THE VIBONACCI NUMBERS

Brian Hayes

A reprint from

American Scientist

the magazine of Sigma Xi, the Scientific Research Society

THE VIBONACCI NUMBERS

Brian Hayes

The Fibonacci numbers 1, 1, 2, 3, 5, 8, 13,... are one of the royal families of mathematics. Like other royalty, they have an ancient pedigree, everybody knows all about them, and no one takes them very seriously anymore. The Fibonaccis and their relations have been thoroughly studied for centuries, so no one would expect much in the way of novelty or innovation to turn up among them. Nevertheless, a whole new branch of the family has just sprouted up. What's more, the new cousins are a highly erratic bunch—sports in the royal lineage.

The pattern in the sequence of Fibonacci numbers is easy to see: Each term (except for the first two) is the sum of the two preceding terms. Expressed as a formula: $f(n) = f(n–2) + f(n–1)$. The new variation on the series changes this formula in only one detail. Instead of always adding two terms to produce the next term, you either add or subtract, depending on the flip of a coin at each stage in the calculation. If the coin comes up *heads,* say, you add as usual, but if the result is *tails,* you subtract $f(n–1)$ from $f(n–2)$. In other words, the formula becomes $f(n) = f(n–2) \pm f(n–1)$, where the symbol "±" signifies that you choose either addition or subtraction randomly and with equal probability.

Here are a few short sequences generated by this random sum-or-difference algorithm:

1, 1, 0, 1, –1, 2, –3, –1, –2, 1, –3, 4, 1, 3, 4, 7, 11
1, 1, 2, –1, 3, 2, 1, 1, 0, 1, –1, 0, –1, 1, –2, –1, –1
1, 1, 2, –1, 3, 2, 1, 3, –2, 1, –3, –2, –5, 3, –8, –5, –13
1,1, 0, 1, –1, 0, –1, 1, –2, 3, 1, 2, –1, 3, –4, 7, –11
1, 1, 0, 1, 1, 2, 3, 5, –2, 7, 5, 12, –7, 19, –26, 45, 19

Certain properties of the Fibonacci sequence continue to show through here, most notably the alternation of two odds and an even in all the series. But the steady growth of the Fibonacci numbers is replaced by fluctuations of increasing amplitude. The way the numbers seem to vibrate between negative and positive values leads me to suggest the name Vibonacci series; I shall designate them by the symbol $v(n)$.

Looking at the fluctuations, you might conclude that the randomness in the formula has wiped out all traces of order. And it is certainly true that the Vibonacci sequence is nondeterministic. Unlike the conventional Fibonacci numbers, where the *n*th term has a single, definite value that needs to be calculated only once, the *n*th term of a Vibonacci series has a distribution of possible values; $v(n)$ is likely to be different every time you compute it. Nevertheless, a great deal of order persists in the randomized sequences. In particular, the absolute value of $v(n)$ grows exponentially as *n* increases. You can measure the growth rate in simple computer experiments. What's more remarkable, the value of the number that determines the growth rate has been pinned down in a mathematical proof.

Breeding Like Rabbits

The Fibonacci numbers were introduced to the world 800 years ago by Leonardo of Pisa, who also made another major contribution to mathematics: It was he who brought Indo-Arabic numerals into European culture. "Fibonacci" was apparently Leonardo's nickname, a shortening of *Filius Bonacci,* or son of Bonacci. Of course Leonardo did not call his sequence the Fibonacci numbers; the name was popularized by the 19th-century French mathematician Edouard Lucas.

The whole matter began with a contrived problem about the breeding of rabbits. Suppose a pair of rabbits breeds once a month and always produces a single pair of offspring, which breeds the following month. Each pair breeds twice and then dies. Starting with a single pair of rabbits under these assumptions, how many pairs will be living after *n* months? The answer is $f(n)$.

