Complexity Science and 21st Century Issues

London School of Economics 25-26 March 2004

Prof. J. Doyne Farmer, McKinsey Research Professor, Santa Fe Institute.

"The Low Intelligence Approach to Economics: Simple Laws Relating Order Flow to Statistical Properties of Markets"

Prof. Farmer presents a model of 'continuous double option' based on zero-intelligence noise traders which can be used to derive simple laws relating order flow to statistical properties of the market, such as volatility and the average 'best bid-best ask' spread that agrees remarkably well with data from the London Stock Exchange. He then discusses the possible ways in which agent intelligence might be added to the model and tested against real data.

Perhaps I should say something about the quirk of fate that brought me to thinking about this subject as a physicist. I spent ten years at Los Alamos and started a complex systems group there, but my background is in dynamical systems and in particular geodynamic and time series forecasting. We were forecasting things like ice ages and sun spots and so on and every so often someone would say to me: 'well haven't you applied this to the stock market?' And as I approached the ten year point and got my little nut dish commemorating my ten years of service I had a mid-life crisis, so instead of buying a red sports car I started a trading company.

We started applying our time series methods to price series and volume series and things like that and created predictive trading models. When we had some models that we thought were doing something reasonable, we thought we'd better find out something about the economics business so we got some books and looked up things like 'portfolio theory', tried it out and found that it didn't work at all. So then we looked at something called a 'capital asset pricing' model and tried that, which also didn't work. And that set me off thinking about what was going on and I discovered that there are things that do work such as 'options pricing' (and that really does work), 'mortgage pricing'; and I'm writing a review paper on critical analysis of efficiency, equilibrium and other ideas in financial economics. My friend John (?) and I work along in parallel. He wrote the entry in the new Palgrave Dictionary of Economics on general equilibrium so he comes from exactly the opposite school from which I do. We've, spent the last five years debating what equilibrium is good for and what it isn't and John also runs a 'hedge' fund based on the exact opposite set of principles from mine. My approach is looking for patterns in data that appear to be statistically significant and trade on without worrying about where they come from. John on the other hand says: 'let's make a model of a typical home owner and when that homeowner will mortgage their house or decide to re-mortgage'. So he creates decision trees which are right out of an economics text book, except for some detuning parameters for irrationality.

We have this debate coming from opposite sides of the spectrum and we're both making successful hedge funds and I'm convinced both work. In fact I'd invest my own money in John's fund to show that at some level equilibrium models can be valid. Hedging is a very useful thing and using portfolio theory for hedging is a very different thing from using portfolio theory to optimise your profits by having writedowns between winners and losers. That doesn't work because of statistical estimation problems.

The most useful things by far for a prediction company are exactly what are not supposed to be useful and those are stock anomalies. Financial newspapers say things like: 'if you invest in January you'll make a better profit' or 'if you buy following a positive earnings announcement the market will continue to go up'. But in my experience 'buy in January' never worked, 'buy on earning announcement' worked until about 1998 and then it evaporated. However there are hundreds of papers in the stock anomaly literature, all published in obscure places because they don't work, but that's what we found to be useful. Whilst most of the papers are bull shit, about 20% of them are interesting and whilst none of them are tradable by themselves, by combining them intelligently you can actually make something work. That experience set me to reading economics and the reason I'm here is because when you're a kid you want to learn the things that are most well known, but when you're grown up you realise that the fun things are those that are least well known. So being a physicist I viewed them as an opportunity.

Economics is hard because it involves people and of course that makes it challenging and interesting at the same time. You have to have some model for the way people behave and you have to have some model for the way they interact strategically. The standard solution for doing that is to assume everybody is a selfish, rationally utility maximising agent and that there's some kind of equilibrium. That's a sensible start and sometimes it even works but there are cases when it doesn't because the people aren't really so rational. The notion of irrationality in economics might be that the peoples' view of the world doesn't incorporate all the information. Then there may be a lack of equilibrium and the point that I think is underrated is that if the models you build are cumbersome, you have to work hard at anything you do to get things out. If you work really hard to some kind of Nash equilibrium then you've probably thrown a lot of useful stuff away, like the dynamics, or the kind of things that I'm going to address here.

