

SEEING AROUND CORNERS

The new science of artificial societies suggests that real ones are both more predictable and more surprising than we thought. Growing long-vanished civilizations and modern-day genocides on computers will probably never enable us to foresee the future in detail—but we might learn to anticipate the kinds of events that lie ahead, and where to look for interventions that might work

BY JONATHAN RAUCH



In about A.D. 1300 the Anasazi people abandoned Long House Valley. To this day the valley, though beautiful in its way, seems touched by desolation. It runs eight miles more or less north to south, on the Navajo reservation in northern Arizona, just west of the broad Black Mesa and half an hour's drive south of Monument Valley. To the west Long House Valley is bounded by gently sloping domes of pink sandstone; to the east are low cliffs of yellow-white sedimentary rock crowned with a mist of windblown juniper. The valley floor is riverless and almost perfectly flat, a sea of blue-gray sagebrush and greasewood in sandy reddish soil carried in by wind and water. Today the valley is home to a modest Navajo farm, a few head of cattle, several electrical transmission towers, and not much else.

Yet it is not hard to imagine the vibrant farming district that this once was. The Anasazi used to cultivate the valley floor and build their settlements on low hills around the valley's perimeter. Remains of their settlements are easy to see, even today. Because the soil is sandy and the wind blows hard, not much stays buried, so if you leave the highway and walk along the edge of the valley (which, by the way, you can't do without a Navajo permit), you frequently happen upon shards of Anasazi pottery, which was eggshell-perfect and luminously painted. On the site of the valley's eponymous Long House—the largest of the ancient settlements—several ancient stone walls remain standing.

Last year I visited the valley with two University of Arizona archaeologists, George Gumerman and Jeffrey Dean, who between them have studied the area for fifty or more years. Every time I picked up a pottery shard, they dated it at a glance. By now they and other archaeologists know a great deal about the Anasazi of Long House Valley: approximately how many lived here, where their dwellings were, how much water was available to them for farming,

and even (though here more guesswork is involved) approximately how much corn each acre of farmland produced. They have built up a whole prehistoric account of the people and their land. But they still do not know what everyone would most like to know, which is what happened to the Anasazi around A.D. 1300.

"Really, we've been sort of spinning our wheels in the last eight to ten years," Gumerman told me during the drive up to the valley. "Even though we were getting more data, we haven't been able to answer that question." Recently, however, they tried something new. Unable to interrogate or observe the real Long House Valley Anasazi, they set about growing artificial ones.

MR. SCHELLING'S NEIGHBORHOOD

Growing artificial societies on computers—*in silico*, so to speak—requires quite a lot of computing power and, still more important, some sophisticated modern programming languages, so the ability to do it is of recent vintage. Moreover, artificial societies do not belong to any one academic discipline, and their roots are, accordingly, difficult to trace. Clearly, however, one pioneer is Thomas C. Schelling, an economist who created a simple artificial neighborhood a generation ago.

Today Schelling is eighty years old. He looks younger than his age and is still active as an academic economist, currently at the University of Maryland. He and his wife, Alice, live in a light-filled house in Bethesda, Maryland, where I went to see him one day not long ago. Schelling is of medium height and slender, with a full head of iron-gray hair, big clear-framed eyeglasses, and a mild, soft-spoken manner. Unlike most other economists I've dealt with, Schelling customarily thinks about everyday questions of collective organization and disorganization, such as lunch-




Figure 1.

room seating and traffic jams. He tends to notice the ways in which complicated social patterns can emerge even when individual people are following very simple rules, and how those patterns can suddenly shift or even reverse as though of their own accord. Years ago, when he taught in a second-floor classroom at Harvard, he noticed that both of the building's two narrow stairwells—one at the front of the building, the other at the rear—were jammed during breaks with students laboriously jostling past one another in both directions. As an experiment, one day he asked his 10:00 A.M. class to begin taking the front stairway up and the back one down. "It took about three days," Schelling told me, "before the nine o'clock class learned you should always come up the front stairs and the eleven o'clock class always came down the back stairs"—without, so far as Schelling knew, any explicit instruction from the ten o'clock class. "I think they just forced the accommodation by changing the traffic pattern," Schelling said.

In the 1960s he grew interested in segregated neighborhoods. It was easy in America, he noticed, to find neighborhoods that were mostly or entirely black or white, and correspondingly difficult to find neighborhoods where neither race made up more than, say, three fourths of the total. "The distribution," he wrote in 1971, "is so U-shaped that it is virtually a choice of two extremes." That might, of course, have been a result of widespread racism, but Schelling suspected otherwise. "I had an intuition," he told me, "that you could get a lot more segregation than would be expected if you put people together and just let them interact."

One day in the late 1960s, on a flight from Chicago to Boston, he found himself with nothing to read and began doodling with pencil and paper. He drew a straight line and then "populated" it with Xs and Os. Then he decreed that each X and O wanted at least two of its six nearest neighbors to be of its own kind, and he began moving them around in ways that would make more of them content with their neighborhood. "It was slow going," he told me, "but by the time I got off the plane in Boston, I knew the results were interesting." When he got home, he and his eldest son, a coin collector, set out copper and zinc pennies (the latter were wartime relics) on a grid that resembled a checkerboard. "We'd look around and find a penny that wanted to move

and figure out where it wanted to move to," he said. "I kept getting results that I found quite striking."

To see what happens in this sort of artificial neighborhood, look at Figure 1, above, which contains a series of stills captured from a Schelling-style computer simulation created for the purposes of this article. (All the illustrations in the article are taken from animated artificial-society simulations that you can view online at the magazine's Web site, at www.theatlantic.com/rauch.) You are looking down on an artificial neighborhood containing two kinds of people, blue and red, with—for simplicity's sake—no blank spaces (that is, every "house" is occupied). The board wraps around, so if a dot exits to the right, it reappears on the left, and if it exits at the top, it re-enters at the bottom.

In the first frame blues and reds are randomly distributed. But they do not stay that way for long, because each agent, each simulated person, is ethnocentric. That is, the agent is happy only if its four nearest neighbors (one at each point of the compass) include at least a certain number of agents of its own color. In the random distribution, of course, many agents are unhappy; and in each of many iterations—in which a computer essentially does what Schelling and his son did as they moved coins around their grid—unhappy agents are allowed to switch places. Very quickly (Frame 2) the reds gravitate to their own neighborhood, and a few seconds later the segregation is complete: reds and blues live in two distinct districts (Frame 3). After that the border between the districts simply shifts a little as reds and blues jockey to move away from the boundary (Frame 4).

