Snakes and Ladders and Intransitivity, or What Mathematicians Do in Their Time Off

Gregory B. Sorkin

For those fortunate enough to be unfamiliar with it, Snakes and Ladders is a children's board game of no skill and no mathematical interest. The game, also known as Chutes and Ladders, uses a board that is in essence a strip of squares, from 1 to 100. A sample board is shown in Figure 1. Each player starts in square 0 (just off the board), and players take turns. On each turn, a player rolls a usual six-sided die and advances by a number of squares equal to the roll of the die. The twist is that some squares ("ladders"), when you land on them, advance you to a later square, and some ("snakes") put you back to an earlier square. The goal is to be first to reach square 100.

A key point is that the players are independent. They could as well play separately, each counting how many moves they took, and compare notes at the end: the one with the smaller number of moves wins. This, along with the fact that no skill is involved (there are no choices to make), is why I disparaged the game as being of no (mathematical) interest; but actually, there are some mathematically interesting aspects.

In particular, as the game is played, both players tend to advance, but there are frequent setbacks. When you are at square i and your opponent is at j, it is natural to wonder who has the advantage. We'll say that square i is "better than" j (and write $i \succ j$) if a player at i is more likely to win than one at j: if the two players bet even odds on the outcome over multiple games, then in the long run, the player at i would come out ahead. Are later squares always better? Probably not, since it's probably better to have the possibility of a long ladder just ahead of you than to be just past it. Does it even make sense to ask what square is better, or does it depend on your opponent's square? Specifically, might the game be "intransitive": is it possible that square i is better than j, and j better than k, but k is better than i, so that $i \succ j \succ k \succ i$?

We will answer these questions. We are not aware of the intransitivity question having been asked before for Snakes and Ladders. Along the way, we'll visit Markov chains, simulation, a paradox of size-biased sampling of geometric random variables, and intransitive dice.

We'll begin with a soupçon of history and some pesky details. Then we will find the expected number of rolls to reach square 100, using simulation (and exploring the size-biased sampling paradox), then Markov chains. In the next section, we find the winning edge between two cells and look for intransitive triples, first using Markov chains and then simulations. We then discuss intransitivity in dice, and conclude by connecting such dice directly to Snakes and Ladders.

History, Details, and the Markov Chain

Snakes and Ladders is widely agreed to derive from an Indian game, called *gyān chaupar* in Hindi, making its way to Victorian England as a side effect of British colonialism. The "infobox" of the Snakes and Ladders Wikipedia page [11] asserts without attribution that the game has been played since the second century CE, while a number of sources credit its origin to the thirteenth century poet and philosopher Sant Dnyaneshwar, but again without citing any basis. The scholarly article [7] cites concrete evidence for the game's having been played in the eighteenth century and says that it "is doubtless much older," but that since board materials are ephemeral, "[u]ntil earlier evidence is available, the origins ... of the game must remain obscure." (See also [8, Conclusions].) The boards vary in size, as do the numbers of snakes and ladders and their positions, depiction, and labeling, but the game play remains the same.

Wherever played, the game was meant to be morally educational. Virtuous ladders, and vices represented by snakes, would bring you toward or away from some version of heaven. Their depiction and labeling would suit the morality of the time and place, a Victorian version, for example, having a ladder of Penitence leading to a square of Grace.

Whether or not morally instructive, the game is a fine illustration of randomness.

An important detail for us is what it means to win. One definition is that if you are the first to finish you win, but that would give an unfair advantage to the first player. (A Markov chain analysis of this version of Snakes and Ladders is given in [1].) In our house we play fair: the game goes in rounds (in each round, player 1, then player 2), and a player wins if they finish in a round and the other player does not. So, if player 1 finishes, player 2 has one last turn: if they also finish, the game is a draw. Either way, we can consider the two players separately and simply count how many rounds it takes for each to finish: in our fair version, the player finishing in an earlier round wins, and if both finish in the same round, it is a draw.

