In Search of the Variable Geometry Firm: An Org Chart that Flexes and Flows

An interview with Joshua Epstein

Matrixed organizations, top-down leadership, centers of excellence, natural work teams. Business libraries and popular bookstore shelves are crammed full of lengthy tomes on the subject of finding the perfect organizational structure. It seems there's no shortage of advice about the perfect recipe for success. What's our fascination with this topic? And why has the holy grail of org charts proved so elusive?

Joshua Epstein is a Senior Fellow in Economic Studies at the Brookings Institution, a Member of the Brookings-Johns Hopkins Center on Social and Economic Dynamics, and a Member of the External Faculty of The Santa Fe Institute. His pioneering work with Robert Axtell in agent-based computational modeling brought us "Sugarscape" – a computer-generated artificial society where natives ("agents") live, work, reproduce, consume, and ultimately thrive or die based on their varied individual attributes – just like humans. The life and times of these societies are charted and examined in Epstein and Axtell's MIT Press book, *Growing Artificial Societies: Social Science from the Bottom Up.*

Epstein has now taken this work in a new direction, exploring what type of agent attributes nurture a flexible, adaptive, and ultimately successful organization over the long-haul. The findings are discussed in "Growing Adaptive Organizations: An Agent-Based Computational Approach" – a research paper that will be included as a chapter in his forthcoming book, *Generative Social Science: Studies in Agent-Based Computational Modeling,* to be published later this year by Princeton University Press. The book will include a CD with animations that illustrate the modeling in action. Readers will be able to watch an adaptive organization grow, an epidemic spread, or a society rise and fall.

Focus recently sat down with Epstein to find out more about what's on the horizon with the organizational green house and the future impact of this work for business.

For the uninitiated, can you tell us what agent-based modeling is all about?
My bumper-sticker line about agent-based modeling is, "If you didn't grow it, you didn't explain it." It's an approach to mathematical social science that starts with the individuals. We begin by representing individuals in the process of interest, whether that's a market, an epidemic, a war – whatever it is. We create individual data structures - cyber people - that have diverse characteristics. Some will be rich, others will not; they will have different appetites, abilities, social networks, memories, preferences, different everything—agents are heterogeneous, partially informed, and boundedly rational.

This is very different from traditional economic models, where all the individual agents are assumed to be rational, fully informed, optimizers. Those models create "homo economicus." Homo sapiens doesn't operate like that. Real people have crummy information, local information. They use heuristics and procedures for making decisions that are *not* rational. They imitate neighbors, do random things, and use rules of thumb. They are going to behave in ways that homo economicus never would. Agent-based modeling acknowledges all this.

Another important attribute of agent-based modeling is that activities unfold on actual spaces where there can be local interaction. In a lot of modeling there is no space. We have real landscapes – in the epidemic model for example there are real homes, hospitals, workplaces, schools. For the organizational model there are different layers of management and locations in an organization. There's much more verisimilitude about the interaction space than with other types of modeling.

So we put these agents in these environments, set them in motion, and watch what unfolds. As with any modeling, the ultimate objective is to better understand the process at hand.

How have you applied agent-based modeling to organizations?
We used the agent-based approach to see if we could "grow" an optimally adaptive organization – in other words, an organization that consistently responds to opportunities, problems, and changes in its environment, better than any other.

However, I didn't want the outcome to be prescriptive along the lines of build a pyramid today, then be flat tomorrow, be matrixed on Wednesday, and divide into three sub units on Friday. That's just a top-down directive. I wanted to find out what corporate culture will generate the optimal sequence of adaptations *from the bottom-up*. This is what makes it an agent-based model—answering the question: What sorts of individual attributes confer adaptiveness on the organization?

So, we modeled the environment as a flow of opportunities or problems that confront an organization. That organization has fundamental resources (agents) deployed across the market to intercept these opportunities as they approach. The basic dilemma is one of labor allocation – ensuring that everyone who's leaning on a rake gets assigned to an outstanding problem.

Within the environment, the agents are endowed with certain characteristics that give them thresholds and tolerances that dictate their willingness to let opportunities pass by, call for help when resources are needed, lend resources when requested from other agents, escalate requests, etc.

