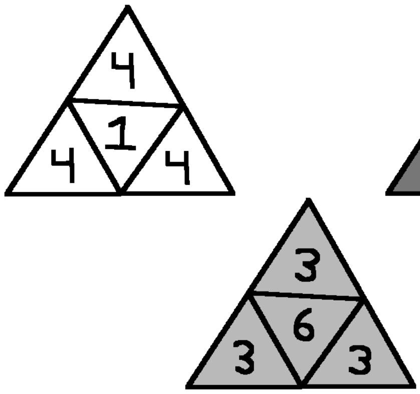

FIGURE 6. The nets for a set of nontransitive tetrahedral dice. These are known as "Tiggemann's Dice".

nontransitive dice for both the octahedral (figure 7) and icosahedral (figure 8) dice, sets that don't rely on using the same faces as the tetrahedral dice.

While we're at it, we might as well look for a set of dodecahedral nontransitive dice. Several sets are known, though arguably the most interesting are known as "Miwin's Intransitive Prime Dodecahedron". These nontransitive dice each beat each other with a probability of $\frac{35}{69}$, and all the face values happen to be prime numbers. If that wasn't enough, for each die, if you add up all the die's face values, you get a sum of 468. None of these extra properties contribute anything to the nontransitivity of the dice, but it's interesting that such a set of dice can exist. A more thorough look at these non-cubic nontransitive dice can be found in Pasciuto [2016].

One final set of dice worth mentioning are the New Grime Dice. Up to now, all the sets of dice we've discussed are meant to be used in two-player games, such as

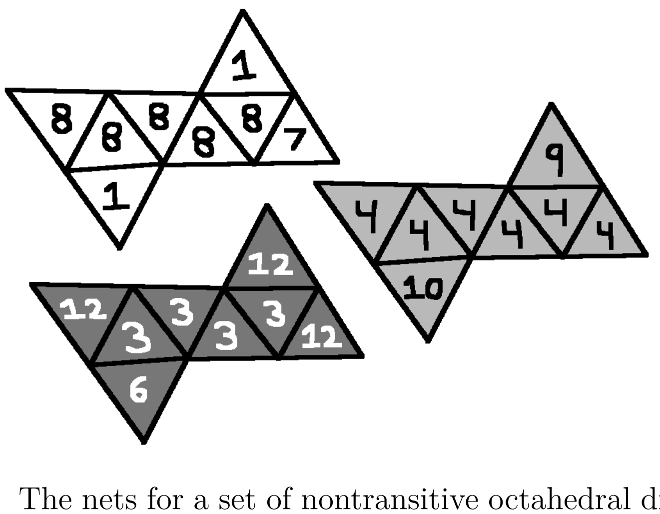

FIGURE 7. The nets for a set of nontransitive octahedral dice. These are known as the "Nichlman Dice".

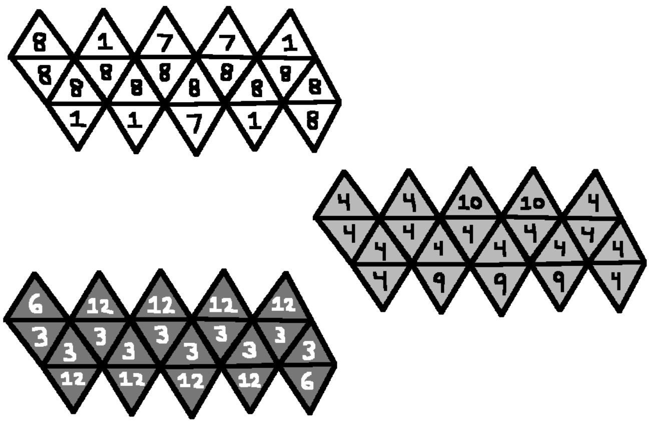

Figure 8. The nets for a set of nontransitive icosahedral dice. These are known as the “Pascannell Dice”.

the dice game between us and our friend. However, Dr. James Grime wondered if it would be possible to construct a set of dice that could be used in a *three*-player game. In other words, Grime wondered if it was possible for our friend to challenge two people at once—letting each person pick their die before they did—and still be able to choose a single die that would have a higher probability of rolling higher than *both* the other two chosen dice.

After some trial and error, Grime came across the New Grime Dice, a set of five cubic dice with *two* distinct nontransitive cycles among the dice, meaning there are two ways to rank the dice using our power relation that exhibit nontransitivity.

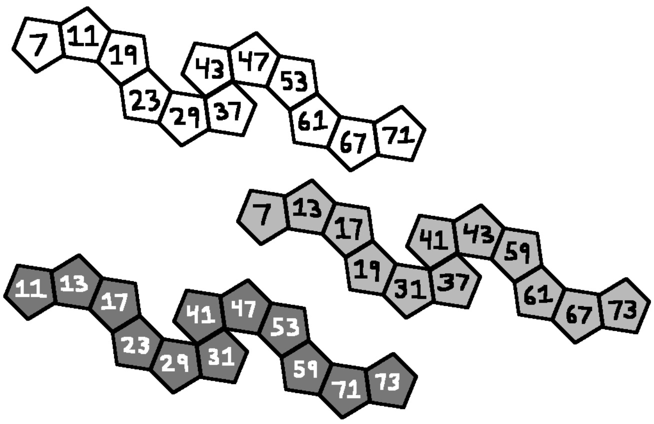

Figure 9. The nets for a set of nontransitive dodecahedral dice. These are known as “Miwin’s Intransitive Prime Dodecahedron”. A more detailed breakdown of these dice can be found at Winkelmann [2015]. Credit to Benjohn, Oolong, and Intelligenti pauca from https://math.stackexchange.com/questions/1634991 for providing the dodecahedron net used.

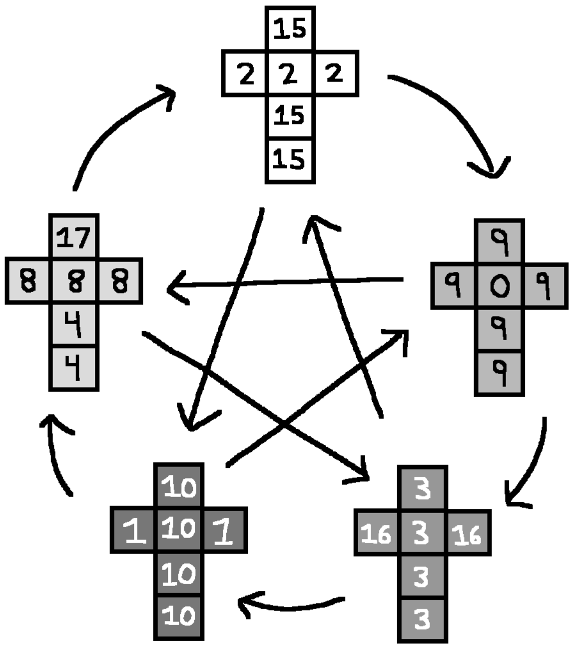

FIGURE 10. The nets for the New Grime Dice. An arrow from die A to B indicates that die A is more likely to roll higher than die B.

Crucially, however, if we decide to play the two-die-per-player variant of the dice game (where both of each player's dice are summed together to get their outcome), one of the nontransitive cycles swaps its ordering, while the other one retains its ordering. Using this order-swapping property, no matter which two dice our friend's opponents choose, our friend can always choose a die that has a greater probability of rolling higher than both... provided our friend is allowed to decide on-the-fly whether each player rolls one or two dice! [19]

According to page 5 of Grime [2017], only nontransitive dice games for two or three players are known to exist. For a four player game to exist—where our friend challenges *three* players simultaneously—it seems that at least nineteen dice are required for such a game (if only one of each die is rolled per person). Other than that, we don't know much more about what a set of nontransitive dice for four or more players would look like.

Grime goes into more detail surrounding the New Grime Dice (and other attempts at creating nontransitive sets of dice) in Grime [2017]. As well, he also casually

[19]To elaborate: sometimes, in the single-die game, the chosen dice of the two opponents leave our friend without a die that beats them both. However, switching to the double-dice game changes some of the probabilities in the exact way needed to allow our friend a choice that beats both opponents. So, provided our friend can decide before each game whether they're to roll one or two dice, they can always choose a die that has a greater probability of rolling higher than both their opponents simultaneously.

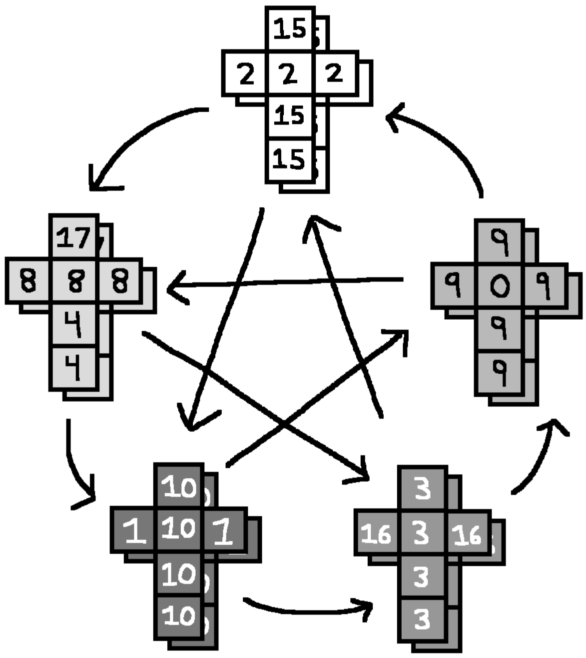

FIGURE 11. The nets for the New Grime Dice. Here, two dice are rolled per die choice, and the outcome is the sum of the two dice. An arrow from die A to B indicates that $A+A$ is more likely to roll higher than $B+B$.

mentions a particularly interesting result concerning the types of setups our friend could create for new dice games.

From page 4 of Grime [2017], it's said that "one can always set up a nontransitive system of m n-sided dice". For our purposes, this means that our friend can *always* find at least one set of m dice, each with n possible outcomes, where a nontransitive power relation occurs between the dice (where there's no single best die to choose to get the highest possible chance of winning). So, no matter the shape of the dice or the number of dice used, our friend can find a dice game where they always have a higher chance of winning!