https://doi.org/10.1007/s00283-022-10250-6

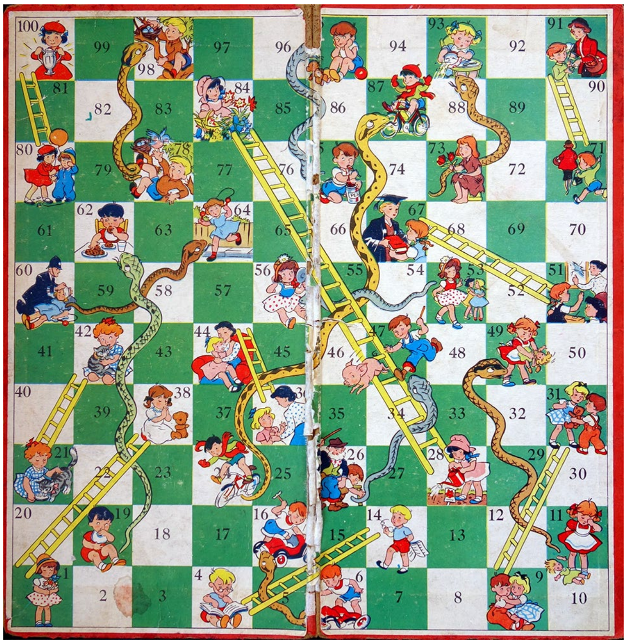

Figure 1. A sample Snakes and Ladders board. This sample, owned by the author, bears no identifying marks or copyright. Squares are numbered in a serpentine pattern, moving to the right through the bottom row, left in the row above it, and so on.

With this fair version, $i \succ j$ means that if one player is in square i and the other in j in the same round, then i wins more often than j (with draws not counting either way): in the long run, i has a winning advantage.

Our main interest is in "intransitivity," i.e., in whether there is a "triangle" (or longer cycle) of squares such that $i \succ j \succ k \succ \cdots \succ i$. The notion of intransitivity is natural in the fair version, where we consider squares i, j, and k all in the same round. (It is less natural in the unfair version, where we must take into account whose move it is. There, perhaps we'd look for a 4-cycle where square i having the move has an advantage over square j without it, which in turn has an advantage over square k with, that over square l without, and that over square i with.)

Let's return to the game's setup and clarify some details. First, an example. In our board, there is a ladder from square 4 to 14. This means that square 4 can never be occupied: if, for example, a player is on square 3 and rolls a 1, they move to square 14.

There are two minor details. One is how you finish. If you overshoot 100, does that count as a finish, do you stay in the same square to try again on the next turn, or do you "reflect" back from 100? We arbitrarily choose the "reflecting" version: for example, from 99, a roll of 3 would bring you 1 step forward to 100, then 2 steps back to 98, where on our board there is a snake, so you'd wind up at 78. A second detail is that sometimes the game is played with the rule that if a player rolls a 6, they are allowed an extra roll in the same turn; it makes no essential difference, and we eschew this complication.

38 1	2 2	3 3	14 4	5 5	6 6	7 7	8 8	31 9	10 10
11 11	12 12	13 13	14 14	15 15	6 16	17 17	18 18	19 19	20 20
42 21	22 22	23 23	24 24	25 25	26 26	27 27	84 28	29 29	30 30
31 31	32 32	33 33	34 34	35 35	44 36	37 37	38 38	39 39	40 40
41 41	42 42	43 43	44 44	45 45	46 46	26 47	48 48	11 49	50 50
67 51	52 52	53 53	54 54	55 55	53 56	57 57	58 58	59 59	60 60
61 61	19 62	63 63	60 64	65 65	66 66	67 67	68 68	69 69	70 70
91 71	72 72	73 73	74 74	75 75	76 76	77 77	78 78	79 79	100 80
81 81	82 82	83 83	84 84	85 85	86 86	24 87	88 88	89 89	90 90
91 91	92 92	73 93	94 94	75 95	96 96	97 97	78 98	99 99	100 100
99 101	78 102	97 103	96 104	75 105					

Figure 2. The Snakes and Ladders board of Figure 1, laid out from top to bottom and left to right. Each square has an index (lower right corner) from 1 to 105, and a value (center) corresponding to a state of the Markov chain. Your previous value, plus your die roll (from 1 to 6), gives an index, whose value is the index of your new square. For example, from square 5, a roll of 4 brings you to index 9 and thus, via a ladder, to state 31 (the value in the square of index 9). The starting state is 0 (not shown). The finishing square of 100 is followed by the squares to which you are redirected (your next state) if you overshoot 100. Squares that are the starting point of a snake or ladder are shown in gray: they are not Markov chain states because it is impossible to wind up in such a square. The occupiable squares, shown in white, have their value equal to their index.

To recapitulate, in essence, the game consists of a set of squares, or "states." From each state, there are six possible next states, the one obtained depending on the roll of the die. Such a game defines a Markov chain. (See Figure 2.) Our board has 82 states, including 0 and 100: squares that are the starting point of a snake or ladder do not appear as states, since it is impossible to wind up in such a square. The winner is the first player to reach a specified state (100 in our case).