Because no two runs begin from the same random starting point, and because each agent's moves affect every subsequent move, no two runs are alike; but this one is typical. When I first looked at it, I thought I must be seeing a model of a community full of racists. I assumed, that is, that each agent wanted to live only among neighbors of its own color. I was wrong. In the simulation I've just described, each agent seeks only two neighbors of its own color. That is, these "people" would all be perfectly happy in an integrated neighborhood, half red, half blue. If they were real, they might well swear that they valued diversity. The realization that their individual preferences lead to a collective outcome indistinguishable from thoroughgoing




Figure 2.

racism might surprise them no less than it surprised me and, many years ago, Thomas Schelling.

In the same connection, look at Figure 2, above. This time the agents seek only one neighbor of their own color. Again the simulation begins with a random distribution (Frame 1). This time sorting proceeds more slowly and less starkly. But it does proceed. About a third of the way through the simulation, discernible ethnic clusters have emerged (Frame 2). As time goes on, the boundaries tend to harden (Frames 3 and 4). Most agents live in areas that are identifiably blue or red. Yet these “people” would be perfectly happy to be in the minority; they want only to avoid being completely alone. Each would no doubt regard itself as a model of tolerance and, noticing the formation of color clusters, might conclude that a lot of other agents must be racists.

Schelling’s model implied that even the simplest of societies could produce outcomes that were simultaneously orderly and unintended: outcomes that were in no sense accidental, but also in no sense deliberate. “The interplay of individual choices, where unorganized segregation is concerned, is a complex system with collective results that bear no close relation to the individual intent,” he wrote in 1969. In other words, even in this extremely crude little world, knowing individuals’ intent does not allow you to foresee the social outcome, and knowing the social outcome does not give you an accurate picture of individuals’ intent. Furthermore, the godlike outside observer—Schelling, or me, or you—is no more able to foresee what will happen than are the agents themselves. The only way to discover what pattern, if any, will emerge from a given set of rules and a particular starting point is to move the pennies around and watch the results.

Schelling moved on to other subjects in the 1970s. A few years later a political scientist named Robert Axelrod (now at the University of Michigan) used a computer simulation to show that cooperation could emerge spontaneously in a world of self-interested actors. His work and Schelling’s work and other dribs and drabs of research hinting at simulated societies were, however, isolated threads; and for the next decade or more the threads remained ungathered.

SUGARSCAPE AND BEYOND

I have office space at The Brookings Institution, which is the oldest of Washington’s think tanks. Since it is one of the more staid places in town, it was probably inevitable that I would notice Joshua Epstein. Epstein is tall and portly, with a wild tuft of graying hair above each ear, a round face, and the sort of exuberant manner that brings to mind a Saint Patrick’s Day parade more readily than a Washington think tank. “No foam!” he roared, grinning, to a Starbucks server one day when we went out for coffee. “Keep your damn foam!” Anyone who notices Epstein is soon likely to encounter Robert Axtell, his collaborator and alter ego. A programming wizard with training in economics and public policy, Axtell is of medium height, quiet, and as understated as Epstein is boisterous. When he speaks, the words spill out so quickly and unemphatically that the listener must mentally insert spaces between them.

Epstein was born in New York City and grew up in Amherst, Massachusetts. His father was a logician and a philosopher of science. Nonetheless, Epstein never managed to finish high school. Instead he got into college on a piano audition and, after composing a series of chamber-music pieces, ended up switching to the study of mathematics and political economy. That led to a Ph.D. in political science in 1981 and then a position at Brookings, plus the realization that he was fascinated by mathematical models. One day in the early 1990s, when he was giving a talk about his model of arms races, he met Axtell, who was then a graduate student. He wound up bringing Axtell to Brookings, in 1992.

Not long after, Epstein attended a conference at the Santa Fe Institute—renowned as a pioneering center for research on “complexity,” the generation of spontaneous order and intricate patterns from seemingly simple rules. At Santa Fe just then a big subject was artificial life, often called A-life. “All of the work was about coral reefs, ecology, growing things that look like trees, growing things that look like flocks of birds, schools of fish, coral, and so on,” Epstein told me. “And I thought, jeez, why don’t we try to use these techniques to grow societies?” Fired up, he returned to Brookings and discussed the idea with Axtell.

There followed the inevitable napkin moment, when the

two of them sat in the cafeteria and sketched out a simple artificial world in which little hunter-gatherer creatures would move around a landscape finding, storing, and consuming the only resource, sugar. When they brought Sugarscape, as they called it, to life with the computer, they were startled to see that almost immediately their rudimentary A-society produced a skewed distribution of sugar that looked very much like the skewed distribution of wealth in human societies, even though nothing about the agents' simple behavioral rules pointed to any such outcome. For several years they built up and elaborated Sugarscape, and discovered that simple rules could produce complex social phenomena that mimicked migrations, epidemics, trade. "Every time we build one of these things, it does some shocking thing," Epstein told me. "You can make it as simple as you want, and it will do something surprising, almost certainly."

Epstein and Axtell then began applying their technique, which they called agent-based modeling, to a variety of problems and questions, and as they did so they quietly inverted a number of tenets of the more conventional varieties of social modeling. In Sugarscape, and in the other artificial societies that followed, Epstein and Axtell made their agents heterogeneous. That is, the artificial people, like real people, were different from one another. Each Sugarscape agent has its own "genetic code": a distinctive combination of metabolic rate (how much sugar each agent needs in order to stay alive), vision (how far the agent can "see" as it hunts for sugar), and so forth. This was a small move that was actually quite radical, and not just because of the daunting computational requirements. In most conventional social-science models people are assumed to be more or less the same: multiple copies of a single representative person. Even in Thomas Schelling's artificial neighborhood all the agents are alike except in color. Moreover, conventional models tend to assume that all their clonelike individuals have complete or near complete knowledge of their world. In Schelling's model unhappy agents, like the modeler himself, could survey the whole scene to find a better situation. In ordinary economic models, by the same token, people all see essentially the same big picture, so if a stock is underpriced, for example, traders will quickly spot the anomaly. Epstein and Axtell instead built models in which agents' vision and knowledge were limited; agents knew only what was going on nearby or what they "heard" from their "friends" (often a unique social network was assigned to every agent). Each agent, therefore, had unique preferences and unique knowledge.

It took me a little while to understand why in some respects this is a whole new ball game. In years of writing on economics I had grown comfortable with the sort of equation-based modeling that is common and, unquestionably, indispensable in the social sciences. The modeler looks at social patterns in the real world and tries to write equations that describe what's going on. The modeler, that is, views the

world from on high and attempts to fit it to regular lines and curves, which are then used to make predictions. A simple and elegant artificial society created by Ross Hammond brought home to me what I had been missing.

