Information • Knowledge • Systems *Management* 1 (1999) 179–183
IOS Press

Would-Be Business Worlds
An Interview with John L. Casti

John L. Casti is author of *Would-Be Worlds: How Simulation is Changing the Frontiers of Science* (Wiley, 1997), as well as many other best-selling science and technology-related books including *The Cambridge Quintet* (Little, Brown, 1998) and *Paradigms Lost* (Morrow, 1989). He is a member of the Santa Fe Institute and on the faculty of the University of Vienna. He received his Ph.D. in mathematics from the University of Southern California.
This interview focuses on Casti's experiences in developing simulations of complex organizational systems. The central elements of these simulations are agent-based models that represent the behavioral mechanisms underlying individual actors in these complex systems. The resulting simulations are used to study the organizational behaviors that emerge from the collective behaviors of these actors.

IKSM: Your books and articles describe a wide variety of fascinating simulation models, ranging from stock markets to insurance markets to supermarkets. What prompted focusing on these particular domains? Was it theoretical interests, specific issues the customers for these models wanted to address, or simply the best opportunities at hand?

Casti: My simulation models of the catastrophe insurance industry and the supermarket were purely targets of opportunity: the people in those industries came to me, following a lecture or reading of my book Would-Be Worlds, and in effect said, "Can you help us?" So the choice of topic was customer-driven. The stock market model is not my own, but one created by Brian Arthur, John Holland and others, in an attempt to construct a laboratory for studying the Efficient Market Hypothesis, as well as other aspects of academic finance.

IKSM: Our readers will be very interested in who should knock on your door, and for what purpose. What type of questions are best answered by simulation approaches? Some problems are tractable analytically. Other types of problems are not tractable at all. Which types of problems match well what simulation approaches can offer?

Casti: In my view, the best type of problems/questions to use agent-based modeling to address are those that have a collection of objects making up the system that are heterogeneous, intelligent and adaptive. For example, drivers in a road-traffic network, shoppers in a supermarket, traders in a financial market or animals in an ecosystem. By these terms, I mean that these agents are intelligent, in the sense that they use rules to decide upon their actions at any time, they are adaptive in the sense that they are ready to change their rule if the old rule isn't working so well anymore, and they are heterogeneous in that the agents don't all use the same rule. Such systems are VERY difficult to model using traditional mathematics, since that mathematics is based on physics, where the "agents" — planets, billiard balls, electrons — always use the same rule. Complex, adaptive systems are very different, as the spectrum of available rules and the fact that the agents change their rules takes these systems outside the domain of any known mathematical structures for characterizing them. So if you are dealing

1389-1995/99/$8.00

with this kind of system — and almost all the systems of daily life are of this sort — then you have very little recourse other than to study them via simulation; there is no mathematics for such structures.

IKSM: Once you have developed such a model, what kinds of questions can you address with it? Further, how do you validate the model so that you can be confident of the answers you generate to these questions? Can you provide examples of how you and/or your colleagues have done this?

Casti: These simulation models can be used to answer a variety of questions, depending on the scale of resolution of the model, the time step in the model, the accuracy of the input data and a lot of other factors. In other words, not all simulations are created equal. But the one thing I'd like to emphasize is that the purpose of these simulations is NOT prediction; it's insight and understanding. People often have a misconception in this regard; they think that all mathematical and computer models — simulations or otherwise — aim at the kind of predictability you get from using Newtonian mechanics to calculate, say, planetary positions. But it's a very rare system, usually found only in physics, in which this kind of predictability is possible. This is especially true in the social and behavioral areas, where the agents making up the system have a choice of what rule to use at any stage — or can even create entirely new rules that never existed before. So I have to spend a lot of time discouraging people from expecting too much from mathematics and computing. After all, mathematics and computing are not magic.