Expected Time to Finish

Let's return now to our questions. It is natural to wonder, first, to what degree being farther along the board is actually helpful, and by how much. For each state (each board square that is not the start of a snake or ladder), what is the expected number of moves from that state to the end state (100), that is, the number of moves it would take, on average, in the limit over an infinite number of games?

Figure 3 shows that indeed it is generally better to be farther advanced along the board—later squares have for

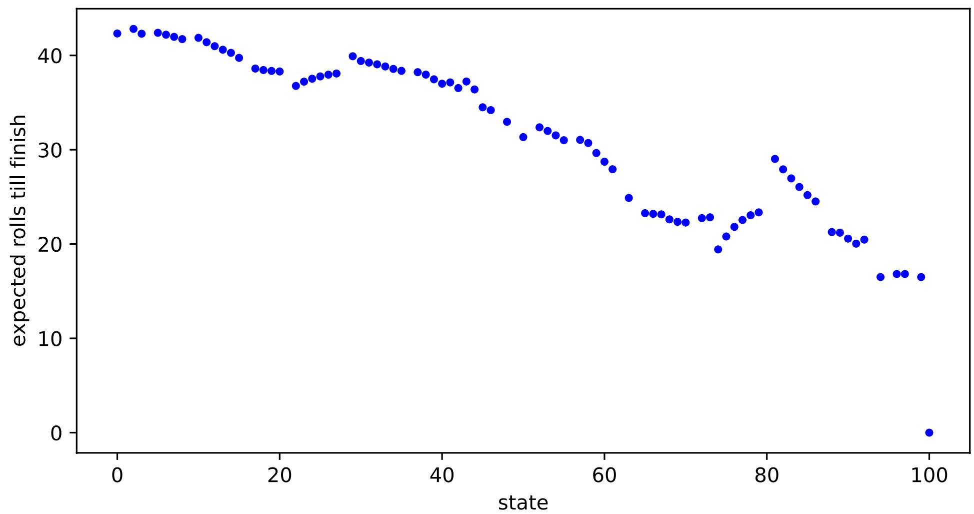

Figure 3. Expected number of moves to finish versus state (square number). The game starts in state 0. The ending state is 100, so from it, the expected time to finish is 0. There are gaps at, for example, positions 1 and 4, because these board squares are not states: in this case, they are ladders to 38 and 14. Being farther along the board is generally helpful, but not consistently so.

the most part a lower expected time (i.e., number of turns) to completion—but there are many exceptions. For instance, the situation is successively worse from squares 22 to 27, because square 28 is a ladder to 84, and being a bit earlier maximizes the chance of landing at that ladder.

The results shown in Figure 3 are drawn from more detailed results (not presented in this paper) giving for each state i, the probability that starting in state i, the game finishes within k moves; in principle this should be done for all k from 0 through ∞, but in practice, the probability that the game has not ended after 1,000 moves is less than 10^{-14} (no matter what square you begin on), so we limited calculation to this number. This information is equivalent to knowing, for each k, the probability that the game ends precisely on the kth roll: the differences in successive "by time k" probabilities are the "at time k" probabilities, and the cumulative sums of the "at k" probabilities are the "by k" probabilities. There are two methods of going about finding this information: simulation of the game and calculation from the Markov transition matrix.

Simulation

Since these questions were just a flicker of curiosity, not a serious research agenda, it was natural to address them by a quick and easy simulation. We can program a computer to start a player in a specified square i, perform a simulated die roll, advance the player accordingly, and continue likewise until the player finishes. Repeating this for square i gives a sample of the game lengths that with a large number of repetitions should be an accurate sample of the true distribution of the duration.

Observing that the game is memoryless, the simulation can be done much more efficiently. Memorylessness means that if we are in square i, the remaining time until the end of the game is independent of what came earlier (though of course dependent on the random future die rolls). Thus, instead of getting just a single duration out of one game simulation, we can get many. Suppose a simulated game visits squares 0, 2, 6, 10, 6, 9, ... , 100. (From square 10, a roll of 6 brings you to square 6 via a snake at 16.) If there were 50 rolls in all, this play gives 51 simulated values: from 0 the game ended in 50 steps, from 2 in 49 steps, and so on, until from 100 it ended in 0 steps. Note that from the first 6, the game ends in 48 steps, and from the second 6 in 46 steps: the simulation can give several remaining-time samples for a single i. All in all, a play of n steps gives n samples (ignoring the final 100), much better than playing a whole game to get just one sample.