For even more examples of nontransitive dice, pages 287 and 294 of Gardner [2001] contain some interesting sets, as well as some variations on the dice game itself (e.g. using cards to determine outcomes instead of dice). Table 2 of Page [1992] also contains more sets of nontransitive dice.

7. Penney's Game

7.1. **Devising New Nontransitive Games.** Our explorations into nontransitive games have certainly broadened our understanding and appreciation of nontransitive relations. Maybe they've also caused us to become more wary of relations which *appear* to be transitive but give no specific reason to believe that they're transitive (see section 8 for some examples). Most of all, though, our exploration has caused us to become envious of our friend. Twice they were able to fool us with nontransitivity!

The only way to reclaim our mathematical prowess is to beat our friend at their own game. However, we can't make use of either of the dice games, both because they're rigged against us and because our friend has a complete understanding of them. To beat our friend, we'll have to create our own nontransitive game, one where the nontransitivity is even less obvious! That way, our friend is sure to be fooled.

In order to create a nontransitive game, we first need to identify what made our friend's game so counter-intuitive. What was it about the dice game that escaped our reasoning? How was our friend able to construct such a biased game without us realising? In both cases, it seems the answer is our assumption of transitivity. However, this raises another question: why did we assume transitivity in the first place?

Transitivity is a fairly common assumption to make in plenty of scenarios. It's implicitly assumed across several different disciplines. As Klimenko comments on page 4370 of Klimenko [2015], many people "have psychological difficulties in accepting a relativistic approach [to ranking], expecting an absolute scale of judgments from 'bad' to 'good', which can be suitable in some cases but excessively simplistic in the others." Klimenko attributes this inclination to the human mind "filling in the gaps" whenever we're tasked with comparing a subset of objects when the full set is partially or completely unknown. We make assumptions about the underlying relationships between objects which are perfectly transitive with regard to the subset we're presented, but lose their transitivity when applied to the entire set.

For example, with our friend's original dice game, the subsets in question are the pairwise "battles" between the dice. In pairs, we can easily rank one die as more powerful than the other based on which one rolls higher in the majority of roll outcomes. However, when we try to apply these ranking to the entire set of Efron dice, the rankings suddenly become nontransitive.

What this implies is that, to make our nontransitive game, as long as we can get our friend to focus on some subset of our game's nontransitive "objects"—whatever they happen to be—we can exploit their brain's natural tendency to create seemingly-transitive relations between them in order to hide the objects' inherent nontransitivity!

Now, with us being the incredibly evil masterminds that we are, we begin idly flipping a coin as we look for inspiration. What sorts of objects could we use, other than dice, that exhibit some sort of nontransitivity when attempting to rank them? Our eyes flick around our room, occasionally stopping on our coin, then moving somewhere else, then moving back to the coin.

However, something odd catches our attention: the patterns our coin is making. Assuming the coin we're flipping is fair (in the same way we discussed for dice), every possible pattern of heads (H) and tails (T) we could flip should be equally likely. Despite this, whenever our gaze flicks back down to our coin, we seem to notice the pattern HT being flipped before TT more times than not. Similarly, the pattern TH seems to be flipped before HH more times than not.

Focusing our attention on the coin, we begin a fresh sequence of flips, recording the sequence and analysing it for patterns. Let's say our recorded sequence of flips ends up looking like

$$H, T, T, T, H, T, T, H, H, T, H, H, T, H, T, T, H, H, T, H, H, T, H, H, H.$$

If we start from the first flip, we notice that HT occurs before TT, and TH occurs before HH. From the second flip, TT occurs before HT, but TH still occurs before HH. If we were to start from all twenty-five of these flips, counting how many times HT occurs before TT, we'd get a count of 17. Similarly, if we start from all twenty-five flips, counting how many times TH occurs before HH, we'd get a count of 18. These seem exceptionally high counts, especially for flips generated from a seemingly-fair coin...

We need some notation to effectively talk about these patterns. If A and B are chosen coin flip patterns (such as HT or TT), then let $A \gtrdot B$ represent the event that pattern A is flipped before pattern B in the way described above. Can we find a way to calculate $\mathrm{P}(A \gtrdot B)$?

As a specific example, let's try and compute $\mathrm{P}(HT \gtrdot TT)$. Assuming our coin is fair, then there's a 50% chance the first flip in our sequence is H and a 50% chance the first flip is T. Graphically, our outcomes can be expressed as an outcome tree, where each leaf node at the bottom of the tree represents a possible flip sequence:

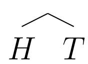

As well, the second flip has a 50% chance to be H and a 50% chance to be T. The outcome of the first flip has no bearing on the outcome of the second (i.e. coin flips are independent), so regardless of whether the first flip was H or T, the second flip is chosen between H or T with equal probability. With our tree, this is represented as:

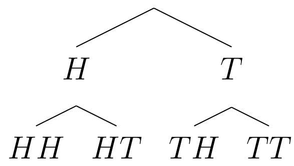

From this tree, we can see that one equally-likely flip sequence is TT, and one is HT. The remaining two possible sequences all end in H. Therefore, if we were to continue adding flips to these sequences, in order for TT to be flipped somewhere in these sequences, a T must be added to the sequence. However, when a T is added to the sequence, the pattern HT will necessarily occur first, since a H will necessarily be the end of the sequence before a T is flipped. Thus, for the remaining two possible sequences, HT will occur before TT:

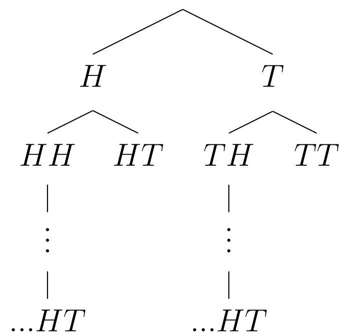

Therefore, since the four outcomes on the second layer of the tree are equally likely, we see that the pattern HT has a 75% chance of being flipped before TT. In our notation, we'll write this as

$$\mathrm{P}(HT \gtrdot TT) = \frac{3}{4}.$$

Using the exact same sort of argument, we get that

$$\mathrm{P}(TH \gtrdot HH) = \frac{3}{4}.$$

Despite our coin being fair, there exist certain sequences of coin flips that occur before other ones *more often*. There's nothing special about sequences of two flips, either. Building an outcome tree for the patterns TTH and HTH, we get:

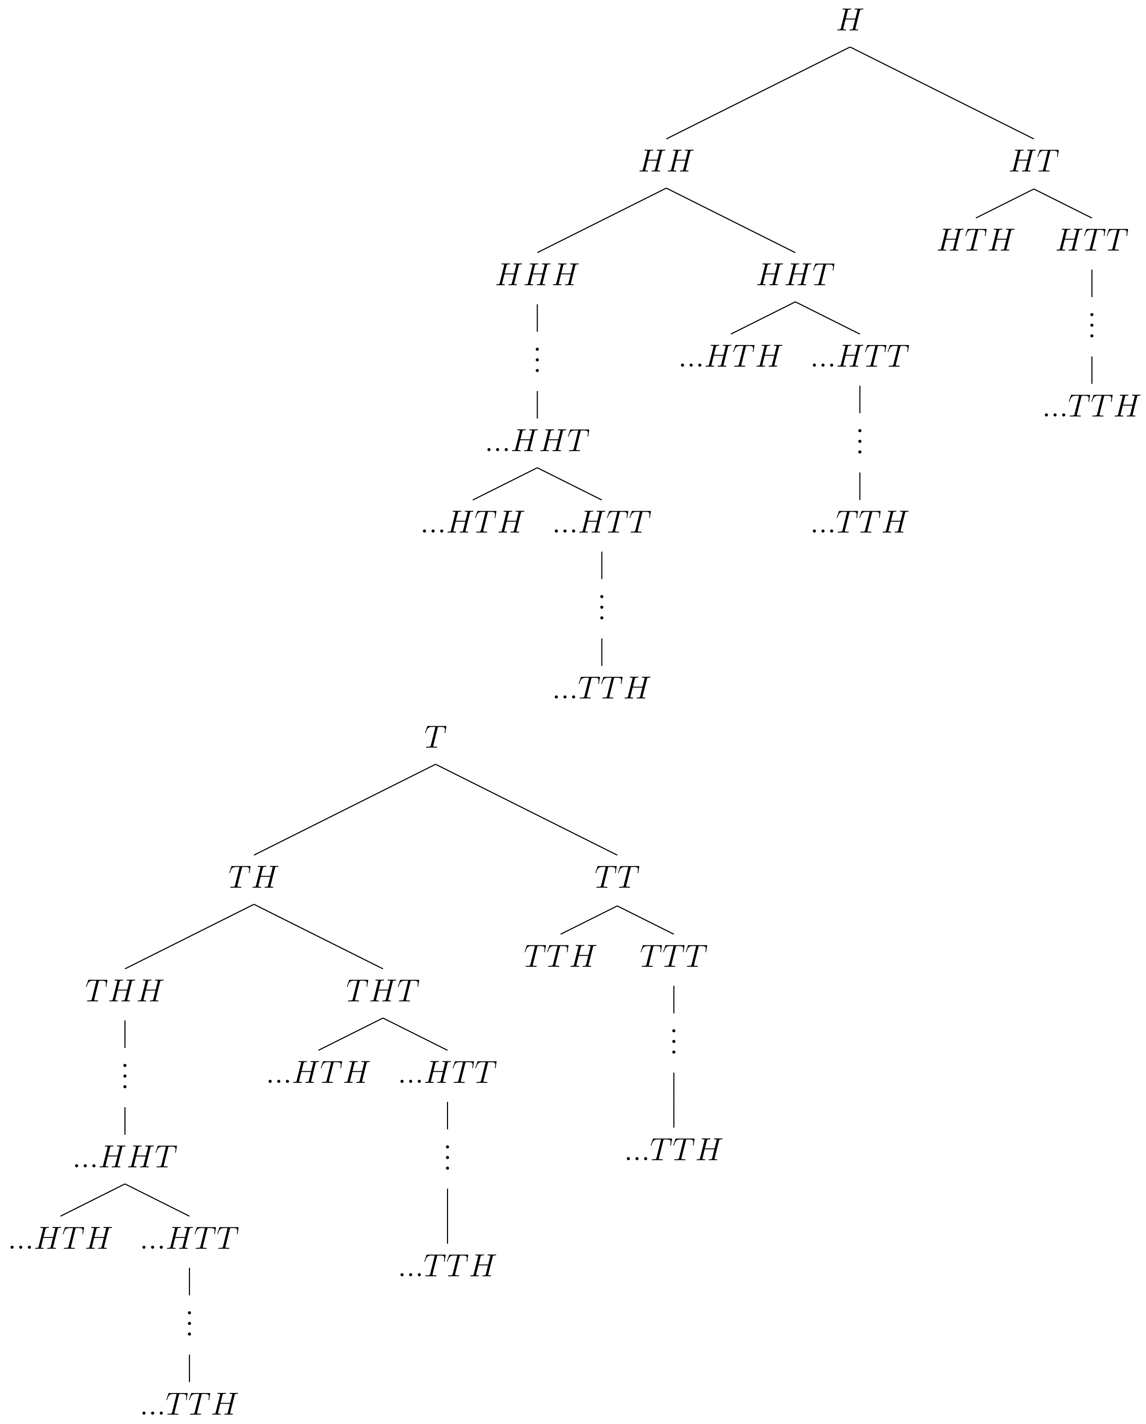

Note that this outcome tree was separated into two: one for when the first flip is H, and one for when the first flip is T. Using these two trees, we see that

$$\mathrm{P}(TTH \gtrdot HTH) = \frac{5}{8}.$$

Indeed, for any two given sequences of coin flips, we can build an outcome tree as we did above and determine which sequence is more likely to occur first.

