The Physics of Finance: Collective Dynamics in a Complex World

Lisa Borland

Integral Development Corporation,
3400 Hillview Ave, Palo Alto, CA USA
lisa@evafunds.com

Abstract. Exploring the dynamics of financial time-series is an exciting and interesting challenge because of the many truly complex interactions that underly the price formation process. In this contribution we describe some of the anomalous statistical features of such time-series and review models of the price dynamics both across time and across the universe of stocks. In particular we discuss a non-Gaussian statistical feedback process of stock returns which we have developed over the past years with the particular application of option pricing. We then discuss a co-operative model for the correlations of stock dynamics which has its roots in the field of synergetics. In all cases numerical simulations and comparisons with real data are presented.

Keywords: finance, fat-tails, long-range memory, statistical feedback, correlation dynamics

1 Introduction

The field of finance is one with a rich history of its own. Yet in recent years it has attracted more and more people with scientific backgrounds, especially from the field of physics. This is not only due to the fact that physicists are particularly well-equipped to tackle many of the challenges of mathematical finance which provides alternative career paths on Wall Street, but also because vast amounts of financial data are now available. Nearly every transaction on a tick-by-tick basis for thousands of stocks and other financial instruments is being recorded electronically as we speak, resulting in something akin to a huge data-base of one of the most complex systems we can imagine. Hardly the result of a controlled experiment in a physics lab, the price formation process of a publicly traded asset is clearly the product of a multitude of evasive interactions. Individuals around the globe post orders to buy or sell a particular stock at a particular price. Transactions are cleared at a certain price at a given time, either by passing through the hands of a specialist on the trading floor, or automatically on the many electronic markets which have sprouted in recent years. Apart from fundamental properties of the company whose stock is being traded, factors such as supply and demand clearly must affect the price of stocks, as well as general trends in the particular industry in question. Stock specific events, such

A. Pelster and G. Wunner (eds.), *Selforganization in Complex Systems: The Past, Present, and Future of Synergetics*, Understanding Complex Systems,
DOI: 10.1007/978-3-319-27635-9_6

as mergers and acquisitions, have a big impact, as do world events, such as wars, terrorist attacks, and natural disasters. A recent example of this are the dramatic events seen in 2007 and 2008 which are perhaps due to fundamental flaws in our credit-based economy. The current European debt crisis is another example.

Stocks are traded for the most part on a central limit order book, such as the New York Stock Exchange. Modeling the intricate dynamics and microstructure of this order book is a field of study which has gotten some traction in recent years [1, 2]. When comparing to physics, that level of description can be seen as the microscopic level. However, it is often more tractable to use a mesoscopic description which aims at describing the price process as a stochastic Langevin equation where the key feature is how to capture the volatility, or noise, that drives the process. This is the most important effect since stock price changes (or returns) form moment to moment are essentially unpredictable so the deterministic part of the equation is less interesting (though of course, if you can predict it ever so slightly then you might become quite wealthy!)

For many years and in a large body of the financial literature, the random nature of price time-series was modeled by most as a simple Brownian motion. The first to propose such a model was Bachelier in his thesis in 1900, which lay largely undiscovered until much later when Black and Scholes wrote their famous paper in 1973 based on a very similar model. They made important contributions in particular to the pricing of options, for which they received the Nobel Price [3]. Options are traded instruments that give the right, not the obligation, to buy a stock at a later date at a certain price, called the strike price. In Black and Scholes work the log price is assumed to follow a Gaussian distribution and even today many trading assumptions and risk control notions are based off of that prior.

However, in more recent years there has been a large body of work which - in quite some detail and statistical accuracy due largely to the vast amount of observations available - has been able to document that the time-series of financial market data show some intriguing statistical properties, which deviate quite substantially from the Gaussian assumption. These features are referred to as stylized facts (cf [4, 5]). It is interesting to note that many of the stylized facts appear to be universal, in the sense that they are exhibited by a vast variety of financial instruments, as different as commodities like wheat, currencies such as the Euro-Dollar rate, and individual stocks. Some are also exhibited over various periods of history (and so can be seen as stationary), others are exhibited on a multiple of time-scales (and so can be seen as self-similar). Other interesting properties pertain to the dynamics and statistics of the cross-section of financial instruments. Ultimately, the goal is to comprehend and model the joint stochastic process of the price formation dynamics of the collection of stocks across time, so investigating various statistical properties both across time and across stocks is essential. The challenge then lies in coming up with a model that captures the dynamics inherent in the data. In addition, it is desirable that such a model is somewhat intuitive, parsimonious, and analytically tractable.

