WIRED

Cracking Wall Street

Suppose you could discern market trends, speed up time to see where those trends were going, then bet on what you discovered. Geeks in suits are doing that today.

By Kevin Kelly

"Tell me about the future," I begged.

I was sitting on a sofa in the guru's office. I trekked to this high mountain outpost a couple of years ago to arrive at one of the planet's power points, the national research labs at Los Alamos, New Mexico. The guru's office was decorated with colorful posters of past conferences that traced the almost mythical career of this high-tech legend: from a maverick physics student who formed an underground band of hippie hackers to break the bank at Las Vegas with a wearable computer, to a principal character in a renegade band of scientists who invented the accelerating science of chaos by studying a dripping faucet, to a founding father of the artificial life movement, to the head of a small lab investigating the new science of complexity in an office kitty-corner to the museum of atomic weapons at Los Alamos -- the office I had trekked to.

The guru, Doyne Farmer, looked like Ichabod Crane in a bolo tie. Tall, bony, probably thirty-something, Doyne (pronounced Doan) was about to embark on his next remarkable adventure. He was starting a company to beat the odds on Wall Street by predicting stock prices with computer simulations. He was going to hack the global economy.

Money is just a type of information, a pattern that, once digitized, becomes subject to persistent programmatic hacking by the mathematically skilled. As the information of money swishes around the planet, it leaves in its wake a history of its flow, and if any of that complex flow can be anticipated, then the hacker who cracks the pattern will become a rich hacker.

Some sort of financial hacking has been around as long as computers. The art reached an apex in the late 1980s when powerful PCs allowed any numskull to click on a computer and perform quick and massive trades under certain favorable conditions. When a heavily computerized stock market collapsed in October 1987, this "programmed trading" was blamed.

In the subsequent years, as the stock market rebounded, the increasing computerization of financial trading was ignored. More and more of the everyday tide of money flowed in digital bits, and more and more sophisticated games could be played with them. The traditional financial gamble of derivatives is the latest to be computerized.

A derivative is sort of a bet on a bet, or a speculation squared. Really complex derivatives may give you, say, the option of buying milk at a certain price in New Zealand while simultaneously selling oil in Taiwan. Third- and fourth-order derivatives -- those betting on an option based on a bet that hinges on another gamble -- up the complexity and incomprehensibility of these financial instruments.

Derivatives are only made possible by the immense number-crunching power of 1990s desktop computers. Yet exotic bets form an increasingly major part of the world economy. The bulk of the approximately US$14 trillion that is entangled in derivatives is three times as much money as is tied up in the ordinary stocks and bonds that these esoteric gambles are derived from. When the stock market shuddered in early 1994, computerized derivatives were blamed.

But I wasn't entreating Doyne Farmer because I was interested in derivatives. "I've been thinking about the future, and I have one question," I told Farmer.

"You want to know if IBM is gonna be up or down," Farmer suggested with a wry smile.

"No. I want to know why the future is so hard to predict."

"Oh, that's simple."

Of course I wouldn't mind knowing if IBM is going up or down. Might be quite handy every now and then. But hacking the global economy big time wasn't going to be any fun unless it could be done regularly, dependably, and on a scale where real money could be wagered. What I wanted to know was, can you really predict something as complex and squirrelly as the stock market? Can you really predict the future at all?

Farmer thinks so. He likes to use a favorite example when explaining the anatomy of a prediction. "Here, catch this!" he says, tossing you a ball. You grab it. "You know how you caught that?" he asks. "By prediction."

Farmer contends you have a model in your head of how baseballs fly. You could predict the trajectory of a high fly using Newton's classic equation f=ma, but your brain doesn't stock up on elementary physics equations. Rather, it builds a model directly from experiential data. A baseball player watches a thousand baseballs come off a bat, and a thousand times lifts his gloved hand, and a thousand times adjusts his guess with his mitt. Without his knowing how, his brain gradually compiles a model of where the ball lands -- a model almost as good as f=ma, but not as generalized. It's based entirely on a collection of hand-eye data from past catches. In the field of logic such a process is known as induction, in contradistinction to the deduction process that leads to f=ma.