Equipped now with some notation and a way to calculate which patterns are more likely to occur first, we crunch some numbers out of curiosity. We see that

$$\begin{aligned} \mathrm{P}(HTT \gtrdot TTH) &= \frac{3}{4}, \\ \mathrm{P}(HHT \gtrdot HTT) &= \frac{2}{3}, \\ \mathrm{P}(THH \gtrdot HHT) &= \frac{3}{4}, \text{ and} \\ \mathrm{P}(TTH \gtrdot THH) &= \frac{2}{3}. \end{aligned}$$

If we use the notation $A \succ B$ to mean that pattern A is more likely to occur before B, then the above probabilities tell us that

$$HTT \succ TTH \succ THH \succ HHT \succ HTT.$$

This situation looks familiar. No matter which of these four patterns we choose, there's always another pattern that's more likely to occur first in a sequence of coin flips. Crunching some more numbers, we see that

$$\begin{aligned} \mathrm{P}(THH \gtrdot HHH) &= \frac{7}{8}, \\ \mathrm{P}(TTH \gtrdot HTH) &= \frac{5}{8}, \\ \mathrm{P}(HHT \gtrdot THT) &= \frac{5}{8}, \text{ and} \\ \mathrm{P}(HTT \gtrdot TTT) &= \frac{7}{8}, \end{aligned}$$

and so

$$THH \succ HHH, \; TTH \succ HTH, \; HHT \succ THT, \quad \text{and } HTT \succ TTT.$$

For each of the four patterns not included in the nontransitive cycle we found, there still exist patterns that are more likely to occur first. So, no matter which three-flip pattern we choose, there always exists another three-flip pattern that has a higher chance of being flipped first.

It seems we've found another nontransitive game. Imagine we approach our friend with a coin, asking them to choose some three-flip sequence of heads or tails. After they make their selection, we choose our own sequence. Then, we flip the coin repeatedly, recording the generated sequence of heads and tails. The person whose sequence is flipped first is the winner. Based on the calculations we performed above, we can be sure that, no matter which sequence our friend chooses, there will always exist a sequence for us which gives us a better chance of winning!

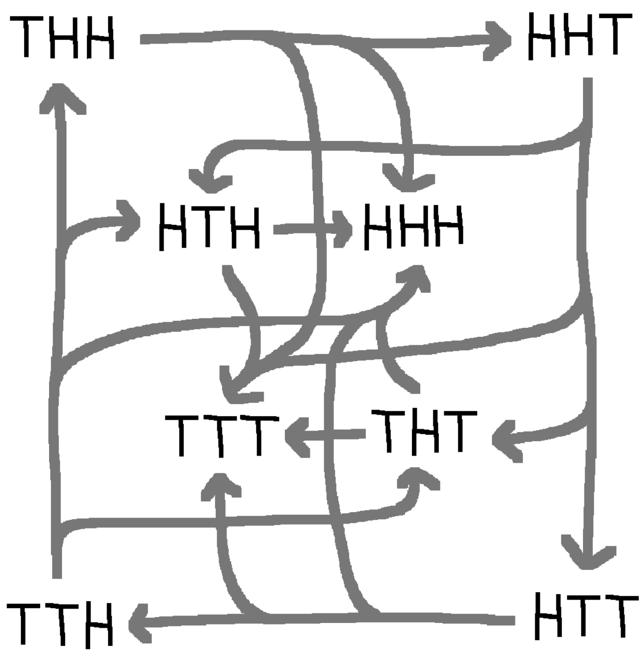

FIGURE 12. A graph showing which three-flip patterns appear first more often than others. An arrow from pattern A to B indicates that pattern A is more likely to occur first when compared to pattern B.

As well, since the idea of a coin flip being fair is so ingrained into our minds, it's likely that our friend will have a hard time recognising the nontransitive nature of these flip sequences. After all, since a coin flip is fair, it *seems* obvious that each sequence should be as likely to be flipped first as any other. What's more, unlike our friend's dice game, the coin we're using is a regular, unadulterated coin, whereas our friend's dice had face values which were altered from the usual one to six in exactly the way needed for nontransitivity to occur. In a sense, the nontransitivity of our new game is even more obscured than it was in the dice game!

The game we've just discovered is known as Penney's Game, named after the mathematician Walter Penney who first described it in 1969. On page 303 of Gardner [2001], Martin Gardner describes the game as "one of the finest of all sucker bets", and it's clear to see why. Without a careful analysis of the game, it *seems* obvious that each possible sequence has equal odds of being flipped first. However, as we're well aware, things *seeming* to be true isn't the same as them being *actually* true! Unlike the dice game, where the fault in our logic comes from assuming that the hierarchy of dice rankings is transitive, the fault in Penney's Game comes from the false assumption that there isn't a hierarchy of rankings at all!

Pages 303-311 of Gardner [2001] delve more in-depth on Penney's Game, considering other ways to analyse flip patterns—such as by using waiting times—as well as considering possible ways to determine which flip patterns give us the highest chance of winning against our friend. In the next subsection, we'll consider one such way to find the best patterns: Conway's Algorithm.

7.2. **Conway's Algorithm.** With our new nontransitive game designed, we now wonder if there's a more efficient way for calculating which sequences of coin flips are more likely to occur first. Certainly, using the brute-force-outcome-tree approach, we can list all possible match-ups of coin flip sequences in a table and calculate which one is the better choice by making trees for all possibilities. For our game, such a table would look something like the following[20]:

	HHH	HHT	HTH	HTT	THH	THT	TTH	TTT
HHH	...	$\frac{1}{2}$	$\frac{2}{5}$	$\frac{2}{5}$	$\frac{1}{8}$	$\frac{5}{12}$	$\frac{3}{10}$	$\frac{1}{2}$
HHT	$\frac{1}{2}$	...	$\frac{2}{3}$	$\frac{2}{3}$	$\frac{1}{4}$	$\frac{5}{8}$	$\frac{1}{2}$	$\frac{7}{10}$
HTH	$\frac{3}{5}$	$\frac{1}{3}$	...	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{3}{8}$	$\frac{7}{12}$
HTT	$\frac{3}{5}$	$\frac{1}{3}$	$\frac{1}{2}$	...	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{3}{4}$	$\frac{7}{8}$
THH	$\frac{7}{8}$	$\frac{3}{4}$	$\frac{1}{2}$	$\frac{1}{2}$	...	$\frac{1}{2}$	$\frac{1}{3}$	$\frac{3}{5}$
THT	$\frac{7}{12}$	$\frac{3}{8}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	...	$\frac{1}{3}$	$\frac{3}{5}$
TTH	$\frac{7}{10}$	$\frac{1}{2}$	$\frac{5}{8}$	$\frac{1}{4}$	$\frac{2}{3}$	$\frac{2}{3}$	...	$\frac{1}{2}$
TTT	$\frac{1}{2}$	$\frac{3}{10}$	$\frac{5}{12}$	$\frac{1}{8}$	$\frac{2}{5}$	$\frac{2}{5}$	$\frac{1}{2}$	...

Here, each cell of the table represents the probability that the sequence in the left column occurs before the sequence in the top row.[21] From this table, we can easily read off which sequence to choose, given the chosen sequence of our friend. For instance, if our friend chooses TTH as their sequence, we should choose HTT, as HTT has a $\frac{3}{4}$ chance of appearing before TTH.

This table works perfectly fine for our particular game, but what if we needed to modify our game in the future? After a while, our friend will definitely figure out that something's amiss with the game. We'll need to introduce some variations on the game to keep them guessing, which means we likely won't be able to keep using the same outcome table.

Thankfully, John Conway has us covered. He devised an algorithm, which we'll call "Conway's Algorithm" for the remainder of this project, that calculates the probability that one sequence of coin flips occurs before another. This algorithm

[20]The table for comparing two coin flip patterns to see which is more likely to occur first was adapted from Gardner [2001].

[21]The ... symbol is used where a particular coin flip sequence is being compared to itself (and so a particular probability value here wouldn't be well-defined). The symbol was adopted from Canada [2019].

works for sequences of *any* length, which makes it a good algorithm to use should we ever consider variants of our original game.

However, there's a slight caveat to Conway's Algorithm. The algorithm itself works perfectly well: give it two coin flip sequences and it'll churn out the relevant probabilities needed for Penney's Game. The caveat is in *understanding* the algorithm. The pmf calculations we performed in section 5.3 may have been a little long-winded, but the intuition behind *why* they worked was quite simple. Joint pmfs are simply functions that give probability information regarding joint probability distributions. Because rolling multiple dice counts as a joint distribution, we could use pmfs to calculate the probabilities we were interested in.

Conway's Algorithm, on the other hand, uses a rather esoteric method for calculating its probabilities. It involves writing sequences on top of each other and sliding them around, tallying up matching symbols and recording the data in binary numbers. Then, these binary numbers are converted to decimal and placed in a seemingly arbitrary ratio which supposedly represents the odds of one sequence occurring before the other. Even Martin Gardner was puzzled by Conway's Algorithm. On page 306 of Gardner [2001], he admits that he "ha[s] no idea why it works. It just cranks out the answer as if by magic, like so many of Conway's other algorithms".

