Speculative Trading Part 2

Evolving Models of Financial Markets

The physics of financial markets is an emerging science. In the second part of a guide to this new field, one of its exponents, **Yi-Cheng Zhang** of Fribourg University, outlines the theory behind some of his own models of the market-place

A new trend has arisen recently: more and more physicists have been attracted to economy-related problems. Evidence of this is the growth in the numbers of papers in physics journals devoted to theoretical and applied issues in economics and finance. In addition, fresh PhDs and seasoned researchers alike are finding careers in finance, new journals are being launched and conferences organized.

Trespassing on the domain of others is a notorious activity of physicists: their insatiable curiosity steadily pushes them into their near and far neighbours' territories such as biology, economics and other natural and social disciplines.

The current predilection seems to be for economics, especially finance. One can probably recognise two reasons for the interest. Firstly, physicists' tools are much in demand and so-called 'technical analysis' has become more and more complex so that an experienced physicist can offer skills that traditionally-trained economists lack. Secondly, such research topics present fundamental, intellectual challenges where the aim is to understand the basic mechanism. In this essay we shall concentrate on this second aspect. Moreover, we shall limit ourselves to modelling mechanisms of financial markets.

What Makes it Interesting

To appreciate why economics is interesting for physicists and what challenges there are, one has to know about the currently accepted theory in the economics field. I will not pretend to summarize the state-of-the-art here; a critical appraisal can be found in the Santa Fe proceedings (The Economy as an Evolving Complex System, ed. P.W. Anderson, K. Arrow and D. Pines, Redwood City, Addison-Wesley 1988). If one looks into the economics literature as a physicist, one may get the strange feeling that the theory is detached from the experiment. On the one hand, the theory is extremely refined and self-consistent with little effort made to compare it with empirical evidence; on the other hand, the experiments (*ie* market traders at work) are extensively performed with little reference to the theory. It is revealing to see how George Soros, one of the top players in global finance, considers the theory. His very success is an embarrassment to the orthodox theory; he considers it to be hardly relevant (*see* Alchemy of Finance, J. Wiley & Sons, New York 1994).

In short, current prevailing theory assumes equilibrium and its descriptions are mostly static. Little is said of how equilibrium can be attained, if it is attainable at all. Such descriptions look like 'mean field' theory of physics. The dynamics of markets can be found nowhere. However, insight gained in statistical mechanics, especially in non-equilibrium processes, may inspire a physicist to have a try in formulating a sort of dynamic theory for some economic processes.

There is no shortage of data. But data of crashes, for example, defy explanation. One is tempted to compare the current state of affairs to thermodynamics before Boltzmann or even Carnot – the framework has not yet been established. One does not know how to put the pieces of empirical law together to form a coherent picture. But let us not carry this comparison too far, since economics provides less precise data, and the fundamental elements are *thinking agents* as compared to the obedient particles in thermodynamics. One can never hope to get a future economic theory as quantitative and predictive as those of physical laws (Traders would die of joy if they could foretell prices as well as we can predict the weather). However, this should not deter us from searching for a framework to understand some basic phenomena qualitatively.

While hoping not to offend our colleagues, we might say that current research in statistical mechanics is somewhat stagnant, in contrast to the exciting times of the 70s and 80s. The 'soft' science of economics presents new challenges, new problems and new ways of thinking; it can teach us some new secrets of how Nature works.

How to build a model

Lacking a general framework one has to search for models in the dark. Before we present our own choice, let us recall that our aim has to be very modest; the best one can have is a sort of paradigm. One also has to keep the model as simple as possible, in order to say something general. One has in mind here the Ising Model which, despite its oversimplification, still offers insight into real magnetism; or the BTW Self-Organized Criticality Model (*see* P. Bak, C. Tang, and K. Wiesenfeld; Self-Organized Criticality *Phys. Rev. Lett.* **60** (1988) 2347), which apparently applies only to ideal sand piles but turns out to offer insight into many natural phenomena.

Let us list a minimal set of ingredients that are indispensable for modelling markets:

- A large number of independent agents participating in a market.
- Each agent has some choice available when making a decision.
- The aggregate activity results in a market price, which is known to all agents.
- Agents use this public price history to make their decisions.