In the early days of astronomy, before the advent of Newton's f=ma, planetary events were predicted on Ptolemy's model of nested circular orbits -- wheels within wheels. Because the central premise upon which Ptolemy's theory was founded (that all heavenly bodies orbited the Earth) was wrong, his model needed mending every time new astronomical observations delivered more exact data for a planet's motions. But wheels-within-wheels was a model amazingly robust to amendments. Each time better data arrived, another layer of wheels inside wheels inside wheels was added to adjust the model. For all its serious faults, this baroque simulation worked and "learned." Ptolemy's simple-minded scheme served well enough to regulate the calendar and make practical celestial predictions for 1,400 years.

An outfielder's empirically based "theory" of missiles is reminiscent of the latter stages of Ptolemaic epicyclic models. If we parsed an outfielder's "theory" we would find it to be incoherent, ad hoc, convoluted, and approximate. But it would also be evolvable. It's a rat's nest of a theory, but it works and improves. If we humans had to wait until each of our minds figured out f=ma (and half of f=ma is worse than nothing), no one would ever catch anything. Even knowing the equation now doesn't help. "You can do the flying baseball problem with f=ma, but you can't do it in the outfield in real time," says Farmer.

"Now catch this!" Farmer says as he releases an inflated balloon. It ricochets around the room in a wild, drunken zoom. No one ever catches it. It's a classic illustration of chaos -- a system with sensitive dependence on initial conditions. Imperceptible changes in the launch can amplify into enormous changes in flight direction. Although the f=ma law still holds sway over the balloon, other forces such as propulsion and air lift push and pull, generating an unpredictable trajectory. In its chaotic dance, the careening balloon mirrors the unpredictable waltz of sunspot cycles, the Ice Age's temperatures, epidemics, the flow of water down a tube, and -- more to the point -- the flux of the stock market.

But is the balloon really unpredictable? If you tried to solve the equations for the balloon's crazy flitter, its path would be nonlinear, therefore almost unsolvable, and therefore unforeseeable. Yet, a teenager reared on Nintendo could learn how to catch the balloon, sometimes. Not infallibly, but at a rate better than chance. After a couple dozen tries, the teenage brain begins to mold a theory -- an intuition, an induction -- based on the data. After a thousand balloon takeoffs, his brain has modeled some aspect of their flight. It cannot predict precisely where a balloon will land, but it detects a direction the missile favors, say, to the rear of the launch or following a certain pattern of loops. Perhaps over time, the balloon catcher hits 10 percent more than chance would dictate.

If balloon catching paid $1,000 a hit, what more do you need? In some games, one doesn't require much information to make a prediction that is useful. While running from lions, or investing in stocks, the tiniest edge over raw luck is significant.

Farmer calls it "gambling with a positive edge." He believes that by gambling with a positive edge he can crack the stock market. "The nice thing about markets is that you don't really have to predict very much to do an awful lot," he says.

Plotted on the gray end-pages of a newspaper, the graphed journey of the stock market as it rises and falls has just two dimensions: time and price. For as long as there has been a stock market, investors have scrutinized that wavering two-dimensional black line in the hopes of discerning some pattern that might predict its course. Even the vaguest, if reliable, hint in direction would lead to a pot of gold. Pricey financial newsletters promoting this or that method for forecasting the chart's future are a perennial fixture in the stock market world. Practitioners are known as chartists.

In the '70s and '80s chartists had modest success in predicting currency markets because, one theory says, the strong role of central banks and treasuries in currency markets constrained the variables so that they could be described in relatively simple linear equations. (In a linear equation, a solution can be expressed in a graph as a straight line.) As more and more chartists exploited the easy linear equations and successfully spotted trends, the market became less profitable. Naturally, forecasters began to look at the wild and woolly places where only chaotic nonlinear equations ruled. In nonlinear systems, the outcome is not proportional to the input. Most complexity in the world -- including all markets -- are nonlinear.

More importantly, all classical economic understanding has been based on the premise that the collective action of a market is not biased toward particular players or patterns but operates consistently and uniformly throughout the marketplace. An unbiased marketplace was seen to be maximally efficient in allotting resources and rewards; any inefficiency in the marketplace would be exploited by players until it disappeared. Indeed, the efficiency of a random market was the cornerstone of modern economic theory.

