different model. I think that intrinsic noise with $\sigma_e^2 > 0$ leads us immediately to a stochastic world, and if $\sigma_e^2 = 0$ but $\sigma_\varepsilon^2 \neq 0$ and not small, as is often the case in economics, distinguishing between iid and low-dimensional white chaos will be extremely difficult.

This leads to the question of whether the real world, such as an actual economy, contains chaos. Chatterjee and Yilmaz take the position that it is ubiquitous, finding examples in "such diverse fields as physiology, geology, ..., economics ..." and "theoretical models of population biology." There are also theoretical models in economics that produce chaos, but that does not imply that it occurs in practice. I would prefer to suggest the opposite view that there is *no* evidence of chaos outside of laboratories. My reason is that there exists no statistical test, that I know of, that has chaos as its null hypothesis. There are plenty of tests that have as a null H_0:iid (or linear) and are designed to have power against chaos. However, as is well known by all statisticians, if one rejects the null a specific alternative hypothesis cannot be accepted. If a null of linearity or iid is rejected, one cannot accept (white) chaos, as nonlinear stochastic models are also appropriate. For example, the test (based on the correlation dimension) by Brock, Dechert and Scheinkman (1987) (the BDS test) that was applied in Brock and Sayers (1988) often finds evidence of nonlinearity but not of chaos. Until a property P can be found that holds *only* for chaos and not for stochastic series, and a test is based on P with chaos as the null, can there be a suggestion that chaos is found in the real world.

Finally, I would suggest that bifurcation and fractional integrated models are irrelevant for the main topic discussed in the articles, but space limitations prevent me from expanding on this point.

In conclusion, I think that scientists working on the area of chaos are doing a disservice to this important area of research by overselling its relevance, by trying to equate it with randomness and by using concepts (such as probability) that are unnecessary and only lead to confusion. The techniques being developed for analysis of chaotic processes, such as the BDS test or estimates of the Lyapunov exponent, or methods of forecasting using $\sigma_e^2 = 0$, are potentially powerful and useful when applied to truly stochastic, real-world series. There is a need for statistical methods to investigate the properties of these techniques in this case, and this, in my opinion, is the natural link between chaos and statistics.

ACKNOWLEDGMENT

This study was supported by NSF Grant SES-89-02950.

Comment: Randomness in Complex Systems

David Griffeath

1. WHAT IS RANDOM?

Professors Berliner, Chatterjee and Yilmaz are to be commended for their thoughtful overviews of the recent explosion in experimental and theoretical research on chaos. They identify a host of challenging statistical questions fundamental to the subject and make timely appeals for the readership of *Statistical Science* to join the fray. Over the past decade, I have tried to track the major currents of chaos, studying many of the articles and books mentioned in the authors' fine reference list. I strongly urge others to peruse those sources and seek out a few.

Berliner and Chatterjee and Yilmaz note that the term "chaos" is not used in a consistent manner by the scientific community; for example, there is no universally accepted mathematical definition. In my experience, the word means so many different things to different people that it threatens to become scientifically dangerous. Apparently, Bernoulli shift, the most basic stochastic process, is deemed chaotic. But how is it distinguished from those strange attractors, delicately perched on the boundary between order and randomness, that have dramatically captured the imagination of both scientists and the general public? The phenomenology of chaos is leaving its mark across a broad spectrum of contemporary culture: from physics to philosophy to recreational computing to textile design. At the hairdresser, I discovered an article in a summer issue of *Gentleman's Quarterly* linking mathematical chaos, Silicon Valley nerds and late-

David Griffeath is Professor of Mathematics, University of Wisconsin, Madison, Wisconsin 53706.

60's psychedelia. (I resist including the precise citation among my references, but curious readers can certainly track it down.) As one colleague recently remarked, "Before long, the Mandelbrot Set will be showing up on toilet paper."

