



**Agent-based simulations  
are allowing companies to build  
silicon versions of themselves,  
piece by piece and person by person.**

# SUPERMODELS TO THE RESCUE

by Mark Buchanan

**Two years ago, the European packaging company** SCA faced a dilemma when one of its largest and most dependable customers asked to purchase an additional 20 million boxes each year. Unable to increase its manufacturing capacity in the required time, SCA had to consider cutting back on supply to other customers. The trouble was, the high-volume customer was a tough price negotiator, and SCA earned a relatively small profit on each box. To meet its demand, SCA would lose business from other customers who seemed more profitable on a per-box basis.

How could the packaging company work through the trade-offs?

John Williams, managing director of SCA Packaging Ltd. in the U.K., found his answer through complexity science. Working with Eurobios, a European

consulting firm with links to the famed Santa Fe Institute and top-heavy with physicists and mathematicians, the company built a detailed computer model of the various cutting, printing, and gluing operations involved in producing SCA's custom-made corrugated boxes. With this, they combined models of the mechanisms for managing demand and capacity, for organizing warehouse usage and avoiding "missed deliveries," and for dealing with unexpected processing line failures. The result was a virtual model of SCA's operations on which the company could run "experiments" and explore the likely consequences of different decisions.

The model quickly turned up a surprise. Another of SCA's large customers — in fact, its largest customer — paid very high prices for its boxes. But the irregularity of its ordering behavior created a hidden cost, as it forced

## Mark Buchanan

(mark.buchanan@wanadoo.fr)  
is the author of *Nexus: Small Worlds and the Ground-breaking Science of Networks* (W.W. Norton, 2002) and *Ubiquity: The Science of History, or Why the World Is Simpler Than We Think* (Random House, 2001). Formerly an editor with *Nature* and *New Scientist*, he holds a Ph.D. in physics from the University of Virginia.

SCA to hold a large inventory. Using the computational model, the company discovered that this customer, in the long run, wasn't nearly as profitable as it appeared; losing its business wouldn't be such a bad thing. So SCA dropped the customer and took on the new demands of the price-sensitive client. In just one factory alone, the packaging company's inventory costs fell by 30 percent, and profits rose by \$200,000 in the next quarter. SCA is currently rolling the model out to its other factories (roughly 100 in total), and extending it to examine supply chain and transportation issues as well.

Effective business leaders have long had to spot problems and opportunities through a forest of obscuring details and distractions. But the complexity of the global business environment routinely overwhelms the analytical capacity of even the most gifted leader. In seeking efficiency through multilayered production processes and extended supply networks, while keeping pace with rapid technological change and shifting consumer demands, contemporary managers face conflicting constraints and impenetrable webs of cause and effect. Ever more frequently, cutting through the complexity is not possible, and executives, lacking real knowledge, are forced instead to rely, however imperfectly, on instinct.

The challenge is only elevated by the growing demand for customization in products and services. Increased pricing transparency, customer mobility, and technology transfer speeds, together with continually improving supply chains, make it ever easier for customers to demand more individualized products and services. Like SCA, firms across industries find themselves faced with the choice of acceding to the demands, or losing the customer to an accommodating competi-

tor. (See "Smart Customization: Profitable Growth Through Tailored Business Streams," by Keith Oliver, Leslie H. Moeller, and Bill Lakenan, *s+b*, Spring 2004.)

In the race to offer customized solutions, most business leaders do not adequately judge the complex trade-offs that affect their bottom lines; too often, customization strategies go awry, offering too much costly service to low-profit customers, or inadequate attention to core customers. The new generation of simulation tools can help companies navigate this unfamiliar terrain, penetrating complexity and allowing them to become "smart customizers."

Done properly, computer simulation represents a kind of "telescope for the mind" — it multiplies our powers of analysis and insight just as a telescope does our powers of vision. With simulations, leaders can surpass their ordinary abilities and discover relationships that the unaided human mind would never grasp.

## Thought Experiments

In science, simulation proved its value long ago as an irreplaceable tool for exploration and discovery. With simulations, researchers have discovered new materials and tested theories of the early universe. The very best models of the human heart now run on supercomputers, and are so accurate that the Food and Drug Administration uses them to test drugs in "virtual experiments" that involve no patients.