We omit from the ingredients two important factors. One, no fundamental news (*ie* economic news from the outside world) reaches the market traders besides news of their own trading activity; and two, agents do not believe in any theory (*ie* traders do not derive their predictions from an established theory, but use some *ad hoc* personal rules; they learn from their own experience, and believe that the price history contains information). In this way we can begin to study the inherent dynamics of a market, in the absence of external influences, even though real economies have both internal and external components.

W. B. Arthur has advocated the so-called 'inductive thinking' approach (*see* Inductive Reasoning and Bounded

Rationality, *Am. Econ. Assoc. Papers and Proc.* **84** (1994) p 406–411) which corresponds to the opinion of a minority in economics. His idea is that since an agent cannot use theory to make a decision, his (or her) only choice is to learn from his own experience, as many a trader would attest to. Our own model is inspired from Arthur's El Farol problem, described in the above paper. We shall illustrate our ideas using two models. The first (the minority model) is intended to reveal the rich intrinsic market dynamics and general issues; the second (the trading model) is an attempt to apply the basic ideas to an artificial market.

The Minority Model

The simplest model we can think of is defined in the form of an evolutionary game. Let us consider a population of N (odd number) players. The game evolves discretely in time steps. At each time step, everybody has to choose to be on either side A or B. The payoff of the game is that, after everyone has chosen sides independently, those who are on the side of the minority win. In the simplest version all winners receive one point. Players make decisions based on the past record, which is common knowledge. But the past record only records which side was the winning side, without the actual attendance number.

The time series (the sides chosen or won) can be represented by a binary sequence, 1 or 0 meaning A or B is the chosen or winning side. Let us assume that our players are quite limited in their analysing power; they can only remember the last M bits of the system's results and can only make their next decision based on these M bits. Each player has a finite set of available strategies, S. A strategy is defined as the next action (to choose either A, B; or rather 1, 0) given a specific past record (of M bits). An example of one strategy is illustrated in *table* 1 for $M = 3$.

There are 8 ($= 2^M$) bits we can assign to the decision. Each configuration of 8 bits corresponds to a distinct strategy, this makes the total number of strategies $2^{2^M} = 256$. This is indeed a fastly increasing number, for $M = 2, 3, 4, 5$ it is 16, 256, 65536, 655362.

We randomly draw S strategies for each player from this pool. All S strategies in a player's bag can collect points depending on whether they would win or not given the past M bits and the actual outcome of the next play. However, these are only *virtual* points as they record the merit of a strategy as if it were used. The player actually puts into play the strategy which has accumulated the most virtual points, and gets a real point only if this strategy used happens to win in the next play. The method of using alternative strategies makes the players adaptive to the market. A player thus tends to maximize his capital (accumulated points) and his performance is judged only on his time averaged capital gain.

Several remarks are in order. By the very definition of minority, agents are not encouraged to form commonly-agreed views on the market. This is like real market trading: bears and bulls live together. In real trading it is often observed that a minority of traders first get into a trend (buying or selling), then the majority get dragged in. When the minority anticipates correctly and gets out of the trend in time, it pockets the profit at the expense of the majority. There are limited resources available for competition. If the players manage to coordinate well, per play they can expect $(N - 1)/2$ points, the maximum gain possible. Since our players are selfish, no explicit coordination is imposed and their fate is left to the market. The important question is whether they can somehow learn to spontaneously cooperate. $S = 1$ simplifies further the model so that, instead of the players, the strategies compete directly. The outcome is a trivial deterministic set of results. The extra layer of complexity at the player's level ensures adaptability.

This binary model is very suitable for numerical experiments. Damien Challet, a graduate student at Fribourg, implemented the game on a computer. (The preliminary results are reported in D. Challet and Y.-C. Zhang, Emergence of Cooperation in an Evolutionary Game, *Physica A* **246** (1997) 407). Note the word 'experiment' instead of 'simulation' is used at the beginning of this paragraph. This is to emphasize that we did not have precise goals at the start – as in an exploratory experiment the players are let loose to play and we observe. But we were rather amazed by the complex, rich consequences of the model.

record	decision
000	1
001	0
010	0
011	1
100	1
101	0
110	1
111	0

Table 1 An example strategy. The decision on which side to choose next is based on the last three winning sides (the record). If the last three wins were 001, then the decision would be 0. The eight decisions together make up the strategy

There are just three parameters in the model: N, M and S. However, there are hidden parameters, which is illustrated by looking at the total number of strategies. At first glance this number seems so huge that even for realistic parameter values, say $M = 10$, this number would be regarded as infinite ($2^{2^M} > 10^{300}$) for all practical purposes. And the number is so large that you would not expect changing it to have any effect on the model. But numerical experiments show that this is not the case. Depending on N the market has distinct behaviour, $M = 10$ can, in fact, be too large or too small for achieving coordination.