There is another c-word that I think more accurately expresses a general framework for the research Berliner and Chatterjee and Yilmaz describe. That word is complexity. Roughly speaking, a complex system is an evolution generated by simple mathematical rules or physical principles that exhibits complicated, unpredictable behavior. Because many complex systems are deterministic, one should add the proviso: unpredictable if one does not know the transition mechanism. Without a doubt, the highest profile examples at present are iterates of nonlinear maps in low-dimensional Euclidean space. Gleick's bestseller and the admirable marketing ability of Mandelbrot, Peitgen, Barnsley and others have secured center stage for this important modeling framework. But many modeling environments are being developed in parallel: quantum and statistical mechanics, percolation, neural nets and artificial life, among others. Cross-fertilization between areas is quite common. To cite just one important example, universality of the Feigenbaum constant was anticipated by the celebrated renormalization group theory of Kadanoff and Wilson. A pervasive signature of complexity is the appearance of phase transitions in parameterized families of models: the system undergoes sudden dramatic changes as parameters vary from one region of the phase diagram to another. Iterates of maps constitute one small part of this big picture.

What is the connection between chaos and statistics? As Professor Berliner has suggested in his provocative Section 2.2, it is almost irresistible to describe the unpredictable trajectories of nonlinear dynamics—sequences of values without apparent rhyme or reason—by the word "random." Indeed, since the popularization of chaos, scientists and lay people are using the r-word quite indiscriminately to describe just about any phenomenon they do not understand. Another short step leads to a distressing "holistic" philosophy of chaos that abandons quantitative analysis altogether. If a butterfly in China can really cause a tornado in the Midwest, maybe we should just let it flow. Life is so complicated!

Trained as a mathematical probabilist, I find this trend rather disconcerting. It is easy to forget that probability theory gives mathematically precise meaning to notions like "a completely random sequence of 0's and 1's" and a "completely random sequence of reals in [0,1]." Almost every such sequence (in the sense of measure theory) matches a huge dossier of character traits, such as the central limit theorem and law of the iterated logarithm. The classical theory of iid sequences can be viewed as a remarkably successful effort to quantify ideal randomness. More recently, researchers in stochastic processes have assembled an enormous menagerie of random evolutions, typically Markovian, incorporating various dependencies and structural features appropriate to the specific context. Diffusions, branching and queueing processes, for example, generalize the notion of randomness. But these are all precise mathematical models with exact quantitative properties inherited from the iid building blocks used in their construction. I think it is only useful to call chaotic trajectories random, insofar as precise numerical connections can be made between the statistics of these trajectories and the statistics of stochastic sample paths. A major objective for the next generation of probabilists and statisticians, in my opinion, should be the discovery and analysis of such connections.

As a purist, the claim at the end of Section 2 in Berliner makes no sense to me. The law of any deterministic process concentrates on a single sample path, so every event has measure 0 or 1. To debunk the probability statement, one could start from a "typical" x_0 and consider the quantity

$$\text{(1)} \qquad \frac{1}{n+1} \sum_{i=10^{100}}^{10^{100}+n} 1_{\{Y_i=1, Y_{i+1}=0\}}$$

for n large (but not necessarily as large as 10^{100}), or perhaps

$$|I|^{-1} \int_I P_x(Y_{10^{100}} = 1, Y_{10^{100}+1} = 0)\, dx,$$

where I is a small interval around x_0 (but not too small!). To show that these quantities are close to 0.25, one uses a conjugacy transformation, discovered by Ulam, that maps the $\alpha = 4$ logistic map to Bernoulli trials. This connection reveals the invariant arc sine law density mentioned in both Berliner and Chatterjee and Yilmaz; it is a beautiful example of precise quantitative interplay between probability theory and the iterates of a (very special) nonlinear map. Motivated by ergodic theory, we might say that a deterministic dynamical system matches a stationary stochastic sequence X_n, if all its limits of time averages such as (1) agree with the corresponding finite dimensional distributions of the random process. From a practical point of view, this seems like a reasonable requirement. But "truly" random sequences like typical Bernoulli 0's and 1's have much more stringent structure; so, it is tempting to try to exclude all aberrations from the ideal. Such attempts lead to the beautiful but decidedly esoteric subject of algorithmic complexity discussed briefly in Chatterjee and Yilmaz. See Cover, Gacs and Gray (1989) for a recent scholarly reference that stresses pioneering contributions of Kolmogorov. Bennett and Gardner (1979) give a colorful popular account focusing on Chaitin's completely random sequence Ω: if only one knew this remarkable sequence explicitly, then most of the world's (mathematical) problems would be solved and we could all retire early.