And what did you find?
The optimal organization turns out, in some of these environments, not to have a fixed structure. There is no one organization chart that consistently appears to be optimal at all. It turns out that there is a whole structural repertoire. The optimal performance suggests an adaptive organization in the technical, strict, rigorous sense: That is to say, the ideal organization does not have a fixed organization chart. Rather, it has an adaptive repertoire of structures.

What do you mean by "adaptive repertoire"?
An adaptive repertoire is actually a space of organization charts – not a fixed wiring diagram. That means local rules and decentralized interactions sometimes, and then very high level intervention other times, or the creation of local pyramids in different chunks of the market. It's like a good jazz band –you pass the lead around. If it's a ten-piece jazz band and someone hogs all the leads, that's not good. In the optimal geometry firm, the locus of decision-making authority moves around.

For example, if you have a CEO who never "goes away," the adaptive repertoire simply doesn't include the ability to dismantle hierarchies when they are no longer needed. If everyone at a certain level has no ability to ask for additional resources when needed, that adaptive repertoire excludes the construction of hierarchies when they are needed.

And this raises all sorts of puzzles about why you see pyramids everywhere. Why is the pyramid so ubiquitous when, statistically, it's not the optimal structure?

Are there any companies that you've seen that seem to understand and/or embrace this adaptive repertoire concept?
I've been told that Alfred Sloan, in his memoirs about his years at General Motors, talked a lot about the same idea: that you shouldn't have a fixed organization chart. My suspicion is that if you look closely at companies that have capitalized well on opportunities – like Microsoft and its decision to move from mainframes to PCs – you might see something like this. There of course are many more examples where failure to be flexible has been costly. If you look at so-called "disruptive technology" stories, the field will be littered with examples like that.

What surprised you about this model?
It all surprised me. One very interesting thing that came up had to do with competition. The model can be extended to introduce many different companies. And one of the questions we explored was how do the different companies divide up the market? You can populate this landscape with multiple adaptive firms and then ask, "How do they co-evolve?" If they are both acting optimally, how do they divvy up the market and both learn to survive?

And what happened?
They divided the market rather nicely. In this case the model worked in a way where if a company didn't take an opportunity, that opportunity would pass by and could be picked up by another company. And what the organizations did was to arrange their resources so that they divvied up the market so they would both survive.

How does all this apply to the real world?
This is the power of the model: The user can say, "here's the time frame I'm worrying about, here's the cost and salary structure, here's what I want to do (for example, "I want to maximize profit, or market share" – you can put in any objective you like), and here's what I expect the environment to look like and how I expect it to change." And then the model does what is called "combinatorial optimization." This means the computer grinds through every possibility with brute force, checks all the possible corporate cultures, and spits out the one that maximizes the objective the user specified, within the environment the user posited.

The model allows you to answer the question: What structural adaptations are required to maximize a specified objective, in a posited environment, over a prescribed time horizon? And in the paper I talk about cases where, with totally standard objectives like profit maximization subject to some market share constraint, the optimal firm for some environments has no fixed structure – it's oscillating.

The tool lets you run this experiment. You could populate the landscape with some company's characteristics that you know, and say, "OK, that's who I'm worrying about. How should I morph myself around that company's activities?" You could also say, "If I want to wipe that company out, what do I do? Maybe I do things in the short term that are quite different than the long term. What is my overall strategy?"

You could dig very deeply into data on actual companies and try to calibrate the model to find out what it is about a company that's making it fail to adapt or succeed in adapting. Companies and managers could actually use this to chart their adaptive course in the environmental dynamics, technical landscapes, regulatory environments, and so on, that they face. In the end, I see it as applied, practical decision-support software.

By plugging in data you could find a sort of an ideal genome for your company. Presumably you'd want to do this many times over a whole range of possible futures.

So what can we do with this information right now?

I always say the first thing to do is think differently. If your brain works differently, your actions will follow. This model invites managers to worry about the organizational diversity with which they're equipped. By that I mean the composition of individuals. And the really distinctive thing about the model is that it's bottom-up. It's all generated from individual information and individual behavior. It really takes very seriously the idea that an organization is ultimately just a bunch of people and takes into account the kind of people they are, how they interact, and their local rules of engagement.

In the end it tells us that it's the variable geometry firm that is optimal, not the fixed structure firm. That's a new way of looking at things.