

A reprint from

American Scientist

the magazine of Sigma Xi, The Scientific Research Society

This reprint is provided for personal and noncommercial use. For any other use, please send a request Brian Hayes by electronic mail to bhayes@amsci.org.

Wagering with Zeno

Brian Hayes

VACATIONING IN ITALY, you wander into the coastal village of Velia, a few hours south of Naples. On the edge of town you notice an archaeological dig. When you go to have a look at the ruins, you learn that the place now called Velia was once the Greek settlement of Elea, home to the philosopher Parmenides and his disciple Zeno. You stroll through the excavated baths and trace the city walls, then climb a steep, cobbled roadway to an arch called the Porta Rosa. Perhaps Zeno formulated his famous paradoxes while pacing these same stones 900,000 days ago. Was there something special about the terrain that led him to imagine arrows frozen in flight and runners who go halfway, then half the remaining half, but never get to the finish line?

That night, Zeno visits you in a dream. He brings along a sack of ancient coins, which come in denominations of 1, $\frac{1}{2}$, $\frac{1}{4}$, $\frac{1}{8}$, $\frac{1}{16}$, and so on. Evidently the Eleatic currency had no smallest unit: For every coin of value $1/2^n$, there is another of value $1/2^{n+1}$. Zeno's bag holds exactly one coin of each denomination.

He teaches you a gambling game. First the coin of value 1 is set aside; it belongs to neither of you but will be flipped to decide the outcome of each round of play. Now the remaining coins are divided in such a way that each of you has a total initial stake of exactly $\frac{1}{2}$. The distinctively Eleatic part of the game is the rule for setting the amount of the wager. Before each coin toss, you and Zeno each count your current

A philosopher who did everything by halves may never win, but he won't go broke

holdings, and the bet is one-half of the lesser of these two amounts. Thus the first wager is $\frac{1}{4}$. Suppose you win that toss. After the bet is paid, you have $\frac{3}{4}$, and Zeno's fortune is reduced to $\frac{1}{4}$; the amount of the next bet is therefore $\frac{1}{8}$. Say Zeno wins this time; then the score stands at $\frac{5}{8}$ for you and $\frac{3}{8}$ for him, and the next amount at stake is $\frac{3}{16}$. If Zeno wins again, he takes the lead, $\frac{9}{16}$ to $\frac{7}{16}$.

In the morning you wake up wondering about this curious game. What is the likely outcome if you continue playing indefinitely? Is one player sure to win eventually, or could the lead be traded back and forth forever?

Can't Win, Can't Lose, Can't Tie

I briefly discussed a version of the Zeno gambling process in an earlier column ("Follow the Money," September–October 2002). Since then I have continued to explore the game, trying to understand its long-term behavior and relate it to other models in probability theory. I've had only partial success, and so what follows is a progress report, presented in the hope that others will build on it.

A few properties of the Zeno game are easy to state. For example, the betting process appears to be fair (assuming that the coin being flipped is unbiased). Each player has the same odds of winning or losing each round, and the amount at risk is the same.

Another way of saying that the game is fair is that the expectation value for each player is $\frac{1}{2}$. If you play many in-

dependent games, you should come out roughly even in the end. But an expectation value of $\frac{1}{2}$ does *not* mean you should expect to go home with half the money at the end of a single game. Indeed, after the first coin toss, the game cannot possibly end in a tie.

But you can never go broke, either—at least not in a finite number of plays. However small your remaining wealth, the wagering rule says you can't risk more than half of it. Of course the same reasoning protects your opponent as well: If you can't lose everything, neither can you win it all.

Here's another observation: In the game-within-a-dream described above, all of the numbers mentioned have a distinctive appearance. They are fractions whose denominator is a power of 2. In other words, they are numbers of the form $m/2^n$, called dyadic rationals. Is this predilection for halves, fourths, eighths, sixteenths, etc., a peculiarity of that one example, or does the pattern carry over to all Zeno games?

The answer comes from an inductive argument. Suppose at some stage of the game your score is a dyadic rational, x , and is less than or equal to $\frac{1}{2}$. Then the amount at stake in the next round of wagering is $x/2$, so that your new tally will be either $x - x/2$ or $x + x/2$. But $x - x/2$ is simply $x/2$, and $x + x/2$ is $3x/2$; both of these numbers are dyadic rationals. A similar (but messier) argument establishes the same result for values of x greater than $\frac{1}{2}$. Thus if your score is ever a dyadic rational, it will remain one for the rest of the game. But the starting value, $\frac{1}{2}$, is itself a dyadic rational, and so the only numbers that can ever arise in the game are fractions of the form $m/2^n$.

