PARADOXES WITH DICE AND ELECTIONS

Marcel Jackson

La Trobe University

Some startling paradoxes can be found with games involving relabeled dice. We look at some of the possibilities and their connection with other seemingly impossible situations arising from rankings.

Elementary probability has the habit of producing some nasty surprises to one's intuition. One of the best and easiest ways of experiencing some of the seeming impossibilities is by way of generalised dice.

Paradox 1 (The Steinhaus-Trybula paradox)

Here is an easy two player "game" (as we will see, it is not so much fun for one of the players). We begin with three standard six sided dice X, Y, Z whose faces have been relabelled in the following unusual way.

	Face value					
X:	1	1	4	4	4	4
Y:	3	3	3	3	3	3
Z:	2	2	2	2	5	5

Table 1

There are two players playing the game, player **U** and player **I**. Each player chooses a (distinct) die and rolls – the player with the higher value is the winner. **I** kindly gives player **U** the first choice of any one of the three dice. Let us say that **U** chooses die X. After some thought, **I** chooses die Z, and then both players roll their dice. What are the possibilities? Because both dice have 6 sides there are $6\times 6=36$ possibilities, however because of repetitions in the labeling of the dice, many of these outcomes are indistinguishable. In Table 2, the columns are labelled by the six faces values of die X and the rows by the face values of die Z. In a row with Z face value z and X face value x, we have put a Z if $z>x$ and an X if $x>z$, in other words, the X's and Z's in the content of the table correspond to possible wins by the respective dice.

	1	1	4	4	4	4
2	Z	Z	X	X	X	X
2	Z	Z	X	X	X	X
2	Z	Z	X	X	X	X
2	Z	Z	X	X	X	X
5	Z	Z	Z	Z	Z	Z
5	Z	Z	Z	Z	Z	Z

Table 2

Evidently, Z wins in 20 out of the 36 possible outcomes. For an alternative analysis, note that dice Z beats dice X if:

- Z comes up 2 and X comes up 1 (happens in $4 \times 2 = 8$ ways); or
- Z rolls a 5 (happens in $2 \times 6 = 12$ ways),

giving a total of 8+12=20 wins out of a possible 36.

So, in all probability, **I** will beat **U**. It seems like **I** made a good choice! The amazing fact is that *no matter which dice is chosen by **U**, there is always a dice **I** can choose that will win in a two way roll with probability more than one half*. The corresponding probabilities are:

- die X beats die Y with probability 2/3,
- die Y beats die Z with probability 2/3,
- die Z beats die X with probability 5/9 (=20/36, as observed above).

This seemingly impossible scenario is the *Steinhaus-Trybula paradox* (which more generally asserts that it is possible to have three independent random variables X, Y, Z such that each of the probabilities $P(X > Y)$, $P(Y > Z)$, and $P(Z > X)$ are greater than one half).

The paradox can be interpreted in many other ways too. For example it (arguably) goes some way toward explaining a common situation in competitive sports and games whereby one team (or player) has a good winning record against a second team which in turn has a good winning record against a third team which in turn has a good winning record against the first team.

For a more rigorous example consider three buses (say, A, B and C) scheduled to pass by a particular bus stop at the same time each day. The Steinhaus-Trybula paradox shows that it is possible that bus A is "on average" more likely to arrive at the bus stop before bus B, bus B is more likely to arrive before bus C, while bus C is more likely to arrive before bus A! For example, the faces of the three dice X, Y, Z of Table 1 could represent possible (independent) random delays for the three buses A, C, B respectively.

Paradox 2.

After their disappointing performance in our first die game, player **U** decides to team up with a new player, player **V**, to beat the devious player **I** in a three way roll-off using a fresh set of dice, labeled as follows:

A:	1	5	5	5	5	7
B:	2	4	4	6	6	6
C:	3	3	3	3	8	8

Table 3

Player **U** insists that **I** goes first this time, and so player **I** chooses die C. This looks good for player **U** and player **V** (who choose, say, dice A and B respectively). They compare their dice to die C as follows: die A will beat C if and only if C rolls a 3 and A rolls 5 or 7 (happens in 4×5 ways), which is with probability 20/36, or 5/9. Similarly, B will beat C with probability 5/9 also. So it appears that die A and die B are both better choices than die C.