Size-Biased Sampling

The latter method, if you look at it from a certain angle, appears wrong. If we start a simulation from i, clearly the duration from that visit of i is what we want. The memorylessness tells us that for later visits to i, the time remaining until the finish is also a valid sample. But paradoxically, those later visits to i obviously have shorter game durations than the starting one. Though this approach seems wrong, in fact, it is right; the mystery lies in size-biased sampling. (The "bus waiting time paradox" is a beautiful example. If buses come on average once an hour but randomly—technically, following a Poisson process of rate 1—the expected time from one bus to the next is one hour. But from the moment you arrive at the bus stop, the expected time to the next bus is an hour, and by symmetry, the expected time since the previous bus is also an hour, giving an expected time of two hours between these two buses. This appears paradoxical.)

Here, intuitively, while it is true that looking at later visits to i would lead to smaller estimates of the game duration (certainly compared to the first visit to i in the same game simulation), countering this is that long games, with more visits to i, are overrepresented in the sampling. It's not obvious that these two effects exactly balance one another, but—trusting to the memorylessness perspective—they must.

To check, we can calculate. From state i, let p be the probability that i is visited again before the end of the game. In that event, let D be the distribution of time until the next visit to i. And let D' be the distribution of time from the final visit of i until the game's end. For a given visit to i, let K be the number of visits to i during the game (including this visit, but no earlier ones, if this was not the first). Conditional upon $K = k$, the length of the game is $\sum_{s=1}^{k} A_s + B$, where $A_s \sim D$ are independent random variables for the revisit durations (each A_s having distribution D), and $B \sim D'$ is the time to get from the final visit of i to the finish. So K tells us everything: if the two methods of simulation result in the same distribution of K, then they give (in the long run) the same sampling of game durations. This, then, is just a question of two ways of sampling the geometric random variable K.

For the first method of simulation, K is just geometrically distributed with parameter p:

$$\mathbb{P}(K = k) = p^{k-1}(1-p)\,.$$

For the second method, of all the visits to i sampled in all the games, we wish to know what fraction of these had exactly k more visits before the game's end (including this visit but no earlier ones). For $k \geq 1$, this is

$$\mathbb{P}(K = k) = \frac{\sum_{t \geq k} p^{t-1}(1-p)}{\sum_{t \geq 1} t \cdot p^{t-1}(1-p)}\,;$$

a game with t visits to i occurs with probability $p^{t-1}(1-p)$, gives one kth to last visit to i iff $t \geq k$, and gives t visits to i in all. It is not hard to check that this expression simplifies to $p^{k-1}(1-p)$. That is, the fraction of i-visits that are kth to last ones in the second simulation approach is the same as the first approach's probability that there are k visits to i, and the two approaches do (as they must) lead to the same result.

The Markov Chain

The Snakes and Ladders Markov chain, like any other, is completely described by its transition matrix A. For states i and j, A_{ij} is the probability of moving from state i to state j in one step. Here, for example, $A_{17,19} = 1/6$: only a die roll of 2 brings us from 17 to 19. To get from i to j in exactly two steps means moving from i to some k in one step and k to j in the next, which happens with probability $\sum_k A_{ik}A_{jk} = (A^2)_{ij}$. Repeating this gives a fundamental property of Markov processes, namely that the probability of getting from i to j in exactly s steps is $(A^s)_{ij}$.

The probability, starting from i, of reaching the final state 100 in s steps is given by $(A^s)_{i,100}$. Specifically, the finishing state is "absorbing": from state 100 there is probability 1 of returning to 100 ($A_{100,100} = 1$) and probability 0 of moving to any other state. In this case, $(A^s)_{i,100}$ represents the probability of being in the finish state at time s, perhaps having reached it earlier. Writing $f_i(s)$ for the probability that the game duration from i is exactly s, we have $f_i(s) = (A^s)_{i,100} - (A^{s-1})_{i,100}$.

Calculating $f_i(s)$ for s from 0 to say 1000 gives, for each i, the distribution of game lengths (the only error being the $< 10^{-14}$ fraction of games that are longer than 1000 rolls). The expected duration of the game from i is simply

$$\sum_{s=0}^{\infty} s \cdot f_i(s)\,.$$

This leads to the results shown in Figure 3.

