ABOVE: © COSMO CONDINA / ALAMY

finance

Economics and Markets as Complex Systems:

A Postcard from the 2007 Complex Systems Summer School

By Daniel Rockmore

Every June, an effervescent and diverse mix of graduate students and postdoctoral fellows, plus a few junior faculty and even a handful of business folk, come to the Santa Fe Institute Complex Systems Summer School (CSSS). The CSSS is a four-week intellectual sprint designed to introduce a new generation of researchers to the ideas and techniques of complex systems. A four-week course can't possibly touch on every aspect of the discipline, so each week is given a theme. The 2007 CSSS began with a week-long introduction to the fundamentals of the subject, outlining its physical roots in chaotic dynamics and the basics of agent-based modeling. Week two provided a deep look at the new science of networks and at ecological systems—all of which was good preparation for the third week's focus on financial markets. Week four touched on allometry and scaling laws in biology—for which SFI has become so well-known—as well as epidemiology and evolution.

In this article we look at some of the ideas explored during "Markets Week."[1] The markets are a fantastic and data-rich example of a complex system. In fact, part of SFI's beginnings can be traced to the bringing together of physicists and economists (at the behest and with backing from then Citicorp CEO John Reed) to understand the global financial market as a complex system.

Beware the "Analogon"

The compressed form of the CSSS lets participants see the interrelations among different complex systems, be they ecologies, economies, societies, or brains. But good science requires more than analogies. Identification must lead to investigation, careful consideration, and brutally honest evaluation to see where an analogy holds, where it is useful, and where it falls apart.

Appropriately, Markets Week began with a cautionary tale of the use of scientific analogy in the study of economics and markets. "Neoclassical economics," the modern paradigm of markets in which price emerges out of the competing pressures of supply and demand, has been interpreted in the light of contemporaneous physical discoveries and theories. SFI Research Professor Eric Smith gave a wonderful thumbnail sketch of this history—which becomes something of a pre-history of econophysics. This dialogue is generally thought to have begun with the French economist Leon Walras (1834–1910). For Walras, the physical forces that could hold an object at rest, the so-called "balance point," were directly comparable to the "forces" of supply and demand that would eventually balance out to create a market equilibrium price.

As Smith pointed out, the analogy has many flaws—chief of which is that it neglects the possibility for a physical system to move forever about an equilibrium point. To the audience's delight, Smith introduced the term "analogon" to describe alluring but incorrect (or even dangerous) scientific analogies.[2]

Although Walras's analogizing of mechanics and markets was a misstep, it sowed the seeds of an interesting idea. The study of markets through the lens of physics and the tools of relatively sophisticated mathematics began to gain adherents, and the metaphor of mechanics was soon replaced by a new physical analogy, in the form of classical thermodynamics. With its laws describing the behavior of a perfect gas—relating temperature, pressure, and volume through equations of state—the analogy would be one of a market driven toward an equilibrium assignment of goods or an equilibrium price, like a gas in a container compressed by a piston. As Smith pointed out, this comparison also falls prey to analogons, chief among which is the path-dependent nature[3] of most economic processes. In recent work with External Faculty Member and New School Professor Duncan Foley, Smith gives the analogy more solid foundations, showing that for "quasi-linear economies"[4] there is a mapping or dictionary that essentially allows one to use classical thermodynamics to draw rigorous or formal conclusions regarding the market of interest. However, while the technical analogy makes sense, it is, Smith says, debatable as to whether or not these quasi-linear economies are a reasonable model of reality. The lesson is to apply a healthy skepticism at all points of the investigation.