Anyway, to validate the behaviors seen in a simulation one does about the same thing you do to validate a mathematical model. You first look to see if the general structure of the simulated behavior agrees with the structure seen in the actually observed behavior of the real system. Of course, one doesn't demand that the model exactly reproduce the real data, since the model has left out many aspects of the real system that may influence the precise way it behaves. What we're interested in is whether the simulated system displays behavior of the same type as the real system.

The second thing you do to validate the model is to ask the opinion of experts. For instance, in a road-traffic model you set up some sort of scenario, then ask experts what they think will happen. Once the model has produced its behavior, it's often the case that the simulation will show things like traffic jams, accidents and the like that the experts didn't anticipate. But they can usually give a good a posteriori explanation for why things happened the way they did — once they've seen it happen. This is the kind of expert judgment that tends to validate the model, since it involves people who have spent their lives working with the real system saying that the model's behavior "feels right".

IKSM: It would seem that emergent, counter-intuitive behaviors would be of most interest — the "I never would have guessed" sort of behaviors. Are you suggesting that experts can usually reverse engineer these behaviors, determine why they arose, and realign their intuitions so that these unexpected behaviors are no longer counter-intuitive? This would seem to offer ample opportunities for people to mislead themselves and jump to conclusions due to the natural human tendency to focus on confirming evidence. Have you observed this phenomena? How might it be counteracted?

Casti: Yes, emergent behavior is really the sine qua non of complex systems. But I don't think it's possible, in general, to "reverse engineer" these phenomena to discover their causes and, in effect, predict surprises. Now in principle this might actually be done. But the problem is that the solution to such "inverse problems" are generally highly non-unique. This means that for a given observed behavior, there are usually LOTS of types of rules that the individual agents making up the system might use, all of which lead to the very same behavior. For example, suppose you observe a stock market crash of a certain magnitude at a particular moment. Then you may well wonder what kinds of trading strategies the investors and traders were using that gave rise to that crash. The answer is that there are usually an infinite number of different trading rules that could all generate that very same crash. You would then have to inject additional information into the situation to pick out the one set of rules that might have been responsible for the particular crash you've observed.

IKSM: A central question is whether decision makers should believe models' counter-intuitive outputs. If they cannot probe the model to find the causal chain of the unexpected outputs, then their only recourse is to probe the real world to find conditions under which the unexpected model outputs might actually occur in the modeled world. How do you approach decision makers' needs to gain confidence in counter-intuitive behaviors.

Casti: Why do you say the decision maker can't probe the model to discover the causal chain leading up to counterintuitive behaviors? There is no reason, in principle, why this can't be done. In fact, this is an excellent reason for having the simulation as it allows this possibility — in contrast to studying the real system, which doesn't.

But there still remains the "Can you trust it?" question. Basically, the way we validate these simulations is twofold. First, we look to see if the model's behavior is at least qualitatively similar to the known behavior of the real system. Then we create various scenarios and put real people who have spent their professional lives working with the system in front of the computer and ask them to anticipate what will happen when the simulation is turned on. Usually, they are unable to predict "surprises". But they are also generally very adept at explaining what occurs after the fact. In short, they are able to trace out the causal connections generating the hard-to-predict behaviors. So this is the validation method, a combination of approximate agreement with whatever data is at hand and expert opinion and judgment.

IKSM: In responding to an earlier question, you said, "I don't think it's possible, in general, to 'reverse engineer' these phenomena to discover their causes and, in effect, predict surprises". Yet, here you seem to say that this is a good approach. There must be some unstated distinction underlying these two responses. Is that the case?

Casti: The resolution of those two seemingly contradictory statements resides in what is possible "in principle" and what can actually be done "in practice". In principle, one could trace all the statements in the program representing the simulation and discover exactly why things unfolded the way they did. But this is about as practical as doing the same thing to track down all the possible Y2K bugs and squash them; it just can't be done in practice. In practice, you can find "some" of the causal paths — but you really have to just run the simulation and see what happens.