Hammond is well over six feet tall and reed thin, with a broad forehead and a pointed chin that make his face a neat triangle. When I met him, last year, he worked as an assistant to Epstein and Axtell (he has since moved on to graduate school at the University of Michigan), but he originally devised his world in 1999, for a senior thesis at Williams College. He decided to make an abstract model of social corruption. He created an artificial world populated with two kinds of agents: citizens and bureaucrats. Each of these agents has his own susceptibility to corruption and his own network of friends. Every time a citizen meets a bureaucrat, the two conduct a transaction. If they collude corruptly, both pocket a nice kickback, whereas if both behave honestly, neither gets payola. If a mismatch occurs, and only one agent is willing to cheat, the honest agent "reports" the corrupt one to an unseen policing authority.

So far the setup is conventional game theory. Less conventional is this: no agent knows exactly how many reports of corruption will land him in jail, or how many other agents are honest or corrupt, or what most other agents are doing. He knows only what has happened recently to himself and his friends. If suddenly many of them land in jail, he will assume that the cops are cracking down and will behave more honestly until the coast looks clearer. (This excludes a sprinkling of George Washingtons—agents who are incorruptibly honest.) The agents, in other words, have varying personalities and limited information, and they display what economists call "bounded rationality"—that is, they make the most rational choices they can based on that *limited* information.

Hammond had no idea what his stipulations would produce. Somewhat surprisingly, he found that within many plausible ranges of corruption payoffs, punishments, and agent characteristics, his artificial society quickly settled down into rampant honesty. But there were some plausible parameters (big payoffs and short jail terms) that produced a truly startling result. To see it, look at Figure 3, on the facing page.

This shows Ross Hammond's little A-society, a world of citizens (bureaucrats are omitted for simplicity's sake) who at any given moment can be either corrupt, honest, or in jail. Schelling's checkerboard represented a physical space; the space in Figure 3, in contrast, is purely abstract. Whether agents are near each other makes no difference. What does matter is whether in any given transaction they behave honestly or corruptly. A corrupt agent is a yellow rectangle, an honest one blue, and a jailed one red. The population at any given moment stretches along a thin horizontal ribbon one rectangle deep, so the window actually portrays society over time. Thus a long vertical blue bar

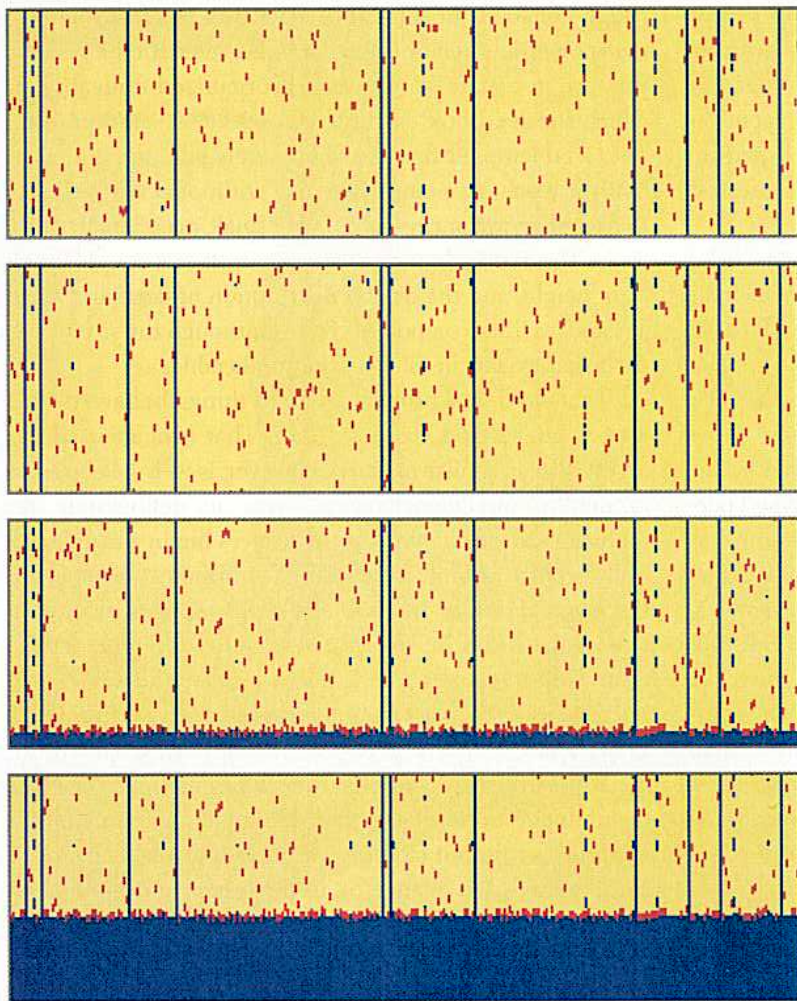


Figure 3.

represents a single agent who is incorruptible (a George Washington), whereas an isolated blue rectangle represents an agent who usually behaves corruptly but on that occasion chooses honesty.

At the top of the first frame, as the agents begin doing business, they are randomly distributed. The field is almost entirely yellow, which means that corruption is the norm. Only occasionally does a yellow agent turn blue—presumably when a bunch of his friends have gone to jail (the friends are not necessarily near him physically, and the social networks are not displayed in this demonstration). Frame 2, captured later, shows more of the same; in this society, clearly, corruption pays and is the norm. Look closely, though, a little more than halfway down Frame 2, and you may notice a vaguely horizontal cluster of reds. Just randomly, in the course of things, there has been a surge of agents going to jail. That turns out to be important for reasons that become clearer when you look at Frame 3, captured later still. Here, just above the bottom of the frame, an unusually large number of agents are again being jailed—and suddenly everyone turns blue. This predominantly corrupt society has become uniformly honest. But for how long? As the last frame shows, honesty is the new

norm. With everybody behaving honestly, there is no payoff for corruption (payoff requires two corrupt dealers), so the A-society stays honest. If the simulation continued running, it would show nothing but blue.

In the jargon, a dynamic system's sudden shift from one kind of behavior to another is typically referred to as "tipping" (and has been since well before the term became a fashionable metaphor for sudden change of whatever sort). Hammond's little world, despite its almost brutal simplicity, had tipped.