In this chapter we shall review a class of models that we have proposed in recent years, aimed at modeling the stock price dynamics in such a way as to capture as many of the statistical properties of real financial data as possible. We also show how these models could be used for important applications such as the pricing of options and other derivative instruments. In the first part we focus on financial time-series, and in the second part we look at cross-sectional dynamics across a universe of stocks. In all cases we shall see that notions of nonlinear cooperative feedback appear to be essential ingredients of this very complex real-world system.

2 Stylized Facts across Time

2.1 Returns

While a random walk with Gaussian noise will show fluctuations that are relatively constant over time, the case is quite different for a time-series of stock returns, usually defined as relative price changes or log price changes. As shown in Figure 1, there appear to be intermittent clusters of higher versus lower magnitude returns. This phenomenon is known as volatility clustering, and we will look at that in more detail further along. The probability distribution of these financial returns are typically fat-tailed. In fact, it has been shown that the distribution of both intra-day and daily returns can be very well fit by a power-law tail of about -3 (referred to by some as the cubic law of finance [6]). This tail index is consistent with that of a Tsallis distribution [7] with index q in the range 1.4-1.5, which fits very well to that of the returns (Figure 2). We shall work a lot with this class of distributions (which are equivalent to Student-t distributions), because this very same distribution fits to returns from a wide variety of financial instruments, such as stocks, currencies and commodities, with much the same q-index. Furthermore, this fat-tailed nature of returns holds for data from several different geographic regions such as North America, Japan and Europe. Because of this rather universal behavior, the non-Gaussian distribution is an important stylized fact. Clearly then, any model of financial data should try to capture at least this important feature. A related stylized fact is constituted by the observation that, as the time lag over which returns are calculated in increased, this power-law behavior of the distribution of returns persists for quite a while, decaying slowly to a Gaussian distribution as the lag increases, becoming indistinguishable from the Gaussian as the lag approaches something on the order of a few months [5].

2.2 Volatility

Volatility is typically defined as the square root of the squared variation of returns. Obviously, this definition is not unique because one can choose to calculate the squared variation over an arbitrary historical period. As it turns out, the statistical properties of volatility will not be too sensitive to this exact choice. What

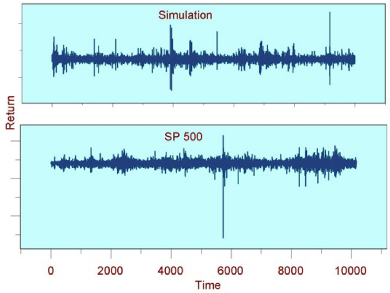

Fig. 1. Market Returns (bottom) and from our multi-time scale model (top)

do we know about volatility? Why do we care? In order to answer these questions, our view is to stand back a bit, and let the data speak. Are there consistent patterns in the statistical properties of the data? Are they universal? If financial data were Gaussian distributed, then the volatility would be the standard deviation of that distribution. In fact, this assumption is still made both by many practitioners and in many theoretical works in mathematical finance, although there is the increasing awareness of the shortfalls of such an assumption. In practice, for example, the devastating effects of the Gaussian assumption were seen in August of 2007 and October of 2008, when investors panicked claiming that 25-sigma events were wiping them out. In fact, these types of statements are only true if the underlying distribution is assumed to be Gaussian. If a heavy tailed distribution such as the Tsallis with q = 1.4 or 1.5 is assumed, then the behavior of the markets as we have seen them in recent years is to be expected, with a probability of about one or two extreme events per decade. Clearly it is extremely important for hedging and risk control purposes to have a richer understanding of volatility.

Data shows that time series of the volatility of the Dow Jones index for the past century exhibits clear periods of lower and higher volatility, typically clustered together. Note that while the distribution of returns is fat-tailed, volatility itself follows a close-to log-normal distribution. Interestingly, the same type of statistics is valid for volatility calculated every 5 minutes intra-day, implying a self-similar structure.The fact that volatility is self-similar on different timescales is a universal feature. The clustering feature that is observed is a signature that there is memory inherent in volatility; if volatility is high, it will persist for a while. This memory can be quantified by looking at the autocorrelation of volatility over increasing time lags. Such an analysis shows that there is a persistence which decays as a power law as the lag increases. There is also the feature of causality in the structure of volatility. In other words, some statistical properties of future volatility conditioned on past volatility are not invariant to reversing