Now let's see what happens in the three way roll-off. Die C wins if:

- C shows a 3 while die A and B show 1 and 2 respectively ($4 \times 1 \times 1 = 4$ possible ways); or
- C produces an 8 ($2 \times 6 \times 6 = 72$ ways).

So C wins in 4+72=76 out of the possible $6 \times 6 \times 6 = 216$ possibilities. Die A wins if:
A is 5 (4 ways), B is 2 or 4 (3 ways) and C is 3 (4 ways); or
A is 7 (1 way), B is anything (6 ways), and C is 3 (4 ways).

This means that A wins in $4 \times 3 \times 4 + 1 \times 6 \times 4 = 72$ out of 216 possibilities. We can find out how often die B wins in a similar way, or alternatively we can subtract the number of ways that A and C can win from the total number, giving $216 - 76 - 72 = 68$ ways. So in fact, die C is the most likely winner in the three way roll, even though it is the least likely to win in a two way roll. Again, **I** made a good choice.

This paradox also has a nice interpretation in terms of public transport. Say that A, B and C are also three buses scheduled to pass by the same stop at the same time. It is possible that

- bus A most likely arrives before bus C,
- bus B most likely arrives before bus C, but
- bus C is the bus most likely to arrive first out of the three!

1. Some notation.

To give other examples of the type described above, it is useful to have dice with a number of sides other than just 6. By a *dice with n sides* (or an *n sided dice*) we mean some solid object with *n* faces and which can be rolled (or thrown or spun) on a flat surface such that in the final resting position, exactly one of the *n* faces is lying against the flat surface, and this face is equally likely to be any of the *n* different possibilities. With regular dice we read the value from the side facing upwards, but for non-standard dice there may conceivably be no such face (in such a case we could choose the face that rests against the surface). An easy construction of an *n*-sided die would be a rod (sharpened at both ends) and whose cross section is a regular *n*-gon – similar to the shape of several common brands of pencil. There are also a number of commercially available games that use dice that are octahedrons and other shapes (the potential possibilities are worthy of an article themselves). Many of the basic examples however can be achieved using just the regular number of sides; it is possible to buy blank (six-sided) dice which can then be labelled using stickers

Each of the faces of one of these generalised dice will be labelled by a whole number (the *value* of the face), and after rolling two dice A and B we say that *A beats* (or *wins over*) *B* if the value of the face shown on A is greater than that shown on B. We will write A>B to indicate that the probability that A wins over B in a two way roll is greater than 1/2. We will refer to this probability as the *dominating probability* of A over B. Dice in which A>B>C>…>A are often called *non-transitive* dice because they demonstrate that the relation > between dice is a non-transitive relation (transitivity of a relation ~ means that if x~y and y~z then x~z).

2. Basic examples and easy investigations.

There are limits to the dominating probabilities that can be achieved for three non-transitive dice and in fact the example in Table 1 is probably the most convincing demonstration of the Steinhaus-Trybula paradox that can be achieved using three regularly shaped (that is, cubic) dice. The original investigation of the Steinhaus-Trybula paradox in [stetry] and [ste] does not mention dice at all, rather it concerns random variables. These articles investigate three random variables X, Y and Z under the restriction that each of the probabilities $P(X > Y)$, $P(Y > Z)$ and $P(Z > X)$ are greater than 1/2. One interesting result found is that at least one of these probabilities must be at most $\frac{\sqrt{5}-1}{2} = 0.61803\ldots$, the golden ratio! (The golden ratio is a ubiquitous mathematical constant that also appears as the "continued fraction" $\frac{1}{1+\frac{1}{1+\frac{1}{\ddots}}}$ and as the limit of the sequence $\frac{1}{1}, \frac{1}{2}, \frac{2}{3}, \frac{3}{5}, \frac{5}{8}, \frac{8}{13}, \ldots,$ where 1,1,2,3,5,8,13,21,… is the sequence of Fibonacci numbers.) The following example from [ste] shows that this bound can actually be obtained. Define *X, Y* and *Z* as follows.