While understanding the algorithm isn't strictly necessary to use it, the mathematician in us craves answers. Section 7.3 details a more intuitive way of understanding the process, adapted from Miller [2019]. For now, we'll focus solely on how to use the algorithm.

Let's say Alice and Bob from figure 1 decided to play Penney's Game, with Alice choosing the sequence $A = HTT$ and Bob choosing the sequence $B = THT$. Conway's Algorithm has us calculate four different numbers—called "leading numbers"—in order to construct a ratio for deciding which sequence is more likely to occur first. We'll call these four numbers $\mathrm{R}_A(A)$, $\mathrm{R}_A(B)$, $\mathrm{R}_B(A)$, and $\mathrm{R}_B(B)$.[22]

Now, we'll walk through calculating one of these four numbers, say $\mathrm{R}_A(B)$. The other three numbers are computed similarly, just using different sequences. First, we write the sequences aligned, with sequence A being written above, while sequence B is written below:

$$\begin{array}{cccc} \textbf{A:} & H & T & T \\ \textbf{B:} & T & H & T \end{array}$$

[22]The notation used here is slightly different from what's typically used when discussing Conway's Algorithm, but it aligns nicely with how the intuition for Conway's Algorithm is presented in section 7.3. In the author's opinion, it's also less confusing.

With this alignment, we check whether all the pairs of flips match each other. In this case, the first flip of A doesn't match with the first flip of B, so the flips don't match each other. Since they don't match, we write a 0 on top of the first column:

$$\begin{array}{rccc} \mathrm{R}_A(B): & {\scriptstyle 0} & & \\ \textbf{A:} & H & T & T \\ \textbf{B:} & T & H & T \end{array}$$

Next, we shift the bottom row one flip to the right and check again whether all the pairs of flips match:

$$\begin{array}{rcccc} \mathrm{R}_A(B): & {\scriptstyle 0} & & & \\ \textbf{A:} & H & T & T & \\ \textbf{B:} & & T & H & T \end{array}$$

In the third column, the T flip in A doesn't match the H flip in B, so again, the flips don't match each other. We write a 0 atop the second column:

$$\begin{array}{rcccc} \mathrm{R}_A(B): & {\scriptstyle 0} & {\scriptstyle 0} & & \\ \textbf{A:} & H & T & T & \\ \textbf{B:} & & T & H & T \end{array}$$

Then, we shift the bottom row one final time, and check again whether the pairs of flips match each other. In this case, there's only one pair of flips to consider, and they're both T flips, so they match. We write a 1 atop the third column to signify this match:

$$\begin{array}{rccccc} \mathrm{R}_A(B): & {\scriptstyle 0} & {\scriptstyle 0} & {\scriptstyle 1} & & \\ \textbf{A:} & H & T & T & & \\ \textbf{B:} & & & T & H & T \end{array}$$

The value of $\mathrm{R}_A(B)$ is then given by converting the binary number we got—001—into decimal. In this case, we get that $\mathrm{R}_A(B) = 1$. If we were to compute the other three numbers using the same method, we'd get that $\mathrm{R}_A(A) = 4$, $\mathrm{R}_B(A) = 2$, and $\mathrm{R}_B(B) = 5$.

Using these four numbers, we can now construct a ratio that represents the odds of sequence B occurring before sequence A. If $B \gtrdot A$ represents the event that sequence B occurs before sequence A, then

$$\begin{aligned} \textbf{Odds of } B \gtrdot A &= \mathrm{R}_A(A) - \mathrm{R}_A(B) : \mathrm{R}_B(B) - \mathrm{R}_B(A) \\ &= 4 - 1 : 5 - 2 \\ &= 3 : 3. \end{aligned}$$

Thus, the odds of sequence B occurring before A are $3 : 3$, meaning it's equally likely for either A or B to occur first.

If, instead, Alice chose the sequence $A = TTT$ and Bob chose the sequence $B = HTT$, then $\mathrm{R}_A(A) = 7$, $\mathrm{R}_A(B) = 0$, $\mathrm{R}_B(A) = 3$, and $\mathrm{R}_B(B) = 4$, and so

$$\begin{aligned}\textbf{Odds of } B \gtrdot A &= \mathrm{R}_A(A) - \mathrm{R}_A(B) : \mathrm{R}_B(B) - \mathrm{R}_B(A)\\ &= 7 - 0 : 4 - 3\\ &= 7 : 1.\end{aligned}$$

In this case, Bob's sequence has far greater odds of occurring first. For every seven times Bob's sequence occurs first, Alice's sequence will only occur first once. It's clear to see that Bob has a severe advantage in this scenario.

Conway's Algorithm works perfectly well with sequences of greater than three flips, it's just that more shifts will have to be done in order to construct each of the four leading numbers. As an example, given the two sequences $A = HHTTTT$ and $B = TTTHHH$, Conway's Algorithm gives $\mathrm{R}_A(B) = 7$.

If we wanted the *probability* that sequence B occurs before sequence A, the leading numbers of Conway's Algorithm can also give us that:

$$\mathrm{P}(B \gtrdot A) = \frac{\mathrm{R}_A(A) - \mathrm{R}_A(B)}{\mathrm{R}_A(A) - \mathrm{R}_A(B) + \mathrm{R}_B(B) - \mathrm{R}_B(A)}.$$

If we wanted, we could verify that this formula gives the exact probabilities that were listed in the above probability table. The point, however, is that Conway's Algorithm eliminates the need for such a table. It also eliminates the need for devising a different way of calculating these probabilities should we choose to change the length of each player's sequence.

7.3. **Intuition Behind Conway's Algorithm.** While Conway's Algorithm provides a handy way for checking which coin flip patterns we should choose against our friend, it doesn't exactly say *why* we should choose one pattern over another. What exactly is Conway's Algorithm computing? How does it know which pattern is better than the other?

The answer to this question isn't exactly obvious. There are many different proofs of Conway's Algorithm,[23] though all we're concerned with is understanding *why* the processes used in Conway's Algorithm should be related to Penney's game. What we need is a slightly different perspective on Penney's Game, one that elucidates how the calculations of Conway's Algorithm are relevant to sequences of coin flips.

[23]For example, pages 401 to 410 of Graham et al. [2011] describe a way to prove Conway's Algorithm using generating functions.

Recall Alice, Bob, and Chris from figure 1. Imagine they've all gone to a casino, and have decided to make some bets on a particular coin-flipping game.[24]

FIGURE 13. In case you needed a visual of Alice, Bob, and Chris at the casino. Credit to Jeremy Chevere for the illustration.

The premise of the game is simple: the betters each bet on a particular side of a coin, either heads (H) or tails (T). The casino then flips a fair coin. If a better's side of the coin appears, their bet is doubled and given back to them. Otherwise, the casino pockets the bet. There's no limit to how many plays each player can play; they can bet on coin flips as much as they want, and, if they're lucky, can win as much as they want.

Now, Chris happens to have a lot of money, and has decided to fund Alice's and Bob's betting. On each coin flip, Chris bets \$1 for both Alice's and Bob's guesses. If one of them wins their bet—therefore earning them double their bet—Chris gives the *entire* winnings to that player.

After a few rounds, Alice and Bob realise that, in its current state, the game is perfectly balanced between them; the casino's coin is a fair coin, and so both players should expect to win about 50% of the time. Since Alice and Bob are constantly

[24]This game (along with its analysis) was adapted from Miller [2019].

determined to prove their superiority to each other, they decide to add some additional structure to their betting in the hopes of adding some strategy, and therefore adding a chance for one player to win more bets than the other.

Each player is to pick a different sequence of three coin flips, say HTH or TTT. Chris will continue betting for them, but the way in which he'll be betting will change depending on each player's chosen coin pattern (explained below). Once one of the two players' chosen sequences appears as a sequence of the casino's coin flips, the winnings of both players' bets are distributed to them by Chris. Alice and Bob will play a few rounds like this simultaneously (and at the same table, so as to bet on the same coin being flipped), and whoever earns the most money at the end will be deemed the superior gambler.

The way in which Chris bets for the two players is a little convoluted under this new set-up. Therefore, an example is in order.

Imagine Alice's chosen coin flip sequence is HTH. On each casino flip, Chris bets \$1 on the coin being H, the first flip in Alice's sequence. If the casino's flip is T, then the casino pockets the \$1 and nothing happens. However, if the casino's flip is H, then Chris wins \$2 from the casino. Instead of immediately giving these winnings back to Alice, however, Chris now bets the \$2 on the next flip being T, the second flip in Alice's sequence. As well, Chris also starts another, separate \$1 bet on the next flip being H. If the flip is H, then Chris' \$2 bet is lost, while the \$1 bet wins him \$2. If the flip is T, then the \$1 bet is lost while the \$2 bet wins him \$4.

For each of the casino's coin flips, Chris *always* bets \$1 on the flip being H, the first flip in Alice's sequence. If any of these bets win, then the entire winnings from that particular bet are then put on a separate bet, with the predicted flip being the second flip in Alice's sequence. If that bet also happens to win, then the winnings of that bet are again put on another bet, with the predicted flip being the third in Alice's sequence. Finally, if this third bet wins, then Chris gives the winnings of the bet to Alice, along with any winnings from "in progress" bet sequences (ones that were bet on the first or second flip in Alice's sequence). The first of the two diagrams on the following pages illustrates how much Alice would win if the casino flips her sequence.

Note that Alice's sequence must exactly occur in the casino's sequence of flips for Alice to receive her winnings from Chris.

While Chris is betting on Alice's sequence, Chris is *also* betting on Bob's sequence in the exact same fashion. Whenever *either* of their sequences are flipped by the casino, Chris distributes both the winning bet sequence to the winning player and the "in progress" winnings to both players. The second diagram on the following pages illustrates how much Bob would win if Alice's sequence is flipped when his chosen sequence of flips was HTT.

The challenge for both players is to choose a sequence of coin flips that's "better" than the other player's sequence—that is, a sequence that's more likely to occur *before* the other player's so as to earn more winnings.[25] With this extra structure established, the link between Alice's and Bob's game and Penney's Game becomes clear.