As political scientist Robert Axelrod of the University of Michigan suggests, simulation can even be seen as a "third way" of doing science. Whereas *deductive science* derives the consequences that follow logically from basic assumptions, *inductive science* gathers empirical facts and tries to generalize to a pattern. *Science by*

*simulation* does neither. Simulations begin with assumptions, yet they explore logic in an empirical or experimental fashion, producing output that reveals the consequences that are likely to unfold from the setting of a certain situation. Simulations are thought experiments.



ver the past decade, firms such as Cisco Systems, Nokia, Capital One, and Boeing have pioneered the use of advanced simulation to model and prototype both product and process innovations. By recasting designs or by altering logistics at the touch of a button, they can advance complex adaptations of key methodologies and goods, work that in the past would have taken months. Simulation-based process reengineering gives these and other firms an edge on their competitors. (See “Here Comes Hyperinnovation,” by Michael Schrage, *s+b*, First Quarter 2001.)

But businesses are also going further with simulations — using them to tame unruly fluctuations in complex production lines; to foresee and predict the consequences of organizational change; and to respond with intelligence and flexibility to myriad market shifts. With a technique called agent-based modeling, business leaders can now model entire business organizations quite effectively, creating full-scale virtual laboratories in which to test their organizations and try strategies “offline” before exposing themselves to the risks of the live competition.

Computer simulation in itself, of course, has been

around for decades, and strategists and operations engineers have come to rely on it when exploring models of manufacturing processes, scheduling problems, and strategic business challenges. The computational approach — based on differential equations, linear programming, and other mathematical methods — has typically employed computers to do traditional mathematical analysis more powerfully. What makes agent-based modeling so different is a determined commitment to modeling organizations and business processes piece by piece from the “bottom up” — not so much analyzing operations as replicating them *in silico*.

These models begin with relatively simple component models of machines, products, employees, and customers, simulating their characteristics, behavioral habits, and aims. The computer then puts all these “agents” together within the realistic structure of the organization and its business environment, and lets them interact. Decision makers can run experiments — reorganizing employees or offering them differing incentives, devoting more effort to one customer rather than another — and can watch as outcomes emerge naturally, unconstrained by the decision makers’ own prejudices about what “ought” to happen. The technique offers insight into the unexpected.

### Committing “Cybermistakes”

Decisions in the real world have consequences. Decisions in a virtual world do not. Hence the first reason for strategic exploration through simulation: Organizations can learn from their mistakes without paying the costs of making them.

Two years ago, a major multinational pharmaceutical company discovered a serious problem in its drug development process. The company, whose associates requested anonymity for themselves and the firm, organized R&D by assigning potential new drugs to individual development teams. Naturally, the success or failure of any specific project could affect a team’s reputation, and this meant that teams were often lured into making “selfish” decisions that were bad for the company; they might keep a project going longer than warranted, for example, to preserve the impression of possible success. Executives found themselves canceling projects in the third phase of clinical trials, with losses that were hundreds of millions of dollars higher than if development had been abandoned at an earlier stage.

How could the therapeutic and financial goals of the company and the aims of individual teams be better

aligned? The company proposed to create a market in which drug development in the early phases of human trials would be contracted out to independent research companies. Management's assumption, largely premised on economics, was that since different companies have different costs of capital and capital requirements, as well as different risk attitudes and resource constraints, a market among such companies might naturally match research tasks to companies willing to "own" the associated risks. These companies would then have clear incentives to judge the merits of any project in realistic terms.

Given the inherent uncertainties of drug development, and the difficulty in coordinating early and later stages of development, the company's executives could not be sure the concept would work. To test the idea, they turned to agent simulation.



ided by Icosystem — a Cambridge, Mass., firm specializing in agent simulations in business, and, like Eurobios, founded by people affiliated with the complexity-research organization Santa Fe Institute — the company constructed a virtual market in which agents representing employees would interact in plausible ways with other agents representing potential contractors, including contractors specializing in managing clinical trials. Within the model, each of the various agents was endowed with some plausible rules by which it would make decisions — about bidding on a particular drug compound, for example. The likelihood of a bid would depend on the company's current resource availabilities and cost of capital, as well as its perception of a compound's commercial potential. These rules reflected real data that the model's developers collected on the capital costs, resource utilization, and costing strategies of independent R&D companies. With the model up and running, and the model showing reasonable results on simple tests, it could then be used to answer the principle question.