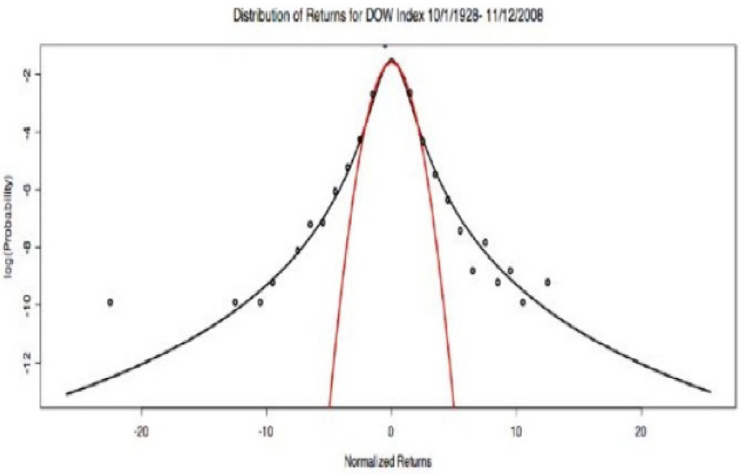

Fig. 2. The distribution of returns is well fit by a Tsallis distribution

the time order of a volatility time series. This reinforces the notion of memory and a dependency of volatility on its own past behavior. Another stylized fact of volatility that exhibits some asymmetry is the so-called leverage effect, which describes the positive correlation between negative returns and volatility; large negative price drops will give rise to subsequently higher volatility.

Some other interesting statistical features relate to the conditional volatility. Specifically, if you look at the probability of observing a volatility of a certain magnitude given that a volatility shock larger than some threshold was just observed, you will see something that translates into an Omori law for volatility. Just as earthquakes are followed by aftershocks, so are volatility shocks followed by other larger than normal shocks.

3 Stock Price Models

Several different models [8–10] have been proposed in an attempt to capture fat tails and volatility clustering which don't exist in the Gaussian Bachelier or Black-Scholes model. Popular approaches include Levy processes, which induce jumps and thus fat tails on short time- scales, but convolve too quickly to the Gaussian distribution as the time-scale increases. Stochastic volatility models, such as the Heston model where the volatility is assumed to follow its own mean-reverting stochastic process, reproduce fat tails, but not the long memory observed in the data. The same holds true for the simplest of Engle's Nobel prize winning GARCH models in which the volatility is essentially an autoregressive function of past returns. Multifractal stochastic volatility models (similar to cascade models of turbulent flow) are another promising candidate [11], reproducing many of the stylized facts, lacking mainly in that they are strictly time reversal symmetric in contrast to empirical evidence. In addition, most of the above mentioned models are difficult if not impossible to deal with analytically. Analytic tractability is desirable for reasons such as efficiently calculating the fair price of options or other financial derivatives which in their own right are traded globally in high volumes and will be discussed in more detail later on. For now we focus on presenting a somewhat realistic model of stock returns themselves which we developed a few years ago [12–14]

3.1 A Non-Gaussian Model of Returns

The standard Black-Scholes stock price model reads

$$dS = \mu S dt + \sigma S d\omega \tag{1}$$

where $d\omega$ represents a zero mean Brownian random noise correlated in time t as

$$< d\omega(t) d\omega(t') >_F = \delta(t - t'), \tag{2}$$

Here, μ represents the rate of return and σ the volatility of log stock returns. This model implies that stock returns follow a lognormal distribution, which is only an approximate description of the actual situation. Furthermore, there is no memory in this process as it is purely Markovian. As already mentioned, it is observed that the power-law statistics of the distributions of real returns are very stable, well-fit by a Tsallis distribution of index $q = 1.4$ for returns taken over time-scales ranging from minutes to weeks, only slowly converging to Gaussian statistics for very long time-scales. Here we review a model of the underlying stock which is consistent with the returns distribution.

Our model bases on the non-Gaussian model [12, 13], where it was proposed that the fluctuations driving stock returns could be modeled by a statistical feedback process, namely:

$$dS = \mu S dt + \sigma S d\Omega \tag{3}$$

where

$$d\Omega = P(\Omega)^{\frac{1-q}{2}} d\omega. \tag{4}$$

In this equation, P corresponds to the probability distribution of Ω, which simultaneously evolves according to the corresponding nonlinear Fokker-Planck equation [15, 16]

$$\frac{\partial P}{\partial t} = \frac{\partial P^{2-q}}{\partial \Omega^2}. \tag{5}$$

The index q will be taken $q \geq 1$. It can be solved exactly yielding

$$P = \frac{1}{Z(t)} (1 - (1 - q)\beta(t)\Omega(t))^{\frac{1}{1-q}} \tag{6}$$

The exact form of the coefficients Z and β are given in [13].

Eq. (6) recovers a Gaussian in the limit $q \to 1$ while exhibiting power law tails for $q > 1$. In that case, our model is exactly equivalent to the Black-Scholes model.