Hammond was astonished, so he ran the simulation again and again. No two runs were the same, because each began from a different random starting point, and no run was predictable in its details, because the agents' interactions, even in so simple a world, were unfathomably complicated. Sometimes the A-society would tip from corrupt to honest almost immediately; sometimes it would tip only after running for hours on end; but always, sooner or later, it tipped. The switch appeared to be inevitable, but its timing and the path taken to reach it were completely unpredictable. What was going on?

Every so often, in the course of random events, a particularly large number of corrupt agents, who happen to have particularly large networks of friends who perhaps themselves have large social networks,

will be arrested. That, Hammond figures, has a double-barreled effect: it leads a lot of agents to notice that many of their friends are under arrest, and it also increases the likelihood that they will encounter an honest agent in the next transaction. Fearing that they will meet their friends' fate, the agents behave more honestly; and in doing so they heighten yet further the odds that a corrupt agent will be nailed, inspiring still more caution about corruption. Soon—in fact, almost instantly—so many agents are behaving honestly that corruption ceases to pay, and everyone turns honest.

"There are plenty of different cities and countries that have gone from a high degree of corruption to a low degree of corruption," Hammond says. His A-society suggests that in such a transition, the fear of being caught may be at least as important as the odds of actually being caught. To test that possibility, Hammond re-ran his simulation, but this time he allowed all the agents to know not just how many of their friends were in jail but how many people were jailed throughout the whole society: in other words, the agents knew the odds of arrest as well as the police did. Sure enough, fully informed agents never got scared enough to reform. Hammond's A-society seemed to have

“grown” a piece of knowledge that many law-enforcement agencies (think of the Internal Revenue Service, with its targeted, high-profile audits) have long intuited—namely, that limited resources are often more effectively spent on fearsome, and fearsomely unpredictable, high-profile sweeps than on uniform and thus easily second-guessed patterns of enforcement.

Hammond also wondered what would happen if he made all the agents alike, instead of giving each a personality marked by a randomly varied proclivity to cheat. What if, say, all agents preferred honesty exactly half the time? The answer was that the A-society never made a transition; it stayed corrupt forever, because everyone “knew” how everyone else would behave. A social model that viewed individuals as multiple copies of the same fully informed person could thus never “see” the social transformation that Hammond found, for the simple reason that without diversity and limited knowledge, *the transformation never happens*. Given that human beings are invariably diverse and that the knowledge at their disposal is invariably limited, it would seem to follow that even societies in which unsophisticated people obey rudimentary rules will produce surprises and discontinuities—events that cannot be foreseen either through intuition or through the more conventional sorts of social science.

GROWING ZIPF’S LAW

Every so often scientists notice a rule or a regularity that makes no particular sense on its face but seems to hold true nonetheless. One such is a curiosity called Zipf’s Law. George Kingsley Zipf was a Harvard linguist who in the 1930s noticed that the distribution of words adhered to a regular statistical pattern. The most common word in English—“the”—appears roughly twice as often in ordinary usage as the second most common word, three times as often as the third most common, ten times as often as the tenth most common, and so on. As an afterthought, Zipf also observed that cities’ sizes followed the same sort of pattern, which became known as a Zipf distribution. Oversimplifying a bit, if you rank cities by population, you find that City No. 10 will have roughly a tenth as many residents as City No. 1, City No. 100 a hundredth as many, and so forth. (Actually the relationship isn’t quite that clean, but mathematically it is strong nonetheless.) Subsequent observers later noticed that this same Zipfian relationship between size and rank applies to many things: for instance, corporations and firms in a modern economy are Zipf-distributed.

Nature is replete with such mysteriously constant statistical relationships. “Power laws,” scientists call them, because the relationship between size and rank is expressed as an exponent. Earthquakes, for instance, follow Zipf-style power laws. Large earthquakes are rare, small ones are common, and the size of each event multiplied by its rank is a

rough constant. In the 1980s scientists began to believe that power-law relationships are characteristic of systems that are in a state known as self-organized criticality, of which the textbook example is a trickle of sand pouring onto a tabletop. At first the sand merely piles up, but eventually it reaches a point where any additional sand is likely to trigger an avalanche—often very small, occasionally quite large. The sand pile now maintains itself at a roughly constant height, and the overall distribution of large and small avalanches follows a power law, even though the size of any particular avalanche is always unpredictable.

That sand and other inanimate things behave in this way is interesting, even striking. That human societies might display similar patterns, however, is weird. People are (generally) intelligent creatures who act deliberately. Yet their cities, for example, sort themselves out in a mathematically regular fashion, a fact that I confirmed by glancing at the *World Almanac*. In 1950 and 1998 the lists of the top twenty-five cities in America were quite different, yet the cities’ relative sizes were almost exactly the same. The biggest city (New York in both years) was about four times as big as the fourth biggest (Los Angeles in 1950, Houston in 1998), which was about three times as big as the sixteenth biggest (New Orleans in 1950, Baltimore in 1998)—not an exact fit, but close. It was as though each city knew its permitted size relative to all the others and modulated its growth to keep the relationships constant. But, obviously, people moving to one city have not the faintest notion how their movements will affect the relative sizes of all cities. What might be going on? One plausible inference is that societies are like sand piles: complex systems whose next perturbation is unpredictable but whose behavior, viewed on a large scale and over time, follows certain patterns—patterns, moreover, that the individual actors in the system (grains of sand, human beings) are quite unaware of generating.

The day I started getting really excited about artificial societies was the day Rob Axtell mentioned that he had created artificial companies and cities, and that the companies and cities both followed Zipf’s Law. According to Axtell, conventional economic theory has yet to produce any accepted explanation for why the size distribution of firms or cities follows a power law. Perhaps, Axtell thought, the trick is not to explain Zipf’s Law but to grow it. He went to his computer and built an artificial world of diverse agents ranging from workaholics to idlers. Axtell’s workers start out self-employed but can organize themselves into firms and job-hop, always in search of whatever combination of money and leisure fits their temperament. When individuals join forces to form companies, their potential productivity rises, because of companies’ efficiency advantages. At the same time, however, as each company grows larger, each agent faces a greater temptation to slack off, collect the paycheck, and let colleagues carry the load.

The resulting universe of A-firms, Axtell found, is like the sand pile, full of avalanches small and large as firms form, prosper, grow lazy, lose talent to hungrier firms, and then shrink or collapse. As in real life, a few A-firms live and thrive for generations, but most are evanescent, and now and then a really big one collapses despite having been stable for years. Sometimes the addition of one slacker too many can push a seemingly solid firm into instability and fission; but you can't be sure in advance which firm will crumble, or when.