What the model showed, though, was this: The market-based idea was a little too simple.

"We found," says physicist and Icosystem founder Eric Bonabeau, "that because of the diversity of players — their different motivations, aversions to risk, cost structures, and so on — the company could not possi-

bly coordinate all of that activity in an open market."

The model also allowed the company to test alternative options by which employees' interests might be better aligned with those of the firm. One idea was to tie employee incentives, in the form of bonuses, to the success of all the company's drug molecules, rather than just to a single project or set of projects. In this way, employees would not be deterred from doing the "killer experiments" that could weed out bad projects early on. By reducing the number of costly development misadventures, the model suggested, this plan could double the value of the company's recently discovered molecules.

### Seeing Around Corners

Twenty years ago, pharmaceutical executives simply could not have said with any confidence what might happen if they managed their R&D within a completely new framework. They would have tried their "market solution," met with expensive failure, and then gone back to the drawing board, most likely only after several years had passed and the organizational wounds had healed. Simulation lets decision makers "see around the corner," and businesses gain a competitive edge by exploiting it.

Several years ago, Nasdaq planned to change the tick size — the basic price increment — of its securities listings, and switch to decimalization from prices listed in fractions. The electronic securities marketplace anticipated that decimalizing and decreasing the tick size would make it easier for the market to discover the accurate price of stocks, because it would let traders express their market views more precisely. The result would be a smaller difference between the bid and ask prices at which traders are willing to buy and sell securities, mak-

## When Nasdaq modeled the actions of its entire marketplace, it discovered that “decimalization” could make the market less efficient, not more.

ing Nasdaq’s pricing more competitive and attracting both more investors and more listing companies.

It sounded like a good idea. But the exchange thought it best to investigate further before going ahead. It developed a model of the activities of traders and market makers, the firms that maintain liquidity for particular securities by their commitment to buy and sell them at the listed prices. The “agents” were programmed to act like participants in the real market, using common strategies, but they were also endowed with artificial intelligence and so were able to adapt and alter their strategies on the fly as they identified trends or patterns in the market.

This may seem almost like magic, but it can be done. One of the most powerful elements of human behavior is humans’ ability to cope with uncertainty by learning from experience. When faced with some altogether novel task, we almost never work out what to do using strict logic. Rather, we try something, and if it doesn’t work, we try something else. We tend to stick with any approach that works, yet, often by accident, we stumble over minor variations of older ideas that work even better. In this way, through an evolutionary process, we learn.

Nasdaq endowed the agents in its model with a similar capacity to learn by trial and error. In the guts of the model, each agent in the market kept tabs on a handful of possible “strategies” by which it might make decisions. Each such strategy was a mathematical recipe for taking reality — in the form of price movements and the actions of other traders in the recent past — and deciding on some specific action. Each agent would monitor its own handful of strategies, seeing which would have earned it the most money in recent trades, and, at any

moment, would use this “best” strategy to make its next decision. In other words, like humans, the agents kept track of successful ideas and used them, while paying less attention to less successful ideas.

Some of the strategies were those identified through interviews with real market participants, but the model also contained an evolutionary mechanism: Every so often, each agent would try out a completely new strategy obtained by introducing minor, random variations into one of its older strategies. In this way, the agents had the potential to discover new ways of behaving that might be superior to anything they had done in the past. They could even discover possible ways of making profits that the model’s designers themselves could not have foreseen.

In this case, this small element of “artificial intelligence” turned out to be crucial to the model’s success. Once the model was working like the real market — reproducing price fluctuations in a mathematically accurate way — the company could use the virtual market as a laboratory. Nasdaq found a surprise. In repeated experiments, reducing the tick size beyond a certain point actually increased the bid–ask spread. As it turned out, so-called parasitic strategies, which make quick profits for individual market makers at the expense of overall market efficiency, grew less risky and more profitable with a smaller tick size. The artificially intelligent market makers naturally learned this and responded, as live market makers surely would in reality, driving up not only the bid–ask spread but also the overall volatility of price fluctuations. (See Exhibit 1.) When Nasdaq went ahead with its plans and actually changed the tick size from 1/16 to 1/100 in 2001, it was able to anticipate this effect. The exchange slowed the introduction of the

## Agent-based modeling helped Munich's Infineon Technologies increase the throughput of its production line.

new system and developed a new “Super Montage” system for displaying buy and sell orders and their execution, which it hoped would ameliorate problems.