Even here, in this highly-constructed set-up, one may be tempted to say that Alice's and Bob's odds of winning should still be equal, as they were for the single coin flips. After all, the casino's coin is a *fair* coin. Does this not imply that the sequences that each player can choose are equally likely to occur?

In fact, each possible sequence of coin flips actually is equally likely to occur, but they're not equally likely to occur *first*.

Consider the situation in which Alice chooses HHH as her sequence of coin flips. For each flip of the casino's coin, there's a $\frac{1}{8}$ chance that the next three coin flips (with the current flip counting as the first) will match Alice's sequence, no matter what Alice's sequence is.[26] We'll say that there's a $\frac{1}{8}$ chance of the casino's flip *starting Alice's sequence*. Roughly speaking, this means that, for every eight coins the casino flips, we'd expect one of them to start Alice's sequence.

So, if Chris bets on Alice's sequence in the same way described above, then Chris would expect to spend about \$10 before Alice's sequence shows up in the sequence of casino flips—\$8 for the eight flips to get Alice's sequence started, and \$2 to finish flipping her sequence. However, using Chris' betting strategy, Alice would win \$14 using HHH as her chosen pattern. So, for every expected \$10 to be spent, \$14 are expected to be gained, which is a net gain for Alice...

How is this possible? If the coin is fair, how can Alice have possibly derived a strategy to rig the odds in her favour? It turns out, actually, that Alice's strategy doesn't rig the odds in her favour.

[25]In this scenario, the player whose sequence occurs before the other's will *always* earn more winnings since a complete betting sequence from Chris will earn them at least \$8, and the maximum a player can earn without a complete betting sequence is \$6.
[26]The $\frac{1}{8}$ comes from the fact that there are 8 possible sequences of three coin flips, and each one is equally likely since the coin is fair.

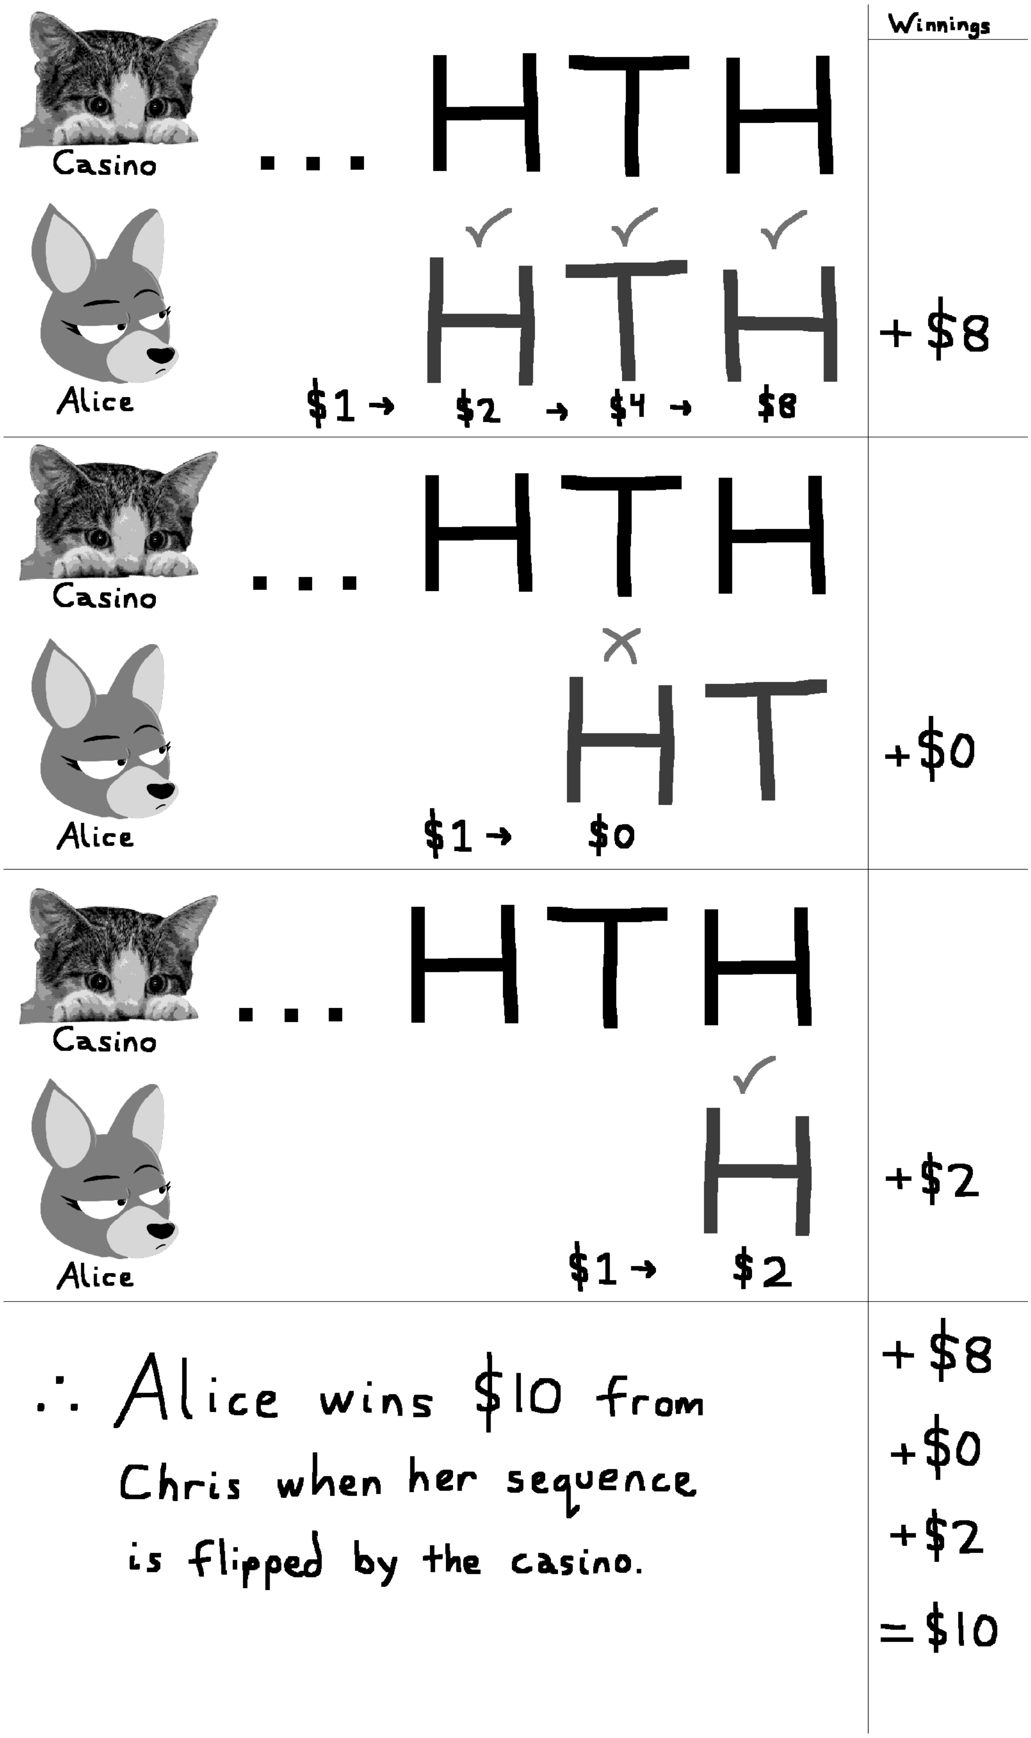
Winnings
Casino
...HTH
Alice
HTH
$1 → $2 → $4 → $8
+$8
Casino
...HTH
Alice
HT
$1 → $0
+$0
Casino
...HTH
Alice
H
$1 → $2
+$2
∴ Alice wins $10 from Chris when her sequence is flipped by the casino.
+$8
+$0
+$2
=$10