Leonardo's fanciful problem is not of any great interest to population biologists, but the Fibonacci numbers are famous for turning up in many other contexts. Their patterns appear in seashells and sunflowers and pinecones; they count the ways of tiling a checkerboard with dominos; they are present in Pascal's Triangle if you know where to look for them. The Fibonacci numbers also played an essential role in the solution to Hilbert's Tenth Problem. They even have their own journal, the *Fibonacci Quarterly,* published since 1963.

One of the most important properties of the Fibonacci numbers was first noted in the 17th

Brian Hayes is a former editor of American Scientist. *Address: 211 Dacian Avenue, Durham, NC 27701. Internet: bhayes@amsci.org.*

century by Johannes Kepler: The ratios of successive Fibonacci numbers, $f(n–1)/f(n–2)$, form a new sequence, beginning 1, 2, 1.5, 1.666,..., 1.6, 1.625. This series converges on a value of 1.618033989..., known as ϕ, or the golden ratio. Over the years the golden ratio has acquired something of a cult following, which I would not want to encourage, and yet it truly is a remarkable number. It has the curious property that if you take its inverse (that is, $1/\phi$) and then add 1, you recover the original number; in other words ϕ is a solution to the equation $1/x = x – 1$. This equation can be rearranged as $x^2 – x – 1 = 0$, for which the quadratic formula gives the solutions $(1 + \sqrt{5})/2$ and $(1 – \sqrt{5})/2$. The first of these numbers is ϕ; the second is $1 – \phi$, or –0.618033989....

Powers of ϕ provide close approximations to the Fibonacci numbers. Since ϕ is the limiting value of the ratio between successive Fibonacci numbers, multiplying any member of the series by ϕ yields an approximation to the next member. Looking at the approximation process through the other end of the telescope, the nth root of the nth Fibonacci number approximates ϕ. A more complex formula, $(\phi^n – (1 – \phi)^n)/\sqrt{5}$, gives the *exact* value of $f(n)$; this strange congeries of irrationals always reduces to an integer, as long as n itself is an integer.

The Fibonacci theme has many variations. Edouard Lucas, the mathematician who named the Fibonacci numbers, has a related series named after him: The Lucas numbers are defined by the same recurrence formula but start with the integers 2, 1 instead of 1, 1. The first few members of the Lucas series are 2, 1, 3, 4, 7, 11, 18, 29, 47, 76. Remarkably, the ratio of successive Lucas numbers also converges to ϕ; indeed, it turns out that you can start a Fibonacci-like series with *any* pair of numbers, and the limit of the growth rate will always be ϕ.

Another variation has come to be known as the Tribonacci series. Instead of adding the previous two terms at each step, you add the previous three. A convenient way to get such a sequence started is to pretend that the initial 1 is preceded by an indefinite list of 0s. With this convention, the Tribonacci sequence begins 1, 1, 2, 4, 7, 13, 24, 44, 81, 149. In this case the growth rate is not ϕ; instead the ratio of successive terms approaches 1.83929.... The analogous Tetrabonacci series, where each term is the sum of the previous four, begins 1, 1, 2, 4, 8, 15, 29, 108, and the ratio of terms converges to 1.92756.... Clearly this process can be generalized to k-bonacci series. You might want to think about what happens in the limiting case where each term is the sum of *all* the previous terms. (The answer is given below.)

Rabbit Cannibalism

If the Fibonacci series describes rabbit breeding, what does the Vibonacci series describe? A cute answer might be the breeding of cannibalistic rabbits—animals that sometimes reproduce normally but at other times consume their own young or their own parents. But the inventor of the series had nothing so whimsical in mind. He was working on a problem in numerical analysis, the branch of mathematics concerned with large-scale computations.

The inventor is Divakar Viswanath, a young mathematician and computer scientist who earned his Ph.D. last year at Cornell University. He has spent the past academic year at the Mathematical Sciences Research Institute in Berkeley; in the fall he will take up a position in the departments of mathematics and computer science at the University of Chicago. His paper on random Fibonacci sequences will be published in *Mathematics of Computation*.