The statistical feedback term P can be seen as capturing the market sentiment. Intuitively, this means that if the market players observe unusually high deviations of Ω (which is essentially equal to the de-trended and normalized log stock price) from the mean, then the effective volatility will be high because in such cases $P(\Omega)$ is small, and the exponent $1-q$ a negative number. Conversely, traders will react more moderately if Ω is close to its more typical or less extreme values. As a result, the model exhibits intermittent behavior consistent with that observed in the effective volatility of markets. To incorporate the subtle effect of skew and the stylized fact known as the leverage effect, we further extend the stock price process as in [14].

4 Option Pricing

Once a model of stock price dynamics is proposed, it can be used for various applications such as estimating and managing risk, or for pricing derivative instruments, such as options. Options are financial derivatives which in their own right are traded globally in high volumes. They fill important financial functions with respect to hedging and risk control, as well as offer purely speculative opportunities. In short, options are financial instruments which depend in some contingent fashion on the underlying stock or other asset class. The simplest example is perhaps the European call option. This is the right (not obligation) to buy a stock at a certain price (called the strike K) at a certain time (called the expiration T) in the future.

Contracts similar to options were exploited already by the Romans and story has it that Thales the Greek mathematician used call options on olives to make a huge profit when he had reason to believe that the harvest would be particularly good. In Holland in the 1600s, tulip options were traded quite a bit by speculators prior to the famous tulip bubble. But it wasn't until 1974 that the fair price of options could be calculated somewhat reliably with the publication of the Nobel-prize winning Black-Scholes formula. This is still the most widely used option pricing model, not because of its accuracy (since it is based on a Gaussian model for stock returns which, as we discussed above, is unrealistic) but rather due to its mathematical tractability (which exists due to the same Gaussian assumptions). In fact, an impressive school of mathematical finance has been developed over the past three decades, and is based largely on notions stemming from the famous Black-Scholes paradigm.

Because real stock returns exhibit fat tails, yet the Black-Scholes pricing formula is based on a Gaussian distribution for returns, the probability that the stock price will expire at strikes far from its current price will be underestimated. Traders seem to correct for this intuitively; for the Black-Scholes model to match empirical option prices, higher volatilities must be used the farther away the strike price is from the current stock price value. A plot of these Black-Scholes implied volatilities as a function of the strike price is thus not constant but instead most typically a convex shape, often referred to as the volatility smile. This way of representing option prices in terms of the Black-Scholes volatility is so widely used that prices are often quoted just in terms of this quantity, most often referred to simply as the vol.

The property that the Black-Scholes volatility exhibits a smile or skew shape that slowly flattens out as the time to expiration increases is the most important stylized fact of options. Many of the more realistic models convolve way too quickly to yield a Gaussian distribution although they might reproduce fat tails over a small time-scale. One of the nice properties of our statistical feedback models is just this slow convolution to a Gaussian. Hence, option pricing based on this model could be quite interesting. Another challenge with many models is related to notions pertaining to the definition of a unique equivalent martingale measure. Defining such a measure was the key ingredient in Black and Scholes seminal work, but many more complicated models that deviate from the Gaussian have often failed. Indeed, we were able to find such a unique equivalent Martingale measure for Eq(1), and this resulted in the ability to obtain closed-form solutions for pricing European call options. The price f of such an option is given by its expectation value in a risk-free (martingale) world as

$$f = E[e^{-rT} \max[S(T) - K, 0]] \tag{7}$$

where r is the risk-free rate. Assuming that S follows Eq(3) we obtain [13, 14]

$$\begin{aligned} f = S_0 \int_{d_1}^{d_2} \exp\left\{\sigma \Omega_T - \frac{\sigma^2}{2}[1-(1-q)(\beta(T)\Omega_T^2)]\right\} \\ P_q(\Omega_T)d\Omega_T - e^{-rT} K \int_{d_1}^{d_2} P_q(\Omega_T)d\Omega_T \end{aligned} \tag{8}$$

with P_q as in Eq. (6) and the explicit form of the coefficients given in [14]. These formulae are analytically tractable and converge to the Black-Scholes equation as $q \to 1$. However, for $q > 1$ they incorporate the effects of fatter tails. We also found option pricing formaulae for the generalized model with skew [14], and further work was done by [17]. Since a value of $q = 1.4$ nicely fits real returns over short to intermediate time horizons, this model is clearly more realistic than the standard Gaussian model. Using that particular value of q as calibrated from the historical returns distribution, fair prices of options can be calculated easily and compared with empirical traded option prices, exhibiting a very good agreement. In particular, while the Black-Scholes equation must use a different value of the volatility for each value of the option strike price in order to reproduce theoretical values which match empirical ones, the $q = 1.4$ model uses just one value of the volatility parameter across all strikes. One can calculate the Black-Scholes implied volatilities corresponding to the theoretical values based on the $q = 1.4$ model, and a comparison of this with the volatility smile observed in the market will reflect how closely the $q = 1.4$ model fits real