Pair Competitions and Intransitivity

What about the probability that a player in state i finishes in fewer rounds than an opponent in state j? For a state i, define $g_i(s) = \sum_{t \leq s} f_i(s)$; this is the probability that the game has finished within time s, starting from i. For i to beat j means that i finishes in some round s by which j has not yet finished, so i beats j with probability

$$Q_{ij} := \sum_{s=0}^{\infty} f_i(s)(1 - g_j(s))\,.$$

Truncating this to a finite sum gives our estimate of the probability Q_{ij} that i beats j. We compute this for all pairs i, j. Specifically, for each s we compute the array of all $f_i(s)(1 - g_j(s))$ (an "outer product" of the vector of all $f_i(s)$ with that of all $g_j(s)$). This calculation is perhaps not as elegant as it could be, but since the outer product array can be computed faster than a single matrix product $A \cdot A^s$ (they are of the same dimension, square over the number of states), it is efficient enough.

Define the "excess" X_{ij} of i over j by $X_{ij} = Q_{ij} - Q_{ji}$, the win probability of i over j versus that of j over i. If on each game the winner received £1, with no money exchanged for a draw, X_{ij} would be i's average winnings playing against j. We won't need it, but the probability of a draw is just $1 - Q_{ij} - Q_{ji}$.

Our original question translates to whether there are states i, j, k such that X_{ij}, X_{jk}, and X_{ki} are all positive. Specifically, let's look for such states where $X_{ij} \geq c$, $X_{jk} \geq c$, and $X_{ki} \geq c$, and c is as large as possible. We found this using a trick something like a matrix product, but that way is no faster than trying all triples of states, so let us not explain it but just assert that we found the best triple.

The result is that states 69, 79, and 73 form such a triangle, each with a winning advantage at least $\frac{1}{2}\%$ over the next in the cycle. Specifically, state 69 has a winning advantage over state 79 of $X_{69,79} \approx 0.0077$, with larger winning advantages of $X_{79,73} \approx 0.0112$ and $X_{73,69} \approx 0.0171$. The respective win probabilities are $Q_{69,79} \approx 0.4970$, $Q_{79,73} \approx 0.4990$ and $Q_{73,69} \approx 0.4930$; they are less than 1/2 because there are also draws. (Each corresponding loss probability is given by $L = Q - X$, and the draw probability by $1 - Q - L$. Draws are relatively rare, under 4% in these three cases, which is unsurprising given that the number of moves to finish ranges from about 20 to 40, and a draw requires a coincidence that the two players

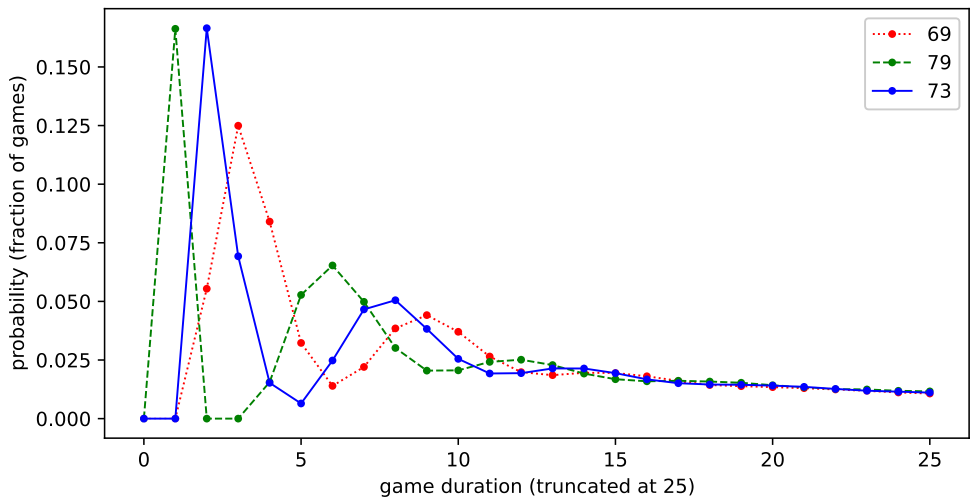

Figure 4. Fraction of simulated games of each duration, starting from squares 69, 79, and 73. For visual clarity, only durations up to 25 are shown; the maximum observed duration over 1 million games per starting square was under 500.

take exactly the same number of moves.) The time-to-finish distribution for each of these three states (e.g., f_{69} in the notation of the previous subsection) is shown in Figure 4.

