Tales of Wild Dice

J. Buhler, A. Gamst, A. Hales*

Martin Gardner played a large role in popularizing nontransitive dice, starting with his December, 1970, column in which he focused on dice due to Efron. Another example, due to Moon and Moser [7], is pictured below. For us, dice are lists (with repetition allowed) of equally likely integer outcomes of a "roll," and we write the Moon/Moser dice as

$$\mathbf{R} = [2, 6, 7], \quad \mathbf{G} = [1, 5, 9], \quad \mathbf{B} = [3, 4, 8].$$

(If you wonder what's on the hidden faces, then we'd tell you that each face is identical to its antipodal face, so that we should have written, e.g., $\mathbf{R} = [2^2, 6^2, 7^2]$; however, this die is equivalent to the simpler form above.)

In addition to being nontransitive (**R** is higher than **G**, on average, **G** is higher than **B**, and yet **B** is higher than **R**), these dice have an even stranger property: the triple $\mathbf{R}^{[2]}$, $\mathbf{G}^{[2]}$, $\mathbf{B}^{[2]}$ is also a nontransitive cycle, but in the opposite direction! (The superscript on $\mathbf{R}^{[2]}$ indicates the die whose roll is the sum of two independent rolls of **R**.) To our knowledge, these curious facts were first noticed by Tom Leighton around 1990, when he was working on notes for a course at MIT (for an amusing discussion of swindles based on these dice, see section 17.3 in the book [6] that grew out of these notes). This idea also appeared in a paper by Allen Schwenk [8], whose title, *Beware Geeks bearing Grifts*, is hard to beat.

Ron Graham asked how far these curious properties could be pushed. He showed that much more exotic outcomes were possible, and his effervescent (and, OK, insistent) personality led to two joint papers: one [3] showing how to produce arbitrary tournaments, in a sense that will be made precise below, and another [2] that shows how to fix a fascinating gap that arises when one tries to argue that these examples are actually explicit. This is an overview of some of the results and techniques, and our real goal is to entice you into reading those papers! We are deeply grateful for Ron's mathematical and non-mathematical friendship, his insights, and the extraordinary opportunity to collaborate with him.

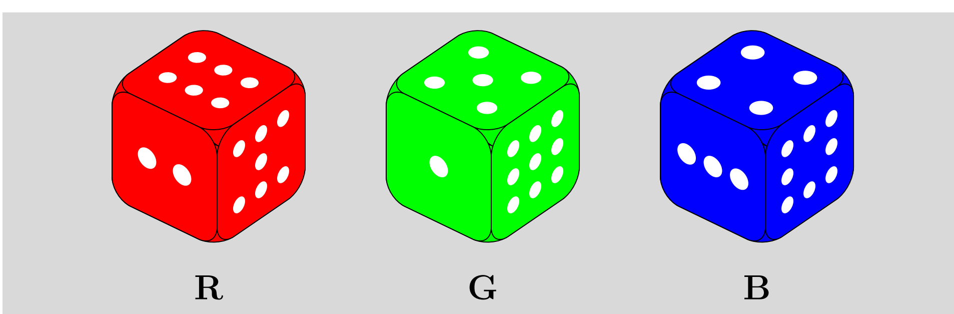

*jpb@reed.edu, acgamst@ucsd.edu, awh@math.ucla.edu

Dice

For us, a *die* (plural: dice) is a finite list of equally likely outcomes (rolls). Dice can be added or multiplied by constants; e.g.,

$$\mathbf{B}+\mathbf{R} = [5,6,9,10^3,11,14,15], \qquad 2\mathbf{G}=[2,10,18]$$
$$\mathbf{B}-\mathbf{R} = [-4,-3^2,-2,1^2,2^2,6], \qquad \mathbf{G}^{[2]}=[2,6^2,10^3,14^2,18].$$

Note that multiplying a roll by 2 is a very different from adding 2 independent rolls. Bracketed exponents indicate repeated addition, and exponents on the values indicate multiplicity (repetition).