The most symmetric version of the Vibonacci series assigns randomly chosen signs to *both* of the preceding terms in the sequence; that is, the recurrence formula is $v(n) = \pm v(n–2) \pm v(n–1)$. But if you are interested mainly in the absolute value of each term—ignoring the sign—the version with just a single random sign yields the same result. The main question in the study of the series is how fast the absolute value of $v(n)$ grows as n
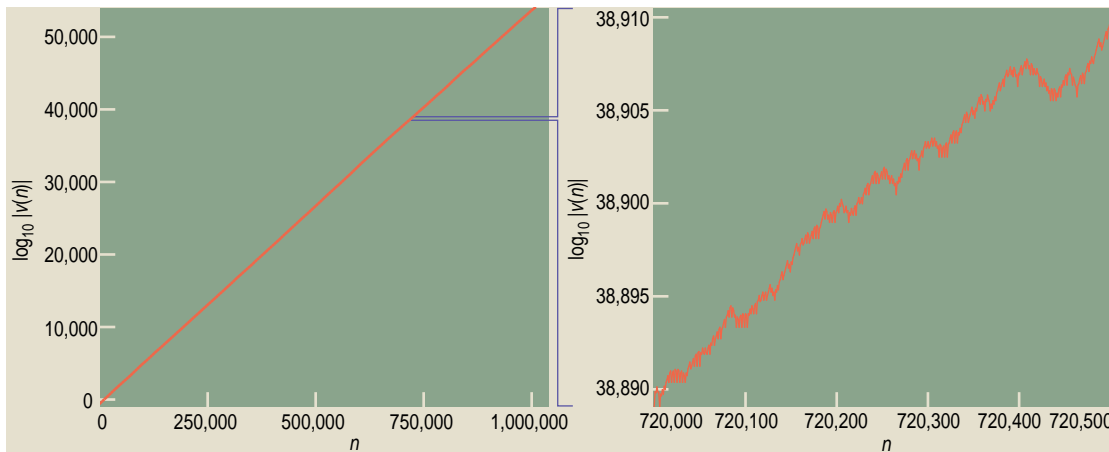

Figure 1. Vibonacci numbers are defined by a random process and at each step can either grow or diminish. Nevertheless, over the range up to $n = 1,000,000$ *(left)*, **the absolute value of $v(n)$ grows at a remarkably steady pace. Seen in finer detail** *(right)*, **the sequence has the jitters. Note that the scale is logarithmic, and $v(1,000,000)$ is greater than $10^{50,000}$.**
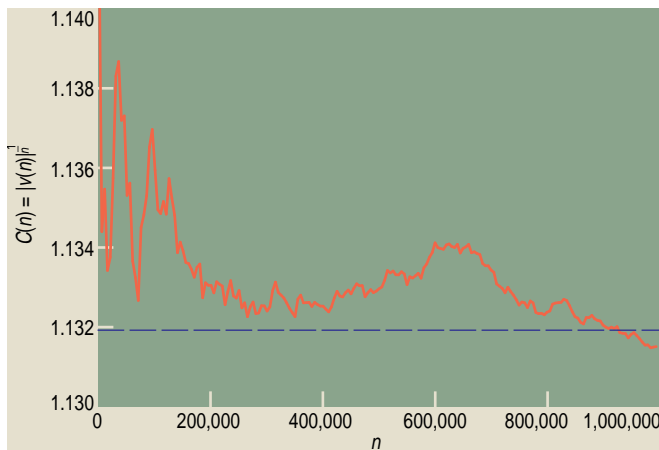
Figure 2. Growth constant of the Vibonacci sequence, designated *C*, is defined as the *n*th root of the absolute value of *v(n)*. The red curve traces the result of a numerical experiment; the dashed blue line indicates the value *C*=1.13198824... calculated by Divakar Viswanath.

increases. In other words, the aim is to find for the Vibonacci numbers a constant *C* that plays the same role as ϕ does for the Fibonacci numbers. This hypothetical growth rate *C* is defined as the *n*th root of $|v(n)|$, where the notation $|x|$ signifies the absolute value of *x*.