$P(X=1):=\frac{3-\sqrt{5}}{2}$ and $P(X=4):=\frac{\sqrt{5}-1}{2}$,
$P(Y=3):=1$,
$P(Z=2):=\frac{\sqrt{5}-1}{2}$ and $P(Z=5):=\frac{3-\sqrt{5}}{2}$.
(note that X cannot equal any other value than 1 or 4, while Y can only equal 3 and Z can only equal 2 or 5; note also the similarity with the three dice in Table 1). Then $P(X>Y)=P(Y>Z)=P(Z>X)=\frac{\sqrt{5}-1}{2}$. In other examples presented in [ste] the three variables each have the same mean and variance.
We cannot achieve the actual golden ratio bound for dice, because the golden ratio is irrational while the probabilities that arise from two-way dice comparisons are always fractions between whole numbers. This is why the example in Table 1 is about the best that can be achieved using three regularly shaped dice. If dice with more than six sides are allowed, the effect can be increased slightly. The following set of 8-sided non-transitive dice with A>B, B>C and C>A have dominating probabilities that are very close to the golden ratio.

A:	1	1	1	4	4	4	4	4
B:	3	3	3	3	3	3	3	3
C:	2	2	2	2	2	5	5	5

Table 4

In fact the dominating probabilities in this example give the closest approximation to the golden ratio that is possible for dice with at most 8 sides [sav].

Exercise. Get students to make some regular octahedrons out of cardboard and mark them according to Table 3. Get them to calculate the probabilities of each die winning over the other (as in Table 2, or the alternative approach that follows Table 2)

Challenging exercise. Is there a set of three non-transitive dice with only 2 sides? With 3 sides?

(The answer to the first question is “no” while the answer to the second question is “yes” – for a hint, look at removing repetitions in Table 1.)

As the golden ratio is only around 0.618, the probabilities with three dice are still not really striking. With 4 regular dice however, it is possible to achieve cyclical probabilities of 2/3. A nice example due to Effron (see [gar]) is as follows

A:	1	1	5	5	5	5
B:	4	4	4	4	4	4
C:	3	3	3	3	7	7
D:	2	2	2	6	6	6

Table 5

In this example one finds that A>B, B>C, C>D and D>A, all with probability 2/3. Some reasonably convincing demonstrations of the Steinhaus-Trybula paradox become possible. Say that 10 rounds of the dice trick are performed using these four dice, with **U** choosing first, and **I** choosing second (and choosing to win). The likelihood of player **I** winning *at least half* of the games is around 92%.

At 20 rounds it is 96%. If **I** wants to win *more than half*, the figures are lowered slightly with only around 91% certainty for 20 games. In fact it is better to play an odd number of games – after just 15 games there is already around 91% certainty of **I** winning more often than **U**. Calculating these probabilities is also a good exercise in the use of the binomial distribution.

Exercise. Get students to mark blank dice according to Table 4. Perform 10, 15 or 20 trials with one student choosing first, and the second student choosing to win.

Exercise. Get students to calculate the probabilities for the four dice of Table 4. There are two comparisons not mentioned above (A versus C and B versus D). What are the probabilities for these? If we are playing to win and are choosing second, is the ever more than one sensible choice?

When the number of generalised dice is increased, the cyclical dominating probabilities can be increased, but at least one of the probabilities must always be less than 3/4 (this is proved in [ste2] and [lic]). While this bound of 3/4 can be approached to any degree of accuracy, it is necessary to also increase the number of sides. If we restrict to sets of the usual six-sided dice, then 2/3 is the best value that can be achieved, and this is already to be found in the four dice in Table 5 (Martin Gardner attributes this to Effron [gar]; a proof is given in [tenfos]).

Bounds for Paradox 2.