Given a set of k dice, we are interested in the results of comparing all pairs of the above dice when each is rolled n times and the results are added. For the **RGB** triple of dice we want to know the 3 pairwise comparisons from the set $\{\mathbf{R}^{[n]}, \mathbf{G}^{[n]}, \mathbf{B}^{[n]}\}$, for all $n \geq 1$. The results are summarized in the figure below for $n \leq 100$. For each n, we label the edges of a triangle whose vertices are the dice with an arrow that points from the winner to the loser. Three different "tournament results" occur. For a few small n, they are nontransitive cycles; for a few other small n, and apparently all $n \geq 9$, the result is the same as for $n = 3$, i.e., **R** beats both **G** and **B**, and and **G** beats **B**.

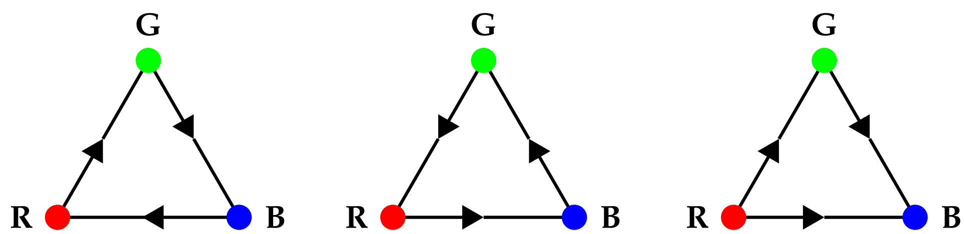

$n = 1, 4$ $\qquad$ $n = 2, 5, 8$ $\qquad$ $n = 3, 6, 7, 9, \ldots, 100$

This means, for $n = 1$, that (loosely speaking) **R** is "better" than **G**, **G** is better than **B** and yet **B** is better than **R**. This is easy to check; e.g., looking at the list for $\mathbf{B}-\mathbf{R}$ above, we see that 5 of the 9 outcomes have **B** winning, so **B** will beat **R** more often than not. For the **RGB** dice, nontransitivity for $n = 1$ is perhaps a bit surprising, but the fact that for $(\mathbf{RGB})^{[2]}$ case (i.e., $\mathbf{R}^{[2]}$, $\mathbf{G}^{[2]}$, $\mathbf{B}^{[2]}$) the tournament is a nontransitive cycle in the reverse direction is doubly surprising, and the fact that $(\mathbf{RGB})^{[3]}$ is neither of the nontransitive cycles is perhaps triply surprising. For the nth powers (as we will call them) for $n \geq 9$ the outcome is always the same tournament, mentioned above. As Schwenk's title hints, these properties offer numerous opportunities for swindles (a.k.a. grifts), at least if wagers, n, and the gradual emergence of the full situation are carefully managed.

Our primary goal in the next section is to construct sets of k dice $D = \{D_1, \ldots, D_k\}$ that exhibit vastly more deranged properties.

For the **RGB** dice, all of the pairwise comparisons seem fair because the means (averages) of the compared dice are equal. Indeed, if the means were unequal, well-known results from Statistics 101 show that, at least when larger powers are taken, dice with higher means will win, in the long run. In other words, for the sake of finding counterintuitive examples it suffices to consider only sets of dice that all have the same mean.

There is an interesting higher order grift that might arise if the players are mathematicians. Imagine that you and an opponent play this game repeatedly (perhaps for the sake of speed and magnifying the small winning margins as n increases, letting a computer roll the dice, of course using a fair random number generator). You spar over who chooses the first die, and what n will be, and after a while you both understand what's going on, e.g., understand the above figure whether or not you cop to that knowledge to your opponent. The size of the bets has been steadily increasing.