So there's a trade off and sometimes it's good to take an alternative approach even if it sounds crazy. Certainly there's no easy fix in dealing with this problem and you have to avoid getting lost in the 'wooliness of bounded rationality'. The ways in which people are non-rational are complex and idiosyncratic, difficult to characterise and if you' re going to stop the research agenda until you have a complete model of human psychology you're probably dead in the water. So I'm going to use an extremely simple agent model, in this case just random behaviour, to understand what market institutions do, because I think we tend not to appreciate what powerful things market institutions or social institutions actually are. Social institutions are things that we have evolved over thousands of years to allow us to not have to think so hard and still do things that are very effective. If we can understand what the market institution really does and how it constrains our behaviour we can win.

Using the random agent model as a benchmark can tell us how agent intelligence really does affect things, because we can gradually introduce simple forms of intelligence and see how the simulation improves. There's a bit of a joke I like to tell which asks: 'how do economists play chess?' and the answer is that they draw to see whether they play white or black and then they spend the rest of the time arguing about whether the Nash equilibrium for white or black results in white or black winning. On the other hand you could ask how two physicists play chess and say they just randomly move the pieces round the board and argue about the scaling laws associated with that.

I’m going to take the latter approach but of course neither approach is a very good model for playing chess. If you move the pieces at random you have a hard time producing a checkmate and it may be better to find something out about the properties of chess so you can build a higher order model. That’s more or less how ’Deep Blue’ works. It investigates one move ahead, two moves ahead, three moves ahead and systematically investigates all possibilities which is roughly like moving at random and having some score function to score how these things go.

Now I want to emphasis that I’m not the first person to think about this approach to economics. A lot of people, both economists and physicists have shown that you can reproduce economic data and get a fair approximation of what’s going on by assuming random agents.

So I want to address the question: 'what drives the changes in prices?' Most of you will know that the standard story goes more or less as follows. You have expectations about future earnings that are driven by new information. We all have a model of the world and we evaluate new information that flows in which alters our expectations of earnings. We then do something that causes a price change and prices are unpredictable simply because new information is random.

A paper by Larry Summers and colleagues looked at the moves on the American stock market (SMP index) from 1946 to 1987 in rank order as percent movement from absolute value. If we take the top 100 companies and rank the size of the move and note the date we can see what kind of news seemed to drive the move. So if we take the newspaper headline of October 19th 1987 it says: 'worry over dollar decline rate deficit', 'fear of US not supporting dollar'. Was this news? I can tell you as a market practitioner I never went through a day without experiencing worry and fear. If the person managing your money isn't that way then you take the money away because if they aren't worried they should be. However there is often real news like: 'Outbreak of Korean war' or 'Eisenhower suffers heart attack' which is affective. So the point is that news, whilst it does seem to have something to do with market movements and does seem to have some correlation with large market movements it doesn't seem to offer a complete explanation. People are quite good at responding to the general direction of news i.e. if there’s an increasing stream of information about Putin having a heart attack most people can guess which way that's going to drive the rouble. On the other hand if Putin has a heart attack the question of where the new value of the rouble should be over the long term instead of the short term is a very hard question.

What about ‘fundamentals’ as a source of information? I'll show you a famous slide by Campbell and Chiller (?), showing how hard it is to make the argument that fundamentals and prices really match up. You can see on the graph that the dashed line represents fundamental values and people have tried many different contents for them. This one's based on dividends and the solid line shows the prices from some DOW Jones index. What you can see is that the two things can be out of line sometimes for decades and be out of line by a factor of two. So there are problems in making prices match fundamentals. Admittedly this may be because people don't understand the right way to assess fundamentals, but the fact that we have no way of doing this means that we cannot test them properly. We would really have to know things like a company’s utility function and have a proper measure for information arrival and these things are difficult to determine.