In such a world you might expect no regularity at all. And yet, Axtell told me, "The first time we turned it on, we got Zipf!" Despite the firms' constant churning, the distribution of large and small firms maintained the same sort of mathematical regularity seen in real life. Axtell and Richard Florida, a professor of regional economic development at Carnegie Mellon University, took the logical next step and built a model of cities, which were assumed to be basically agglomerations of firms. Same result: with no tuning or tweaking, the artificial cities unfailingly lined themselves up in a Zipf distribution and then, as a group, preserved that distribution even as particular cities grew and shrank in what looked to the naked eye like random turmoil. "All of a sudden," Florida told me, "I looked at Rob's model and it dawned on me. This creates the city *system*." The artificial cities and their artificial residents were all unknowingly locked in a competition for talent, but they could retain only so much of it before they lost ground relative to other clusters of talent. Richard Florida, to whom the Zipf distribution of cities had previously seemed a mere curiosity, infers that the Zipf relationship is much more than a pretty anomaly or a statistical parlor trick. It bespeaks the higher-order patterns into which human beings, and thus societies, unconsciously arrange themselves.

ARTIFICIAL GENOCIDE

If societies can order themselves systematically but unconsciously, it stands to reason that they can also disorder themselves systematically but unconsciously. As societies, the Balkans, Rwanda, Indonesia, and South Central Los Angeles have little in common, yet all have experienced, in recent memory, sudden and shocking transformations from a tense but seemingly sustainable communal peace to communal disorder and violence. Obviously, riots in America are in no way morally comparable to genocide in Rwanda, but what is striking in all these cases is the abruptness with which seemingly law-abiding and peaceable people turned into looters or killers. Scholars often use the metaphor of contagion in talking and thinking about mass violence, because the violence seems to spread so quickly from person to person and neighborhood to neighborhood. Yet sociologists who have studied mass behavior have learned that people in crowds and groups usually remain rational, retain their individuality, and exercise




Figure 4.

their good judgment; that is, they remain very much themselves. The illusion that some larger collective mind, or some sort of infectious hysteria, has seized control is just that: an illusion. Somehow, when communal violence takes hold, individuals make choices, presumably responding to local incentives or conditions, that make the whole society seem to have suddenly decided to turn savage. Might it be that rampant violence is no more the result of mass hysteria than the rampant segregation in Thomas Schelling's artificial neighborhood is the result of mass racism?

Figure 4, above, shows Joshua Epstein's artificial society containing two kinds of people, blues and greens. As usual in Epstein's models, each agent has his own personality—the relevant traits being, in this case, the agent's degree of privation or discontent, his level of ethnic hostility, and his willingness to risk arrest when the police are around. Also as usual, agents can "see" what is going on only in their immediate neighborhoods, not across the whole society. The agents' environment is one of ethnic tension between blues and greens; the higher the tension, the more likely it is that the agents will, in Epstein's term, "go active"—which in real life could mean looting a neighbor's store or seizing his house, but which in the current instance will mean killing him. When an agent turns red, his discontent or hatred has overcome his fear of arrest, and he has killed one randomly selected neighbor of the other color. Those are the rules. They are very simple rules.

In Figure 4 none of the agents are red. There is not enough ethnic tension to inspire them to go active, so they coexist peacefully, and indeed fill up the screen as their populations grow (they can procreate). Between Frames 1 and 2 all that happens is that blues and greens move around and occupy previously empty spaces. The situation looks safe and stable, but it is not. In Figure 5, on the following two pages, ethnic tension has increased only slightly, but that increment has shifted the society into a radically different state. In Frame 1 the randomly distributed agents have set about killing one another, so their world is awash with red dots. Shortly afterward, only a few seconds into the simulation, the population has thinned dramatically (Frame 2), with most of the agents who live in ethnically mixed zones having been

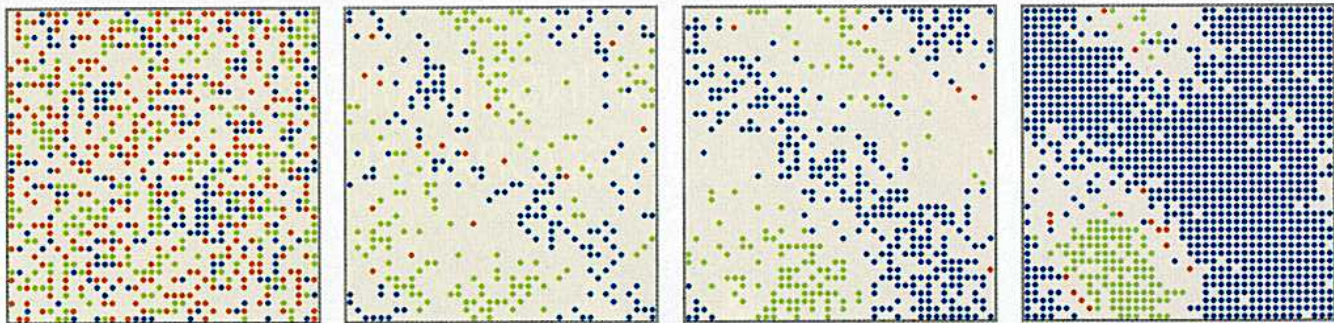


Figure 5.




Figure 6.

picked off. By Frame 3 blues and greens have separated, with violence flaring along the borders and blues predominating.

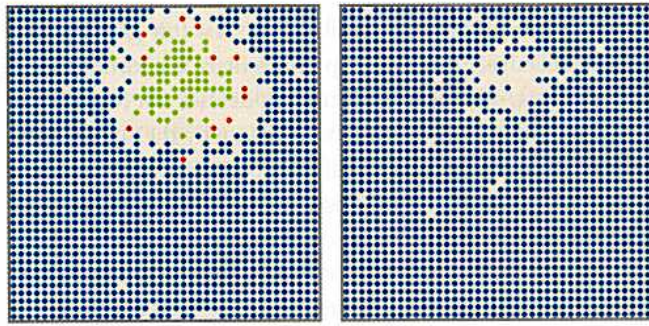
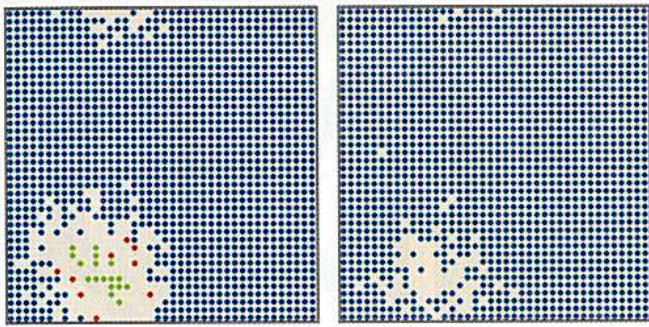
Epstein has run this simulation countless times from different random starting points, and it turns out that neither color enjoys an inherent advantage: blues and greens are equally likely to prevail, with the outcome depending on random local events that tilt the balance one way or the other. No two runs are quite alike. But all are the same in one respect: once a side has attained the upper hand, its greater numbers allow it to annihilate the other side sooner or later. In Frame 4 greens are confined to a single ethnic enclave (the bottom of the frame wraps around to join the top), where they huddle in beleaguered solidarity as blues continue to nibble at them. The rest of the story, in Frames 5 and 6, speaks for itself.