But with the advent of cheap, industrial-strength computers, mathematicians have been able to understand certain aspects of nonlinearity and are now questioning the premise of random markets. While traditional economists elaborate ever more complicated explanations of how markets cannot be predicted, money -- big money -- is being made by math jockeys who extract reliable patterns out of the nonlinearity of financial prices and then bet on them. The computer nerds who decipher these esoteric methods, and who have invented computerized derivatives, are called "rocket scientists" or "quants," short for quantitative analysts. These geeks in suits, working in the basements of trading companies, are the capitalist hackers of the '90s.

The aim of the quants is to use sophisticated computer programs and mathematical formulae to play with market data and to try to predict the future a bit. Since all that is required to play the game is a Unix workstation and a math degree, a lot of the action takes place off the big city trading floors. At Wall Street Analytics, based at the center of Silicon Valley in Palo Alto, California, rocket scientists devised a program (list price $50,000), running on PCs, to compute scenarios of mortgage-pool investments. By ingeniously regrouping pools of mortgages and then extrapolating the profits under different "what if" conditions (much like a spreadsheet but with massive criss-crossing feedback of the variables), the software is able to isolate crummy mortgages (nicknamed "toxic waste") from more profitable ones. The computer plays out hundreds of scenarios of different "slices" of the mortgage pool in rapidly fluctuating financial environments. One moment a particular mortgage is toxic, the next it's golden. The quants use the computer to sift through this seething mass of interlinked considerations to select the most profitable buy at any given moment.

Given the propeller-headed math needed to build this incredibly complex forecast, it is not surprising that Wall Street Analytics was started by ex-physicists. "Investing is increasingly becoming dominated by mathematicians, electrical engineers, and programmers," says Adrian Cooper, vice president of Wall Street Analysts. "As a grad student in physics you spend five years banging your head against an abstract problem that no one can solve. This prepares you for working with intractable problems on Wall Street." And the pay can pop the eyes of a postdoc. Salaries on Wall Street can be three times those of physicists. According to *Science* magazine, "Two of the four students who received doctorates in theoretical physics from Harvard went off to jobs on Wall Street.... Of the twenty or so students who received theoretical physics doctorates over the last five years from Stanford University, only two or three [classmates say] are still in physics; they can name eight or nine who are working in finance."

I sought out Doyne Farmer, former mathematical physicist (he banged his head against the insoluble problem of systems with infinite numbers of dimensions while in graduate school), because he is currently one of the financial world's hottest rocket scientists. Together with colleagues from his earlier mathematical adventures, Farmer moved his filing cabinets from Los Alamos and set up an office in a small, four-room house in adobe-adorned Santa Fe. Using their own custom workstations and software, these ex-physicists are directly wired into the financial trading markets in Chicago. They are building the Supercollider of Finance.

Farmer is betting that they can take what they learn from dissecting the trajectories of colliding bits of digitized money and use this information to play the market.

The two-dimensional chart of stock market prices over time hides a hodge-podge of forces driving the price line. A true graph would include an axis for every influence and would thus become an unpicturable, thousand-armed monster.

Mathematicians struggle with ways to tame these monsters, which they call "high-dimensional" systems. A mere 100 variables create a humongous swarm of possibilities. Because each behavior impinges upon the 99 others, it is impossible to examine one parameter without examining the whole interacting swarm at once. Even a simple three-variable model of weather, say, touches back upon itself in strange loops, breeding chaos and making any kind of linear prediction unlikely. (The failure to predict weather led to the discovery of chaos theory in the first place.)

Pop wisdom says that chaos theory proves that these high-dimensional complex systems -- such as the weather, the economy, army ants, and, of course, stock prices -- are intrinsically, no-way-around-it unpredictable. So ironclad is the popular assumption that in common perception any design for predicting the outcome of a complex system is considered naive or mad.

But chaos theory is vastly misunderstood. It has another face. Farmer suggests chaos is like a hit (vinyl) record with two sides. The lyrics to the hit side go: "By the laws of chaos, initial order can unravel into raw unpredictability. You can't predict far." But the flip side goes: "By the laws of chaos, things that look completely disordered may be predictable over the short term. You can predict short."