At one point, your opponent offers you the following game. She names an n, you pick whether to chose first or second, and then, rather than the grubby rolling of dice and doing arithmetic, or even simulating that using a computer, it is agreed that whoever chooses the first die will win if they can give a reasonably concise proof that they win (in the probabilistic sense of having a greater probability of rolling a higher number). If you've taken Probability 301 then you know that the computation comes down to comparing the median of the difference of nth powers of two of the **RGB** dice to the mean, which is 0. As n gets large, the median and mean are close (by the Central Limit Theorem), but their difference is governed by the "skewness" of the distribution (roughly, which way it leans away from being a symmetric normal distribution). This primarily depends on the third moment $\sum \mathbf{Pr}(X = x)x^3$ where X is the difference of the nth powers of two of the dice. Moreover, you dimly recall that the winning margin goes to 0 as a function of n something like $c/\sqrt{n}$. Making the choice, and giving a proof, is trivial for $n \leq 8$ and is easy for large n because of the theorem below, which you will have to cite. The stakes are of course tripled, and since you know the full story and have a sure thing, you accept.

Your opponent says that n will be 10^{24}. Which, in case you are counting (and use American terminology) is one septillion. This is a bit unnerving, but of course you know to choose **R**. Your proof begins by quoting the following theorem, which you cleverly stored on your phone during the break.

Theorem 1. *Let X be a die (an integer-valued random variable) with span 1 and mean 0. Then for n going to infinity,*

$$\mathbf{Pr}(X > 0) - \mathbf{Pr}(X < 0) = \frac{c}{\sqrt{n}} + o\left(\frac{1}{\sqrt{n}}\right), \quad c = \frac{-\mu_3(X)}{3\sqrt{2\pi\mu_2(X)}}.$$

You (tediously) continue your proof by explaining that the "little-oh" term $o(1/\sqrt{n})$ means that for all positive ε there is an $n_0 = n_0(\varepsilon)$ such that the term is less than $\varepsilon/\sqrt{n}$ for $n \geq n_0$. And that the jth central moment of a die X is

$$\mu_j(X) = \sum_x \mathbf{Pr}(X = x)\, x^j.$$

Finally, the *span* of the (integer-valued) X is the largest m such that the values of X are contained in a single congruence class $a + m\mathbf{Z}$ modulo m. Any element of the congruence class is called the *shift* of the die. For instance the values of **G**, namely, 1, 5, and 9, are all 1 mod 4 and contained in the congruence class $1 + 4\mathbf{Z}$ modulo 4; moreover, this is clearly the largest possible m with that property, since two of the rolls differ by 4, so the span of **G** is 4. The shift is only defined modulo m, so the shift could be said to be 5, 1, or -3. However, all of the pairwise differences of the three dice have values that differ by 1, so their span is 1 and you are relieved that the above theorem applies.

You finish by asserting that **R** is clearly going to win — the moments are easy to compute, and the chosen n is (insanely) large.

You are horrified when your opponent points out that this is not a proof because you haven't yet named an explicit function $n_0(\varepsilon)$ and proved that it works. As unlikely as it seems given the outcomes up to $n = 100$, you have not proven that the third moment term dominates the error term for $n = 10^{24}$.

You may now be in a spot of trouble (especially because the stakes have been increasing). Stay tuned for useful remarks in the last section.

In the meantime, you might take some solace from having noticed that there is an intuitive reason for the minus sign in front of the third moment term in the Theorem—if the third moment is positive, then there has to be more probability mass that is negative in order to balance the die so that its mean is 0. So the median will be negative.

Tournaments

Let $D = \{D_1, \ldots, D_k\}$ be a set of k dice. The result of all $k(k-1)/2$ pairwise comparisons between the k dice can be recorded as an antisymmetric matrix whose the entry in row i and column j is 1 if D_i beats D_j in the long run, 0 if the contest is fair (i.e., a probabilisitic tie), or -1 if D_j beats D_i in the long run. This matrix is a *tournament* if there are no ties, and a *partial tournament* if there are ties. For instance, the three tournaments in the figure above could be represented (less elegantly) as the following matrices:

$$\begin{bmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{bmatrix}, \qquad \begin{bmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ 1 & 1 & 0 \end{bmatrix}, \qquad \begin{bmatrix} 0 & 1 & 1 \\ -1 & 0 & 1 \\ -1 & -1 & 0 \end{bmatrix}.$$

If X is a die then, in order to cope with ties, and reduce the outcome to ± 1 if there is a winner (and 0 for a tie), it is convenient to define the "positivity" of a die X as $w(X) = \operatorname{sgn}(\mathbf{Pr}(X > 0) - \mathbf{Pr}(X < 0))$, where $\operatorname{sgn}(x)$, for a real number x, is 1, 0, or -1 according to whether x is positive, zero, or negative.