A famous theory that was used for things like option pricing was that prices describe a random walk and this was due to Bachelier (?) around 1900, five years before Einstein introduced the random walk model to describe the photoelectric effect. Even so, one of the things that isn't known in this very simple model is the diffusion rate, i.e. the volatility with which prices jump around and move in a direction. What does that depend on and why do price changes have fat tails? In other words a normal random walk is a good starting model but it's more complicated than it seems. Large steps in the random walk are more common than they ought to be and we've already seen that out of the twelve largest moves, five were in one month and two were in another month. That's not very likely to happen at random. If we plot volatility over more than 100 year span where volatility is the standard deviation of daily prices month by month, we can see periods like the Depression where things were extraordinarily volatile and we can see other periods where prices were much less volatile.

We need to get some idea, maybe just of the intermediate drivers, of what is going on and my goal in this task is to understand factors such as risk, price volatility and liquidity. Liquidity here means the impact of trading on prices. If someone suddenly trades a million shares how much will it push prices around? What is the effect of transaction costs? How different are the prices for buying and selling which gives an idea of the friction in a market? I suggest we can do this by investigating the mechanics of the 'continuous double option'. 'Double option' is the price adjustment in the orders book between 'to buy' and 'to sell' and assumed to take place simultaneously. It's also continuous because transactions can be made anytime.

The best buying price is actually called the 'best bid' and the best selling price is called the 'best ask' and the difference between them is called the spread. We'll also have two kinds of order: 'market order' and 'limit order'. Market orders result in an immediate transaction; limit orders show a degree of patience. I mean if you submit an order and say; 'I want 1000 shares and I don't care what the price is', then it will have an immediate effect. On the other hand if you give an order to buy or sell and specify a price then that's a limit order.

We can represent the volume of sell and buy orders on a continuously updated graph above and below the line respectively and show the effect on the best bid or best ask price. So the program here is to investigate how prices are being formed at a microscopic scale and then build back up. OK, so market orders arrive and cross with existing limit orders and transactions are made. We also have cancellations of orders and each transaction eventually results in a price change. This is the model and it's not just some change in expectations but the mechanical order arrival and transaction that does it.

My collaborators, Julie Laurie, Eric Smith, Marcus Daniels, Lazlo Jumo and I built the simplest model we could think of in which buy limit orders rain down randomly from above and sell limit orders rain up. Market orders don't arrive at random and we have different rates as we do for cancellations. The model dynamic is that the boundary conditions for order placement respond to price formation and the key thing that we're capturing, that's very hard to capture in equilibrium models, is the feedback between order placement and price formation. This is actually an extreme version of the notion that prices are important since prices in this model are completely informative because all the traders do is respond to where prices are formed by changing the boundary condition of where orders are placed. It's a very ad hoc and arbitrary distribution for order placement. By plotting the distribution of orders that are eventually executed actually gives you a reasonably looking graph that's peaked towards the centre which is what the real data looks like.

It's a simple model but it's quite hard to solve because it's really a queuing problem with two queues which are dynamic with dynamic boundaries influencing each other. The buying limit order divided by the sell market order process determines the best bid which influences the boundary condition for the sell market order/buy market order process and that's all coupled. So as an analogy to a physics model it's actually quite hard and one of the ways that we solved it was with sol/gel transitions involving evaporation deposition problems. This is a kind of simpler version of the problem here.

Without going into any of the mathematics let me give you an overview of some of the things we have done to this model. One of the things was to use a physics engineering trick called dimensional (?) analysis which allows us to both simplify the analysis of the model and make first cut predictions. This is exciting because normally we're brought up to think of physics in terms of length, time, mass, charge. These are the sort of sacred things that everything is built out of and knowing that the world is built out of those things we can guess formulae and so on. We'll do the same thing here except the fundamental things are price shares and times. We have a different notion of time here, because economic time in this model is determined by the cancellation rate so any time interval you look at divided by the cancellation gives you a rescaled time that simplifies things.