Check by Simulation …

These results were checked against simulations. Starting from 0, we simulated 100,000 games, with a total of about 4.4 million die rolls; the maximum game length observed was under 500 rolls. Following the "efficient" simulation approach above, each of the three states in question was visited at least 25,000 times. Compared with the calculated winning advantages above, namely 0.0077, 0.0112, and 0.0171, the simulated ones were about 0.0090, 0.0096, and 0.0172, not very precise agreement. A second simulation gave similar results: 0.0081, 0.0127, 0.0146. Each simulation's time is dominated by the game play, taking under three minutes in an inefficient implementation. Simulations with 10,000 games are even less accurate, often showing negative values rather than the positive ones desired. A simulation with one million (1M) games takes five minutes (in a quickly improved implementation) and gives results wonderfully close to the calculated ones: 0.0072, 0.0115, and 0.0171 (each off by at most 0.0005 from the Markov chain calculation).

This simulation method has the advantage that it generates data for all squares. However, it has some implementation complications and thus possible errors.

… and Check It Twice

To be sure, I also quickly tried the very simplest "dumb" simulation: simulating just the time of a single game started from a given state. I simulated 1M games from each of these three states and computed the fraction of them in which player 1 beat player 2. This gave an estimate of 0.0067, 0.0100, 0.0177. This is not terribly accurate, but about what you should expect statistically: each player wins about half the time (draws are rare), which would mean a number of wins distributed as $B(1\mathrm{M}, 1/2)$; this has standard deviation $\sqrt{1/2 \cdot (1-1/2) \cdot 1\mathrm{M}} = 500$, leading to typical errors of $500/1\mathrm{M} = 0.0005$ in each probability estimate.

Other than simplicity, an advantage of this method is that we are assured that each square is represented by 1M samples. In the method in which 1M games were played and each contributed as many samples as its length, altogether some 44M samples were generated, but that works out to only about 500,000 per state, and square 69 was underrepresented, with only about 261,000 samples.

On the other hand, the "dumb" method's comparison of the starting squares "game by game" is inefficient, since we could just as well compare square 69's game #1 with square 79's game #17, in principle, 1M times 1M comparisons for each pair of starting squares. This can be done efficiently, since the game length was always under 500, so for each starting square, a histogram of the 1M duration samples is quite compact. This gave a better estimate, 0.0074, 0.0111, 0.0176, using exactly the same simulation data. The cruder method does have one advantage, though: the 10^6 "games between" player 1 and player 2 are all independent, so it is straightforward to estimate the variance in the winning-advantage estimate. The cleverer method here makes 10^{12} comparisons from these 10^6 games, but they are not independent, and the variance of its estimate, while smaller, is harder to estimate.

Intransitive Dice

Despite having been open to the theoretical possibility of there being intransitivity in Snakes and Ladders, I was struck by it; I did not know anything else like it. I expected that something related must be known, and did an internet search, but not knowing the right keywords, I did not quickly find anything. I reached out to a couple of colleagues and got an almost immediate response: “intransitive dice.”

Intransitive Random Variables and Dice

The roots of this sort of intransitivity go back to a “paradox of three random variables” [9], namely that it is possible for three independent random variables X, Y, and Z to have

$$\mathbb{P}(X > Y) > \frac{1}{2}, \quad \mathbb{P}(Y > Z) > \frac{1}{2}, \quad \mathbb{P}(Z > X) > \frac{1}{2}. \tag{1}$$

The paper [9] even makes reference to practical applications in which the random variables represent the breaking strengths of iron bars.

The concept was popularized in a “Mathematical Games” column by Martin Gardner [4], based on intransitive dice invented sometime earlier by Bradley Efron. Here, the random variables are the rolls of dice, and the relation $X \succ Y$ that X beats Y may be defined, as in (1), that $\mathbb{P}(X > Y) > 1/2$ or (analogous to our Snakes and Ladders interpretation) that $\mathbb{P}(X > Y) > \mathbb{P}(Y > X)$. There is, as usual, a nice Wikipedia article [10] on the topic of intransitive dice.

Examples of Intransitive Dice

One example in the Wikipedia article features three six-sided dice, die A labeled $2, 2, 6, 6, 7, 7$, die B labeled $1, 1, 5, 5, 9, 9$, and die C labeled $3, 3, 4, 4, 8, 8$. Setting aside the conventionality and physical practicality of six sides, we can as well think of these as three-sided dice, die A labeled $2, 6, 7$, die B labeled $1, 5, 9$, and die C labeled $3, 4, 8$. All face values are distinct, so there are no ties. Die A beats B 5/9 of the time, for an advantage of 1/9; B beats C with the same advantage; and C beats A with the same advantage.