Fundamentally, $w(X)$ measures whether the median of X is above or below 0. For a set $D = \{D_i\}$ of k dice, define the tournament on their nth powers setting the element in the ith row and jth column to be

$$T_n(D)_{ij} = w(D_i^{[n]} - D_j^{[n]}).$$

(This is only a partial tournament if some of the off-diagonal entries are 0.)

Suppose that there are $k = 3$ dice. Then there are 3 pairwise contests of interest, and $8 = 2^3$ possible ways for those contests to be decided (ignoring ties), and therefore 8 possible tournaments. Only 3 of the 8 possible tournaments showed up in the contests between the nth powers of **R**, **G**, and **B**. Wouldn't it be cool, or at least more nontransitive, if there was another set of dice A, B, C such that all 8 of the possible tournaments occurred as a "dominance tournament" on $A^{[n]}$, $B^{[n]}$, and $C^{[n]}$, for some value of n?

The next theorem says that such a set exists. Worse, it says that such a set that realizes all possible k-person tournaments exists for *any* $k > 3$. Worse yet, there will be no "limiting tournament" as in the case of the **RGB** dice (i.e., a tournament that is the result for all sufficiently n), because the dice are so exquisitely balanced that each tournament not only occurs, but occurs infinitely often.

Theorem 2. *For every $k > 2$ there is a set of k dice $\{D_i\}$ such that for any $k \times k$ tournament matrix T on k players there are infinitely many n such that $T = T_n(D)$.*

A proof can be found in [3]. For the sake of giving a sense of what is going on, we consider the example of $k = 5$ dice.

First, there are $2^{10} = 1024$ possible tournaments (!). The 10 edges of the complete graph on 5 dice (illustrated below) can be oriented in 1024 ways, and each of those tournaments on $D = \{D_i\}$ can be realized as the tournament graph on the nth powers, for infinitely many n.

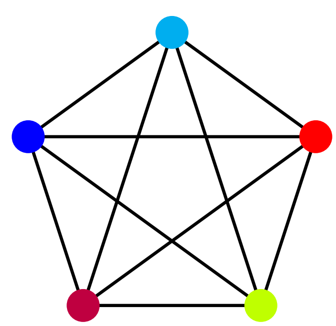

This sounds like a tall order. Curiously, it turns out that our only recourse is to go back to Theorem 1 and look at the omitted case: dice with spans larger than 1. Also it turns out that the third moment term is just an annoyance, and it simplifies things to just require that the third moment is always 0.

Suppose that D_i is a collection of k dice that have mean and third moment equal to 0, with shifts a_i, spans m_i respectively. It isn't hard to show that mean and third moment of $D_i^{[n]}$ are 0, and its shift and span are na_i and and m_i. And to show that the span of a sum of two die is the gcd of their spans, and the shift of the sum is the sum of the shifts.

If x is a real number, its fractional part, written $\{x\}$, is x minus the largest integer less than or equal to x. It is convenient, for the sake of stating the improved version of Theorem 1, to define

$$\langle x \rangle = \begin{cases} 1 & \text{if } 0 < \{x\} < 1/2 \\ -1 & \text{if } 1/2 < \{x\} < 1 \\ 0 & \text{otherwise.} \end{cases}$$

Note that the "otherwise" case occurs precisely when x is an integer or half-integer. Carefully generalizing Theorem 1 for spans larger than 1 leads to a very concrete description of the (matrix of) tournaments $T_n(D)$ for sets of dice $\{D_i\}$ that that have $\mu_1 = \mu_3 = 0$, at least for large enough n.