The NASDAQ (National Association of Securities Dealers Automated Quotations), the largest electronic screen-based equity securities trading market in the U.S., lists approximately 3,200 companies

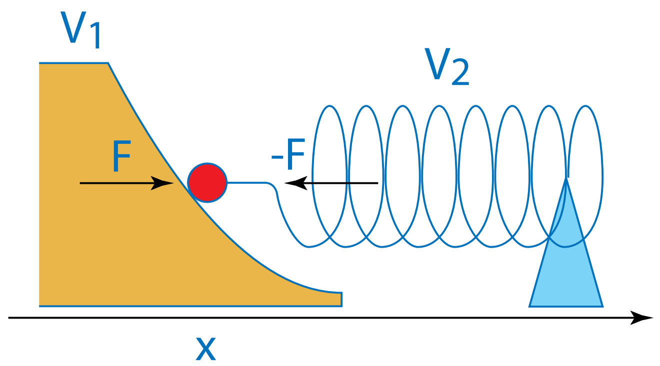

Figure 1: A cartoon comparison of forces on a ball rolling down a hill that will result in its equilibrium position at the bottom versus the "forces" of supply and demand that result in an equilibrium price for a good.

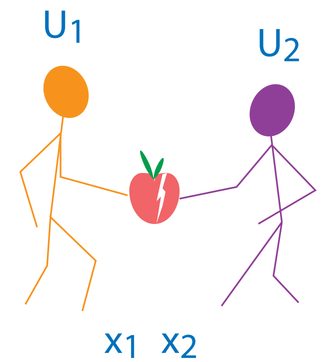

Nevertheless, the thermodynamics analogy is alluring, and closer to what you might call the complex systems ethos, of an emergent structure born of relatively simple interactions among a multitude of individuals. In thermodynamics, these are the trillions upon trillions of particles of varying energies, bouncing around some container according to the laws of classical mechanics, collectively giving rise to temperature and pressure. In the markets, it is a variety of many competing participants, from individuals to institutions, working in the marketplace, which gives rise to price.

Of course, particles in a box have no agency, and the degree of agency assigned to market participants is a critical component of any market model or simulation. The spectrum ranges from the "zero intelligence" model, which mirrors the random motions of particles in a box, to the traditional "homo economicus"—the fully rational and omniscient actor capable of optimizing behavior in light of total information. An interesting interpretation of "rationality" is to consider how much of market dynamics derives from institutional constraints as opposed to individual preference. In this view, zero intelligence means a dynamics completely driven by the rules of the marketplace, while homo economicus is a way to get at a model in which purely internal ("personal") concerns create social allocations without institutions. Presumably, reality lies somewhere between these two extremes, but even zero intelligence can provide a fair amount of insight into some aspects of price dynamics. In particular, Smith showed how even a rule-based marketplace accompanied by the inevitable mismatches that occur in the random arrival of buyers and sellers with random ranges of requirements and constraints (the "continuous double auction") is able to replicate many of the characteristic statistics of the price dynamics of the actual marketplace while simultaneously exposing various dependencies of the dynamics of order flow.[5] We may be closer to zero intelligence than we know...

The dynamics and market effects of order flow were also in part the subject of the lectures of Smith's colleague (and sometimes co-author), SFI Research Professor J. Doyne Farmer. Farmer discussed the interplay of order flow and market impact—the effect that a trade has on the price of the thing traded. He also explored what a science of finance and markets might look like, pointing out where current theory (one that puts the paradigm of equilibrium at the center of all considerations) falls down and explaining the importance of simulations based on more realistic assumptions.[6]

Both zero intelligence and full rationality have the advantage of enabling (in many cases) analytic solutions for their models. At intermediate levels of rationality, mathematical proofs are relatively rare. Instead, we can turn to the world of agent-based simulations and artificial markets. In these *in silico* markets, a large population of participants (the agents) trades a generally small number of assets. The agents have a range of strategies—behavioral rules dependent on their risk aversion and the state of the market—which is some simple function of market observables. These agents generally have "bounded rationality": they know something about the global situation, know something about history, and have some simple update rule for making trading decisions and evolving their strategies. In outline, this is similar to the many program trading formats widely used today—they are constrained by partial market information and have hard temporal constraints for executing their strategies. The agents' (and programs') strategies reflect traits such as degree of risk aversion, market depth, and memory of past market behavior. The fun (i.e., surprise) in these models lies in the strategies' evolution—poorly performing strategies go the way of the dodo, while productive strategies persist and possibly adapt, finding trading niches. Questions abound: Do various populations and update rules generate a market equilibrium? How sensitive are the dynamics to the changes in parameters? Are there recognizable regimes of behavior in the parameter space?