This line of argument actually yields a slightly stronger result. For a score $x < \frac{1}{2}$, the net effect of the gambling transaction is to multiply x by either $\frac{1}{2}$ or $\frac{3}{2}$. In either case, the denominator is doubled; as the game proceeds, the de-

Brian Hayes is Senior Writer for American Scientist. A collection of his essays, Group Theory in the Bedroom, and Other Mathematical Diversions, was published in April by Hill and Wang. Additional material related to the "Computing Science" column appears in Hayes's Weblog at <http://bit-player.org>. Address: 211 Dacian Avenue, Durham, NC 27701. Internet: brian@bit-player.org

nominator increases monotonically. An important consequence is that the entire numerical process is *nonrecurrent*: In the course of a game you'll never see the same number twice. This is one reason the game can't end in a tie: After the first flip of the coin, the score can never find its way back to $\frac{1}{2}$.

Walking with Zeno

The evolution of a Zeno wagering game corresponds to a special kind of random walk. A player's gains and losses are represented by the movement of a walker along the interval between 0 and 1. The walker starts at the position $x = \frac{1}{2}$. Each flip of the coin determines whether the next step is to the left (toward 0) or the right (toward 1). The length of the step is half the distance to whichever of these boundaries is nearer. In other words, the step length is $\frac{1}{2} \min(x, 1-x)$.

The upper illustration on page 196 shows a few trajectories constructed according to these rules. One feature of note is an apparent tendency for paths

to flee the middle of the interval and linger near the edges. It's not hard to understand this behavior, at least in a qualitative way. Whenever the walker is near the center, it is moving with higher velocity (that is, taking larger steps per unit time), and so it doesn't stay long in this neighborhood. Out at the periphery, the walker moves very slowly, and so it takes a long time to escape. It's as if the walker were moving over a landscape that's smoothly paved in the middle but becomes a sticky mire near the edges.

A plausible hypothesis suggests that a typical random walk will spend more and more time near the end points of the interval as the walk proceeds, coming arbitrarily close to 0 and 1. To test this idea you might follow a walk for many thousands of steps, but that process is computationally challenging. If you represent the walker's position by means of a floating-point number, the program will usually report that the walker has reached either 0.0 or 1.0 after just a few hundred steps. This

outcome would surprise Zeno! The problem is that floating-point formats have only finite precision, and very small values are rounded to zero.

A remedy for the round-off problem is exact rational arithmetic, but this becomes cumbersome. Here is the unwieldy numerical value of a game score after 150 steps:

2854495385411827052653424041061904510840082661
2854495385411919762116571938898990272765493248

The numerator and denominator both have 45 digits, and they differ by less than one part in a trillion.

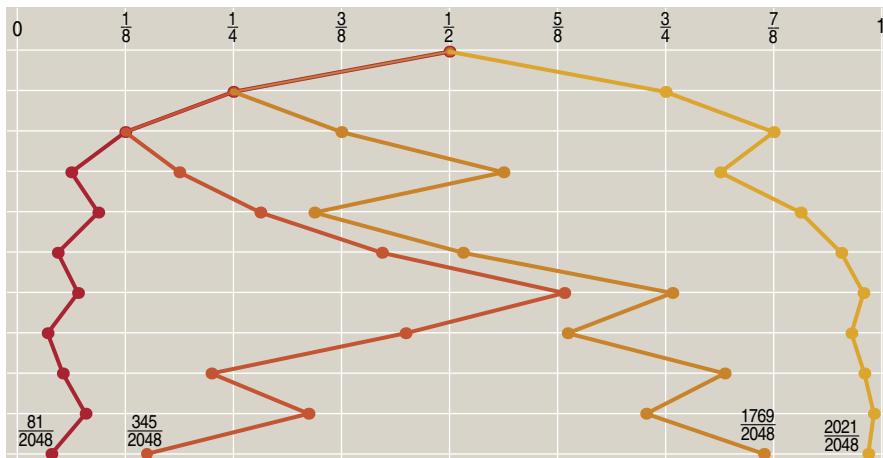
Zeno's Favorite Numbers

An alternative to tracing a few very long games is to gather statistics on the outcome of many shorter games. The lower illustration on page 196 gives the observed frequencies of various outcomes for games of length one through six, based on samples of several thousand trials.