Theorem 3. *Suppose that D_i are k dice with mean and third moment equal to 0, and with shifts a_i and spans m_i. Define*

$$d_{ij}^{[n]} = \frac{n(a_i - a_j)}{\gcd(m_i, m_j)}.$$

Then the ij entry of the tournament matrix $T_n(D)$ is

$$T_n(D)_{ij} = \operatorname{sgn}\left(\frac{c\,\langle d_{ij}^{[n]}\rangle}{\sqrt{n}} + o\left(\frac{1}{\sqrt{n}}\right)\right),\ c = 1/(3\sqrt{2\pi}).$$

Note that if n is large enough the argument of sgn has the same sign as its $\langle\ \rangle$ term, unless that term is 0.

This shows that we can compute tournaments matrices for large n once we know the shifts and signs, but still begs the question of how we can specify the D_i in any straightforward and general way. The following Lemma gives a very practical answer!

Lemma 1. *Given a positive integer m and any integer a there is a die X with shift a, span m, mean 0, and third moment 0.*

If you guess to look at what you can do if X has only 3 values (with repetitions allowed), the proof is a rather unenlightening exercise in high school algebra, though it can be made more appealing by using some simple linear algebra on 2×3 matrices. We also expect that the only chance for more elegant dice (fewer repeats) is to use more values.

Via this lemma, the following table specifies five dice with mean and third moment 0, with the indicated shifts and spans.

a_i	2^4	2^3	2^2	2^1	2^0
m_i	2^5	2^9	2^{12}	2^{14}	2^{15}

Our goal is to prove Theorem 2 which, loosely speaking, asserts the existence of dice that are so wild that their powers have all possible tournaments. To do this we have to show that every tournament matrix T on 5 contestants is of the form $T_n(D)$, for the set D just described, and some positive integer n.

Once all of the results are unraveled, this comes down to showing that the 10-tuple ν of all $(a_i - a_j)/\gcd(m_i, m_j)$ from the above table satisfies the conclusion of the following Lemma, which is something of a "mod-2 equidistribution" law. In our application of the lemma, K is 10, since ν has 10 components. For the sake of stating the Lemma, we let H_K be the set of all elements x in $[0,1]^K$ which have no coordinate x_i equal to 0, 1/2 or 1. This is a disjoint union of 2^K open cubes C of side length 1/2. The Lemma asserts that for every C the integer multiples $\{n\nu\}$ of ν in 10-space intersect some translate $u + C$ of the half-cube C by an *integer* 10-tuple u.

Lemma 2. *With the above notation, there is a vector* ν *such that for every open cube* C *in* H_K *there are infinitely many integers* n *such that* $\langle n\nu\rangle$ *lies in* C*, where* $\langle u\rangle$ *denotes the result of applying the bracket operator to every component of* u*.*

The reader might enjoy either or both of the (nontrivial!) exercises of proving the Lemma, or of showing that the 10-dimensional ν arising from the 5 dice above satisfies the Lemma.

This finishes the outline of the proof of Theorem 2.

Comments.

A number of remarks are in order.

#1: Linguistically, "intransitive" (meaning definitely not transitive) is probably a better term than "nontransitive" (meaning not necessarily transitive), but we stick with the latter as it has become thoroughly ingrained in the mathematical literature on the topic.

#2: Although there may have been related ideas that arose earlier, the specific idea of nontransitive dice seems to have first appeared in work of Steinhaus and Trybula [9] in the late 1950s. A number of interesting references to results on nontransitive dice can be found in the bibliographies of the more recent papers in the bibliography below. Also, there is a lot of information online, including, of course, the Wikipedia article on Intransitive Dice as well as web pages by James Grime, Oskar Deventer, and no doubt others.