ALL DIAGRAMS COURTESY OF DANIEL ROCKMORE AND ERIC SMITH

The CSSS was treated to four lectures on agent-based markets by SFI External Professor and Brandeis Professor of Economics, Blake LeBaron. LeBaron was one of the developers of the Santa Fe Institute Artificial Stock Market, which was among the early agent-based markets. LeBaron went over a good deal of the basics on how these markets are constructed, and discussed some more recent work in which he and his colleagues have built artificial markets with heterogeneous "investors." These models generate market dynamics that replicate much of the behavior usually associated with the time series of a given stock price, such as distributions with "fat tails"[7] the "persistence" of trading volume,[8] the so-called "long memory" of volatility (meaning that price fluctuations persist), and correlations between volatility and volume. In particular, LeBaron finds connections between persistence and the modification of individual strategies via imitation of known successful strategies in an order-driven market.

Markets—A Top-Down Networks View

Agent-based models are "bottom-up": the system's laws of interaction are specified, with the hope that this seed will grow into the dynamics of the whole system. The art is not to simplify too much—remember Einstein's admonition: "Make everything as simple as possible, but not simpler." Alternatively, we might also take the system as it is, consider it from the "top-down" and try to characterize the emergent phenomenon of interest. Once again, this involves choices.

Understandings of a complex system are achieved when top-down and bottom-up approaches meet. In the case of an individual equity, replicating various properties of the price series (long tails, long memory, correlated volatility) in simulation hints that the model may be on the right track. A variety of basic statistical tools can measure the distance between simulation and reality. But what happens when we turn our attention to an actual market—say, a collection of assets whose behaviors are necessarily correlated—where strategies must account for the interrelation of a wide range of complicated assets. From 10,000 feet, a view of "the market" might look something like the cartoon in Figure 2—in which capital flows between the different "submarkets" of equities, bonds, etc.

The market is a network of large markets, which are in turn themselves networks of markets and so on and so on, until we arrive at some sort of tradable "atom." Any bottom-up model would ultimately need to reproduce this kind of hierarchical structure, a structure that seems to be both characteristic of financial markets, but also a general feature of many complex systems.

How to articulate this structure in a mathematical way? This was the topic of three lectures by Greg Leibon of Dartmouth College. Leibon showed how tools from statistical learning can analyze the correlation structure of the equities market, or the "network of equities," and reveal the hierarchical structure described above. The core object of study was the weighted network of equities, in which two equities are connected by an edge whose strength reflects the degree of correlation in the time series of normalized daily

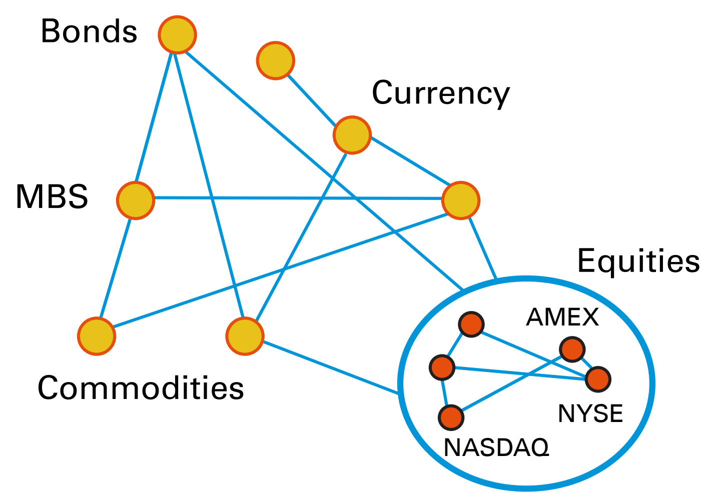

Figure 2: Model of the market network in which capital flows between the different "submarkets."