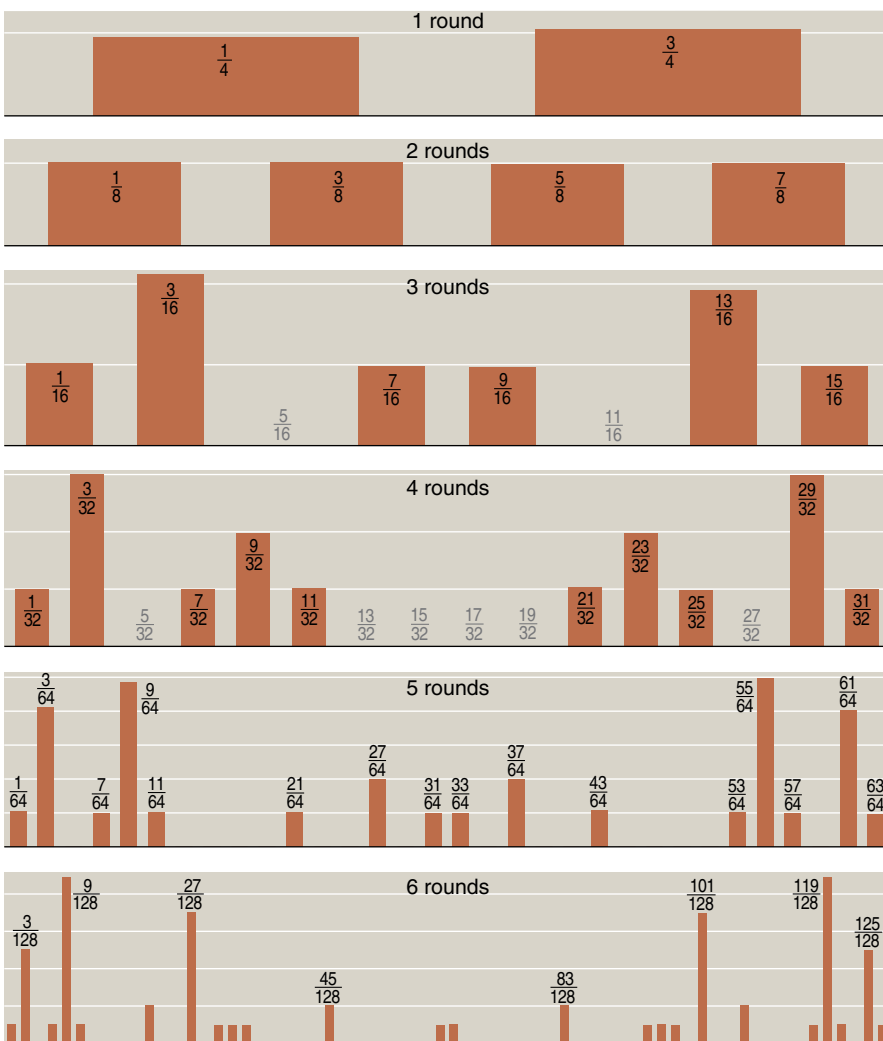
Games of length one (a single coin toss) can have only two possible out-



A stone gate called the Porta Rosa once connected two quarters of the ancient city of Elea, near what is now the town of Velia in southern Italy. Elea was the home of Zeno, the philosopher famous for his paradoxes of multitude and motion. The seemingly redundant "eyebrow" directly above the main arch of the Porta Rosa may have had a structural role, relieving stresses in the stonework. But followers of Zeno's thinking could tell a different story. If Zeno had designed the gate, he would have placed the primary arch halfway between the ground and the top of the wall, then a second arch halfway between the first arch and the top, a third arch halfway again, and so on. Perhaps the repairs near the top of the wall testify to the impracticality of this scheme? (Photograph by the author.)



In Zeno's wagering game, two players start with a stake of $\frac{1}{2}$ each and decide a series of bets by flipping a fair coin. The size of each bet is half the current wealth of the poorer player. The game is equivalent to a random walk on the interval from 0 to 1, taking steps whose size is equal to half the distance to the nearer boundary. The trajectories of four such walks are shown here.