#3: In aiming at tournaments, we skipped over the interesting case of n-cycles. The method of construction of the 5×5 matrix (extracted verbatim from [1])

7	8	9	10	25
4	5	6	23	24
2	3	20	21	22
1	16	17	18	19
11	12	13	14	15

is clear. A bit of careful thinking about the dice formed from its rows should convince you that it is easy to construct nontransitive cycles of any length. Indeed, rows of this matrix form a nontransitive cycle of length 5: each row is better than the one below it, and the bottom is better than the top! One question that has been around since the 1960s is: what is the highest possible success probability around the cycle? I.e., given n, what is the largest p such that there is a nontransitive n-cycle with $\mathbf{Pr}(D_{i+1} > D_i) \geq p$ for all i? A definitive treatment appeared in the Monthly recently [5], and it confirms that the largest p is

$$p = 1 - \frac{1}{4\cos^2(\pi/(n+2))} = \frac{3}{4} - O(1/n^2).$$

(Strictly speaking, for dice in our sense, this is the supremum of all such p, but this value can actually be attained if irrational probabilities are allowed.) In other words, the best winning probability around a cycle approaches 3/4 from below, as n goes to infinity.

#4: The dice in the proofs might not be aesthetically pleasing or "practical" since our goal was only to push the envelope on what was known to be possible. There are lots of opportunities to do better. One obvious open question is to find a reasonable set of 3 dice that realize all 8 tournaments. One measure of the size of a set of dice is the least common multiple L of the spans of its dice (though we think that this is not exactly the same as "practical"). For $n = 3$, the smallest possible value of L is 10 and, perhaps slightly surprisingly, a set of 3 dice exist with $L = 10$. For $n = 4$, easy arguments show that the optimal L satisfies $64 < L \leq 512$ (the upper bound coming from the construction above). In fact $L = 68$ is possible (and almost certainly best possible). This seems to suggest that wild dice exist more or less as soon as there is any room for them to exist.

#5: A full proof of Theorem 3 can be found in [3] where the probability in question is initially expressed as a contour integral.

#6: Explicit estimates error estimates for Theorem 3, i.e., with overt functions $n_0(\varepsilon)$ that you might need (in the circumstances mentioned above), can be found in [2]. The quest for explicit error estimates arises in number theory, probability, and elsewhere. In the case of [2] (first order lattice Edgeworth error estimates, in the lingo) common wisdom was probably that precise estimates were hard to find, likely to be uninteresting because they would be unrealistically large, and not terribly useful because computers can compute values for large n, leaving no doubt as to what actually happens in, say, dice tournaments. Although a variety of ideas were needed, the estimates in [2] turned out to be unexpectedly good at giving realistic estimates even in fairly pathological cases. Be that as it may, computations do give a good sense of what is going on, and you sort of have to be a mathematician (at least at heart) to enjoy this quest for explicit estimates.

References

[1] I. I. Bogdanov, Нетранзитивные рулетки (Nontransitive Roulette), Mat. Prosvesch. 3 (2010), 240–255 (Russian).

[2] Joe Buhler, Anthony Gamst, Ron Graham, and Alfred Hales, *Explicit error bounds for lattice Edgeworth expansions*, Cambridge: Cambridge University Press, 2018.

[3] Joe Buhler, Ron Graham, and Alfred Hales, *Maximally nontransitive dice*, Am. Math. Mon. **125** (2018), no. 5, 387–399.

[4] James Grime, *The bizarre world of nontransitive dice: games for two or more players*, Coll. Math. J. **48** (2017), no. 1, 2–9.

[5] Andrzej Komisarski, *Nontransitive random variables and nontransitive dice*, Am. Math. Mon. **128**, no. 5, 423–434.

[6] Eric Lehman, Tom F. Leighton, and Albert R. Meyer, *Mathematics for Computer Science*, MIT OpenCourseWare.

[7] J. Moon and L. Moser, *Generating oriented graphs by means of team comparisons*, Pac. J. Math. **48** (1967), no. 21, 531–535.

[8] Allen Schwenk, *Beware of geeks bearing grifts*, Math. Horiz. **48** (2000), no. 7, 10–13.

[9] Hugo Steinhaus and S. Trybula, *On a paradox in applied probabilities*, Bull. Acad. Pol. Sci., Sér. Sci. Math. Astron. Phys. **7** (1959), 67–69.

[10] Pavle Vuksanovic and A. J. Hildebrand, *On cyclic and nontransitive probabilities*, Involve **14** (2021), no. 2, 327–348.