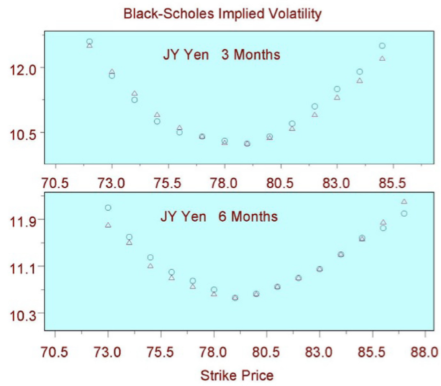

Fig. 3. Implied Black-Scholes volalities from the $q = 1.4$ model mactch extremely well to real traded values

prices. Excellent agreement over several expirations (time horizons) can be seen in Figure 3.

5 The Multi Time-Scale Model

Although very successful for pricing options, the statistical feedback model is still not entirely realistic. The main reason is that there is one single characteristic time in that model, and in particular the effective volatility at each time is related to the conditional probability of observing an outcome of the process at time t given what was observed at time $t = 0$. For option pricing this is perfectly reasonable because one is interested in the probability of the price reaching a certain value at some time in the future, based entirely on ones knowledge now. But this is a shortcoming as a model of real stock returns. In particular, in real markets, traders drive the price of the stock based on their own trading horizon. There are traders who react to each tick the stock makes, ranging to those reacting to what they believe is relevant on the horizon of a year or more, and of course, there is the entire spectrum in-between. Therefore, an optimal model of real price movements should attempt to capture this existence of multiple time-scales and long-range memory. We shall now show how we extend the above model to include multiple time-scales.

The effective volatility term $P^{1/(1-q)}$ in the statistical feedback model Eq(3) can be rewritten in the form

$$\sqrt{N'(1-(1-q)\beta'(\Omega(t)-\Omega(0))^2)} \tag{9}$$

by inserting the expression for P explicitly. where N' and β' are time-dependent constants. In this form it is clear that the volatility depends on the change in log-price at time t to the initial time t_0, since Ω is essentially the log of the price. A natural extension is to generalize to include feedback over not just one but many time-scales, as discussed in [18]). The random return is constructed as the product of a time dependent volatility σ_i and a random variable ω_i of zero mean and unit variance:

$$dy_i = \mu dt + \sigma_i d\omega_i \tag{10}$$

where μ is the average drift, which we will set to zero in the sequel, meaning that we measure all returns relative to the average drift. The volatility is written to include feedback over multiple time-scales $(i-j)$.

$$\sigma_i^k = \sigma_0 \sqrt{1 + g \sum_{j=1}^{\infty} \frac{1}{(i-j)^\gamma} (y_i^k - y_j^k)^2} \tag{11}$$

where i corresponds to time. The parameter g is a coupling constant that controls the strength of the feedback, σ_0 is the baseline volatility and γ is a factor that determines the decay rate of memory in the system which can also be seen as the relative importance between short-term and long-term traders. This model is well-defined as soon as $\sum_{j=0}^{\infty} \frac{g}{i-j}^\gamma < 1$ or approximately that $\frac{g}{1-\gamma} < 1$ and $\gamma > 1$ [18]. The model can be calibrated to real stock data yielding $\gamma = 1.15$ and $g = 0.12$, and it is seen that all of the main stylized facts of stock returns are reproduced. Interestingly, the model also shows multi-fractal scaling akin to that seen in turbulent systems, although there is no explicit multi-fractality injected in the model. These results are discussed in [18]. In Figure(1) a simulation of that model is shown. The process describes real data very well, with obvious periods of lower and higher volatility clustering together.

6 Statistical Signatures across Stocks : Self-orgnaization of correlation

Up until now we have been focused on understanding and modeling the dynamics of stock returns across time. However, in order to grasp properties of the full joint stochastic process driving markets we turn our attention to the cross-sectional dynamics, across a universe of stocks at any given point in time. Along these lines, there have in recent years been several studies which focus on exploring the structure and dynamics of correlations across the different stocks that comprise the market [19–25]. In particular, these models and analysis explore correlations during market stress or times of bubbles, panic and crashes. This is important because it is under such extreme scenarios that investors are at most risk; their usual models and world-views might break down and the results could be devastating as seen in August 2007, October 2008, or during the so called Flash Crash of May 2010. Some of our contributions in this area [24, 25] has been to see if there are any particular cross-sectional statistical signatures in these periods of market panic. To get a grasp on the cross-sectional distribution of stock returns at a given time point, we can look at moments such as the mean, standard deviation, skew and kurtosis. We then look at these as a