2. RANDOM NUMBERS

Coming back down to earth, suppose one wants to implement Monte Carlo simulation of a specific, messy stochastic process on a computer. I have never understood why analog sources of random noise are seldom used. Maybe it is difficult to design a real-world noise source that conforms to all our expectations. If so, maybe that should tell us something. In any case the usual approach involves a deterministic digital pseudo-random number algorithm, in essence a chaotic iterated map. Surprisingly, neither survey article mentions these complex dynamics that statisticians have been using for decades. Most are based on linear congruences. Moreover they are truly unpredictable because their code is typically undocumented and buried deep within a compiler. The staggering range of complex behaviors now being discovered and classified should make us more wary then ever about the ability of chaotic systems to capture ideal statistical independence. To illustrate this point, let me describe four specific algorithms for ostensible randomness, two pretty terrible, two not so bad.

First, consider the rule $x_n = \text{rnd}^n(x)$, where rnd is Microsoft's random number generator for the BASIC interpreter that was bundled with the first generation of IBM PC's and XT's in the early 1980s. If we form the two-dimensional system $(x_n, y_n) = (\text{rnd}^{2n}(x), \text{rnd}^{2n+1}(x))$, scale $[0, 1]^2$ to fill a 320×200 array, and plot the first 10,000 points, the result is shown in Figure 1. Evidently, this generator had serious pair correlation problems. It was bad enough to fail almost any standard statistical test, but visualization delivers the verdict dramatically. For the record, rnd was upgraded and stripes disappeared when the IBM AT was released.

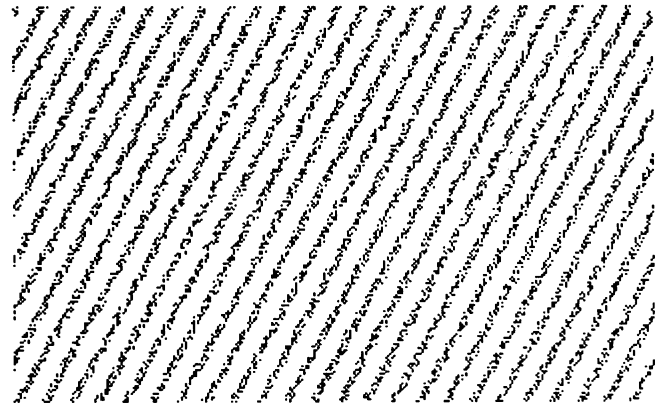

FIG. 1. *The first MS BASIC random number generator.*

Those were the early days of microcomputers, you say. Consider next the bit generator random(2) for the current version of Turbo PASCAL, a compiler that has sold hundreds of thousands of copies. Moving across each row, from top to bottom of a 256×240 array, color a pixel black whenever a bit is 1 but leave the pixel white when the bit is 0. Figure 2 shows the result. The pathology here is more subtle; I suspect that many simple statistical tests might miss the glitches made evident by visualization. In fact, if we change the row size to anything other than a power of 2, the complex patterns disappear from view. Presumably, the folks at Borland are using an algorithm based on arithmetic mod 2^p for some p. This choice is especially convenient for binary shift operations, but notoriously ill-advised.

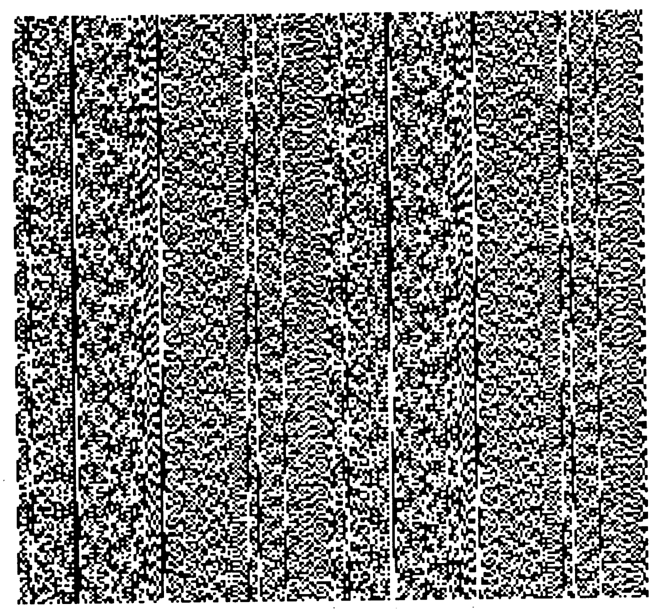

FIG. 2. *The current TURBO PASCAL random(2) generator.*

In light of such horror stories, Park and Miller (1988) have proposed a so-called minimal standard 32-bit generator $x_n = f^n(x)$, where

$$f(x) = 16807x \bmod 2147483647. \tag{2}$$

They urge their colleagues to adopt this algorithm instead of using those supplied with compilers or attempting to invent their own. They argue that (2) has the blessing of D. Knuth, the undisputed authority on design of pseudo-random number generators, and that it is "demonstrably random," whatever that means. I have no doubt that their rule is far superior to the vast majority of generators in current use. But there is something unsettling about the notion of any officially sanctioned standard. What if all Monte Carlo simulations during the next century are based on (2) and then someone uncovers a terrible flaw?