We're also able to use mean field theory which is a fancy word for an approximation technique we use in physics. One of the big differences between physics and economics is that in physics we don't prove theorems, we just learn to approximate things. In the simulations we came up with scaling relations like the ones that Geoff West showed you which we could test against real data. The data set that we used comes from the London Stock Exchange, and shows you how many shares are sitting there in the limit order book at each price level. So it's a completely transparent market where everything has been electronic since 1998 so we can see every action by every trader on every stock making a data base of 350 million events over a four year period.

This means we can actually measure the order flows directly. Here's an example of the data that we get. This is Shell Oil Company plotting times on this axis at 2002 and price on that axis and the colour code represent the supply and demand and the volume stored at each price level Somebody places an order 'buy at this price level' and it sits on this order to buy for months. Eventually it's executed here and black means not many shares, white means not any shares , blue means a few, red means a whole lot and green or yellow means something in the middle. So you can see that most of the supply and demand is concentrated in this random walking curve which is where the transactions have actually taken place. That's the interface from our point of view. Things above the curve are orders to sell and it's the collision between buying and selling in this concentrated region of supply and demand that's actually making the price diffuse around.

Just for the heck of it I'll show you same thing on a daily basis. You can see the boiling region in the centre where things are happening very fast and as you move away you can see the fine scales are rolling out enabling you to see certain regularities in this picture right away. You start looking at this data and you automatically start seeing certain regularities and what we're trying to do is characterise those regularities.

OK, just to review the model that I mentioned a moment ago. It has five parameters. The limit order rate, the market order rate and the order cancellation rate and these turn out to be the really dominant things. The fourth parameter is typical order size and the fifth is discreteness in the same way that a beach with fine grains of sand is different from a beach with pebbles. All of these can be decomposed into the three fundamental quantities of number of shares, price and time. Bearing these parameters in mind we make predictions that the typical size of the spread (distance between the best bid and the best ask) should be the ratio of the market order rate to the limit order rate times some very simple function, non additional function of the following non dimensional combination of parameters; order size times cancellation rate divided by market order rate.

$\mu / \alpha f(\sigma\delta / \mu)$ where

α = limit order rate
μ = market order rate
δ = order cancellation rate
σ = typical order size

We can derive that from field theory and our graph shows the predicted values of the spread against the actual value of the spread. We took 50 stocks and measured the average values of these parameters over a two year period so it's not a prediction like saying the stock tomorrow will have this or that spread. It's more like 'the ideal Gas Law' where if you know, the temperature and volume of the gas you can make a prediction about the pressure. You can see the parameters for the 50 different stocks and you can see that they roughly cluster along a line, the slope of which is very close to 1, as it should be. This is on a $\log_n.\log_n$ scale, but the basic scaling relation is captured. Similarly price diffusion rate or volatility should scale as some non trivial ratio of market order rate, cancellation rate, order size and limit order rate and we see the data clustering along the line. The slope is actually a little greater than 1 indicating the scaling relation isn't quite right but it's in the ball park.

Another thing of interest to anyone who trades in a financial market is 'market impact' which tells you when you try to trade how much you're going to move prices. In this case we define it as the instantaneous change in price caused by the price immediately after the order arrives minus the price before it arrives. When an order arrives we start with the best ask price and see a shift in the ask price and that's the market impact. If we plot this by taking the order size and units like pounds or dollars and look at the price shift we see a bunch of different curves for each stock. In the model we have to plot the same data in non-dimensional units so we take the number of shares and we multiply by the order cancellation rate divided by marker order rate which gives us a non dimensional quantity relating the number of shares to the price. Similarly if we take the price shift and multiply it by the limit order rate over the market order rate then we see the data for the stocks collapse onto the same curve illustrating that we now understand something about what is influencing price formation.