It is not immediately obvious that $|v(n)|$ should be expected to grow at all in the long run. With equal numbers of random additions and subtractions, you might guess that the series would hover around some fixed average, so that the *n*th root of $|v(n)|$ would converge to a value of 1, signifying no growth. Or, conversely, the Vibonacci numbers might bounce around so chaotically that the growth rate would never converge on any stable value; the limit of the *n*th root of $|v(n)|$ might simply not exist.

The question of whether a limiting growth rate exists was settled 40 years ago by Harry Furstenberg and Harry Kesten, then both at Princeton University. For a broad class of random processes, including the one I'm calling the Vibonacci series, they showed that a limit does exist, given a few mild assumptions. Three years later Furstenberg proved that the growth rate is "almost surely" greater than 1. The "almost surely" disclaimer is needed because of the probabilistic nature of the system. Some sequences do fail to grow, and you cannot absolutely exclude the possibility of stumbling onto one. For example, a carefully chosen pattern of alternating plus and minus signs generates the cyclic Vibonacci series 1, 1, 0, 1, 1, 0,.... But such exceptional sequences are unlikely; indeed, in the limit as *n* goes to infinity, they have probability zero. They exist, but you have no chance of ever finding them. Thus Furstenberg's "almost surely" result is actually quite a strong statement. It implies not merely that almost all Vibonacci sequences grow but that any individual sequence grows with probability 1, if it is allowed to continue long enough.

Unfortunately, apart from showing that *C* must exist and be greater than 1, Furstenberg's theo-

rem gives no information about the magnitude of *C*. The value 1 is merely a lower bound. There is a complementary upper bound: The value of C cannot be greater than 1.618..., since that is the growth rate of the ordinary Fibonacci series, with all additions and no subtractions.

Numerical experiments provide estimates of *C*. For small *n*, it's easy to enumerate all possible *n*-step Vibonacci seqences and calculate their growth rates by taking the *n*th root of the final term. Since each of these series is equally likely, the arithmetic average is an estimate of *C*. For example, there are four Vibonacci series for *n* = 4, namely 1, 1, 0, –1; 1, 1, 0, 1; 1, 1, 2, 1; and 1, 1, 2, 3. Thus the final terms are –1, 1, 1, 3, and the average of the fourth roots of their absolute values is about 1.08. For *n* = 20 the corresponding estimate of *C* is about 1.18. But tracing out the entire tree of Vibonacci sequences becomes impractical for large *n*. At *n* = 20 there are already half a million branches to be tabulated.

Random sampling gives approximations to *C* for much larger values of *n*. Figure 1 shows the outcome of a single computer run generating Vibonacci numbers up to $v(10^6)$. There is no question that this particular sequence is growing exponentially; it attains heights of greater than $10^{50,000}$. The value of *C* calculated from the series exhibits distinctive fluctuations, which appear to diminish in amplitude as *n* increases, but which also tend toward longer wavelengths at higher *n*. Convergence is slow. The growth rate appears to be somewhere near 1.13, but from these data it would be difficult to estimate *C* with greater precision.

C = 1.13198824...

Viswanath's approach to determining *C* is indirect, and indeed it takes a detour through some areas of the mathematical landscape that might seem to have no connection with Fibonacci-like series. But by transposing the problem into another realm, where more powerful tools can be brought to bear, he is able not merely to produce an empirical estimate of *C* but to prove that the value of the constant must lie in a certain narrow interval. The proof is not a simple one, and here I shall present only a hasty sketch; the details are in Viswanath's paper.