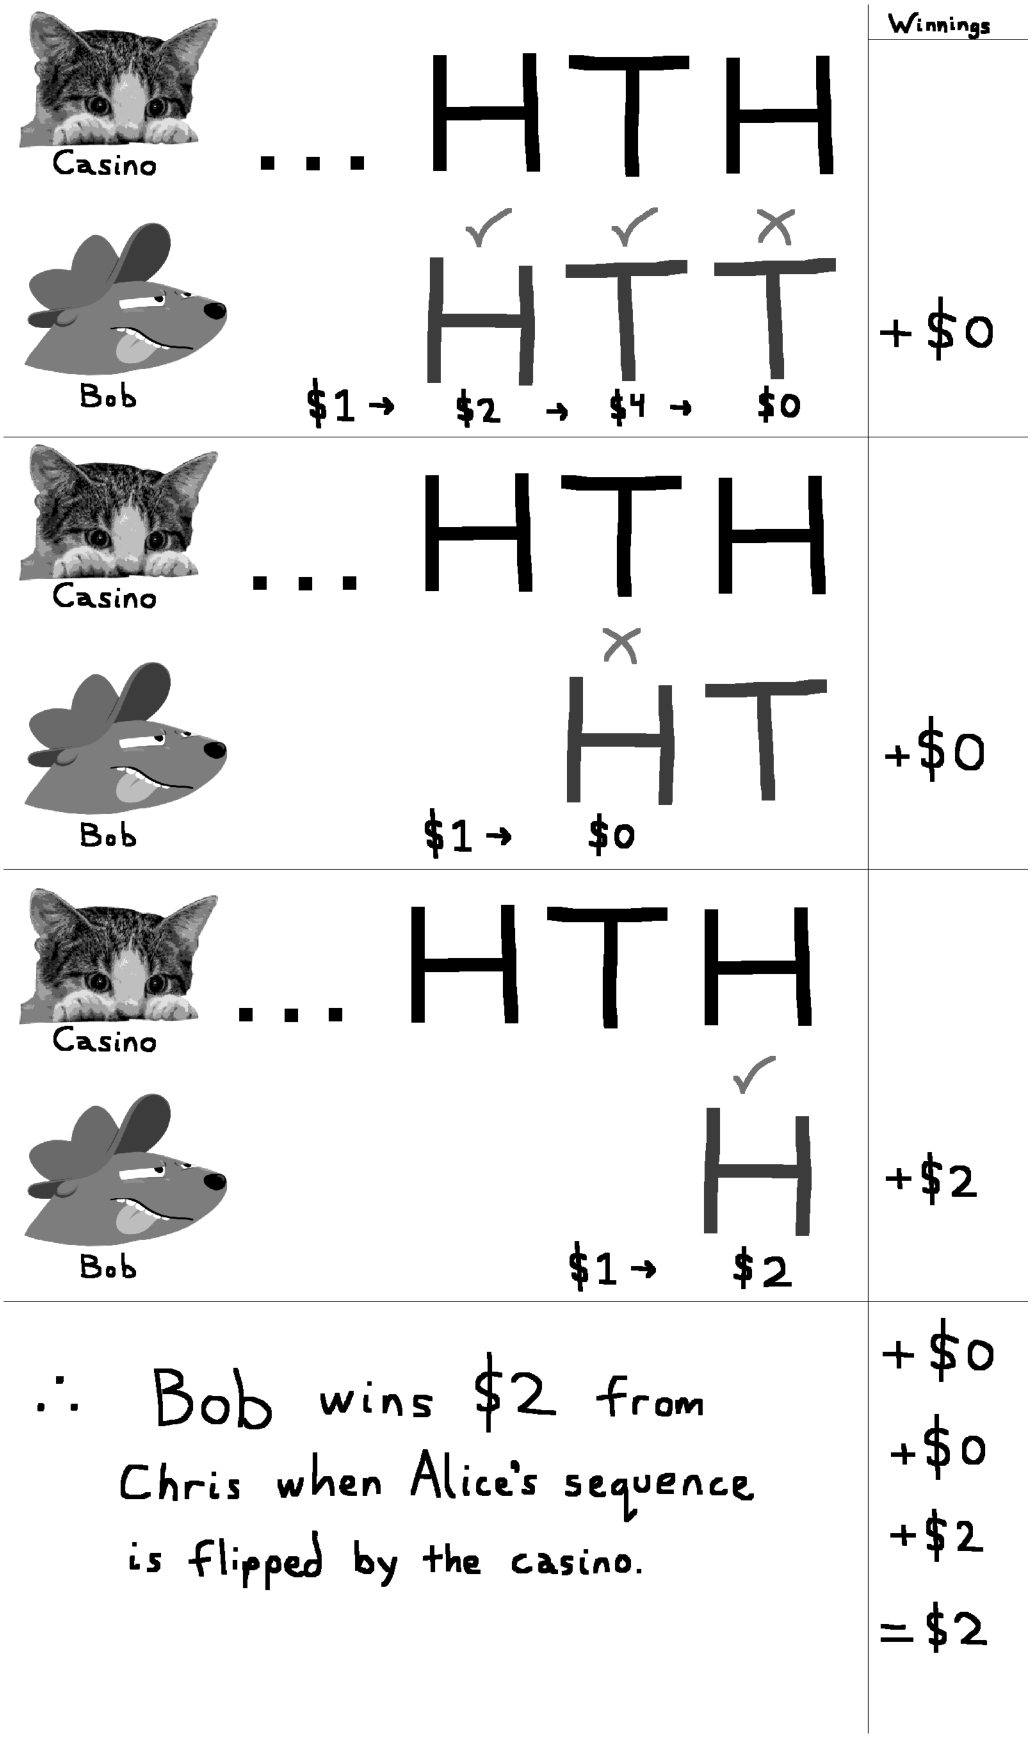
Winnings
Casino
... HTH
Bob
HTT
$1 → $2 → $4 → $0
+$0
Casino
... HTH
Bob
HT
$1 → $0
+$0
Casino
... HTH
Bob
H
$1 → $2
+$2
∴ Bob wins $2 from Chris when Alice's sequence is flipped by the casino.
+$0
+$0
+$2
=$2

Recall that, when Alice's sequence occurs in the casino's flips, Chris returns *all* winnings to Alice—even those from "in progress" betting sequences—and begins betting afresh. This has the effect of starting a *new* sequence of casino coin flips, where previous flips are no longer of any significance. If we refer to the probability table constructed in section 7.2, we see that HHH is less likely to occur first when compared to almost every other sequence. Therefore, every time a new sequence of casino flips starts, the chance of the casino's flips starting Alice's sequence is *less* than $\frac{1}{8}$, meaning we can't expect to make \$14 for every \$10 spent like we predicted above. In fact, the chance of Alice's sequence being started has decreased by the exact amount required to make Alice's expected profit \$0.

Returning to the two player case with both Alice and Bob playing at the same table, let's consider the expected winnings of both players. Consider the case where Alice bets on the pattern HTT and Bob bets on the pattern THT. Since Alice and Bob are solely concerned with earning more money than the other, let's look at how much *more* money each player wins, comparatively, when their pattern occurs first in the casino's flips. We'll refer to this amount as the *lucre* for each player, just to ease in description. If Alice's pattern occurs first, then by Chris' betting rules, Alice wins \$8, while Bob wins \$2. Therefore, Alice's lucre is \$6 since she wins \$6 more than Bob. Similarly, if Bob's pattern occurs first, then Bob wins \$10, while Alice wins \$4. Therefore, Bob's lucre is also \$6. We can represent this info with the ratio $6 : 6$, where the left number represents Alice's lucre, while the right represents Bob's lucre.

Being the astute observers that we are, we notice that the ratio $6 : 6$ is the same ratio that Conway's Algorithm gives us when comparing the sequences HTT and THT, except with an extra factor of 2 present on both sides (which thankfully doesn't affect the calculated odds). However, now that we've constructed this ratio from Alice and Bob's convoluted betting game, we can intuitively understand what the ratio is telling us: because both patterns give the same expected lucre, each pattern must be equally likely to occur first, since neither pattern rigs the betting returns in their favour. If we check the probability table from section 7.2, we see that, indeed, the patterns HTT and THT are equally likely to occur first.

If, instead, Alice bets on the pattern TTT while Bob bets on the pattern HTT, then Alice's lucre becomes \$14 while Bob's lucre becomes \$2, giving us a lucre ratio of $14 : 2$. This time, it seems Alice has massively rigged the betting returns in her favour: for every \$14 she gains over Bob when her pattern occurs, she'll only make \$2 less than Bob if Bob's pattern occurs. If both their patterns were equally likely to occur first, Alice would have a huge advantage! However, since the coin is fair, there's no pattern that either player can choose so as to give them a monetary advantage. Therefore, it must be the case that Alice's pattern is *less likely* to occur first so as to ensure she isn't given a monetary advantage over Bob. Again referring to the table in section 7.2, this is exactly what we see.

Symbolically, our findings can be summarised as follows: let $\mathrm{R}_X(Y)$ represent how much money pattern Y would win if pattern X occurs first in the sequence of casino flips. Then, if Alice's chosen pattern is represented as A, and Bob's chosen pattern is represented as B, then Alice's lucre can be represented as

$$\mathrm{R}_A(A) - \mathrm{R}_A(B),$$

while Bob's lucre can be represented as

$$\mathrm{R}_B(B) - \mathrm{R}_B(A).$$

Our lucre ratio is then represented as

$$\mathrm{R}_A(A) - \mathrm{R}_A(B) : \mathrm{R}_B(B) - \mathrm{R}_B(A).$$

From our reasoning above, whichever side has the lower number represents the pattern that's more likely to occur first. If we compare this form with Conway's Algorithm, we see that both are equivalent (except for the extra factor of 2 mentioned previously). So, using our hypothetical gambling scenario with Alice and Bob, we were able to derive an equivalent representation of Conway's Algorithm.

Of course, none of this is a rigorous proof that Conway's Algorithm is correct.[27] However, after having considered Alice's and Bob's betting game, the intuition behind Conway's Algorithm should now be clear. The leading numbers used in the algorithm are just a tool used to quantify the supposed "advantage" one pattern could have over another. In Alice's and Bob's case, those leading numbers correspond to expected profits for each player. Unfortunately for Alice and Bob, those same expected profits also ensure that neither can expect to have a financial advantage over the other. It seems they'll have to prove their worth to each other some other way.

8. Nontransitivity in the Real World

After having become familiar with a handful of nontransitive games, we may wonder whether there are other, less artificial examples of nontransitivity. For the original dice game introduced in section 3, it initially seemed obvious that our "power" relationship between the dice should be transitive, and yet it wasn't. Are there

[27] Again, see pages 401 to 410 of Graham et al. [2011] for a more formal proof of Conway's Algorithm.

any situations in our lives where we *assume* a specific relationship to be transitive without having any real reason for assuming so?

It turns out, perhaps unnervingly, that there are several such instances where transitivity is assumed.

One of the most important places where we tend to rank things are in elections. There are many, many different election systems around the world, but the general process is always the same: rank a set of candidates from most favourable to least favourable, and choose the most favourable candidates for whatever position the election is for.