In spite of the fact that in Russia the translation of 'random agents' led to the belief that model traders were more or less complete idiots, I don't want to argue that you can understand everything about a market by assuming that traders do nothing but flip coins. They do process and incorporate information and it's important to try to understand how that works. There are a lot of regularities that the model I have just presented to you doesn't capture and one is that volatility is heavily articulated in time. A big price move today doesn't say anything about the size of a price move tomorrow per se, but it does say something about the probability. Actually if you take the relationship function and plot it on the log/log scale you see that the tail is a straight line, meaning there's a power law function tail. It's well known that price fluctuations have fat tails and it's a probability that you'll see a price return greater than a certain threshold. This is something that Mandelbrot discovered back in the 60s that raised a lot of interest in the physics community and there must be 15-20 theories as to why this is so. I'll just show you a picture of what that data looked like. This is actually an average over ten different stocks and we just plotted the log of the probability that a return was greater than the threshold against the size of that threshold and you can see it seems to be roughly a straight line.

These are things that we would like to understand and now we work the problem back from the other end. Instead of starting with the model we work back from the data asking what the key thing is that's driving these fat tail phenomena. My collaborators were two graduate students at the Santa Fe Institute and the first thing we did was look at the data to see what might drive the return at the individual event level. We took each price change triggered by market orders and asked what actually triggered that price change. So, the market order hits the book and causes the bid offer to move some way. We then segmented the data into five groups based on the size of the order hit and looked at which of the five groups of order size caused an event to happen. What we see is a very similar distribution for all five groups and no real tendency for larger orders to cause larger price changes, which is a puzzle because it seems intuitive that larger orders penetrating the book would cause larger changes. Looking at some time frames of big events in Astra Zeneca stock we can see the best ask price sitting there and a buy market order arrives and knocks that out and we see the subsequent event. So there's a large price change and it occurred because there's a big gap in the limit order book. The distribution of price moves in the book precisely reflects the distribution of gap sizes in the limit orders. The pattern is set by how many buyers and sellers there are, but it's actually not the number, it's the way in which they cluster their orders in the book (in physics jargon it's a finite size effect) and depends on the fact that there are a finite number of orders in the book and the granularity.

What we're trying to do is build a better empirical model for real order flow working the problem backwards.We notice quite a lot of fascinating regularities in what's going on inside this limit order book. For example, a $\log_n.\log_n$ plot of the probability that somebody places an order at a distance from the corresponding best price gives a straight line and indicates another power law. In fact buy and sell orders lay over each other across almost the whole range. A typical British stock of about twenty pounds, has a tick size of about a pence and measuring ticks we see agreement all the way out to about 2000 curves given the asymmetry between buying and selling (?).

Another fascinating regularity is shown by making orders + 1 if they're buy orders and -1 if they're sell orders and then taking an order correlation function of the series. Again we see a linear region in the curve. Furthermore the exponent of this power law is less than 1, which means that the order correlation integral is infinite illustrating that it has a long memory i.e. events in the long past have a substantial influence on the future.

The model has its limitations because it contains arbitrary parameters. The random orders falling out of the sky are, in economic jargon, 'random liquidity demanders' (these refer to the guys that randomly place market orders or say: 'gee, I need some money now never mind what's going on with the stock). This gives rise to arbitrage in prices and asymmetries in the order book. What we do want to see is the effect of adding intelligence a bit at a time. We could draw an analogy that markets are a lot like biology. You have variation in that people are constantly trying out new strategies. The ones that are successful get selected and propagated through time either because the same person goes on using them or somebody moves to a new firm. Trading firms represent a very diverse ecology and people use very different strategies in different places and the role that diversity plays is poorly studied. We should be able to characterise typical strategies in markets and maybe explain why those strategies exist. The reason for their existence may be really important in understanding what goes on in a financial ecology. So we're building a model where we deal with liquidity demand as food for arbitrageurs who are like predators in an ecology of arbitrage.

We must also take account of market makers and agents who exploit order imbalances and tactical trading to see how the statistical properties of the market are affected. Prediction companies, for example, have to understand market impact, and understand what the spread really means. Designing an automated market maker is a very active topic in the financial community these days and regulators might ask whether it's possible to reduce volatility by putting incentives on patient order flow and charging for impatient order flow. If we change the order flow by x amount can we predict how much it will change volatility and spread? We can also maybe use this to detect poor markets by seeing which markets lie above or below the scaling law curves.