### Penetrating the Confusion

In addition to helping companies resolve conflicts and improve processes, agent-based simulation also represents an engine for innovation and creativity, as leaders can discover and grab hold of opportunities that even the cleverest analyst would never see.

Two years ago, Post Danmark, the Danish postal service, undertook a project to improve the efficiency of its “last mile” postal delivery routes. The planning of postal delivery routes faces countless conflicting constraints involving the physical placement and capacity of roads, patterns of mail demand, union work rules, and delivery time requirements for different categories of mail, such as priority post.

Working with Eurobios, Post Danmark developed an agent simulation to search for better delivery routes, looking for solutions that would lower costs or speed delivery time. The model quickly turned up routes that reduced the distance postal workers had to travel on the existing routes by 15 to 20 percent. The solution even forced the postal service’s planners to consider ideas they previously never would have entertained; for example, the simulation showed that optimum delivery routing required some paths to cross, an idea that human planners had dismissed as “obviously” unsuitable.

The model also helped settle another issue. Post Danmark offers priority postal delivery that is guaranteed to arrive before 10:00 A.M.; executives were toying with the idea of offering a 9:30 A.M. delivery to give customers better service. But to meet the earlier deadline,

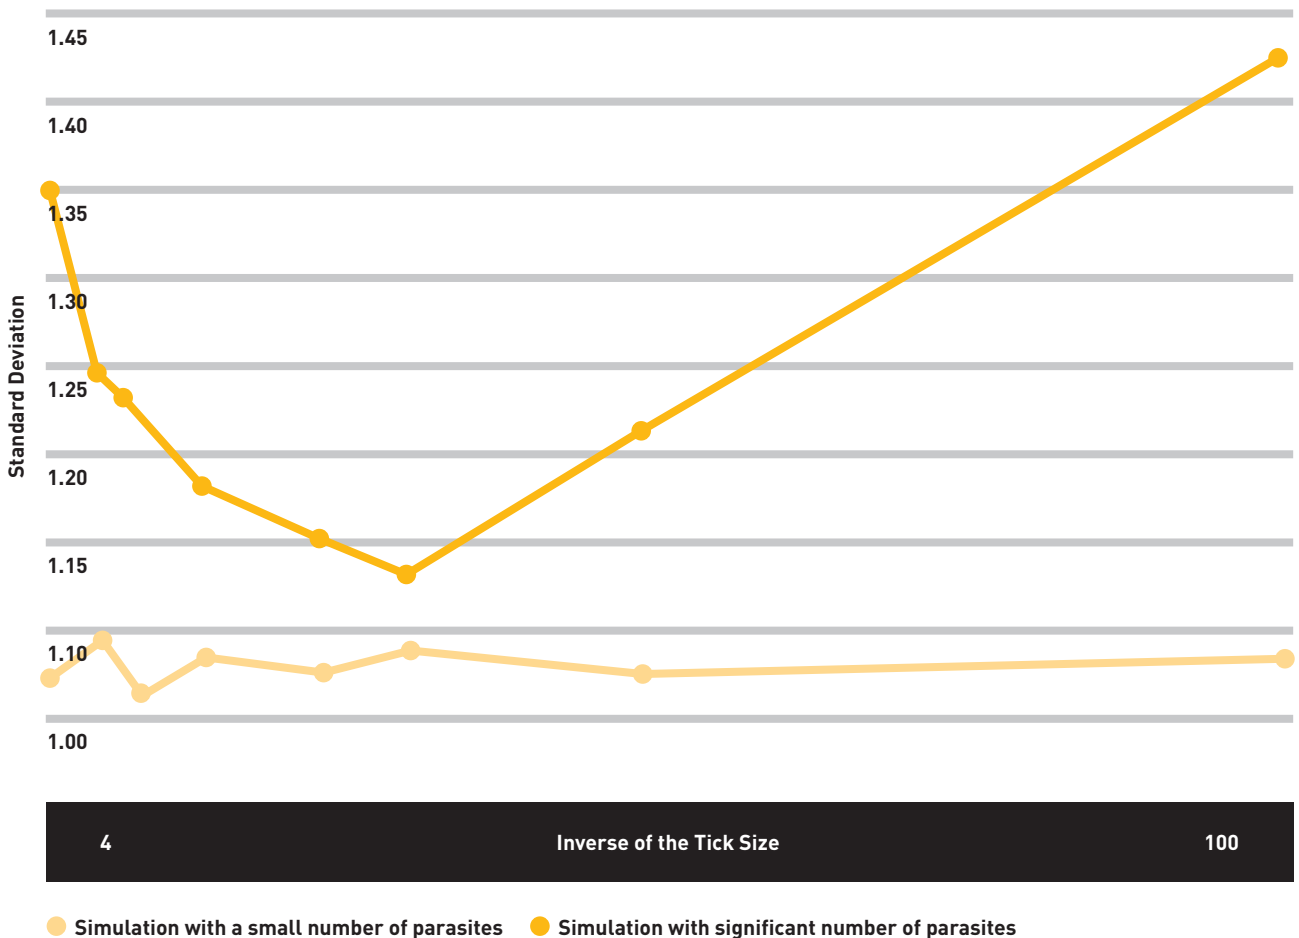
the model showed, Post Danmark would have had to add a significant number of extra postal carriers, and even then would likely have failed to meet its basic standards of performance. A change of one half-hour would seem minimal, yet the model revealed it would push the delivery into a complexity crisis, undermining its ability to cope.