The first step is to recast the Vibonacci process in terms of matrices. The ordinary Fibonacci series can be defined by the matrix equation:

$$\begin{pmatrix} f(n-1) \\ f(n) \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} f(n-2) \\ f(n-1) \end{pmatrix}$$

Applying the rules for matrix multiplication confirms that the equation has the expected behavior. In the product on the left side of the equation the first row is given by the sum $0 \times f(n-2) + 1 \times f(n-1)$, which of course is just *f(n–1)*; the second row of the product is $1 \times f(n-2) + 1 \times f(n-1)$, which is the definition of *f(n)*.

Repeating the matrix multiplication generates successive terms of the Fibonacci sequence. For

example, given the initial terms 1, 1, three matrix multiplications produce the terms 3, 5:

$$\begin{pmatrix} 3 \\ 5 \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}\begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}\begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}\begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

Continuing in the same way generates every Fibonacci number in turn, and only those numbers.

The Vibonacci sequence can also be redefined as a matrix product, the only difference being that there are two matrices to choose from:

$$\begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 0 & 1 \\ 1 & -1 \end{pmatrix}$$

At each step in generating the series, one of these matrices is selected at random with probability $1/2$. Whenever the second matrix happens to be picked, the minus sign in the lower right corner has the effect of subtracting $v(n-1)$ from $v(n-2)$, rather than adding the two terms, just as in the more direct definition of the Vibonacci series.

Why bother with this elaborate reformulation of the problem if it merely reproduces the same result? Because the study of products of random matrices offers a handy toolkit of useful methods. It allows the Vibonacci process to be viewed in a geometrical context. Suppose that any two adjacent Vibonacci numbers, $v(n-1)$ and $v(n-2)$, represent the coordinates of a point in the x,y plane. Drawing a line from the origin at 0,0 to this point defines a direction in the plane, specified by an angle θ or a slope m. The slope is simply y/x, so it is given directly by the coordinates of the point. Now multiply the pair of coordinates by one of the 2×2 Vibonacci matrices. What happens to the point? It is mapped into a new point—with coordinates $v(n)$ and $v(n-1)$—which defines a new direction from the origin. Specifically, multiplying by one of the Vibonacci matrices causes a rotation to a new slope of either $(1 + 1/m)$ or $(1/m-1)$.

These transformations are easier to understand through a brief example. Starting with a slope of $m = 1$ (equivalent to an angle of 45 degrees), the $m \rightarrow (1 + 1/m)$ transformation yields a new slope of 2 (about 63 degrees); applying the $m \rightarrow (1/m-1)$ transformation rotates the line to a slope of $-\frac{1}{2}$ (about 333 degrees). If you continue to iterate this process, always taking the last value of m and replacing it with a random choice of either $(1 + 1/m)$ or $(1/m-1)$, you create a kind of random walk through the space of possible slopes. The progression of slopes seems to have nothing in common with the Vibonacci series, and yet the connection through products of random matrices shows they are really just two manifestations of the same process. (And the connection isn't as mysterious as it might seem. Note that repeating the transformation $m \rightarrow (1 + 1/m)$ yields a series of numbers that converges to ϕ for any starting m. Hidden in the mapping $m \rightarrow (1 + 1/m)$ is the equation $1/x = x-1$ that defines ϕ.)

The slope m in these formulas can take on any value along the real number line, from negative to positive infinity, but not all slopes are equally likely. The key to understanding the random walk—and also to calculating the growth rate of the Vibonacci series—is identifying the probability distribution that determines the likelihood of every possible slope. Viswanath's main contribution was finding a way to estimate this distribution to any desired degree of accuracy.

The probability distribution is a peculiar one—not at all like the smooth Gaussian curve that describes so many random processes. Instead it has multiple spiky peaks and deep canyons, and if you look closer, you find that the spikes have spicules, and the canyons are creased by smaller canyons. Thus the distribution appears to be a fractal landscape that cannot be described by any continuous function. Viswanath sidestepped this problem by constructing an ingenious discrete partitioning of the real number line. The structure is called the Stern-Brocot tree, after the mathematician Moriz Stern and the watchmaker Achille Brocot, who discovered it more than a century before Viswanath did.