Let's consider a population of voters where each voter has their own particular ranking for each candidate. For instance, if three major candidates are running for office—Mr. Bat (🦇), Ms. Email (✉), and Dr. Skull (☠)—then one voter's candidate ranking may look like

$$\text{🦇} \succ \text{✉} \succ \text{☠},$$

where this voter prefers Mr. Bat to Ms. Email, and prefers Ms. Email to Dr. Skull. Let's imagine three hypothetical voters, say A, B, and C. Their candidate rankings are given as follows:

$$\begin{array}{rccccc} A: & \text{🦇} & \succ & \text{✉} & \succ & \text{☠}, \\ B: & \text{☠} & \succ & \text{🦇} & \succ & \text{✉}, \\ C: & \text{✉} & \succ & \text{☠} & \succ & \text{🦇}. \end{array}$$

Clearly, if each voter voted solely for their top candidate, there would be a tie. However, let's say that voters A, B, and C are instead voting on a *subset* of possible candidates. Maybe Mr. Bat and Ms. Email are facing off at a debate, and our voters must choose which candidate they prefer from the two that are debating. In this case, the candidate rankings for each voter change slightly since Dr. Skull isn't a potential candidate in this scenario:

$$\begin{array}{rccccc} A: & \text{🦇} & \succ & \text{✉} & \succ & \cancel{\text{☠}}, \\ B: & \cancel{\text{☠}} & \succ & \text{🦇} & \succ & \text{✉}, \\ C: & \text{✉} & \succ & \cancel{\text{☠}} & \succ & \text{🦇}, \end{array} \quad \Longrightarrow \quad \begin{array}{rccc} A: & \text{🦇} & \succ & \text{✉}, \\ B: & \text{🦇} & \succ & \text{✉}, \\ C: & \text{✉} & \succ & \text{🦇}. \end{array}$$

When ranking exclusively Mr. Bat and Ms. Email, Mr. Bat would win the public vote 2:1. Similarly, if Ms. Email and Dr. Skull were debating,

$$\begin{array}{rccccc} A: & \cancel{\text{🦇}} & \succ & \text{✉} & \succ & \text{☠}, \\ B: & \text{☠} & \succ & \cancel{\text{🦇}} & \succ & \text{✉}, \\ C: & \text{✉} & \succ & \text{☠} & \succ & \cancel{\text{🦇}}, \end{array} \quad \Longrightarrow \quad \begin{array}{rccc} A: & \text{✉} & \succ & \text{☠}, \\ B: & \text{☠} & \succ & \text{✉}, \\ C: & \text{✉} & \succ & \text{☠}, \end{array}$$

so Ms. Email would win the public vote 2:1. Finally, if Dr. Skull and Mr. Bat debated,

A : 🦇 $\succ$ ✉ $\succ$ ☠, $\Longrightarrow$ A : 🦇 $\succ$ ☠,
B : ☠ $\succ$ 🦇 $\succ$ ✉,   B : ☠ $\succ$ 🦇,
C : ✉ $\succ$ ☠ $\succ$ 🦇,   C : ☠ $\succ$ 🦇,

so Dr. Skull would win the public vote 2:1. Summarising, we see that, for pairwise comparisons, 🦇 $\succ$ ✉, ✉ $\succ$ ☠, and ☠ $\succ$ 🦇. This means

🦇 $\succ$ ✉ $\succ$ ☠ $\succ$ 🦇.

Once again, nontransitivity appears. The order in which candidates are paired for debates completely determines which candidate can win the public's support! This situation is known as the Condorcet Paradox, named after Marquis de Condorcet who first described such a situation in the 1700s. Whether one uses a single-vote plurality system or a ranked ballot, the Condorcet Paradox can occur.[28]

The Condorcet Paradox has pretty striking implications for elections, mainly that the candidate who manages to receive the most votes isn't necessarily the obvious, undisputed preference of the entire voting body. As we saw above, the way in which an election is conducted can have as much an effect on the election's outcome as the popularity of the candidates themselves.

In fact, in the most extreme case, the way in which the election is conducted can *entirely* decide the outcome of the election! The first two theorems of Saari [1995] (pages 42 and 45) say that, given the right voter preferences, it's possible to have *any* election behaviour we want for any subset of candidates we choose.[29] In essence, what this means is that the voters' preferences applied to a particular subset of candidates in an election *cannot* necessarily be used to determine anything about voter preferences applied to a different subset. For instance, if, after tallying up the preferences of an entire voting population, we found that

🦇 $\succ$ ✉ $\succ$ ☠,

[28]For more on the Condorcet Paradox, see pages 298-299 of Gardner [2001] for a brief history of the paradox and its consequences for creating voting systems.

[29]The theorems make use of an analogy between the structures of voting preferences and chaotic dynamical systems (such as the orbits generated by repeatedly applying Newton's Method to points on a polynomial curve). Because of this, the exact wording of the theorems cannot be replicated here unless a considerable amount of time is spent building up the ideas and terminology used within Saari [1995]. Instead, we choose to paraphrase the theorems in a way that's compatible with this project. The interested reader is encouraged to read Saari [1995] for the motivation leading to these theorems.

it's entirely possible that, when tallying up the voters' preferences on subsets of these candidates, we find that

☒ ≻ 🦇, ☠ ≻ 🦇, and ☠ ≻ ☒,

even though these preferences seem contradictory. Underlying all this confusion is, once again, the false assumption that the preferences of voters should be transitive. Our examples show, however, that under certain conditions, voter preferences can indeed be nontransitive.

These examples only deal with three candidates total, so there are only so many ways to create apparent "paradoxes" in these rankings. However, the theorems in Saari [1995] apply just as well to scenarios with over three candidates. In these cases, because a greater number of candidate subsets are possible, the election behaviour can get much more erratic and unpredictable.

Moving away from elections, the apparent nontransitivity of preferences has implications for *any* area of our lives where things need to be ranked based on some criteria. As an example, say we work for some sort of marketing firm, and we need to determine how a particular population would rank a set of three subscription services.[30] It's (naively) assumed by our managers that there should exist a single, "most attractive" subscription service, and it's this subscription we're tasked to find. Let A represent the percentage of our population that prefers the first subscription, B the corresponding percentage for the second subscription, and likewise for C. When asking consumers to rank between only the first two subscriptions, we see that

$$A = 0.68, \quad B = 0.32.$$

However, when asking consumers to rank all three subscriptions, we find that

$$A = 0.16, \quad B = 0.84, \quad C = 0.$$

Once again, preferences applied to one subset of objects don't necessarily give any information about preferences applied to other subsets. Our consumers' preferences for the first and second subscriptions completely swap when we introduce a third subscription for them to consider. How do we determine which subscription to report to our managers as the "most attractive"?

Compare this situation to our friend's dice game. The power ranking of the dice changes depending on how many dice each player is rolling. Is it not reasonable, then, that other rankings are subject to change depending on how many objects are being considered?

[30]This example was adopted from Klimenko [2015]. The example, however, is based off of real data!

Indeed, Klimenko argues in page 4385 of Klimenko [2015] that "a real buyer in the real world is likely to have a[n] intransitive set of preferences", where adding additional objects to a ranking "provides us with additional information that makes [different options] more attractive".

This human tendency towards nontransitivity can even affect our ability to make emergency decisions. Consider the following safety regulations:[31]

During an emergency situation:

(1) **Leadership:** if a manager is present, they should lead and organise site personnel to contain the cause of emergency.

(2) **Safety:**

  (a) the manager and personnel stay on site during emergency if there is no immediate danger to personnel, but

  (b) personnel must be evacuated whenever there is a significant danger to personnel.

Following these regulations, a manager has several factors to consider when deciding on what to do in the case of an emergency. Thus, it's likely that a manager may want to compare several possible actions to see which one seems the most desirable. Consider the following three actions a manager could take in the case of an emergency fire:

A) Evacuating personnel and abandoning the site.

B) Organising personnel to monitor the situation on site.

C) Organising personnel to contain and extinguish the fire.

If regulation 1 is considered the most important, then action C should be chosen to B as the manager is expected to take action to contain the source of the emergency. If one considers regulation 2a as the most important, then action B should be chosen over A since monitoring a fire won't cause immediate harm to anyone on site. If regulation 2b is considered the most important, then action A should be chosen over C since attempting to quell a fire could potentially cause harm to on-site personnel. From these considerations, we find that

$$\mathrm{C} \succ \mathrm{B} \succ \mathrm{A} \succ \mathrm{C},$$

[31]These example safety regulations are adapted from Klimenko [2015].

where the notation $X \succ Y$ represents action X being favourable to action Y. From the relation above, we see no single course of action is obviously better than the rest. Nontransitivity rears its head once again. Not only does nontransitivity prevent us from easily ranking sets of objects, but it also prevents us from easily deciding on an action given a set list of possibilities. Section 4 of Klimenko [2015] provides another great example of nontransitivity manifesting in human behaviour, this time with respect to choosing actions in light of particular laws.

However paradoxical it may seem, nontransitivity is an innate part of human behaviour and decision making, as the previous examples demonstrate. Unnervingly, the natural world is just as prone to nontransitivity as we are.

Consider the set of three circles in figure 14. If we look at the black ring, we notice that it's completely under the slate (dark grey) ring. Similarly, the slate ring is completely under the grey (light grey) ring. However, unexpectedly, the grey ring is completely under the black ring. Symbolically, if we use the "$\prec$" symbol to represent one ring being under another, then we have that

$$\text{Black} \prec \text{Slate} \prec \text{Grey} \prec \text{Black}.$$

Even physical space can display nontransitivity.

This configuration of rings is known as the Borromean rings. It represents the simplest possible *Brunnian* link, where if any one of the rings are removed, the other two rings can be pulled free of each other. However, with all three rings present, the rings are linked together. If we were to try and make the Borromean rings using perfectly circular rings as in the figure above,[32] we'd find that it's not actually

[32]Just pretend they're perfectly circular. The author is incapable of drawing perfect circles with a mouse.

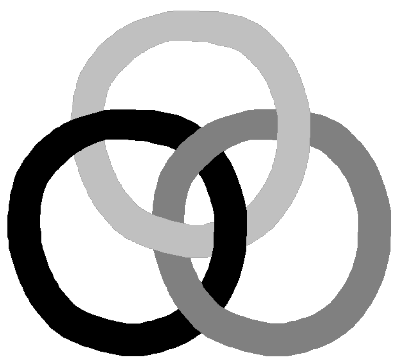

FIGURE 14. A representation of the Borromean rings.

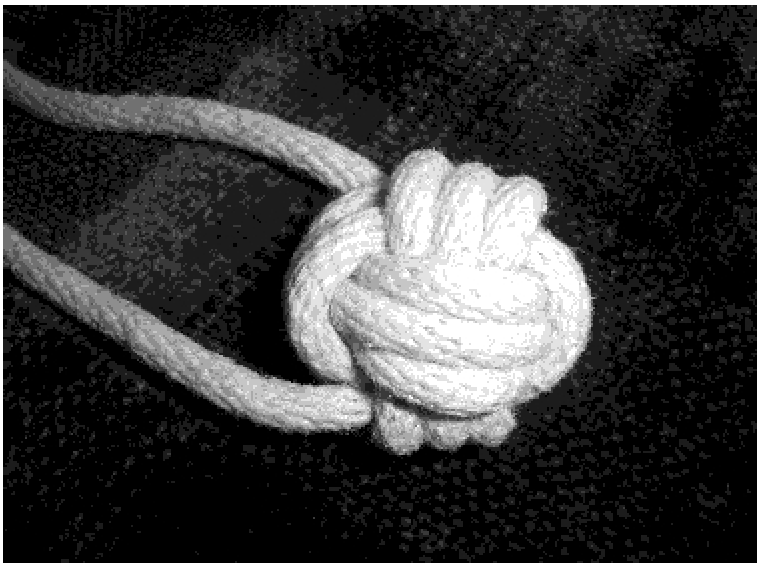

FIGURE 15. "Knot Monkey Fist" by Tortillovsky. Licensed under CC BY-SA 3.0 (http://creativecommons.org/licenses/by-sa/3.0/), via Wikimedia Commons. Desaturated and posterised to ten colours.

possible. As stated in Lamb [2016], in order to construct a set of loops that mimic the Borromean rings in our three-dimensional space, we must use other shapes, such as ellipses or warped circles. One common way to construct the Borromean rings is by tying a monkey's fist knot. Without any of the three triple-rope rings in a monkey's fist, the entire knot falls apart. A structure similar to the Borromean rings is the braid. For the usual three-strand braid people tie their hair into, when comparing any two strands, one will always be above the other. However, all three strands together keep the braid interlinked.