close prices over 15 years.[9]

Leibon's analysis treats the equities network as a spring model, in which pulling or pushing on one piece can send shocks to other pieces. Given this mechanical toy composed of thousands of masses (equities) and almost 2.5 million springs (correlations), the goal is to try to uncover collections of equities that generally move as one and then to try to understand the degrees to which subsets of the market are in phase or out.

A similar problem confronts a chemist or molecular biologist trying to understand (predict) the properties of a complicated molecule. These fields use the mathematics of "spectral analysis"—a subject that derives its name from studying the response of a substance to various kinds of radiation—to reveal the properties of molecules from the (slightly modified) network of connections that describe them. In a similar fashion, the same mathematics can reveal the structure of the equities network. In both cases, the "large" modes of the molecule—the perturbations that cause the most dramatic resonances—indicate the existence of large neighborhoods or clusters that effectively move together. Finding the large modes lets Leibon know how many big clusters to look for. What is of interest here is that "neighborhoods" of the market emerge in an unsupervised fashion, corresponding to labeled sectors such as technology, energy, and so on. Also striking is that the analysis can automatically recognize some known dynamics of the market, such as the cyclic flow of capital between sectors.

While some of the clusters reflect or replicate a traditional view of the market, comprising mainly equities from a single "traditional" sector, other clusters are highly mixed, suggesting that new kinds of hybrid sectors have emerged or evolved in our era of highly diversified, multinational, publicly traded entities. A new structure can suggest new markets, new derivative products, new opportunities for investment, as well as new strategies for ameliorating risk.

Each cluster can then be viewed as its own submarket, ready to be analyzed just like the original market. The process is then iterated until one reaches a submarket that is effectively completely decorrelated. The result is a more highly textured and geometric view of the correlation structure of the market. It suggests a "null model" of the correlation structure of the market that is much more interesting and sensible than a strawman "random" assortment of correlations or correlation matrices.

By analyzing the correlation structure of the market as a network, market analysis has been turned into a geometry problem (Leibon trained as a geometer and topologist). And the visual form of this network reduces the data of thousands of 1,000-dimensional time series to a three-dimensional representation that contains a good deal of the original information.[10]

Identifying this clustered and hierarchical structure is in a sense a data-driven form of equities classification. In Leibon's final lecture, he showed how "Bayesian nets" and more generally "graphical models," probabilistic tools for modeling complicated (i.e., highly non-independent) multivariate distributions, could then use this understanding as training data to solve the problem of classifying unknown or new data streams as (behaving like) members of a given sector (traditional or not). These tools can also be turned upside down to generate time series that behave like a member of a given sector. While such a "generative model" still doesn't account for all other kinds of market dependencies and interactions, it still provides a more sensible model of behavior for price series movement, given sector information. This is a first step toward some sort of meso-scale artificial market working at a level above participants and strategies.

These lectures were the material that made up Markets Week, but of course the subject matter was discussed in other weeks too. The compressed timescale of the CSSS causes the subjects to wash over one another, making participants think about possible connections and relations. For example, discussions of niche development, innovation, and evolution carried over to discussions of emergence in markets. Could the new work in food-web dynamics and extinction patterns give insights into risk management—or vice versa? Do the methodologies of functional neuroimaging and the hierarchical structure of neuroanatomy map onto the measurements and structures of economies and markets? Analogies or analogons? Only time and a lot of careful thought will tell. ◀

Daniel Rockmore is John G. Kemeny Parents Professor of Mathematics at Dartmouth College where he is also Chair of the Department of Mathematics and Professor of Computer Science. He is an External Professor at SFI and Director of the Complex Systems Summer School. This article benefited from the helpful comments of Robert Savell and Eric Smith.