The "reverse engineering" aspect, incidentally, refers more to trying to understand what rules the agents are using, given that you observe some collective effect — emergent behavior — that results from the interaction of the decisions coming out of those rules. Unfortunately, this is not only an ill-posed problem, mathematically; it's also one whose solution is generally highly non-unique. So, in general, there are LOTS of different rules that the agents might be using, ALL of which will lead to the same observed behavior.

IKSM: You've made a pretty convincing case for the role and value of such models. If someone wants to build a model of their business, perhaps in terms of interactions with customers, where should they start? Do they need particular software tools? Do they need to hire a modeling specialist or simulation consultant?

Casti: The place to start is by elucidating carefully what question(s) you want the simulation to be able to address. This determines the level of resolution of the simulation, who the agents are, what kind of interactions take place and so forth. The second step is to make sure you have expert advice available from people who understand the system. This advice is crucial in giving the agents realistic rules of action to choose among, as well as in setting the way the agent's decisions interact with each other to form the collective behaviors which are generally all we really see of complex systems.

At this point, formal modeling and programming can begin. It may be necessary to engage the services of a consultant or expert if the technical skills to do this are not at hand. The issue of what software to use in actually programming the simulation is a tricky one. Every single simulation I've participated in was programmed in a different language — C++, Java, Excel, MathCad, etc. So I certainly cannot say you need this or that package; generally, it comes down to the way the simulation is to be used (e.g., is it necessary that it be very portable and accessible to people who are not very computer literate), the size of the database and the sheer volume of agents and interactions (e.g., do you need a supercomputer to run it), and the tastes of the programmer(s) (e.g., what languages do they like to program in?).

IKSM: In your experience, are these models mainly used to explore "big" questions and gain new insights? Or, are they used more on an ongoing basis for process optimization of sorts? How does the primary use of such models affect model "maintenance" in terms of updates, version control, etc.?

Casti: In general, my feeling is that these simulations are used primarily to address "big" questions and to gain insight, not to answer day-to-day operational questions. But I think this is mostly due to the fact that this type of analysis and modeling is of relatively recent vintage, and so managers and business decision makers are just not yet familiar with the potential — and the limitations — of such tools. I believe that once the simulations have been "vetted", in the sense of being calibrated against real-world events and behaviors, their use as an everyday tool will increase. And once they are calibrated against the real world, we will have to take greater interest in things like model updates, new versions to represent shifting circumstances and the like.

IKSM: Where are simulation models of complex organizational systems headed? Are there likely breakthroughs on the horizon? What types of breakthroughs would make a substantial difference? What are the limits to growth/applicability in developing such models?

Casti: In my opinion, the real advances in agent-based models in business and organizations will come from two directions. The first will be increased popularity in the use of such models, resulting in greater efforts to obtain the input data needed for them, especially the psychological factors underlying individual decision making. The second area of major advances will be technological. Greatly enhanced computing power will enable us to look at these organizational systems in new ways. For example, the ability to quickly run through thousands of scenarios involving millions of agents will provide an almost unlimited playground for testing hypotheses about how, say, stock markets or ecosystems really work. In another direction, this computing capability will enable us to couple virtual reality technology to agent-based models, giving us the capability to actually "experience" what various possible organization structures might be like instead of just observing them. A first-person perspective, if you like, rather than a third-person one. So these are the areas in which I see advancement taking place in the short-run.

IKSM: What types of domains are you attempting to model now, or in upcoming work? What domains would you like to model if money, cooperation, etc. were no object?

Casti: If money and computing resources were no object, I'd like to try tackling some simulations in the area of national economies. Also the political arena seems appealing. For instance, trying to trace the movement of terrorists, drug traffickers, and the like might be possible. Finally, there is the world's financial network. How does money move around the globe? Where are the soft spots in the network that hackers might attack? Questions like that should be approachable via this methodology.

IKSM: Thank you for the wealth of insights you have provided. These insights will be invaluable for our readers who are tempted to construct their own would-be worlds.