Thus, even the space we reside in is prone to nontransitivity. Not only that, but the creatures that inhabit that space also can't seem to resist nontransitivity— human or otherwise.

Consider the Lotka-Volterra population model. If we have two species we'd like to model the population of, and one of the species is assumed to prey on the other, we can model the change of these two populations over time using a system of differential equations. If R represents the number of rabbits in an environment, and F represents the number of foxes in the environment, then the Lotka-Volterra population model says that the two quantities should change according to the following two differential equations:

$$\frac{\mathrm{d}R}{\mathrm{d}t} = a_r R - bRF,$$
$$\frac{\mathrm{d}F}{\mathrm{d}t} = cRF - a_f F,$$

where a_f, a_r, b, and c are all positive constants adjusted to match whatever scenario is being modelled.

From these two equations, we can gleam that the rabbit population benefits from the environment. The more rabbits there are, the more rabbits that are begat, so without any external forces, the rabbits are able to take as much as they need from the environment in order to thrive (as represented by the term $a_r R$ in the equation controlling the change in rabbit population). As well, foxes benefit from the rabbit population. The more rabbits there are, the more foxes that are begat (as represented by the term cRF). However, the foxes do *not* benefit from the environment like the rabbits do. The more foxes there are, the *less* foxes that are begat (as represented by the terms $-bRF$ and $-a_f F$). So, in a slight abuse of language, we could say that the environment benefits from the foxes, not the other way around.

If the symbol "$\succ$" is used to represent one object benefitting from another, then the above analysis of the Lotka-Volterra model shows that

$$\text{Environment} \succ \text{Foxes} \succ \text{Rabbits} \succ \text{Environment.}$$

Yet again, nontransitivity arises,[33] this time in an environment completely separate from human constructs.

This is far from the only example. For instance, it's been observed that the mating habits of particular species of side-blotched lizards form a nontransitive cycle: a habit is preferred when compared to a specific habit, but is subpar when compared to another. As Kristin Leutwyler relays in Leutwyler [2000], "each morph successfully used a different tactic to exploit weaknesses of another strategy and a morph's success depended on the close proximity of a vulnerable alternative strategy", meaning the nature of the competition in the lizards' environment "promote[d] conditions that favor[ed] each morph, and thus preserve[d] all three strategies of the rock-paper-scissors cycle in the long term". In other words, attempting to rank the lizards' mating strategies from best to worst results in the same sort of nontransitive cycle we see in Rock, Paper, Scissors.

From these examples, it's now clear to us that nontransitivity pervades the entirety of our world. Not only does it appear in several of our games, institutions, and thoughts, but also in the processes that surround us, free from any artificial constructs we may impose. It seems, then, that assuming transitivity is not only wrong in some cases, but entirely irrational. There are many things in our world that

[33] Section 5 of Klimenko [2015] goes into further detail regarding the nontransitivity that arises in this population model.

evade a simple placement on a hierarchy. In the case of our friend's games, a more careful analysis is needed—one other than determining the overall "best" die—to fully appreciate the structures at play. For the nontransitive scenarios explored above, a more nuanced approach is required—one other than placing things in a hierarchy—in order to properly understand what's going on.

9. Conclusion

Starting with a seemingly innocuous analysis of a strange dice game, we discovered that its underlying nontransitive structure shares many properties with the structure of coin flip sequences, election ordering, and even population dynamics. Instead of having a hierarchy-like structure, these scenarios have a structure akin to Rock, Paper, Scissors, where the "power" of individual objects is not an absolute measure, but rather a relative measure on pairs or groups of objects.

The apparent strangeness of nontransitive structure almost always comes from a false assumption of transitivity. For our friend's dice game, the assumption of transitivity gave our friend a strategical advantage against us. For other scenarios, such as those described in section 8, assuming transitivity completely throws off our understanding and intuition.

Understanding nontransitivity in the context of our friend's games may seem a bit separated from reality. After all, the games are *meant* to fool us. They're created specifically to force nontransitive cycles to occur. However, understanding nontransitivity in the context of consequence-free games[34] helps better equip us for the cases where nontransitivity *does* have consequences. While the process of analysing dice rolls is quite different from the process of observing lizard species, we've seen that the relationships between pairs of dice and between pairs of lizard species are surprisingly similar, and thus understanding one scenario should give us tools (or, at the very least, ideas) for working with the other.

Ultimately, though, just being aware of nontransitivity is what proves to be the most useful. By understanding that certain relations need not be transitive, we can begin to explore other possible structures for any given situation. As Klimenko states on page 4366 of Klimenko [2015], "[b]oth transitive and [non]transitive effects are common and need to be investigated irrespective of what we tend to call 'rational' or 'irrational'". After having seen nontransitivity in several places throughout this project, we can hopefully see that it's a structure worth considering. As well, we've hopefully come to appreciate nontransitivity as an innate part of many systems and environments, and not as some niche academic curiosity.

[34]Well, the games are only consequence-free if we ignore the loss of our dignity to our friend.

10. Image Credits

The following is a list of image credits for the public domain images used throughout this project. All non-public domain images used are credited with the appropriate attribution where used. All images created by the author, and any edited images used in this project, were created in GIMP 2.10.34.

- "A Kangaroo on a Grass Field in an Enclosure" by 卓浩 虞 from Pexels.
- "Close-Up Shot of an Otter" by Ellie Burgin from Pexels.
- "Selective Focus Photography of Brown Squirrel" by Rod Dion from Pexels.
- "Woman Raising Her Hand" by EKATERINA BOLOVTSOVA from Pexels.
- "Hand Peace Sign" by Kevin Malik from Pexels.
- "Close-up of a Clenched Fist" by Kevin Malik from Pexels.
- "Orange Cube" by Magda Ehlers from Pexels.
- "Alice and Bob Bet It All" by Jeremy Chevere. This image was commissioned for use in this project.
- "Cute Kitten hiding behind a Pillow" by Francesco Ungaro from Pexels.

References

Highest scoring words, 2023. URL `https://www.scrabblewizard.com/highest-scoring-words/`.

Statistics Canada, Sep 2019. URL `https://www.statcan.gc.ca/en/concepts/definitions/guide-symbol`.

Walter W. Funkenbusch and Donald G. Saari. Preferences among preferences or nested cyclic stochastic inequalities. *CONGRESSUS NUMERANTIUM*, 39(1): 419–432, 1983.

Martin Gardner. *The Colossal Book of Mathematics: Classic Puzzles, Paradoxes and Problems*. Norton, New York, 2001.

Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. *Concrete Mathematics*. Addison Wesley, Boston, 2011.

James Grime. The bizarre world of nontransitive dice: Games for two or more players. *The College Mathematics Journal*, 48(1):2–9, 2017. doi: 10.4169/college.math.j.48.1.2. URL `https://doi.org/10.4169/college.math.j.48.1.2`.

Alexander Klimenko. Intransitivity in theory and in the real world. *Entropy*, 17 (12):4364–4412, Jun 2015. ISSN 1099-4300. doi: 10.3390/e17064364. URL http://dx.doi.org/10.3390/e17064364.

Evelyn Lamb. A few of my favorite spaces: Borromean rings, Sep 2016. URL https://blogs.scientificamerican.com/roots-of-unity/a-few-of-my-favorite-spaces-borromean-rings/.

Kristin Leutwyler. Mating lizards play a game of rock-paper-scissors, Dec 2000. URL https://www.scientificamerican.com/article/mating-lizards-play-a-gam/.

Joshua B. Miller. Penney's game odds from no-arbitrage, 2019. URL https://doi.org/10.48550/arXiv.1904.09888.

Robert A. Page. Casino dice game, July 1992. URL https://www.freepatentsonline.com/5133559.html.

Nicholas Pasciuto. The mystery of the non-transitive grime dice. *Undergraduate Review*, 12(1):107–115, 2016. URL https://vc.bridgew.edu/undergrad_rev/vol12/iss1/1/.

Jr. Richard P. Savage. The paradox of nontransitive dice. *The American Mathematical Monthly*, 101(5):429–436, May 1994. doi: 10.2307/2974903. URL https://www.jstor.org/stable/2974903.

Donald G. Saari. A chaotic exploration of aggregation paradoxes. *SIAM Review*, 37 (1):37–52, 1995. ISSN 00361445. URL http://www.jstor.org/stable/2132752.

Stanislaw Czeslaw Trybula. On the paradox of n random variables. *Applicationes Mathematicae*, 8:143–156, 1965. URL https://api.semanticscholar.org/CorpusID:117782454.

Zalman Usiskin. Max-Min Probabilities in the Voting Paradox. *The Annals of Mathematical Statistics*, 35(2):857–862, 1964. doi: 10.1214/aoms/1177703585. URL https://doi.org/10.1214/aoms/1177703585.

Pavle Vuksanovic and A. J. Hildebrand. On cyclic and nontransitive probabilities. *Involve, a Journal of Mathematics*, 14(2):327–348, apr 2021. doi: 10.2140/involve.2021.14.327. URL https://doi.org/10.2140%2Finvolve.2021.14.327.

Eric W Weisstein. Probability axioms, Oct 2023. URL https://mathworld.wolfram.com/ProbabilityAxioms.html.

Michael Winkelmann, Mar 2015. URL http://www.miwin.com/Miwins%20Dodekaeder/Miwins_Dodekaeder.html.