A smaller example, still with all face values distinct (no ties), uses dice with three, three, and two faces: $A = 2, 3, 8$, $B = 1, 6, 7$, and $C = 4, 5$. Then A beats B w.p. (with probability) 5/9 (if A is 8 or B is 1); B beats C w.p. 2/3 (if B is 6 or 7); and C beats A w.p. 2/3 (if A is 2 or 3).

If we allow ties and insist only on a winning advantage (not necessarily winning w.p. greater than 1/2), there is an example, again with fair dice with 2, 3, and 3 faces, but using only 7 distinct values: $A = 2, 3, 7$, $B = 1, 4, 6$, $C = 3, 4$. Here A beats B w.p. 5/9 (when A is 7 or B is 1) and never draws, for a winning advantage of 1/9; B beats C only w.p. 1/2 (when B is 6, or B is 5 and C is 3), but draws w.p. 1/6 and loses w.p. 1/3, for a winning advantage of 1/6; C beats A symmetrically to how B beats C.

If we allow “unfair” dice, with arbitrary probabilities, a particularly understandable example has dice with 1, 2, and 2 sides. We can take $B = 2$ (deterministically), $A = 1, 3$ with probabilities 1/3, 2/3 respectively, and $C = 1, 4$ with probabilities 5/9, 4/9. Clearly, A beats B w.p. 2/3, and B beats C w.p. 5/9. Finally, C beats A w.p. 4/9 (whenever $C = 4$), while A beats C w.p. $2/3 \times 5/9 = 10/27$ (only when $A = 3$ and $C = 1$), giving this case a winning advantage of 2/27.

There can be no smaller example: it would have to have dice with 1, 1, and 2 sides. Without loss of generality, a one-sided die A must have a larger value than the one-sided die B, in which case, since B wins over C, A must also win over C.

Recent Research on Intransitive Dice

There is some recent serious research on intransitive dice. Conrey et al. [2] seek to understand how common intransitivity is. Specifically, they define a random n-sided die to be one whose face values $(a_1, \dots, a_n)$ are drawn from $1, \dots, n$ and average to $(n+1)/2$. They conjectured that for three such dice, knowing that $A \succ B$ and $B \succ C$ says essentially nothing about the probability that $A \succ C$.

This was proved in the collaborative project Polymath 13 [5, 6], which showed that for three random dice, the eight possible tournaments (which of A, B, and C beats which others) are asymptotically equally likely. (It was also conjectured in [2] and proved in [6] that ties of the sort $\mathbb{P}(X > Y) = \mathbb{P}(Y > X)$ are rare.) Cornacchia and Hązła [3] show that for four such dice, the equivalent conjecture is false: not all tournaments are asymptotically equally likely. Specifically, on four dice there are $2^6 = 64$ possible tournaments, of which $4! = 24$ are transitive, a 3/8 fraction, but the probability that a tournament is transitive is, asymptotically, strictly above 3/8.

A Transparent Snakes and Ladders Example

An anonymous referee suggested constructing a minimal example of a Snakes and Ladders board exhibiting intransitivity. We leave this problem open in its literal sense, but we produce a board for which the intransitivity is obvious, not dependent on hard to verify calculations.

To be faithful to the spirit of this challenge, we will use a standard die, with six sides and labels 1 to 6. However, we will in essence simulate three 6-sided dice, A, B, and C, with faces labeled

$$A: 2,2,2,2,3,4; \quad B: 1,1,1,4,4,4; \quad C: 1,2,3,3,3,3. \tag{2}$$

for player A	1 1	29 2	29 3	29 4	29 5	36 6	43 7
for player B	8 8	22 9	22 10	22 11	43 12	43 13	43 14
for player C	15 15	22 16	29 17	36 18	36 19	36 20	36 21
“3 to go”	22 22	29 23	29 24	29 25	29 26	29 27	29 28
“2 to go”	29 29	36 30	36 31	36 32	36 33	36 34	36 35
“1 to go”	36 36	43 37	43 38	43 39	43 40	43 41	43 42
“done!”	**43** 43						

Figure 5 The complete board, annotated.

It is easy to check that they are intransitive.

We will arrange that a player with simulated roll 4 finishes immediately, a player with a 3 will finish on the next roll, while a 2 requires two more rolls, and a 3 one more. (It is thus convenient to have the narrowest range of labels possible, and complete enumeration shows that (2) is best possible: there is no set of three intransitive dice using only the labels 1 to 3.)