In other words, the character of chaos carries both good news and bad news. The bad news is that very little, if anything, is predictable far into the future. The good news -- the flip side of chaos -- is that in the short term, more may be more predictable than it first seems. Both the long-term, unpredictable nature of the high-dimensional systems and the short-term, predictable nature of low-dimensional systems derive from the fact that "chaos" is not the same thing as "randomness." "There is order in chaos," Farmer says.

Farmer should know. He was an original pioneer into the dark frontier of chaos before it jelled into a scientific theory and faddish field of study. In the hip California town of Santa Cruz in the 1970s, Doyne Farmer and friend Norm Packard co-founded a commune of nerd hippies who practiced collective science. They shared a house, meals, cooking, and credit on scientific papers.

As the "Chaos Cabal," the band investigated the weird physics of dripping faucets and other seemingly random generating devices. Farmer in particular was obsessed with the roulette wheel. He was convinced that there must be hidden order in the apparently random spinning of the wheel. If one could discern secret order among the spin-ning chaos, then . . . why, one could get rich . . . very rich.

In 1977, long before the birth of commercial microcomputers such as the Apple, the Chaos Cabal built a set of handcrafted, programmable, tiny microcomputers into the bottoms of three ordinary leather shoes. The computers were keyboarded with toes; their function was to predict the toss of a roulette ball. The home-brew computers ran code devised by Farmer based on the group's study of a purchased secondhand Las Vegas roulette wheel set up in one of the commune's crowded bedrooms. Farmer's computer algorithm was based not on the mathematics of roulette but on the physics of the wheel. In essence, the cabal's code simulated the entire rotating roulette wheel and bouncing ball inside the chip in the shoe. And it did this in a minuscule 4 Kbytes of memory, in an era when computers were behemoths demanding 24-hour air conditioning and an attendant priesthood.

On more than one occasion, the science commune played out the flip side of chaos in a scene like this: Wired up at the casino, one person (usually Farmer) wore a pair of magic shoes to calibrate the roulette operator's flick of the wheel, the speed of the bouncing ball, and the tilt of the wheel's wobble. Nearby, a cabal cohort wore the third magic shoe that was linked to Farmer's by radio signals and placed the actual bet on the table.

Earlier, using his toes, Farmer had tuned his algorithm to the idiosyncrasies of a particular wheel in the casino. Now, in the mere fifteen seconds or so between the drop of the ball and its decisive stop, his shoe-computer simulated the full chaotic run of the ball. About a million times faster than it took the real ball to land in a numbered cup, Farmer's prediction machinery buzzed out in small taps the ball's future destination on his right big toe. Typing with his left big toe, Farmer transmitted that information to his partner, who "heard" it on the bottom of his feet, and then, with a poker face, pushed the chips onto the predetermined squares before the ball stopped.

When everything worked, the chips won. The system never predicted the exact winning number: the cabal members were realists. Their prediction machinery forecast a small neighborhood of numbers -- one octave section of the wheel -- as the bettable destination of the ball. The gambling partner spread the bets over this neighborhood as the ball finished spinning. Out of the bunch, one won. While the companion bets lost, the neighborhood as a whole would win often enough to beat the odds. And make money.

The group sold the quasi-legal system to other gamblers because of unreliability in the hardware. But Farmer learned three important things about predicting the future from this adventure:

First, you can milk underlying patterns inherent in chaotic systems to make good predictions.

Second, you don't need to look very far ahead to make a useful prediction.

And third, even a little bit of information about the future can be valuable.

With these lessons firmly in mind, Farmer, together with five other physicists (one of them a former Chaos Cabal member) engineered a start-up company to crack every gambler's dream: Wall Street. They would use high-powered computers. They would stuff them with experimental nonlinear dynamics and other esoteric rocket-scientist formulas. They would hack the financial world with the help of banks; it would be very legal, too.

They would . . . (drum roll, please) . . . predict the future. With a bit of bravado, the old gang hung out their new shingle: the Prediction Company.