Epstein then added a third element, one that might be of special interest to the United Nations: cops, or, if you prefer, peacekeepers. In Figure 6, above, cops are represented by black dots. Like other agents, they can “see” only in their immediate vicinity. Their rule is to look around for active agents and put them in jail. The less hotheaded agents will behave peaceably when a cop is nearby, so as to avoid arrest. The result is a markedly different situation.

In Frame 1 agents and cops are scattered randomly, and the bolder agents (in red) are setting upon their victims. When they commit murder near a cop, the agents go to jail. Even so, the cops are initially overwhelmed by the sheer quantity of violence, and in Frames 2 and 3 an enclave of embattled greens forms, just as before. Now, however, there is an important difference: the enclave is stable. Once it has

dwindled to a certain size, the cops are able to contain the violence by making arrests along the border. As long as the cops stay in place, the enclave is safe. But what if the cops are withdrawn? The result is exactly the same as what happened when peacekeepers abandoned enclaves in Bosnia and Rwanda. In Frame 4 the cops have all departed. Again, Frames 5 and 6 speak for themselves.

I don’t think I’m alone in finding this artificial genocide eerie. The outcome, of course, is chilling; but what is at least as spooky is that such complicated—to say nothing of familiar—social patterns can be produced by mindless packets of data following a few almost ridiculously simple rules. If I showed you these illustrations and told you they represented genocide, you might well assume you were seeing a schematic diagram of an actual event. Moreover, the model is designed without any element of imitation or communication, so mass hysteria or organized effort is literally impossible. No agent is knowingly copying his peers or following the crowd; none is consciously organizing a self-protective enclave. All the agents are separately and individually reacting “rationally”—according to rules, in any case—to local conditions that the agents themselves are rapidly altering. As hotheads begin to go active, the odds that any one misbehaving agent will be arrested decline, emboldening more-timid agents nearby to act up, reducing the odds of arrest still further, emboldening more agents, and so on. As in real life, the violence, once begun, can spread rapidly as cops are overwhelmed in one neighborhood after another. Although the agents are atomized and disorganized, the violence is communal and coherent. It has form and direction and even a sort of malevolent logic.



At a Brookings conference last year, where Epstein presented his artificial genocide, Alison Des Forges was in attendance. Des Forges, a senior adviser to Human Rights Watch Africa, is one of the world's leading authorities on the Rwandan genocide of 1994. After the session I asked her what she made of Epstein's demonstration. Neither she nor anyone else, Epstein included, believes that an array of little dots explains the Rwandan cataclysm or any other real-world event; the very notion is silly. What the simulation did suggest to Des Forges is that disparate social breakdowns, in widely separated parts of the world, may have common dynamics—linking Rwanda, for instance, to other horrors far away. She also told me that Epstein's demonstration reminded her of Hutu killers' attack on Tutsis who had gathered on a Rwandan hilltop: the torches, the fires, the killing working its way up the hill.

CYBER-ANASAZI

In 1994 Epstein went back to the Santa Fe Institute, this time to lecture on Sugarscape. He told me, "I came to a run in the Sugarscape that we called the Protohistory, which was really this made-up toy history of civilization, where it starts with some little soup of agents and they go to peaks on the Sugarscape and coalesce into tribes and have lots of kids and this forces them down in between the peaks and they smash into the other tribe and they have all this assimilation and combat and all this other stuff. And I showed that toy history to this typically unlikely Santa Fe collection of archaeologists and biologists and physicists, and I said, 'Does this remind anyone of anything real?' And a hand shot up, and it was George Gumerman's hand. I had

never met George. And he said, 'It reminds me of the Anasazi.' I said, 'What the heck is that?' And he told me the story of this tribe that flourished in the Southwest and suddenly vanished. And why did they suddenly vanish? I thought, That's a fascinating question."

The greatest challenge for A-society researchers is to show that their wind-up worlds bear on anything real. Epstein asked Gumerman if he had data on the Anasazi, and Gumerman replied that there were lots of data, data covering a span of centuries and recording, year by year, environmental conditions, settlement patterns, demographic trends, and more. "I thought, jeez," Epstein says, "if there's actual data, maybe we can actually reconstruct this civilization computationally. I came back all excited and told Rob. We built this terrain in a computer and we literally animated this entire history, looking down on it as if it were a movie. We said, Okay, that's what really happened. Let's try to grow that in an agent-based model. Let's create little cyber-Anasazi and see if we can equip them with rules for farming, moving, mating, under which you just leave them alone with the environment changing as it truly did, and see if they reproduce—grow—the true, observed history."

Gumerman and Jeffrey Dean (and several other scholars who joined in the effort) were equally interested, for reasons of their own. Some scholars believed that drought and other environmental problems caused the Anasazi to leave; others blamed marauders or internecine warfare or disease or culture, as well as drought. The argument had waxed and waned ever since the 1920s. "We've thought the environment was important," Gumerman told me, "and other archaeologists said they didn't think it was that important, and that's been the level of argument until now." The prospect of growing artificial Anasazi in cyberspace suggested a new way to get some traction on the question.

So they created a computerized replica of the Long House Valley environment from A.D. 800 to A.D. 1350 and populated it with agents—in this case, digital farmers. Each agent represents a household and is given a set of what the scholars believed to be realistic attributes: family size, life-spans, nutritional needs, and so on. Every year each artificial household harvests the corn on its land during the growing season and draws down its stocks in the winter. If a household's land produces enough corn to feed the family, the family stays and farms the same land again the next year; if the yield is insufficient, the family moves to the nearest available plot that looks promising and tries again; if the family still cannot eke out sustenance, it is removed from the simulation. I have simplified the parameters, which allow for the formation of new households, the birth of children, and so on. Still, the rules are fairly straightforward, basically directing the artificial Anasazi to follow the harvest and to leave or die off if the land fails to support them.