The statistics of the Zeno game show that not all outcomes are equally likely, except in the very shortest games. The charts record the observed frequency of all possible end states for games that last from one round (that is, a single coin toss) to six rounds. The only numbers that can appear in the game are dyadic rationals—fractions of the form $m/2^n$ —but only a subset of these numbers are actually observed. Among the numbers present, some are much more popular than others; especially common are those with a power of 3 in the numerator.

comes, namely $\frac{1}{4}$ and $\frac{3}{4}$, and these events are equally likely. Two-round games must end with a value of $\frac{1}{8}$, $\frac{3}{8}$, $\frac{5}{8}$ or $\frac{7}{8}$, and again all four choices have the same probability.

Things get interesting with games of three or more rounds. After the third coin toss, the score of the gambler (or the position of the random walker) must be a fraction that, when expressed in lowest terms, has a denominator of 16. There are eight such fractions, but only six of them ever turn up as results of Zeno games; $\frac{5}{16}$ and $\frac{11}{16}$ are simply not observed. Among the six values that *do* occur, two of them ($\frac{3}{16}$ and $\frac{13}{16}$) are twice as common as the others.

Going on to four-round games, the pattern gets more peculiar. In this case all game values must be fractions with a denominator of 32. Of the 16 possibilities, only 10 are actually observed, and a few of these are two or three times more frequent than others. The likeliest game outcomes are $\frac{3}{32}$ and $\frac{9}{32}$ (along with the symmetrically related values $\frac{29}{32}$ and $\frac{23}{32}$, which are equal to $1 - \frac{3}{32}$ and $1 - \frac{9}{32}$). The differences in frequency are much too large to be an effect of statistical noise.

As the number of wagering rounds increases further, the patterns become even more pronounced. Wide gaps in the frequency distribution turn the graph into a snaggle-tooth smile. And certain numbers are dramatically more popular than the rest. For games of length six, only 24 of 64 possible outcomes are observed, and much of the probability is concentrated in just three values (and their symmetrical counterparts). The three favored fractions are $\frac{9}{128}$, $\frac{27}{128}$ and $\frac{3}{128}$. Why does the Zeno process favor these particular numbers? The powers of 2 in the denominator have already been explained, but why do the most common game outcomes all have powers of 3 in the numerator? It can't be an accident.

Climbing Zeno's Tree

In an effort to puzzle out these patterns, I tried constructing the tree of all possible outcomes for games of a given depth. In other words, I listed the two available moves from the starting state $x = \frac{1}{2}$, then for each of these positions I wrote down the two possible outcomes of another coin flip, and so on. (Eventually I wrote a program to do the calculations and draw the tree.)

Before plunging into the intricacies of the Zeno tree, it's worth pausing to

look at a simpler model—a random walk in which the step size is halved with every step. Starting at $x = \frac{1}{2}$, the walker goes either left or right a distance of $\frac{1}{4}$, then $\frac{1}{8}$, then $\frac{1}{16}$, etc. Tabulating all possible walks of this kind yields a binary tree that fans out to reach all the dyadic rationals. The branches bifurcate symmetrically, and each branch is completely isolated from all the others. The descendants of two neighboring nodes will come arbitrarily close but never touch. In a wagering game based on this rule, winning the first round is enough to put you in the lead forever. Even if you lose every subsequent bet, your share of the wealth can never fall below $\frac{1}{2}$.

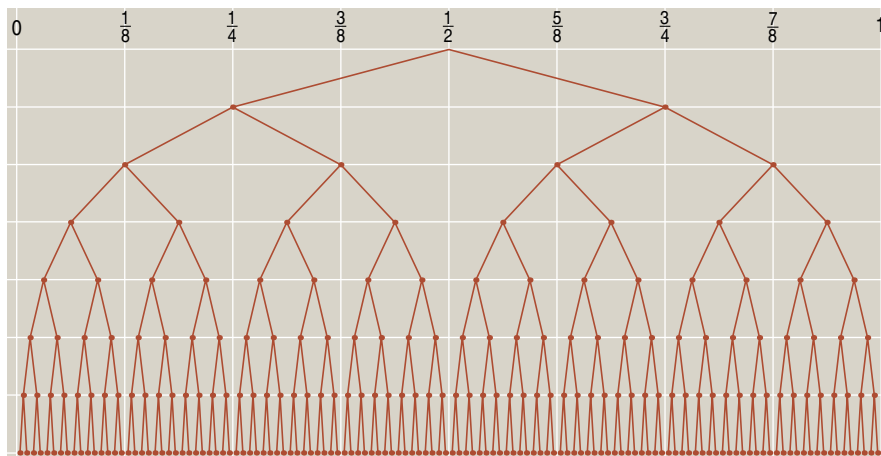
The Zeno game tree starts out exactly like this tidy binary tree—the first three levels are identical—but then things start to get weird. In the lower levels of the tree there's a lot of disorder, with various gaps between adjacent nodes, and many crossing branches. Even more remarkable, many paths through the tree bifurcate and then immediately come together again. (Strictly speaking, a structure with loops of this kind isn't a tree at all, but it will do no harm to continue using the term.)