The basic idea of the tree is to divide the set of real numbers into progressively finer—but not necessarily equal—intervals. Writing zero as $0/1$ and representing infinity by the notation $1/0$, all positive real numbers lie in the interval $[0/1, 1/0]$. This half of the number line is broken down into the subintervals $[0/1, 1/1]$ and $[1/1, 1/0]$, extending from zero to 1 and from 1 to infinity respectively. The left-hand subnode then splits into $[0/1, 1/2]$ and $[1/2, 1/1]$, and the right-hand subnode into $[1/1, 2/1]$ and $[2/1, 1/0]$. The general rule is that an interval $[a/b, c/d]$ splits into $[a/b, (a+c)/(b+d)]$ and $[(a+b)/(c+d), c/d]$. All the numbers in the negative half of the real line are classified in the same way, and the two halves are joined by a special root node labeled $[-1/0, 1/0]$.

The Stern-Brocot tree has a place in this story because there is a direct correspondence between the random walk among slopes m described above and paths traced out through the branches of the tree. Any such a path can be written as a sequence of *left* and *right* commands, giving directions for how to get from the root of the tree to a specific interior node. For example, from the
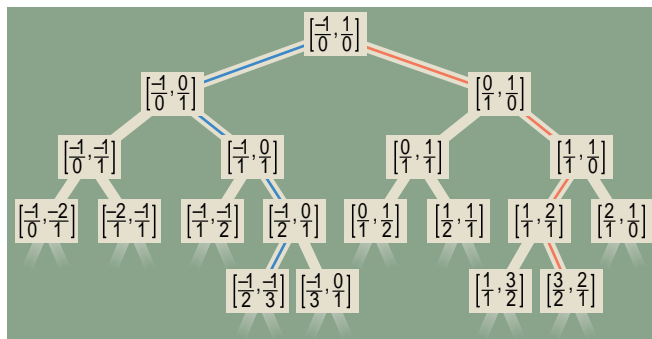

Figure 3. Stern-Brocot tree was the key to Viswanath's calculation of the constant C. Paths through the tree reveal the probability distribution of the Vibonacci sequence; in the section of the tree shown here, a single Vibonacci transformation maps the red path into the blue path.
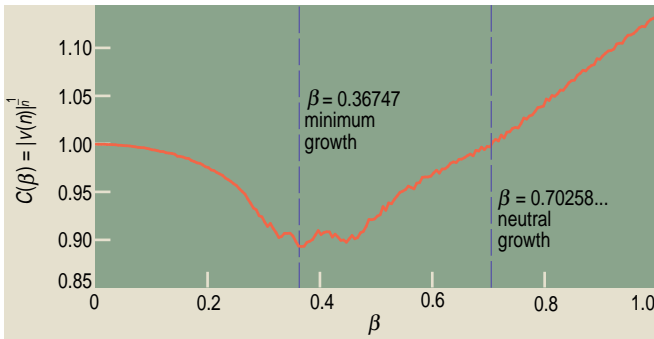
Figure 4. Variant series defined by the equation $v(n) = v(n-1) \pm \beta v(n-2)$
can either grow or dwindle, depending on β**. The minimum point and
the neutral point in the curve were determined accurately by Mark
Embree and Lloyd N. Trefethen of the University of Oxford.**

[0/1, 1/0] node, a path of *right, left, left* brings
you to the [1/1, 3/2] node. Suppose the current
value of the slope m lies somewhere within this
interval. After a transition from m to $1/m$, the
new slope must lie in the [2/3, 1/1] node, which
is the mirror image of the original node; to reach
it, you follow the opposite sequence of turns,
left, right, right. Mapping the transformations
$m \to (1 + 1/m)$ and $m \to (1/m - 1)$ into tree paths
is only a little more complicated. Through this
mapping, the intervals of the Stern-Brocot tree
yield up an approximation to the probability dis-
tribution of the Vibonacci series.