How can such a large number (10^{300}) still be relevant in this model? This apparently large number *is* irrelevant in the model; a much smaller (but hidden) parameter is actually responsible for the dynamics. A more refined analysis of the strategy space shows this.

Boolean 'Genetic' Space

It is instructive to represent the strategy space on a 2^M-dimensional Boolean hypercube. The $N_{tot} = 2^{2^M}$ distinct strategies are on the points (corners) of the hypercube. (This has a striking similarity with the construction of S. Kauffman's Boolean NK network, *see* The Origins of Order, Oxford University Press, New York 1993). Consider two neighbouring strategies which differ only in one bit (*ie* differ in one decision).

We say that the Hamming distance between the two strategies is one. The two strategies almost always predict the same outcome when acting on the past record: out of 2^M possibilities there is only one exception. Therefore, distinct strategies can be highly correlated. And players using correlated strategies tend to obtain the same decision, thus hindering their chance of finding the minority side.

Among the N_{tot} strategies there is a huge redundancy. If two strategies are uncorrelated, their decision outcomes should match with 1/2 probability. This is possible if their Hamming distance is 1/2 of the maximal value (2^M). We are thus led to count mutually uncorrelated strategies. This count will provide a crucial measure of diversity (or independence) of the N_{tot} strategies.

We recall now some useful basic properties of a hypercube. There is a subset of 2^M pairs of points (out of N_{tot}) in which, for each pair, the Hamming distance is maximum (2^M). They are anti-correlation pairs in the sense that the two strategies of a pair always predict opposite actions. All other Hamming distances among these $N_0 = 2^{M+1}$ points are 2^{M-1}, *ie* mutually independent. In other words, these N_0 points are composed of two groups of 2^M points each. Within the group the strategies are completely independent. Some of the cross-group links can be anti-correlated.

It is this reduced number N_0 which plays an important role in the model. If the number of strategies in the population (NS) is larger than N_0 then the players have to use strategies which are positively correlated. The herd effect (of most agents doing the same thing) is unavoidable despite the adaptability of the players, and will result in fluctuations larger than random chance would warrant, which leads to a waste of the limited resources. On the other hand, if $NS << N_0$ then the 'independent' subset of strategies indeed appears to be independent. NS cannot sample enough of the N_0 strategies, the anti-correlation Hamming distances are hardly represented (they appear with probability proportional to NS/N_0). The players in this case will appear as if they were using random strategies, independent from each other. Most interesting is the critical case when $NS \sim N_0$, *ie* when the reduced set is more or less covered by the population. The majority of their mutual Hamming distances implies independence, but a small part (about the square-root of the total) implies anti-correlation. This means that the players behave almost independently, the small number of anti-correlation Hamming distances help them to obtain opposite decisions. This is beneficial for everybody since coordinated avoidance makes the players more evenly distributed on the two sides, and the limited resources are better exploited. There are, therefore, three distinct phases:

1. The overcrowded phase, $NS > N_0$, where positive correlation is inevitable. The herd effect makes it worse for everybody.

2. A cooperative or critical region, $NS \sim N_0$ where the strategies used by N players are mostly independent (plus some anti-correlation links which achieve mutual avoidance). The resource is better shared.

3. A random region $NS << N_0$ where the anti-correlation is hardly present, and strategies appear to be independent.

A recent numerical study is in qualitative agreement with the above analytical result (*see* R. Savit, R. Manuca and R. Riolo, Adaptive Competition, Market Efficiency, Phase Transitions and Spin-Glasses *Michigan Univ. Pre-print*, December 1997).