Advanced management of strategies for adaptive customization — the generation of profitable growth via “smart customization” — lies at the high end of business complexity. But even at the low end, in day-to-day operations and the routine organization of procedure, simulations can show companies surprising ways to improve their performance. In some cases, the right steps may even seem obviously wrong. Consider the experience of Munich-based Infineon Technologies, a major manufacturer of semiconductor chips, which recently discovered a way to increase the overall throughput of its production line.

For years, operations managers had been striving for higher line speeds, trying to push more chips through faster. This seemed the only way forward. But a detailed simulation of the process suggested that slower line speeds might be better. At high speeds, the expensive devices for handling the “wafers” on which chips are manufactured were unable to orchestrate the complex movement of thousands of wafers while still respecting the tight tolerances involved. Costly stoppages were the result. Slower, it turned out, was actually faster, as the line ran more smoothly.

Simulations helped the same pharmaceutical company that cured its “selfish team” syndrome to discover further opportunities to improve operational efficiency. The firm had organized research and development by

Exhibit 1: Nasdaq's Decimalization Simulation



Source: Eurobios

groups with expertise in 22 distinct “functions,” from biochemistry to marketing, and each project drew on these functions when needed. But in moving past R&D and into the clinical trials required to bring a medical

product to the market, all 22 functions had to be present at a series of meetings. Coordinating this activity so that each function could bring its expertise to bear on important strategic decisions led to serious delays.

With the computer model, the company found that breaking trials into smaller and more manageable pieces would demand fewer such coordination activities and improve productivity, in principle, by as much as 80 percent. That was two years ago. Having followed the strategy, the company has in fact seen a 50 percent rise in productivity.

### Practical Matters

New-generation simulations are easy to overuse. Even with the decline in the price of computing power, the cost of modeling a relatively simple problem — such as SCA's box manufacturing process — can run beyond \$100,000; more complex models like Nasdaq's, which need to capture human behavior and learning, can cost \$500,000 or more to develop. Moreover, companies should not reject basic strategic and operational thinking in favor of computer simulations any more than elementary school students should ignore their multiplication tables because calculators are available. Agent-based models should be a tool for better thinking, rather than a replacement for thinking.

In rough terms, there is a threshold for determining when agent-based modeling and other related forms of simulation should be used, a threshold known as “high computational complexity.” The idea of computational complexity is that some systems or problems just cannot be simplified. Any model that can accurately mimic the system has to be roughly as complex as the original system itself. Complexity theorists believe that systems made up of many interacting parts are often of this type — including ecosystems, the global climate, and business and government organizations. And although very recent mathematical research suggests that some of these

systems may have “hidden” structure that would enable their simplification and prediction on a more theoretical basis, the scientific tools for understanding this structure may take decades to develop. Fortunately, the vast processing power of the computer offers a way to meet computational complexity head on, and defeat it, by creating models that run in a world where time is greatly compressed and where history can be repeated many times.



oing it well, of course, requires more than a little practice. An ongoing project at Argonne National Laboratory in Chicago illustrates the process nicely.