My favorite candidate for a deterministic random sequence of 0's and 1's is generated by Wolfram's *rule 30*: the one-dimensional cellular automaton (CA) with update rule

$$(3)\quad x_{n+1}(i) = x_n(i-1)\mathrm{XOR}\big(x_n(i)\mathrm{OR}\, x_n(i+1)\big).$$

Starting from a doubly infinite (symmetric) Bernoulli sequence $x_0(i)$; $i \in \mathbb{Z}$, one can prove that the "time slice at the origin" ($x_n(0)$; $n \geq 0$) is also iid. This is a special property enjoyed by a very few CA rules. Its proof relies on a clever application of XOR to both sides of (3), a trick that Wolfram (1986) attributes to Milnor. Of course, the "PASCAL's triangle mode 2" CA

$$(4)\qquad x_{n+1}(i) = x_n(i-1)\ \mathrm{XOR}\ x_n(i+1),$$

which is linear, satisfies the same property; at first glance, it seems that one is merely trading a random source for random output. The ergodic theorem implies that both rules should generate random time slices starting from almost every x_0. The difference between (3) and (4), due to the nonlinear OR operation in rule 30, is that finite configurations of 1's appear to behave typically in the former case, whereas they certainly do not in the latter. Starting from a single 1 at the origin, linear rule (4) generates all 0's at the origin. One checkerboard subgraph has all 0's, while the other makes a Sierpinski lattice, the discrete version of a fractal known as Sierpinski's gasket. In contrast, rule 30 started from a single 1 generates a stream of bits at 0 that performs quite well on all the standard statistical tests for randomness. Details and lots of pictures appear in Wolfram (1986). Even though the problems posed there seem beyond the reach of rigorous analysis for the time being, I think Wolfram's paper is one of the best examples of experimental mathematics in recent years.

3. A FEW WORDS ABOUT RANDOM CELLULAR AUTOMATA

Professors Chatterjee and Yilmaz have alluded briefly to cellular automata, an area that shares many themes and objectives with chaotic iterates of maps. Since that is my field of expertise, I would like to make some remarks about recent exciting developments, both theoretical and experimental.

My current work with complex systems focuses on the ability of CA dynamics to self-organize out of random initial conditions ("primordial soup"), an interest that evolved from more than a decade of research in the area of probability known as interacting particle systems (IPS). These are Markov processes, typically in continuous time, the states of which are lattice-valued configurations. For example, consider an infinite collection of continuous time random walks moving independently around the d-dimensional integer lattice, except for the times when a particle tries to jump to a site already occupied by another. If the two particles merge at such jump times, we have coalescing random walks; if such jumps are suppressed, we get simple exclusion. Originally motivated by statistical physics, IPS has been one of the most vital areas of probability over the past 25 years. There is now a rich collection of rigorous results for additive and cancellative particle systems, such as coalescing random walks and simple exclusion, the stochastic analogs of (4). This family of systems amenable to theorems and proofs also includes Harris' contact process, an important spatial contagion dynamic that exhibits phase transition, and the voter model, a selectively neutral competition process that clusters for $d = 1, 2$, but has stable equilibria for $d \geq 3$. Authoritative references are Liggett (1985) and Durrett (1991).