We can understand a lot of what's going on in the Zeno tree by looking at two specific fragments. First is the loop formed below the node at $\frac{1}{4}$. One path descending from this node goes left to $\frac{1}{8}$ and then right to $\frac{3}{16}$; another path goes right to $\frac{3}{8}$ and then left to meet up at $\frac{3}{16}$. The fact that these two routes converge on precisely the same point is not some miraculous numerical coincidence; it's just a matter of working out the arithmetic:


$$x - \frac{x}{2} + \frac{x - \frac{x}{2}}{2} = x + \frac{x}{2} - \frac{x + \frac{x}{2}}{2}$$

When the dust settles, this identity comes down to the proposition that a half of three halves is three quarters, and likewise three halves of a half is three quarters.

The second interesting spot in the tree is right next door—the pair of paths descending from the node at $\frac{3}{8}$. The branch to the left, as we've already seen, goes to $\frac{3}{16}$, and then the next turn right on this path lands at $\frac{5}{32}$. The alternative route heads right from $\frac{3}{8}$ to $\frac{5}{16}$; however, on the next leftward bend this path fails to rejoin its partner at $\frac{5}{32}$. Instead the branch stops short at $\frac{11}{32}$. The reason is that this path crosses



A symmetrical binary tree represents all possible trajectories of a game similar to but simpler than the Zeno game. In this variant, the size of the wager—or the length of a step in a random walk—is cut in half after every round. All the dyadic rationals are present in the tree, and at each level they all have the same probability. Paths within the tree never cross or touch. The structure is called a *binary* tree because each node has two children; each child has just one parent node.



All possible trajectories of the Zeno game form a binary tree, but the structure is chaotic. Many branches cross, and there are also mergers, where a single node is a child of two sibling parents. Because of such rejoinings, many nodes can be reached by more than one path from the tree's root. The probability that a game will end at a particular node is proportional to the number of pathways reaching it, indicated here by the size of the dots marking the nodes.

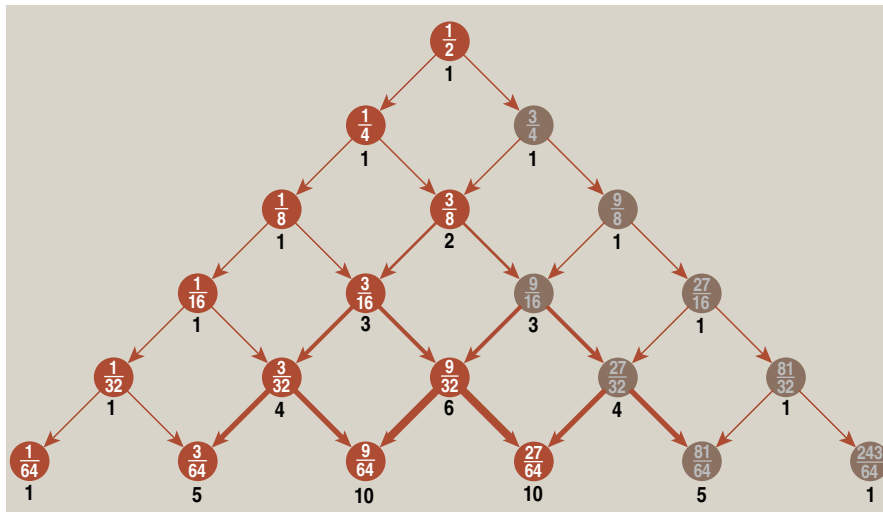
the midline of the tree at $x = \frac{1}{2}$, and for points to the right of this line distance is measured not from 0 but from 1. The two branches fail to meet because the equation no longer holds:

$$x - \frac{x}{2} + \frac{x - \frac{x}{2}}{2} \neq x + \frac{x}{2} - \frac{1 - (x + \frac{x}{2})}{2}$$