So to conclude. We found some partial solutions to some classic problems by dividing the problem into two parts. Rather than trying to do things from first principles we first tried to understand how order flow affects the market. We now need to determine what gives rise to order flow and we haven't even begun to study that. Nevertheless it seems that whenever you can divide a problem into pieces it's a good idea. We have the advantage that in the model we created the parameters are all directly measurable so we can go out and test the theory in a fairly unambiguous way which is one of the major problems with most equilibrium models for markets. They result in theories that are very difficult to test properly and difficult to falsify. Part of what we said is that it's liquidity that's really the important driver of price formation and liquidity is modulated. The key parameter is high market order rate which indicates impatience and a high limit order rate indicates patient behaviour. Likewise high order cancellation rate would indicate impatience and low would indicate patience. So it seems that whatever information arrival is doing and however it's affecting the market, that a key aspect is whether it's making people more or less patient.

I hope I have convinced you that it's important to understand the dynamics of the market institution itself. In a sense what we've done is throw random stuff at this black box which is the continuous double auction algorithm and look at its response function. It's very much the way an engineer would go about trying to understand the components in an engineering system. As a start it makes sense to look at zero intelligence models and add the intelligence incrementally to try to see or pin down exactly which statistical properties of the market depend on the actual intelligence or other aspects of agent behaviour.

Questioner 1: Is the model sensitive to a distinction between institutional investors and individual investors? Do you try to go to actual markets for information on that to make your models?

Answer 1: Yes, we're looking at a 350 million event data set for the London Stock Exchange and there is a kind of population law associated with that because if the book is neither building up nor shrinking then for every order that goes in an order has to be going out and so the conservation laws on average are that the number of limit orders equals the number of market orders plus the number of cancellations. So in fact the number of limit orders is pegged on 50% and the number of market orders and cancellations tend to trade off. Now we see variations in the size of those orders and in the density with which those orders are concentrated. When the market gets volatile the orders are spread out across a large area and when less volatile, they tend to be concentrated more at a simple price. This is one of the key elements we're trying to capture in a more refined model. The other feature this data set has, which we've just begun to exploit, is that we actually have codes for the participants and this was why I was arm wrestling with the London Stock Exchange because they stopped providing that information. We have about four years of data and we can actually track over long periods in some cases who the participants are. We don't know which are institutional and which are not but we can actually see the individual behaviour of the participant codes and it's quite fascinating because what we see is that the different participants in the market give really big effects. So there's a fascinating level of diversity in the data and everywhere you look you see something amazing. We're trying to document this heterogeneity and make use of it in the model. One thing we have shown is the long memory: some of the agents order flow by themselves displaying the long memory, other agents don't show it at all. And that gives a clue about what must be causing long memory but what it is we don't know.

Questioner 2: Might this technique be useful for detecting insider trading? That would be like the reverse of what you're trying to do.

Answer 2: A symptom of insider trading is that you should see a price movement before it is announced so that's a most useful way of spotting insider trading and this initiated a dialogue with the New York Stock Exchange (NYSE).The London Stock Exchange (LSE) is sort of a free fall market with no designated market maker. In contrast at the NYSE, the US government gives a monopoly to different players called specialists who have the regulatory obligation to make an 'orderly market'. So for example, they're not supposed to let the price move by a few ticks at a time because that would inhibit orders. Now if we could get access to the kind of order flow of data that we have for the LSE we could make those same plots that I showed you for the spread and for the volatility and we could see who's below the curve and who's above it. The problem is unless you have access to a benchmark like that you might say: 'gee those shares are awfully high' and they'll say 'well that's because my volume is low or my volatility is high because it's this kind of stock and this is….' and some story like that. But if you could provide a benchmark that scales out some of the excuses that might be quite useful.