1 Copies of the slides of almost all lectures from the 2007 CSSS as well as some related readings can be obtained at http://www.santafe.edu/events/workshops/index.php/CSSS_2007_Santa_Fe-Readings

2 Smith and his co-author SFI External Professor Duncan Foley attribute "analogon" to a 1949 paper of J. Lisman. Coupled with "ridiculogram," Mark Newman's name for the indiscriminate use of obfuscating network diagrams, we now have a growing vocabulary for modes of bad science.

3 In this context, path dependence for an economic process can be thought of as the idea that the equilibrium that is achieved (e.g., price) would depend on the manner (process) by which the price comes to be determined. For example, historical "accidents" could very well play a part. In the physics of an ideal gas, it would be the fact that work (for example) is a path-dependent function of the state variables pressure, temperature, and volume. The analogons arise when we try to find analogies for the state variables in an economic process and then "derive" path dependence—or the lack thereof.

4 A quasi-linear economy is one in which the utility function of the participants (how much they value any particular assignment of "goods," including money) breaks up into a sum of the cash on hand plus a utility function of the non-cash goods alone.

5 "Order flow" refers to the continual influx of buy and sell orders into the marketplace for a particular equity.

6 See also Farmer, J. D., D. E. Smith, and M. Shubik. "Is Economics the Next Physical Science?" *Physics Today* 58(9) (2005): 37-42

7 The notion of "fat tails" makes reference to an underlying probability distribution and the fact that there is a relatively large probability for the occurrence of large deviations from the mean, in contrast to normally distributed (bell-curve) phenomena. With respect to the markets, this is reflected in high volatility, so that the existence of fat tails makes risk management much more difficult.

8 Persistence (or hysteresis) refers to the ability to see the effects of moving large amounts of a stock well beyond the time of the actual sale or purchase—like a slowdown in traffic that persists well past the time of an accident.

9 A direct analysis of the simple time series of close prices results in all sorts of spurious identifications of correlation or the lack thereof. A standard first normalization is to consider the series of relative changes. In addition, we transform the data in order to take out market effects—that is, general effects due to the movement of capital between markets.

10 This is the content of a paper in preparation: G. Leibon, S. Pauls, D. Rockmore, and R. Savell, "Hierarchical structures in the equities market."

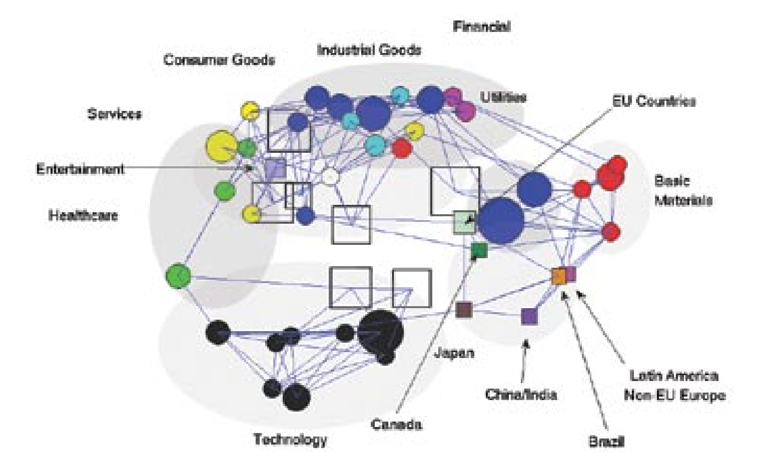

Figure 3: Network structure of the equities market after cluster analysis and dimension reduction. Labels reflect dominant sector or classification (e.g., country). Node size is proportional to cluster size. Unfilled squares indicate highly mixed (i.e., not easily identifiable) clusters. Connections are made for the (relatively) closest cluster centroids.