The guys in the Prediction Company figure that looking ahead a few days into the financial market future is all that is needed to make big bucks. Prediction machinery need not see like a prophet to be of use. It needs only to detect limited patterns -- almost any pattern -- out of a camouflage of randomness and complexity.

According to Farmer, there are two kinds of complexity: inherent and apparent. Inherent complexity is the "true" complexity of chaotic systems. It leads to dark unpredictability. The other kind of complexity is the complement of chaos -- apparent complexity obscuring exploitable order.

Farmer draws a square in the air. Going up the square increases apparent complexity; going across the square increases inherent complexity. "Physics normally works down here," Farmer says, pointing to the bottom corner of low complexity for both sorts, home of the easy problems. "Out there," pointing to the opposite upper corner, "it's all hard. But we are now sliding up to here, where it gets interesting -- where the apparent complexity is high, but the true complexity is still low. Up here, complex problems have something in them you can predict. And those are exactly the ones we are looking for in the stock market."

With crude computer tools that take advantage of the flip side of chaos, the Prediction Company hopes to knock off the easy problems in financial markets.

"We are using every method we can find," says partner Norman Packard, a former Chaos Cabalist. The idea is to throw proven pattern-finding strategies of any stripe at the data and "keep pounding on them" to optimize the algorithms. Find the merest hint of a pattern, and then exploit the daylights out of it. The mind-set here is that of a gambler: any positive edge is an advantage.

Farmer and Packard's motivating faith that chaos possesses a flip side firm enough to bank on is based on their own experience. Nothing overcomes doubts like the tangible money they won from their experiments with the Las Vegas roulette wheel.

In addition to experience, Farmer and Packard place a lot of faith in the well-respected theories they invented during their years in chaos research. Now they are testing their wildest, most controversial theory yet. They believe, against the unbelief of most economists, that "the market is not random." They hold that certain regions of otherwise complicated financial phenomena can be predicted accurately. Packard calls these areas "pockets of predictability" or "local predictability." In other words, the distribution of unpredictability is not uniform throughout systems. Most of the time, most of a complex system may not be forecastable, but some small part of it may be for short times. In hindsight, Packard believes local predictability is what allowed the Santa Cruz Chaos Cabal to make money forecasting the approximate path of a roulette ball.

If there are pockets of predictability, they will surely be buried under a haystack of gross unpredictability. The signal of local predictability can be masked by a swirling mess of noise from a thousand other variables. The Prediction Company's thirteen quants and rocket scientists use a mixture of old and new, high-tech and low-tech search techniques to scan this combinatorial haystack. Their software examines the mathematically high-dimensional space of financial data and searches for local regions -- any local region -- that might match low-dimensional patterns they can predict. They search the financial cosmos for hints of order, any order.

They do this in real time, or what might be called hyperreal time. Just as the simulated bouncing roulette ball in the shoe-computer comes to rest before the real ball does, the Prediction Company's simulated financial patterns are played out faster than they happen on Wall Street. They reenact a simplified portion of the stock market in a computer. When they detect the beginnings of a wave of unfolding local order, they simulate it faster than real life and then bet on where they think the wave will approximately end.

David Berreby, writing in the March 1993 issue of *Discover*, puts the search for pockets of predictability in terms of a lovely metaphor: "Looking at market chaos is like looking at a raging white-water river filled with wildly tossing waves and unpredictably swirling eddies. But suddenly, in one part of the river, you spot a familiar swirl of current, and for the next five or ten seconds you know the direction the water will move in that section of the river."

Sure, you can't predict where the water will go a half-mile downstream, but for five seconds -- or five hours on Wall Street -- you can predict the unfolding show. That's all you really need to be useful (or rich). Find any pattern and exploit it. The Prediction Company's algorithms grab a fleeting bit of order and exploit this ephemeral archetype to make money. Farmer and Packard emphasize that while economists are obliged by their profession to unearth the cause of such patterns, gamblers are not bound so. The exact reason why a pattern forms is not important for the Prediction Company's purposes. In inductive models -- the kind the Prediction Company constructs -- the abstracted causes of events are not needed, just as they aren't needed for an outfielder's internalized ballistic notions or for a dog to catch a tossed stick.