The State of Illinois is legally committed to deregulating its electricity market in the year 2007. Given the recent disaster associated with deregulation in California — with the power “shortage” during the summer of 2000, later traced to the market manipulation of energy traders at Enron and elsewhere — the Illinois Commerce Commission would like to identify potential trouble beforehand and take wise steps to avoid it.

At the request of the state, Charles Macal, a researcher at Argonne National Laboratory, and his colleagues have developed an agent-based model, in an effort to ensure that this deregulation comes with no surprises. They have focused on good modeling practice from the outset. The team first carried out a “participatory simulation,” with knowledgeable people playing the roles of the agents that would ultimately appear in

the computer model. “This helped greatly to identify likely strategies the agents would use in the real world, and the kinds of information the agents would want to use in their decisions,” says Dr. Macal.

Based on their observations, the team built a model with agents to represent companies that generate, consume, transmit, and distribute electrical power, as well as individual consumers and regulators. These agents explore various market strategies, and on the basis of their experience, adapt their behavior as time goes on, constantly searching for new strategies that perform better. The agents also learn how they can potentially influence the market for their own benefit. To make the model credible, Dr. Macal and his team found they also had to model the underlying electric grid, following the flow of electricity through each and every one of the 2,000 physical nodes in the Illinois system.

With a model of this complexity, validation through extensive tests is absolutely crucial. But it is not enough just to look at the totality of what the computer spits out. Agent-based models often yield surprises and explore realms of behavior where no one knows quite what to expect. So there is no way to confirm a model’s validity by comparing its output to “known” results. To validate their model, Dr. Macal and his colleagues instead scrutinized the model’s guts, checking as many factors as possible to be sure it was handling the details correctly — that is, reproducing the behavior of individual agents accurately, whether these agents represented power producers or industry regulators. After an exhaustive study of each component under many conditions, they had the confidence to use the model as a tool for exploration and discovery.

Dr. Macal and the team have now used their model

to identify some potential problems with the proposed deregulation. In particular, they have found that the distributed nature of the electrical network makes it quite conceivable that some companies would be able to engineer geographical “pockets” in which they would effectively have monopoly power and could set prices as they wished. This, of course, is precisely the kind of situation that effective deregulation seeks to avoid. The Argonne project is currently working to establish the credibility of this finding, varying many tiny details within the model to be sure that the problem is not a spurious artifact, but persists under all reasonable assumptions. If the preliminary findings stand up, then Dr. Macal and his colleagues will employ the model in another way — to begin exploring options for mitigating the problem, by maintaining regulation over specific regions of the network or through other mechanisms. +

Reprint No. 05106

---

## Resources

---

Keith Oliver, Leslie H. Moeller, and Bill Lakenan, “Smart Customization: Profitable Growth Through Tailored Business Streams,” *s+b*, Spring 2004; [www.strategy-business.com/press/article/04104](http://www.strategy-business.com/press/article/04104)

Michael Schrage, “Here Comes Hyperinnovation,” *s+b*, First Quarter 2001; [www.strategy-business.com/press/article/10900](http://www.strategy-business.com/press/article/10900)

Eric Bonabeau, “Agent-Based Modeling: Methods and Techniques for Simulating Human Systems,” *PNAS*, Vol. 99 (May 14, 2002), 7280–7287

Joshua Epstein and Robert Axtell, *Growing Artificial Societies: Social Science from the Bottom Up* (MIT Press, 1996)

Navot Israeli and Nigel Goldenfeld, “On Computational Irreducibility and the Predictability of Complex Physical Systems,” unpublished; [www.arxiv.org/abs/nlin.CG/0309047](http://www.arxiv.org/abs/nlin.CG/0309047)

Proceedings of the Agent 2004 Conference on Social Dynamics: Interaction, Reflexivity, and Emergence (University of Chicago Press, forthcoming); <http://agent2004.anl.gov/>

---