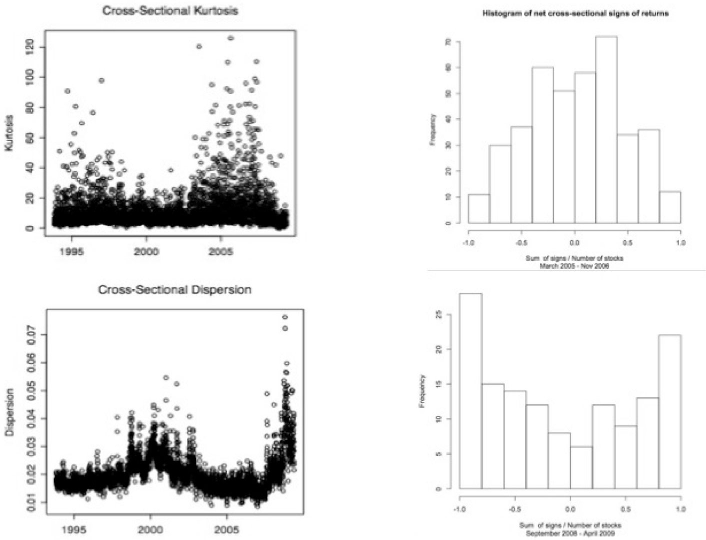

Fig. 4. Statistical signatures in normal and panic times

function of time. The standard deviation of returns is widely referred to as dispersion. Calculated across a universe of 1500 US stocks and plotted out for the time period 1993 - 2009, it is striking to see that the dispersion gets relatively big during the time periods defined as panic according to the discussion above. However, the more striking discovery is to plot out the cross-sectional kurtosis alongside the dispersion, or together with market returns (Figure (5)). Even by eye it is quite clear that there is a strong negative correlation between the two quantities, which is in fact about -25%. In times of panic, dispersion is high yet excess kurtosis practically vanishes. In more normal times, the dispersion is lower but the cross-sectional excess kurtosis is typically very high.

We want a model that can explain all of these findings, namely to preserve the fat-tailed time series properties of stocks, but which gives rise to the remarkable reduction in kurtosis and increase in correlations, cross-sectionally, that are characteristic of market panic. In such times, dispersion is high yet kurtosis is low, which implies that the data are more Gaussian in times of panic. This can be explained partially by the fact that the volatilities of the individual stocks are higher yet more alike in times of panic, a statement that is borne out by the data [24]. This might be one effect contributing to our findings, but we believe that the behavior of cross-sectional correlations is what drives the statistical signatures that we found. As a proxy for the collective behavior of all stocks in the market, we define the following quantity

$$s = \frac{s_{up} - s_{down}}{s_{up} + s_{down}} \tag{12}$$

where s_{up} is the number of stocks that have positive returns over a given interval, and $s_{down}$is the number of stocks that have negative moves on that same interval (for example a day). If $s = 0$ then roughly the same number of stocks moved up as down, and the assumption is that the stocks had little co-movement and so were uncorrelated. If all stocks move together either up or down, though, the value of s will be +1 or -1 and the stocks will have high correlation. So, the following picture emerges: If $s = 0$ there is no correlation, and we are in a disordered state. However if $s \neq 0$ then there is correlation and we are in an ordered state. We will now make a leap and borrow some terminology from physics. We shall call s the order parameter. It is a macroscopic parameter that tells us whether there is order and correlation in the system, or not. In physics, in particular in the field of non-equilibrium thermodynamics and synergetics [26], the concept of the order parameter is often used to describe systems that exhibit spontaneous self-organization. Examples range from chemical kinetics to laser dynamics, from fluid dynamics to biological systems; from collective behavior in both the animal and human world to cloud formation. To illustrate the concept, let us look at an example which should be familiar and intuitive to most, namely magnetism.

In a ferromagnetic system, the total magnetic moment depends on the orientation of the individual magnetic spins comprising the system. It is proportional to the quantity

$$m = \frac{m_{up} - m_{down}}{m_{up} + m_{down}} \tag{13}$$

where m_{up} and m_{down} denote the number of spins lined up and down respectively. The distribution of possible outcomes of this macroscopic quantity is given by

$$P(m) = N \exp(F(m, T)) \tag{14}$$

where T is the temperature, N is a normalization factor and F is the free energy of the system. Depending on the value of T, the magnetic system will either be in an ordered or disorderd state. Assuming one can perform a Taylor expansion of F and invoking symmetry arguments, the corresponding Langevin equation takes the form