Rather than worry about the dim relationships between causes and effects in these massively swarmy systems crowded with circular causality, Farmer says, "The key question to ask in beating the stock market is, what patterns should you pay attention to?" Which ones disguise order? Learning to recognize order, not causes, is the key.

Before a model is used to bet with, Farmer and Packard test it with backcasting. In backcasting techniques (commonly used by professional futurists) a model is built withholding the most recent data from the human managing the model. Once the system finds order in past data, say from the '80s, it is fed the record of the last several years. If it can accurately predict the 1993 outcome, based on what it found in the '80s, then the pattern seeker has won its wings. Farmer: "The system makes twenty models. We run them each through a sieve of diagnostic statistics. Then the six of us will get together to select the one to run live." Each round of model building may take days on the company's computers. But once local order is detected, a prediction based on it can be spun in milliseconds.

For the final step -- running it live with bundles of real money in its fists -- one of the PhDs still has to hit the Enter button. This act thrusts the algorithm into the big-league world of very fast, mind-boggling big bucks. Cut loose from theory, running on automatic, the fleshed-out algorithm hums under the murmurs of its creators in Santa Fe: "Trade, sucker, trade!"

"If we can earn 5 percent better than what the market does, then our investors will make money," Packard says. He clarifies that number by explaining that they can predict 55 percent of market moves, that is, 5 percent more than by random guessing, but that when they do guess right their result can be 200 percent better. The fat-cat financial backers who invest in the Prediction Company (currently O'Connor & Associates which, pending approval, will become part of a new subsidiary of Swiss Bank Corporation) get exclusive use of the algorithms in exchange for payments according to the performance of the predictions. "We have competitors," Packard states with a smile. "I know of four other companies with the same thing in mind" -- capturing patterns in chaos with nonlinear dynamics and predicting from them. "Two of them are up and going. Some involve friends."

One competitor trading real money is Citibank. Since 1990, British mathematician Andrew Colin has been evolving trading algorithms. His forecasting program randomly generates several hundred hypotheses of which parameters influence currency data, and then tests the hundreds against the last five years of data. The most likely influences are sent to a computer neural net, which juggles the weight of each influence to better fit the data rewarding the best combinations in order to produce better guesses. The neural net system keeps feeding the results back in so that the system can hone its guess in a type of learning. When a model fits the past data, it is sent out into the future. In 1992 *The Economist* said, "After two years of experiments, Dr. Colin reckons his computer can make returns of 25 percent a year on its notional dealing capital.... That is several times more than most human traders hope to make."

Midland Bank in London has eight rocket scientists working on prediction machinery. In their scheme, computers breed algorithms. However, just as at the Prediction Company, humans evaluate them before "hitting the Return button." Midland Bank's computers were trading real money by late 1993.

Investors like to ask Farmer how he can prove he can make money in markets with the advantage of only a small bit of information. As an "existence proof" Farmer points to people such as George Soros, who earn millions year after year trading currencies and whatnot on Wall Street. Successful traders, sniffs Farmer, "are pooh-poohed by the academics as being extremely lucky -- but the evidence goes the other way." Human traders unconsciously learn how to spot patterns of local predictability streaking through the ocean of random data. The traders make millions of dollars because they detect patterns (which they cannot articulate), then make an internal model (of which they are unconscious), in order to make predictions (for which they are rewarded or punished, sharpening the feedback loop). They have no more idea of what their model or theory is than of how they catch fly balls. They just do. Yet both kinds of models were empirically constructed in the same inductive Ptolemaic way. And that's how the Prediction Company employs computers to build models of the flow of money -- from the data up.

"Our predictions are not as certain as a physicist's prediction of the motion of planets," Farmer says, "but they don't have to be. They are less certain than the laws of motion, but much more than just random."

But the laws of economics may be more certain than some think. In the early '70s economists Fischer Black and Myron Scholes came up with an equation which explains why derivatives work. The Black-Scholes equation was quickly adopted by Wall Street quants to calculate the value of options in complex derivative dealings. A lot of them being ex-physicists, they noticed that the structure of the equation was parallel to the equation used in physics to describe the dissipation of heat. Manipulating this heavy economic math was just like hauling heavy physics math. Indeed, Chicago-based O'Connor & Associates, the well-heeled firm that funded the Prediction Company's start-up, made multimillions from quants running the Black-Scholes equation in the derivatives market.