There are also some basic restrictions on probabilities for the second paradox. Say that (as in Paradox 2) we have three dice A, B and C where A and B both win a two-way roll against die C with probability greater than one half, but C is the most likely winner in a three-way roll with A and B. Then the probability that A (or B) wins over C in a two way roll cannot be as large as 2/3 and the probability that C wins in the three-way roll cannot be as large as 1/2. (We will not prove these facts here, although the proofs are not difficult). These bounds cannot be further reduced, since they can be approximated to arbitrary accuracy using variations of the example in Table 3. For example, let n be any whole number and construct three dice A, B, C each with $2n+1$ faces and by labelling:

- the faces of A with the numbers $1,3,5,\ldots,4n+1$;
- the faces of B with the numbers $2,4,\ldots,4n+2$;
- the first $n+1$ faces of C with 0 and the remaining n faces with $4n+3$.

Then in a three-way roll with A and B, die C wins with probability $n/(2n+1)$ (which tends toward 1/2 as n tends to infinity), while A and B win with probability approaching just 1/4. Both dice A and B win in two way rolls against A with probability $(n+1)/(2n+1)$, which is always greater than 1/2.

Challenging exercise. After some investigation of their own, a good student might be able to at least speculate the above-mentioned bounds for Paradox 2. For example, a problem for investigation might be to find how high the probability that C wins over A and B can be made while still keeping A>C and B>C.

3. ELECTIONS AND RANKING.

One of the nicest connections with dice paradoxes comes from the theory of rankings. To see the kind of paradox one can encounter here, consider a situation where three sale items A, B and C have been ranked by three customers as follows (this and numerous other interesting examples of non-transitivity can be found in [gar2]):

- Customer 1: A>B>C;
- Customer 2: B>C>A;
- Customer 3: C>A>B.

Here 2 out of 3 customers prefer A to B, 2 out of 3 prefer B to C and 2 out of 3 prefer C to A! Alternatively the items A, B and C could represent candidates in an election, and the customers can represent voters (obviously we can also allow for more than three voters!). This certainly has a similar feel to the dice paradox 1 although as we noted, we cannot get such high proportions when comparing three dice.

In fact all of the dice paradoxes can be reinterpreted as voting paradoxes. To see this, we consider the three dice of Table 1. If we compare all three dice at once, there are 216 possible combinations, each of which gives a ranking of the three dice. We can think of the dice as corresponding to candidates, and the 216 comparisons as individual ballots. Under this interpretation, the proportion of preferences for one candidate over another coincides exactly with the proportion of wins the corresponding dice make over each other. So, for the dice in Table 1, we find that 2/3 of the 216 ballots show a preference for candidate X over Y, 2/3 for Y over X and 5/9 for candidate Z over X. There is of course nothing special about the three dice of Table 1 – all of the dice examples translate in this way. For Paradox 2 (seee Table 3), we get a situation where the majority of voters found candidate A preferable to candidate C, and also a majority found candidate B preferable to candidate C, yet it was C who was ranked first the most often!

The study of rankings and elections goes back at least to the investigations of the Marquis de Condorcet in 1785. Perhaps the most interesting contribution from our perspective is the work of Donald Saari, who has shown that every conceivable paradox is possible in rankings. An example similar to the ranking paradox we translated from Paradox 2 is given in [saa]. Say that a group of 15 friends have voted for three beverages wine, beer, spirits as follows (in this example it is implicit that the relation > is transitive):

- six have ranked spirits>wine>beer;
- five have ranked wine>beer>spirits;
- four have ranked beer>wine>spirits.

On first preferences, we have spirits more preferable than wine, which is in turn more preferable than beer. Upon reaching the bottle shop, it is found that there is no beer (surprising though this may seem, it is not the paradox). While previously, spirits appeared the most preferable, we see that if beer is removed from the list of preferences, there are actually 9 people preferring wine to spirits and only 6 preferring spirits to wine!