This mechanism effectively divides the tree into three vertical zones. For all points to the left of $\frac{1}{3}$ and to the right of $\frac{2}{3}$, a node's two children are both on the same side of the midline and thus are governed by the same rule for cal-

culating step lengths. As a result, these zones form a fairly regular latticelike structure, made up of diamond-shaped panes that grow narrower toward the periphery. In the middle zone, by contrast, all nodes have one child on each side of the midline, where different rules apply. The result is chaos.

The structure of this tree begins to illuminate some of the observations made about the statistical distribution of Zeno-game outcomes. For example, a simple counting argument explains the existence of gaps in the distribution. Every node of a binary



A simplified and regularized version of the Zeno game tree shows how junction points affect the number of paths through the tree and hence the probability of reaching a node. The tree corresponds to a random walk in which the step length is always one-half of x . The number of paths reaching each node (indicated by arrow thicknesses and by numbers in black) creates the pattern known as Pascal's triangle, with the likeliest nodes in the middle. The path counts in the real Zeno tree are different because the grayed nodes at the right are not present in that tree.

tree has two descending links, which is just enough to reach all the dyadic rationals, whose population doubles at each level of the tree. But when links from two or more nodes converge on the same child node, then other nodes must be missing from the tree. (If I have two children and my spouse has two children, that doesn't necessarily mean we have four children.)

Counting can also explain why some game endings are more common than others—such as why $\frac{3}{16}$ turns up twice as often as $\frac{1}{16}$. Think of playing a Zeno game as tracing a pathway through the tree, starting at the root (level 0) and continuing down to some final node. At each interior node, the path turns left or right with equal probability. In the simple binary tree, this procedure yields the same probability for all the nodes at a given level. In particular, at the third level each of the eight nodes is reached with probability $\frac{1}{8}$. In the Zeno tree, however, two different paths both lead to the third-level node at $\frac{3}{16}$. Since there are two ways of getting to this node, the probability of landing there is doubled.

The path-counting analysis can be understood more clearly if we suppress some of the clutter and disorder in the Zeno tree. Suppose a random walk always takes steps of size $\frac{1}{2}$, adopting the rule that governs the left half of the Zeno tree but applying it everywhere. The corresponding tree has a uniform lattice structure in which it's easy to

count paths (see illustration above). Every interior node has two parents, and the number of ways of reaching the node is simply the sum of the number of ways of reaching the parents. Hence the path counts correspond to the entries in Pascal's triangle.

There's something else notable about this tidied-up version of the Zeno tree: All the numbers that appear in the lattice take a very specific form: $3^m/2^n$. Thus the numerators are drawn from the series 1, 3, 9, 27, 81, ..., and the denominators follow the now-familiar progression 2, 4, 8, 16, 32, ... These are the very numbers that appear with higher-than-average frequency in the Zeno game, and it's clear why. The highest path counts run down the middle of the lattice, supporting the observation that numbers such as $\frac{3}{32}$, $\frac{9}{64}$ and $\frac{27}{256}$ are extraordinarily popular.

Of course the real Zeno tree is not so tidy; the lattice is torn apart between any two nodes on opposite sides of the midline. As a result of these trans-midline events, factors other than 3 can enter the numerator; they form shadow lattices of their own, visible in some parts of the tree. Thus the probabilities calculated from the simple lattice are at best approximations.

Midline Crises

Those are all the answers I can offer on the Zeno wagering game, but I certainly haven't run out of questions. Here are three more.

First, how dense is the set of numbers included in the Zeno tree, measured as a proportion of all the numbers that *might* be present? At level n there are 2^n dyadic rationals in all; what fraction of them are Zeno-tree numbers, and how does that fraction evolve as n increases? At the first three levels of the tree the fraction is 1: All the dyadic rationals are included. Then there is a steep linear descent as the fraction goes from $\frac{3}{4}$ to $\frac{5}{8}$ to $\frac{1}{2}$ to $\frac{3}{8}$. This series obviously cannot continue or the tree will disappear in three more steps. And indeed the slope begins to flatten out: The next four elements of the series are $\frac{9}{32}$, $\frac{25}{128}$, $\frac{19}{128}$ and $\frac{7}{64}$. At this point each level of the tree includes only about a tenth of the dyadic rationals at that level. It seems a reasonable guess that the density will approach zero as n tends to infinity.