Darwinism and Evolution

Between the over-competitive and random phases lies the critical region, $NS \sim N_0$, where cooperation emerges. This reminds us of Kauffman's often quoted motto "adaptation toward the edge of chaos". The competing population has different phases of collective behaviour, depending on the choice of N, M and S. We see from the above discussion that adaptation gives an evolutionary pressure to players' strategies, and they tend to be as diversified as possible. In the $NS > N_0$ region this pressure, or tendency, is frustrated since the strategy space (the hypercube) is too crowded. To relieve this pressure, one may want to change S or M for fixed N. Changing S is not very efficient. Increasing M is highly effective because of the exponential dependence.

Many an investor figures out his chances by simulation, and is then attracted to markets if he believes he has an edge. But often the subsequent participation can be disappointing

However, we would like to let the population evolve to this cooperative region without fine tuning the parameters. To this end we introduce another level of adaptation by adding Darwinism to this game of competition. The worst player is periodically eliminated from the game; the best player is allowed to produce an offspring. The new born player is a clone copy of his parent, but with his virtual capital reset to zero. And a small mutation rate is assumed to ensure genetic diversity – one of his inherited strategies is replaced. The total population is kept constant without loss of generality. The newly drawn strategy can have its M changed by one unit. This permits the players to find the best suitable M. The population can nevertheless reach a stationary state, clearly with inhomogeneous M. In short, the population is able to evolve without outside intervention and self-organize itself to find the critical region $NS \sim N_0$ which is beneficial for everybody.

Another interesting variant is to let the strategies fight directly without the players as intermediaries ($S = 1$). In order to have non-trivial dynamics going on, we allow a large number (N_1) of randomly drawn strategies to participate in the game passively, only a small number of strategies of the total ($N_a << N_1$) actively play. The passive strategies observe the system's posting board and count their virtual gain (as if they had taken part in the game). If one strategy's simulated gain is above a certain threshold, say a benchmark average of the actual game, it is bullish enough to come to the fore by becoming an active strategy. The worst strategies recede to the status of passive observers. Each strategy has two balance sheets, one for the accumulated real gain during its active playing periods; another for simulated gain counted from the last time it was relegated to the observer pool. In this variant the Darwinistic evolutionary pressure is also operative. Positively correlated strategies may do well as observers, but do poorly in actual play. Therefore, their active appearance is bound to be brief. Most interesting is to let the number of active strategies N_a free, everyone is allowed to actively play, and the market should decide how many actually do. Clearly, inhomogeneous M is better and mutation should also be operative in order to keep an adequate diversity. The system is like a grand-canonical ensemble. A large pool of observers is also a realistic feature of markets. Many an investor figures out his chances by simulation, and is then attracted to markets if he believes that he has an edge. But often the subsequent participation can be disappointing.

A Prototype Trading Model

The minority model, though very rich, still lacks some of the most basic features of a real market; the price, for example, which is determined by the aggregate supply and demand. We want to keep the above general ingredients to build a more specific trading model. In a trading model players have to decide when to buy and sell, just like taking sides A or B in the minority model. However, the payoff is not a fixed rule, it depends on the price movement which is in turn determined by the players' trading. In the modern market of stocks, currencies, and commodities, trading patterns are becoming more and more global. Market-moving information is available to everybody. However, not all the participants interpret the information in the same way and react at the same time delay. In fact, every participant has a certain fixed framework for facing external events. And it is well known that the global market is far from being at equilibrium, the collective behaviour of the market can occasionally have violent bursts.

Let us define our model more precisely. Each player is initially given the same amount of capital in two forms: cash and stock. There is only one stock in this model. All trading consists of switching back and forth between cash and this stock. Each player has a strategy that makes recommendations for buying or selling a certain amount of stock at the next time step. This depends solely on the price history. Player *i*'s strategy is an arbitrary non-linear function, $F_i(p_t, p_{t-1}, \ldots)$, positive (or negative) values suggest the amount of buying (or selling). The aggregate trading decides the price at the next time tick, p_{t+1}, using the law of supply and demand. Darwinism is also implemented here.