Viswanath tabulated all paths through the
Stern-Brocot tree to a depth of 28 levels, where
the tree has more than 50 million nodes. In this
way he was able to calculate the value of C to
eight decimal places. His result is $C = 1.13198824....$
It has the same "almost surely" status as Fursten-
berg's proof; that is, every instance of the Vi-
bonacci series will grow at this rate with proba-
bility 1 if it is continued long enough.

An unusual feature of Viswanath's proof is
that it relies on floating-point computer arith-
metic. Computer-aided proofs are no longer a
novelty in mathematics, but few of them adopt
floating-point methods because the results are
not exact. Viswanath includes an error analysis
showing that any arithmetic inaccuracies are
much smaller than the uncertainty introduced by
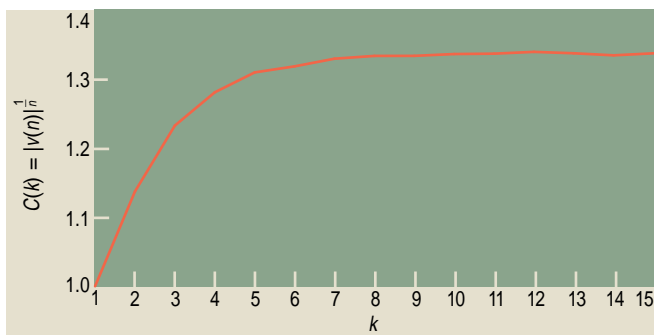truncating the Stern-Brocot tree at a finite depth.

Figure 5. The k**-bonacci series calculates each term by giving random
signs to the previous** k **terms and then taking the sum. As** k **increases,
the growth rate appears to converge on a value of about** $C=1.32$**.**

Variations and Generalizations

All the well-known variations on the Fibonacci
series also have analogues in the randomized
world of the Vibonacci numbers.

It was noted above that the Fibonacci recur-
rence has the same growth rate starting from any
two initial terms. Does the Vibonacci series share
this property? Numerical experiments quickly af-
firm that it does.

A randomized variant of the Tribonacci series
defines each term as the sum of the previous
three terms with randomly chosen signs. Not
surprisingly, the absolute value of the three-term
series grows faster than that of the two-term Vi-
bonacci recurrence. Experiments suggest that the
random Tribonacci growth constant is about 1.22.
For the random-sign analogue of the Tetrabonac-
ci series, with four terms included in each sum,
the growth rate is roughly 1.27. It appears that
the growth rate continues increasing slowly as
more terms are included.

What happens in the limiting case, where *all*
earlier terms are gathered up into the sum? In the
pure Fibonacci version, without randomness, the
first few terms of this series are 1, 1, 2, 4, 8, 16,
32,.... These are the successive powers of 2, and so
the exponential growth rate is exactly 2. The ran-
dom analogue of this series is actually the problem
that got Viswanath started on his investigation of
the Vibonacci numbers. Working with Lloyd N.
Trefethen, his dissertation supervisor at Cornell,
he was studying the series $r(n) = \pm r(1) \pm r(2) \pm ...$
$\pm r(n-1)$, where each term is calculated by giving
random signs to all the preceding terms before
taking the sum. Finding the asymptotic growth
rate of this sequence would settle an outstanding
question in the theory of random matrices, pro-
posed by Trefethen. Viswanath was unable to find
a rigorous solution and so turned to the two-term
Vibonacci sum as a simpler model. Numerical ex-
periments suggest that the growth rate for the ran-
dom sum of all terms is about 1.32.