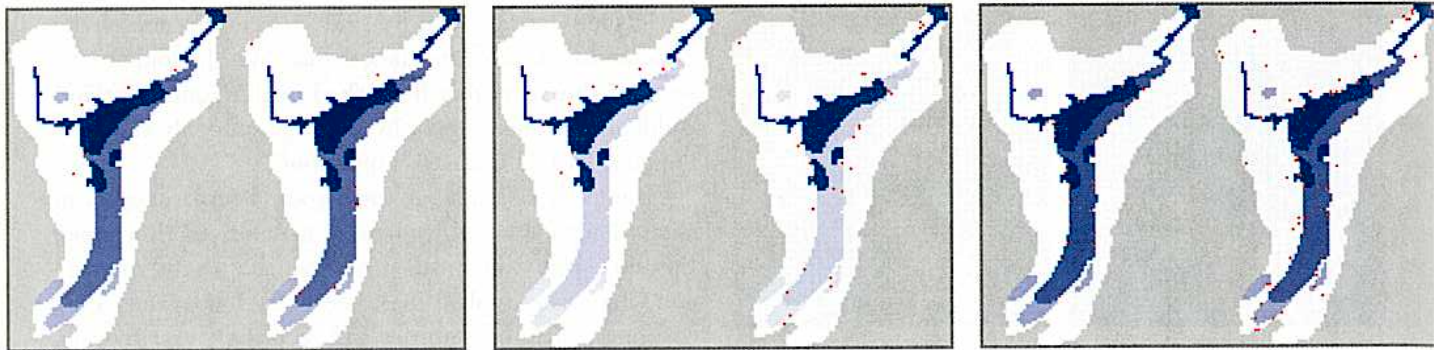


Figure 7.

To see what happens, look at Figure 7, above. You are looking down, as if from a helicopter, on paired images of Long House Valley starting in the year 800. Within the valley blue zones represent places where water is available for farming (darker blue means more water). In both images the red circles represent Anasazi settlements. But—the crucial difference—the right-hand image shows where real Anasazi settlements were, whereas the left-hand one shows where cyber-Anasazi settled.

As always, no two simulations are alike; but once again, this one is pretty typical. In the first frame, as the simulation begins, both the real and the artificial populations are sparse, but the settlements' locations have little in common—to be expected, since this simulation begins randomly. In Frames 2 and 3 (A.D. 855 and A.D. 1021) the real Anasazi population grows and spreads to farmland in the south of the valley; the artificial population also grows and spreads, but with a considerable lag, and the cyber-settlements are more likely than real ones to cling to the edges of fertile zones. Nonetheless, by 1130 (Frame 4) the real and artificial populations look strikingly similar, except that the artificial farmers appear to have overlooked some desirable land in the extreme south. By 1257 (Frame 5) the real population is well along in its decline, and the virtual one continues to track it. (Note that reality and simulation agree that by this point the southern portion of the valley supports only one family, though they disagree about where that family lived.) But in Frame 6, at the end of the period, real history and cyber-history have diverged: the real Anasazi have vanished, whereas several families hang on in the simulation.

What does all this tell us? Nothing for certain; but it suggests two things. First, environmental conditions alone can indeed explain much of what is known about Anasazi population and settlement patterns. Differences between reality and simulation are many; still, given the relative simplicity of the rules and the fact that all but environmental factors are excluded, what is remarkable is how much the simulation manages to look like the real thing. But, second, environmental hardship does not, at least in this model, explain the final disappearance. A steep decline, yes; but a small population could have stayed. Perhaps some unknown force drove

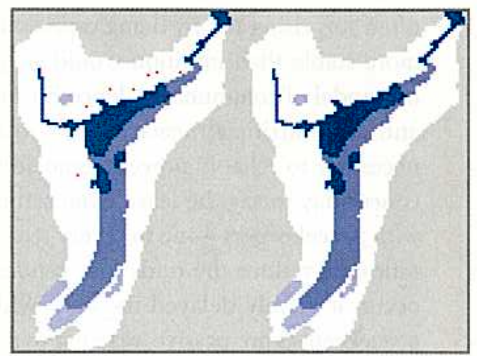
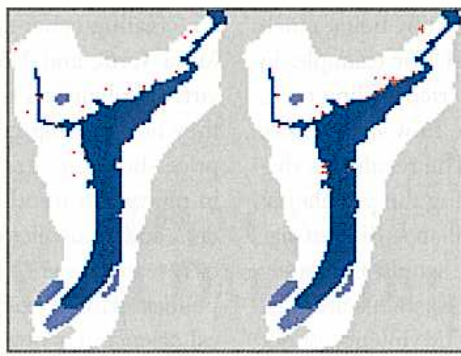
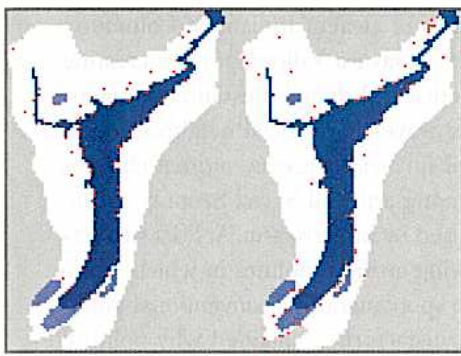
them out; or perhaps, more likely, the last few gave up and chose collectively to leave; or perhaps there is a turning point that this first, still relatively crude model has not found.

Even if the modelers fail to explain why the Anasazi left, they will have shown that artificial societies can come within hailing distance of replicating, in a general but suggestive way, the large trends of real societies, and even some of the smaller trends. In Long House Valley, Gumerman and Dean led me up a sandstone slope to the site of the ancient Long House settlement. Gumerman planted himself in the midst of the ruin and put his arms out and shouted, over an icy morning wind that lashed the valley in early spring, "It boggles the mind. More than half the simulations produce the biggest site right here—where the biggest site actually was."

LEARNING FROM LUMPINESS

“There is no such thing as society,” Margaret Thatcher famously said in 1987. “There are individual men and women, and there are families.” If all she meant was that in a liberal democracy the individual is sovereign, then she was right. But if she also meant that, as some conservatives believe, the notion of a capital-S Society is a collectivist fiction or a sneaky euphemism for the nanny state, then it appears that she was demonstrably wrong; and the artificial societies I have shown you are the demonstrations. They are, it is true, almost laughably simple by comparison with real people and real societies, but that is exactly the point. If even the crudest toy societies take on a life and a logic of their own, then it must be a safe bet that real societies, too, have their own biographies. Intuition tells us that it is meaningful to speak of Society as something greater than and distinct from the sum of individuals and families, just as it is meaningful to speak of the mind as something greater than and distinct from the sum of brain cells. Intuition appears to be correct.