$$\frac{dm}{dt} = -\frac{a}{2}m - \frac{b}{4}m^3 + W_t \tag{15}$$

where W_t is thermal noise. The coefficient a can be written as

$$a = \alpha(T - T_c) \tag{16}$$

where T_c is the so-called critical temperature. One can envision these dynamics as motion in a potential well V given by $V(m) = -F(m)$. If $T > T_c$, the only minimum is the trivial one at $m = 0$. However, for $T < T_c$ there are two real roots appearing, yielding non-zero values of m. Clearly, m can be positive or negative, depending on which minima is reached by the system. This is referred to as symmetry breaking. Due to the noise, the dynamics can also drive m from one minimum to the other. Because the value of T determines whether the system is in the disordered state ($m = 0$) or the ordered state ($m \neq 0$), it is called the control parameter. The probability distribution of the system in the disordered state will be a unimodal one, while the probability distribution of m in the ordered state will be bimodal. As T passes from above to below T_c, or vice-versa, there is clearly a phase transition: the state of the system is drastically altered. In this type of symmetric system, the phase transition is referred to as a second order one.

In the current setting, we make an analogy between the variable s and the magnetic moment m. We have observed rather drastic changes in the cross-sectional distribution of stocks in the times of panic versus more normal market conditions [24]. Histograms of s in both periods show that in normal times, s is unimodal, and in panic times we obtain a bimodal distribution consistent with the frame-work of a phase transition leading to self-organization in panic times (see Figure 4).

We postulate that the dynamics of s be given by

$$\frac{ds}{dt} = -\frac{a}{2}s - \frac{b}{4}s^3 + W_t. \tag{17}$$

We propose that

$$a = \sigma_c - \sigma_0 \tag{18}$$

where W_t is a Gaussian noise term and σ_0 corresponds to the baseline volatility level of stocks. This volatility is assumed constant across all instruments, and essentially measures the general uncertainty in the environment, so in this sense acts much as the temperature in the magnetic system. Note that it is the feedback effects in the system which induce stock-specific variations in volatility over time, and can largely explain most of the excess volatility observed in stock time-series, whereas the parameter σ_0 is not driving the stock-specific dynamics, but simply describes a "global" level of risk. The quantity σ_c would correspond to a critical level of uncertainty, below which the market is in a normal phase, and above which we have the onset of panic. Much as in the case of ferromagnetism, where the control parameter T can be tuned externally above or below the critical temperature, in our model the uncertainty level σ_0 captures the external environment. In a sense it represents the general perception of risk in the public mind. Our hypothesis is then that financial markets appear to exhibit a phase transition from the disordered to ordered state, after crossing a critical level of risk perception. Putting the dynamics together, we have the multi time-scale feedback process for each stock k

$$dy_i^k = \sigma_i^k d\omega_i^k \tag{19}$$

with $k = 1 \cdots N$, and the volatility of each stock k given by Eq(11) The random variables ω_i^k are drawn from a Gaussian distribution, uncorrelated in time such that $< \omega_i^k \omega_{i+\tau}^k >= \delta(i-(i+\tau))\tau$, yet amongst themselves at a given time point

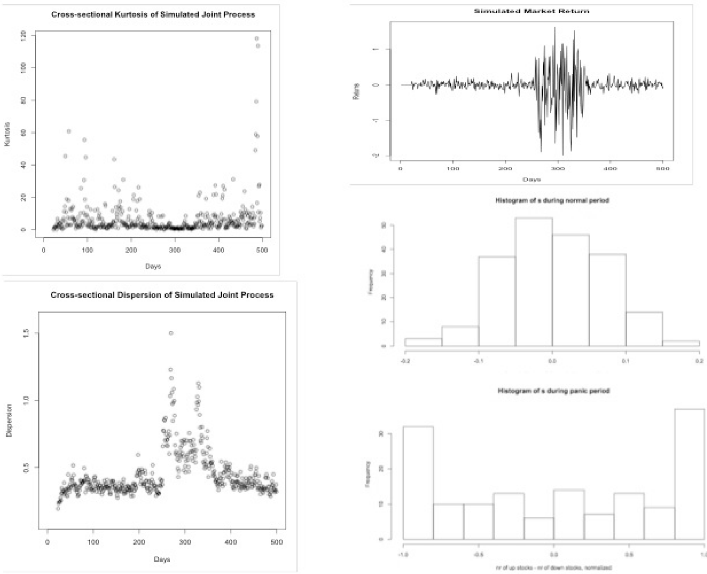

Fig. 5. Results of the simulation of the joint process

i across stocks k, they are correlated with correlation $|s|$. The macroscopic order parameter s is therefore just a signature of the cross-stock correlations, whose dynamic behavior manifests itself in the order parameter equation

$$\frac{ds}{dt} = -\frac{a}{2}s - \frac{b}{4}s^3 + W_t. \tag{20}$$

The coefficients must always be such that $|s| \leq 1$ which can be imposed by running the dynamics of s on a real valued variable $\hat{s}$ such that $s = \tanh(\hat{s})$.