Another way of generalizing the Vibonacci
process is to consider what happens when the
probabilities of choosing plus or minus are not
equal. Intuition suggests that any bias in the prob-
abilities, favoring either plus or minus, ought to
cause faster growth in the absolute value of $v(n)$.
In other words, the sequence with equal probabil-
ities should have the minimum growth rate, with
C increasing toward ϕ in the extreme cases of all-
plus or all-minus sequences. Experiments support
these inferences, but the exact behavior of the se-
ries with skewed probabilities is unknown. The
Stern-Brocot proof works only when plus and mi-
nus are chosen with equal probability.

Perhaps the most interesting Vibonacci varia-
tion has been studied by Mark Embree and Tre-
fethen, who is now at the of the University of Ox-
ford. Instead of skewing the probabilities, they
scale one of the two terms in the random sum by
an adjustable factor, which they denote β. That is,
the recurrence relation is $v(n) = v(n-1) \pm \beta v(n-2)$.

If you try a value of $\beta = \frac{1}{2}$ in this equation, you will immediately see a dramatic change. The sequence no longer grows exponentially; on the contrary, it rapidly dwindles away. In other words, the exponential growth rate is less than 1; numerical evidence gives a value of about 0.929.

Setting $\beta = 1$ recovers the original Vibonacci sequence, where of course the growth rate is known to be 1.13198824.... If the series decays for $\beta = \frac{1}{2}$ and grows for $\beta = 1$, there must be some intermediate value of β where it is "neutrally buoyant," neither rising nor falling on average. Embree and Trefethen have searched for this point of equilibrium, designated β^*, and they find the closest approximation at $\beta^* = 0.70258....$ Computer runs at this setting develop large and erratic fluctuations, but they seem not to veer off into unbounded growth or decay.

Searching for the value of β that minimizes the growth rate (or in other words maximizes the decay rate) turned up more surprises. The minimum cannot be at $\beta = 0$, because that is another neutral point, where $v(n) = 1$ for all n. Embree and Trefethen found the minimum at $\beta = 0.36747....$ In the course of the search they discovered that the curve recording the variation of C as a function of β is not a smooth one. The dips and humps in the curve appear to have a fractal structure, similar to itself at all scales of magnification.

A final question: What is the meaning of numbers such as $C = 1.13198824...$ and $\beta^* = 0.70258...$? Where do they come from? The number ϕ is given by a simple analytic expression, $(1 + \sqrt{5})/2$. Is there any similar formula for C or for β^*? Probably not. Embree and Trefethen point out that C is 0.4 percent greater than the fourth root of ϕ; it is even closer to four-fifths of $\sqrt{2}$, but these numerical coincidences are surely meaningless. Numbers are so plentiful that you can always find relations among them if you look hard enough, but C and β^* will probably have to stand on their own as new constants of nature, or of mathematics. Embree and Trefethen suggest calling 1.13198824... Viswanath's constant.

Bibliography

Bougerol, Philippe, and Jean Lacroix. 1985. *Products of Random Matrices with Applications to Schrödinger Operators*. Boston: Birkhäuser.

Embree, Mark, and Lloyd N. Trefethen. 1997. Growth and decay of random Fibonacci sequences. *Proceedings of the Royal Society of London A*. (To appear.) http://www.comlab.ox.ac.uk/oucl/users/nick.trefethen/embree.ps

Furstenberg, H., and H. Kesten. 1960. Products of random matrices. *The Annals of Mathematical Statistics*. 31:457–469.

Furstenberg, Harry. 1963. Noncommuting random products. *Transactions of the American Mathematical Society*. 108:377–428.

Gardner, Martin. 1994. Notes of a fringe-watcher: The cult of the golden ratio. *Skeptical Inquirer* 18:243–247.

Gould, H. W. 1981. A history of the Fibonacci Q-matrix and a higher-dimensional problem. *The Fibonacci Quarterly* 19:250–257.

Viswanath, Divakar. 1998. Random Fibonacci sequences and the number 1.13198824. *Mathematics of Computation*. (To appear.) http://msri.org/people/members/divakar/papers/rfib.ps