Exercise. Get each of the members of a class to rank some objects (eg food items, some football teams), say A, B, C, D, and E (if you use too many objects it will be too much work collating the results). These should preferably be objects for which there are will be fairly even preferences, so that there will be plenty of variety in the rankings. Examine the rankings for paradoxical cycles. To do this, you will need to first add up how many people prefer object A to B, then A to C and so on, until each pairing is considered. Write A, B, C, D, E in a circle on a piece of paper. Put an arrow from one letter to the other if the corresponding object was found to be preferable. There are non-transitive preferences (a voting paradox) if there is a cycle (a path of arrows that comes back to meet itself) somewhere in this diagram.

4. THREE OR MORE PLAYERS.

We have seen one three player trick above – Paradox 2. A more natural generalisation of non-transitive dice to three player games is as follows. Is it possible to have a set of dice, so that no matter which two dice are chosen by two opponents, it is always possible to choose a dice that will be the most likely winner in a three-way roll? The answer is yes, although it seems very difficult to find examples. The following example was found in a tent while avoiding bad weather on New Years Day 2000.

We need at least five dice A, B, C, D, E (we won't prove here that at least 5 are needed, but the proof is easy – just try to do it with four dice and see what goes wrong!):

A:	1	10	15	18	18	18
B:	0	12	14	14	17	21
C:	4	6	13	13	19	23
D:	3	5	7	9	22	22
E:	2	8	11	16	20	20

Table 6

It is quite tedious to calculate the various probabilities for each of the 10 possible three-way rolls, and furthermore, the winning probabilities are not extreme: in one case the best choice has only 1/216 advantage over one of its opponents. The following table gives the appropriate winning choice and the probability for each of the 10 pairs:

Pair:	A,B	B,C	C,D	D,E	E,A	A,C	B,D	C,E	D,A	E,B
Is beaten by:	E	A	B	C	D	E	A	B	C	D
Winning prob. (.../216)	76	78	76	79	77	76	80	76	79	77

Table 7

A similar 3 player variant but with two-way rolls, is as follows (we call it the *multi-player problem*): is there a set of dice such that no matter which two dice are chosen by two opponents, there is a third die that can be chosen that will more likely than not, beat either opponent in a two way roll? Again, the answer is "yes", but examples are very elusive. It turns out that there is a simple construction to produce any desired combination of two way dominances. While this construction is surprisingly easy, it also requires the introduction of a great deal of new notation and so will not be given here (it will appear in a forthcoming article). The solution it produces for the multi-player problem involves 7 dice with up to 16 sides each. The following string of letters gives the order that the faces of the 7 required dice A, B, C, D, E, F, G should be labeled by, say, the numbers 1,2,3,4,… (note that A appears 16 times, so 16 faces are required on this dice, while the other dice require only 12 sides):

ABACBDCEDFEGFAGCAECGEBGDBFDAFADDGGCCFFBBEEAAAEABEFBCFGCDGADA
FFDDBBGGEECCAAGGFFEEDDCCBBAA.

More complicated still is the 4 player version, which requires 19 dice with up 36 sides each!

References.

[gar] M. Gardner. The paradox of the nontransitive dice and the elusive principal of indifference, *Scientific American*, **223** No. 6 (1970), 110–114.

[gar2] M. Gardner. On the paradoxical situations that arise from non-transitive relations, *Scientific American*, **231** No. 4 (1974), 120–125.

[lic] C. Li-chien. On the minimum probability of cyclic random inequalities, *Scientia Sinica* **10** (1961), 499–504.

[saa] D. Saari. A chaotic exploration of aggregation paradoxes, *SIAM Review*, **37** (1995), 37–52.

[sav] R. P. Savage. The paradox of nontransitive dice, *Amer. Math. Monthly* **101** (1994), 429–436.

[stetry] H. Steinhaus and S. Trybula. On a paradox in applied probabilities, *Bull. Acad. Polon. Sci.* **7** (1959), 67–69.

[ste] S. Trybula. On the paradox of three random variables, *Zastos. Mat.* **5** (1961), 321–330.

[ste2] S. Trybula. On the paradox of *n* random variables, *Zastos. Mat.* **8** (1965), 143–154.

[tenfos] R. L. Tenney and C. C. Foster. Non-transitive dominance, *Math. Magazine*, **49** (1976), 115–120.