Says Farmer, "If we are successful on a broad basis in what we are doing, it will demonstrate that machines are better forecasters than people, and that algorithms are better economists than Milton Friedman. Already, traders are hesitant about this stuff. They feel threatened by it."

The hard part is keeping it simple. Says Farmer, "The more complex the problem is, the simpler the models that you end up having to use. It's easy to fit the data perfectly, but if you do that, you invariably end up just fitting to the flukes. The key is to generalize."

Prediction machinery is ultimately theory-making machinery -- devices for generating abstractions and generalizations. Prediction machinery chews on the mess of seemingly random chicken-scratch data produced by complex and living things. If there is a sufficiently large stream of data over time, the device can discern a small bit of pattern. Slowly, the technology shapes an internal ad hoc model of how the data might be produced. The apparatus shuns "overfitting" the pattern on specific data and leans to the fuzzy fit of a somewhat imprecise generalization. Once it has a general fit -- a theory -- it can make a prediction. In fact, prediction is the whole point of theories. "Prediction is the most useful, the most tangible, and, in many respects, the most important consequence of having a scientific theory," Farmer declares. Manufacturing a theory is a creative act that human minds excel at, although, ironically, we have no theory of how we do it. Farmer calls this mysterious general-pattern-finding ability "intuition." It's the exact technology "lucky" Wall Street traders use.

Prediction machinery is found in biology, too. Dogs don't do math, yet dogs can be trained to predictively calculate the path of a Frisbee and catch it precisely. Intelligence and smartness in general is fundamentally prediction machinery.

Farmer confessed to a private gathering of business CEOs, "Predicting markets is not my long-term goal. Frankly, I'm the kind of guy who has a hard time opening to the financial page of the *Wall Street Journal*. "For an unrepentant ex-hippie, that's no surprise. Farmer sees himself working for five years on the problem of predicting the stock market, scoring big time, and then moving on to more interesting problems -- such as real artificial life, artificial evolution, and artificial intelligence. Financial forecasting, like roulette, is just another hard problem. "We are interested in this because our dream is to produce prediction machinery that will allow us to predict lots of different things" -- weather, global climate, epidemics -- "anything generating a lot of data we don't understand well."

"Ultimately," says Farmer, "we hope to imbue computers with a crude form of intuition."

So what happens when everyone is net surfing the world economy? It'll get more difficult to stay ahead, and that will prompt many novices to drop out, leaving the field to those with the biggest iron, the newest schemes, and the keenest, most nimble insights. The only thing the Prediction Company has over its upcoming competitors is a two-year lead.

By late 1993, the Prediction Company reported success in predicting markets with "computerized intuition" in live trading. They refused, however, to comment about their performance. Their agreement with their investors prohibits them from talking about this, as much as the talkative ex-physicists are dying to do so. They were told: "Do not say anything about performance -- anything!" They won't say how much money they trade, or exactly where. (Farmer asked me to let him vet these closing paragraphs because "if I inadvertently tell you something I shouldn't, it could blow our contract with the Swiss Bank Corporation. You know the Swiss -- they are secretive to the point of paranoia.") Indeed, at his request, I removed the hints of what kind of markets they track that slipped out during his conversations with me.

He would say, though, that the studies they have done prove "by rigorous scientific standards" that financial markets can in fact be beaten. "We really have found statistically significant patterns in financial data. There are pockets of predictability. We have learned a lot, and I would love to be able to describe it all to the world, to write a technical book laying out the knowledge we have accumulated on how to extract the weak signals that exist in financial markets and trade on them -- a kind of Theory of Financial Prediction. Maybe call it How to Beat the Market. But given the quantity of money involved, I'm sure our partners will never let us do that, and I have just enough belief in efficient markets myself that I probably agree with them."

This story is adapted from Kevin Kelly's new book "Out of Control: The Rise of Neo-Biological Civilization," published by Addison-Wesely in June 1994. "Out of Control" is about how machines are becoming biological so that we can manage their increasing complexity.

Kevin Kelly (kk@well.com) is executive editor of Wired.