The results of this model are quite encouraging (*see* G. Caldarelli, M. Marsili and Y.-C. Zhang, A Prototype Model of Stock Exchange, *Europhys. Lett.* 40 (1997) 479). Despite the simplicity and the arbitrariness of the strategies, an extremely rich price history is created. A sample of P_t is shown in *figure* 1, which shows fluctuations of all sizes. During long runs, depending on the parameters, crashes occasionally occur with no warning. New features appear here with respect to the minority model. Even though the same self-organized structure is used here, the system does not appear to reach equilibrium, there is hardly any limit to the range of the price fluctuations. This is due partly to the continuous strategy space, as well as the law of supply and demand.

We have discussed two models of self-organized systems, in which players compete in a common market-place using the results produced by their own activities. We argue that this general scenario should also be present in real markets.

Open Questions and Perspectives

As one would expect from the early stages of any emergent scientific discipline – it has been baptised 'econophysics' – many different models have been proposed. We feel the current need is to learn how to ask the right questions about economic processes. By asking the right questions and by trying to answer them, we have to explore many seemingly isolated models and empirical laws to be able to set up a workable framework. Already, at the level of the simplest minority model, we see that many interdisciplinary subjects have without intention been touched upon, including self-organized criticality, population and Darwinism, ecology, information science, glasses and spin-glasses (*see* A Random Walk in Search of the Glass Transition in this issue), Kauffman's NK model and auto-catalysis, game theory (Prisoner's dilemma pits two players against each other, the minority game is a natural generalisation). Many of these relations deserve further study. Besides some relevance to economics, continuing this exercise of model building and playing is certainly rewarding.

The author would like to thank G. Caldarelli, D. Challet and M. Marsili for a fruitful collaboration

Thoughts on what econophysicists should be doing, written by Marcel Ausloos, appear on page 70, and are followed by a reply from Yi-Cheng Zhang

Past issue
See the last issue of *Europhysics News* (29 1) which carried an article on the Caldarelli-Marsili-Zhang model

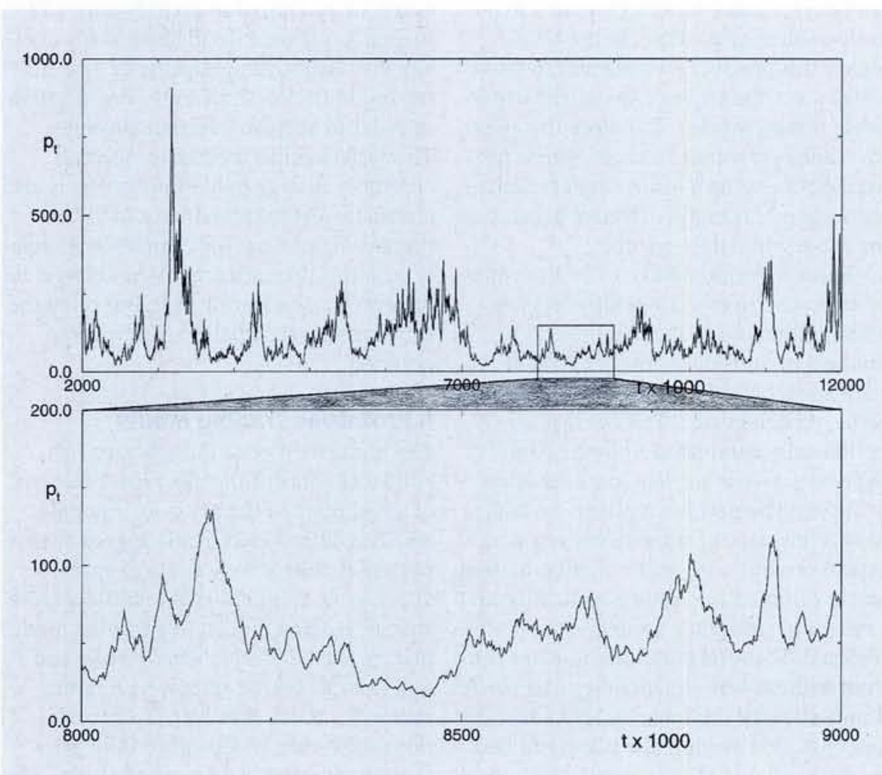

Fig 1 Price history of an artificial trading model. Fluctuations in the price look the same no matter what timescale is used, which is also true of real market data. In fact, the statistical properties of this artificial price index are in reasonable agreement with real data (*see* Mantegna and Stanley, *Nature* **379** (1995) 46)