About 5 years ago, together with my colleague Bob Fisch, I discovered a self-organized equilibrium of random spirals in a simple two-dimensional multi-type IPS called the cyclic particle system (CPS). Its dynamics, revealed by the Cellular Automaton Machine (CAM) of Toffoli and Margolus (1987), are reminiscent of the B–Z reaction mentioned in the introduction of Chatterjee and Yilmaz. Unfortunately, we were able to prove very little about the CPS due to its nonlinear transition mechanism. However, CAM visualization suggested a near determinism to its periodic waves, as if the Markov process were a small perturbation of a CA. Conventional wisdom maintains that stochastic processes are easier to analyze than deterministic ones, because randomness erases memory and induces smoother, averaged behavior. In some highly degenerate instances, though, nonlinear dynamics are stable with respect to noise, and the random component only muddies the water. Our CPS is a case in point. There is a corresponding cyclic cellular automaton (CCA) with very similar qualitative features starting from random initial states. It is actually easier to analyze rigorously, because one can use logic to trace back its trajectories to the random initial state at time 0, about which one knows a great deal from classic probability and percolation theory. In other words, there is a hybrid area of dynamical systems called deterministic dynamics from random initial states. Because the process starts randomly, it is random at all times, (i.e., a stochastic process that avoids the sticky issue of identifying "typical" x_0). Once the process gets going, there is nothing truly unpredictable about its evolution.

I believe that deterministic dynamics from random initial states have much to contribute to the theory of chaos and complex systems. For iterates of maps, the monograph by Lasota and Mackey (1985) is a beginning. For cellular automata, we have now identified a three-parameter universe of CCA rules, and a corresponding universe of Greenberg–Hastings (GH) rules, that each exhibit an elaborate phase diagram featuring several varieties of rather spectacular self-organization, all starting from complete randomness. Experimental findings are reported in Fisch, Gravner and Griffeath, (1991); companion freeware (Fisch and Griffeath) for IBM compatibles with VGA graphics is available by E-mail request to griffeath@math.wisc.edu. Most gratifying is the fact that such CA models for excitable media also succumb to rigorous mathematical analysis. When these complex systems are subjected to a suitable scaling limit, several numerical cutoffs can be calculated either precisely or with surprising numerical accuracy, and certain geometrical features (such as the asymptotic shape of spiral wave fronts) can also be computed exactly. For a popular account of the very latest developments, see Durrett (1988).

Let me conclude my discussion with some comments about CA experimentation and visualization. There is a widespread distrust of CA modeling that rivals the skepticism about chaos in general. To some extent, this is a reaction to all the pop-culture hype, but it is also rooted in the historical dominance of continuous mathematics. Computers are changing tradition very rapidly as the world converts to digital. Parallel processing is still in its infancy; basic research on CA dynamics will help identify organizational principles of digital dynamics, programming and control. There will always be a role for the venerable and powerful methodology of partial differential equations; but now that we have the appropriate tools to study them, CA rules should begin to enjoy a comparable status. Naturally, hybrids will evolve as they prove useful or illuminating. Already there is substantial experimental work on coupled lattice maps, an amalgam of iterates of maps and cellular automata.

Dynamic visualization is an invaluable tool for the study of complexity. Observation of a complex system can tell you the answer to a basic question like, "What does this rule do?" At the very least, you avoid wasting time and energy because you have the wrong movie in your head. Not infrequently, by watching an evolution—especially in real-time—one identifies key structural properties of the model. On a very good day that experience might even lead to a proof of something interesting. For readers who would like to experiment with CA visualization, the best way to start is with Rudy Rucker's CA Lab (1989). His software is accompanied by a delightful and well-written 264-page book on the history, philosophy, theory and use of cellular automata.

When experimentalists want to incorporate randomness into a CA rule, say either on the CAM or in CA Lab, they cannot waste time flipping coins or even using the minimal standard (2). Instead, they tap an invisible background plane of lattice gas (i.e., billiard ball) particles or some nonlinear chaotic CA that looks like a TV set tuned to a nonexistent channel. Our vision is remarkably good at finding patterns in dynamic noise; so, if the noise source passes the eyeball test, it is good enough. They find that complex systems comprised of many locally interacting components are remarkably robust with respect to random input; basic features are typically impervious to fine statistical details. For those of us raised on the ideology that chaotic deterministic maps are a bastardization of ideal randomness, for use in Monte Carlo simulation only when all else fails, witnessing this systemic insensitivity can be quite enlightening. Perhaps independent trials are merely the probabilist's fantasy, a mathematically tractable prototype that captures the essential aspects of a vast domain of complex deterministic dynamics. Ultimately, I feel that this perspective lends greater significance to probability and statistics as a whole.

ACKNOWLEDGMENTS

I thank my colleagues Charlie Bennett, Ted Cox, Rick Durrett and Bob Fisch for many stimulating discussions over the years related to issues of randomness and complexity.