That, however, should not provide a lot of comfort to liberals and progressives. They like the idea of Society because it is not an It but an Us, a group project. For them, Society can be built like a house, or guided like a child, by a community of enlightened activists and politicians who use their own intuition as a blueprint. Artificial societies suggest that real ones do not behave so manageably. Their logic is



their own, and they can be influenced but not directed, understood but not anticipated. Not even the Olympian modeler, who writes the code and looks down from on high, can do more than guess at the effect of any particular rule as it ricochets through a world of diverse actors. The diversity of individuals guarantees that society will never be remotely as malleable or as predictable as any person.

Assimilating this style of thinking took me a while, but then I began seeing human society as both more complicated and less strange than before. Many of the seminal changes in American life have been characterized by the sorts of abrupt discontinuities and emergent patterns that also characterize artificial societies. Why, after twenty-five years of rapid growth, did productivity in America suddenly shift to a dramatically lower gear in the early 1970s? That event, probably more than any other, shaped the discontents of the 1970s and the political and social changes that followed, yet conventional economics still has not mustered an accepted explanation. Why did the homicide rate in New York City, after more than a century of relative stability at a remarkably low level, quadruple after 1960? Why did the rate of violent crime in America as a whole triple from 1965 to 1980? Why did the percentage of children born out of wedlock quadruple from 1965 to 1990? Why did crack use explode in the 1980s and then collapse in the 1990s? If we think of societies in terms of straight lines and smooth curves, such landslides and reversals seem mystifying, bizarre; if we think in terms of sand piles and teeming cyber-agents, it seems surprising if avalanches do *not* happen.

Washington, D.C., is a place deeply committed to linearity. Want to cut crime in half? Then double the number of cops or the length of prison sentences. That is how both Washington and the human brain are wired to think. Yet in recent years many people even in Washington have come to understand that something is amiss with straight-line or smooth-curve thinking. In fact, the notion of unintended consequences has become almost a cliché. Policy measures sometimes work more or less as expected, but often they misfire, or backfire. So far the trouble has been that the idea of unintended consequences, important and well founded though it may be, is an intellectual dead end. Just what is one supposed to do about it? One cannot very well

never do anything (which, in any case, would have unintended consequences of its own), and one also cannot foresee the unforeseeable. And so Washington shuffles along neurotically in a state of befuddled enlightenment, well aware of the law of unintended consequences but helpless to cope with it.

It is at least possible that with the development of artificial societies, we have an inkling of an instrument that can peer into the black box of unintended consequences. That is not to say that A-societies will ever predict exact events and detailed outcomes in real societies; on the contrary, a fundamental lesson of A-societies seems to be that the only way to forecast the future is to live it. However, A-societies may at least suggest the *kinds* of surprises that could pop up. We won't know when we will be blindsided, but we may well learn which direction we are most likely to be hit from.

Moreover, A-societies may also eventually suggest where to look for the sorts of small interventions that can have large, discontinuous consequences. "It may be that you could learn of minimally costly interventions that will give you a more satisfactory outcome," Thomas Schelling told me—interventions not unlike his trick of reordering the traffic flow in Harvard's stairwells by changing the behavior of a single class. I used to think that the notion of government funding for late-night basketball was silly, or at best symbolic. In fact it may be exactly the right approach, because pulling a few influential boys off the streets and out of trouble might halt a chain reaction among their impressionable peers. It now seems to me that programs like President Clinton's effort to hire 100,000 additional police officers and spread them in a uniform film across every jurisdiction are the gestural, brain-dead ones, because they ignore the world's lumpiness. Increasingly, cops themselves are coming to the same conclusion. More than a few cities have learned (or relearned) that pre-emptively concentrating their efforts on key areas and offenders can dramatically reduce crime across an entire city at comparatively little cost.

The flip side of learning to find small interventions with large returns, and at least as important, is learning to avoid large interventions with small returns. In the stretches between avalanches and other discontinuities, A-societies are

often surprising not by being capricious but by being much more stable than intuition would suggest. For example, in his model of communal violence Epstein tried adding more and more artificial peacekeepers to see how many were necessary to reliably prevent genocide. The result was disconcerting, to say the least. Even saturating the population with peacekeepers—one for every ten civilians—did not significantly reduce the odds that genocide would ultimately occur; it merely delayed the end. Why? Epstein's artificial peacekeepers are passive, reacting to nearby violence rather than striking pre-emptively; eventually a rash of clustered killings will always overwhelm their ability to respond, at which point the violence quickly gets out of hand. Epstein concludes that simply throwing forces at an ethnic conflict is no answer; intervention needs to anticipate trouble. That, of course, would not have come as news to the reactive and largely ineffective peacekeeping forces in, say, Rwanda, Bosnia, or Sierra Leone. In Rwanda frustrated peacekeepers pleaded for permission to seize arms caches and intimidate extremists before large-scale killing could begin. Their pleas were denied, at a cost apparent in Figure 6. (See "Bystanders to Genocide," by Samantha Power, September 2001 *Atlantic*.)

The science of artificial societies is in its infancy. Whether toy genocides will truly be relevant to real ones remains an open question. But the field is burgeoning, and a lot is going on, some of which will bear fruit. Researchers

are creating cyber-models of ancient Indians of Colorado's Mesa Verde and Mexico's Oaxaca Valley; they are creating virtual Polynesian societies and digital mesolithic foragers; they are growing crime waves in artificial neighborhoods, price shocks in artificial financial markets, sudden changes in retirement trends among artificial Social Security recipients, and epidemics caused by bioterrorism. At least two sets of researchers are growing artificial polities in which stable political parties emerge spontaneously (conventional political science has never satisfactorily explained why political parties appear to be a feature of every democracy). To me, the early results of this work suggest that social engineering can never be as effective as liberals hope, but also that it need not be as clumsy as conservatives insist.

Today's universities and think tanks are full of analysts who use multivariate equations to model the effects of changes in tax rates or welfare rules or gun laws or farm subsidies; I can easily envision a time, not long from now, when many of those same analysts will test policy changes not on paper but on artificial Americas that live and grow within computers all over the country, like so many bacterial cultures or fruit-fly populations. The rise and refinement of artificial societies is not going to be a magic mirror, but it promises some hope of seeing, however dimly, around the next corner. ■

Computer animations of the artificial societies discussed in this article can be viewed online, at www.theatlantic.com/rauch.