Second, how much structure can we find in the Zeno tree? To make this question more concrete, suppose you want to determine which node, on a certain level of the tree, gives the closest approximation to some specified value on the real number line. With the standard binary tree this is easy: You can list a series of left and right turns that describe a path from the root to the closest node. It's not clear to me how to find such a path in the Zeno tree, except by exhaustive search.

Finally there's this big question: If you allow a game to continue arbitrarily long, will one player gain an advantage and hold onto it indefinitely, or will the lead change hands repeatedly? In terms of the random walk, will the walker get stuck on one side, or will the walk cross back and forth over the midline? Of course it's always possible for the walker to get back to midline; all it takes is a sufficiently long sequence of steps in the right direction. But it doesn't necessarily follow that such an event will have a probability greater than zero when the number of steps grows without limit.

For what it's worth, here's what the experimental evidence suggests. In random walks of 10,000 steps, roughly half exhibit no midline crossings at all: The walker sets out either left or right initially and never gets back to the other side. For those walks that *do* cross the midline, the last observed crossing is usually within the first 10 steps; the latest I have seen is step 101. But trials of 10,000 steps, or even 10 million, prove nothing.

In Search of Zeno's Urn

It's entirely possible that these questions have already been answered somewhere in the literature on probability models. Knowledgeable friends have pointed me to two areas in particular: urn models and reinforced random walks. I have not yet found any discussion of a model that I can recognize as being isomorphic to the Zeno game, but I've found some fascinating reading along the way.

Urn models are the 18th-century ancestor of the Powerball lottery: You mix up a bunch of balls in a container and draw them out one by one. Of particular relevance is a class of models studied by George Pólya, Bernard Freedman, David A. Freedman and a number of others, including recently Robin Pemantle. In one version, an urn initially holds one white ball and one black ball. Then each time a ball is drawn at random from the urn, that ball is put back in along with an extra ball of the same color. When the process is repeated many times, will the ratio of black to white balls settle down to some fixed value, or will it keep fluctuating? The answer is that in any single long run, the ratio will tend to a stable value, but the value

itself is a random variable. If you try the experiment again, it will come out differently.

The reinforced random walk was invented by Persi Diaconis and has also been studied by Pemantle. The idea is to watch a random walker stepping from node to node through a network. Visiting a node increases the probability that the same node will be chosen the next time the walker is nearby. The question asked is whether the walker can range freely throughout the graph forever or will get trapped in some local neighborhood. The answer seems to depend on the details. On a one-dimensional lattice (like the line of integers), the walker gets stuck in a five-node region. But a variation that associates probabilities with the links between nodes rather than the nodes themselves allows the walker to escape confinement and visit every node infinitely often. On two-dimensional infinite lattices the fate of the walker is unknown.

It is unclear to me whether the Zeno game belongs to the same family as these models. The Zeno mechanism includes no explicit reinforcement of probabilities; on the other hand, it does have an indirect form of positive

feedback, because movement away from the center reduces the step size and thereby makes it harder to move back. Perhaps the Zeno process can be reformulated in a way that will make the correspondence with known work clearer. But the analysis of all such models is a subtle art. Every time I think I'm getting close to an answer, it seems the problem has moved on a little further down the road.

Bibliography

- Freedman, David A. 1965. Bernard Friedman's urn. *Annals of Mathematical Statistics* 36:956–970.
- Friedman, Bernard. 1949. A simple urn model. *Communications on Pure and Applied Mathematics* 2:59–70.
- Hayes, Brian. 2002. Follow the money. *American Scientist* 90:400–405.
- Johnson, Norman L., and Samuel Kotz. 1977. *Urn Models and Their Application: An Approach to Modern Discrete Probability Theory*. New York: Wiley.
- Krapivsky, P. L., and S. Redner. 2004. Random walk with shrinking steps. *American Journal of Physics* 72:591–598.
- Pemantle, Robin. 2007. A survey of random processes with reinforcement. *Probability Surveys* 4:1–79.
- Steinsaltz, David. 1997. Zeno's walk: A random walk with refinements. *Probability Theory and Related Fields* 107:99–121.