Our board is illustrated in Figure 5. Intransitivity occurs with player A starting in cell 1, B in cell 8, and C in cell 15. Cell 43 is the winning cell. Cell 36's meaning is "1 roll to go": the board construction means it always takes exactly 1 roll to finish from this cell. Likewise, cell 29 means "2 to go," and cell 22 means "3 to go."

As in Figure 2, the value shown in the cell's center is where this cell takes you via a ladder (our construction has no snakes), or the cell's own number if you can remain there. In contrast to Figure 2, here the gray cells are those that *can* be occupied (1, 8, 15, 22, 29, 36, 43); the rest are all ladders.

Player A's rolls of 1–4, 5, and 6, corresponding to imagined die labels of 2 (four times), 3, and 4, respectively bring it to cells 29 ("2 more rolls to go"), 36 ("1 to go"), and 43 (done). Any roll of 1 to 6 starting from the "3 to go" cell 22 brings a player to the "2 to go" cell 29, and likewise starting from the "1 to go" and "2 to go" cells 29 and 36.

The board is not very small, but it directly translates a set of intransitive dice to a Snakes and Ladders board with intransitivity among cells 1, 8, and 15: the six cells after each of these simulate the faces of the corresponding die.

A few details. As it stands, players B and C cannot come to their starting squares. If desired, this can be rectified by adding six squares at the start of the board, with ladders to the starting squares for A, A, B, B, C, and C (the squares currently numbered 1, 1, 8, 8, 15, and 15), with all squares renumbered appropriately. This randomly gives a player the role A, B, or C.

The board can be made slightly smaller, at the expense of clarity. The last of the "3 to go" cells, cell 28, could itself be used as the base of the "2 to go" cells, with all of the "3 to go" ones having ladders to 28; the same could be done for the "1 to go" and "2 to go" cells. Also, the three dice and the faces within them could be reordered so that some faces can be used by two dice. Specifically, the dice of (2) could be laid out on the board as follows:

C	1	3	3	A|3	2	3	2	2	B|2	4	1	1	1	4	4

Another potential improvement would replace the "x to go" trick by simply positioning the players appropriately (depending on their simulated rolls) and letting them race to the finish. However, this blurs the outcomes—a player in the rear can get lucky and overtake one more advanced—and it is challenging to preserve the intransitivity.

Gregory B. Sorkin, Department of Mathematics,
The London School of Economics and Political Science,
Houghton Street, London, WC2A 2AE, England.
E-mail: g.b.sorkin@lse.ac.uk

Acknowledgments

Thanks to my colleagues Bernhard von Stengel for the pointer to "intransitive dice" and Neil Olver for his encouragement to write this up. Thanks also to an anonymous referee for suggesting that I look for a smallest example.

References

[1] Daniel Audet. Probabilités et espérances dans le jeu de serpents et échelles à deux joueurs. *Bulletin AMQ* LII:4 (2012), 28–40.

[2] B. Conrey, J. Gabbard, K. Grant, A. Liu, and K. Morrison. Intransitive dice. *Mathematics Magazine* 89:2 (2016), 133–143.

[3] E. Cornacchia and J. Hązła. Intransitive dice tournament is not quasirandom. Available at arxiv:2011.10067, 2020.

[4] M. Gardner. The paradox of the nontransitive dice. *Scientific American* 222:6 (1970), 110–111.

[5] Polymath. Polymath 13—a success! Available at https://polymathprojects.org/2017/08/22/, 2017.

[6] D. H. J. Polymath. The probability that a random triple of dice is transitive. Available at https://gowers.files.wordpress.com/2017/07/polymath131.pdf, July (2017).

[7] Andrew Topsfield. The Indian game of Snakes and Ladders. *Artibus Asiae* 46:3 (1985), 203–226.

[8] Andrew Topsfield. Snakes and Ladders in India: some further discoveries. *Artibus Asiae* 66:1 (2006), 143–179.

[9] S. Trybuła. On the paradox of three random variables *Applicationes Mathematicae* 5:4 (1961), 331–332.

[10] Wikipedia contributors. Intransitive dice. *Wikipedia*. Available at https://en.wikipedia.org/w/index.php?title=Intransitive_dice&oldid=1017474621, 2021 (accessed May 18, 2021).

[11] Wikipedia contributors. Snakes and Ladders. *Wikipedia*. Available at https://en.wikipedia.org/w/index.php?title=Snakes_and_ladders&oldid=1022691178, 2021 (Online; accessed May 18, 2021).