What do we expect to see from this model? We have already seen that across time, it models very well many properties of real financial time-series. Across stocks, if $\sigma_0 < \sigma_c$, correlations fluctuate around $s = 0$ and we expect to see a unimodal distribution of s. The cross-sectional kurtosis should be rather high since there is no mechanism to cause either stocks or stock volatilities to have any co-movement at all, so at each time point it is as if the cross-sectional returns are drawn from a Gaussian process with stochastic volatility, yielding a fat-tailed distribution as the superposition. Then as the market crashes with $\sigma_0 > \sigma_c$, the system enters a phase transition. The order parameter s becomes $s \neq 0$ and the system enters the ordered phase with high co-movement. Because the random variables ω_i^k are now correlated across stocks, cross-sectional returns will be more similar and the distribution will have lower kurtosis. Additionally, due to the fact that the phase transition is triggered by an external shock in volatility, all stocks will tend to have higher volatilities and higher cross-sectional dispersion.

Simulations of this model for the joint stochastic process of stocks were performed and we refer the reader to [24] for details on the implimentation. Here we provide a summary. The baseline volatility was assumed to be $\sigma = 0.2$, and at a certain time a volatility shock $\sigma_{shock} = 0.6$ (consistent with levels observed in the VIX volatility index in late 2008) was applied to the system (see Figure 5). This induced the phase transition from the disordered state where correlation among stocks are relatively low, centered around zero, to a highly ordered state where the correlations are different from zero. We found that the main features of financial markets are captured within this framework. The order parameter s goes from 0 (the disordered state) to $s \approx 0.8$ (the ordered state) at the time of the volatility shock. When the shock subsides, it returns to the disordered state again. As expected, the market volatility rises when s is in the ordered state, which corresponds to the panic phase. In addition, the cross-sectional dispersion rises during the market panic, while the cross-sectional kurtosis drops close to zero. The correlation between the two quantities is in this example -17%, consistent with empirical observations that also showed a strong negative correlation. Histograms corresponding to the distribution of the order parameter s in the normal market phase as well as in the panic phase are in excellent agreement with the empirical observations of the real market data, namely unimodal in the normal phase, and clearly bimodal during the panic time. Encouraged by these findings we extended this joint stochastic model to include skew and explored its self-similar properties on different time-scales, as presented in [25]. Further studies of interest in a similar vein were done by others [19].

7 Conclusion

We have reviewed the anomalous features inherent in financial data across time and across stocks, which reflect the very complex and nonlinear interactions leading to price formation in financial markets. Our approach has been to propose models that intuitively capture the dynamics that could be at play, and to verify them in that they can reproduce the observed statistical signatures and stylized facts, also in derivative markets such as options. Common to all our studies is that cooperative effects, memory and nonlinear feedback appear at play, so many interesting techniques for the field of statistical physics and synergetics can be applied. Continuously evolving and changing, the field of finance will certainly continue to pose interesting challenges for scientists and practitioners alike for years to come.

References

1. Bouchaud, J.-P., Gefen, Y., Potters, M., Wyart, M.: Fluctuations and response in Financial markets: The subtle nature of random price changes. Quantitative Finance **4**(2), 176–190 (2004)
2. Farmer, J., Gerig, A., Lillo, F., Mike, S.: Market efficiency and the long-memory of supply and demand: Is price impact variable and permanent or fixed and temporary? Quantitative Finance **6**(2), 107–112 (2006)

3. Black, F., Scholes, M.: The Pricing of Options and Corporate Liabilities. Journal of Political Economy **81**, 637–659 (1973)
4. Bouchaud, J.-P., Potters, M.: Theory of Financial Risks and Derivative Pricing. Cambridge University Press (2004)
5. Gopikrishnan, P., Plerou, V., Nunes Amaral, L.A., Meyer, M., Stanley, H.E.: Scaling of the distribution of fluctuations of financial market indices. Phys. Rev. E **60**, 5305 (1999)
6. Gabaix, X., Gopikrishnan, P., Plerou, V., Stanley, H.E.: A theory of power-law distributions in Financial market Fluctuations. Nature **423**, 267 (2003)
7. Tsallis:, C.: J. Stat. Phys. **52**, 479 (1988); Curado, E.M.F., Tsallis, C.: J. Phys. A **24**, L69 (1991); **24**, 3187 (1991); **25**, 1019 (1992)
8. Heston, S.L.: A closed-form solution for options with stochastic volatility with applications to bond and currency options. Rev. of Fin. Studies **6**, 327–343 (1993)
9. Cont, R., Tankov, P.: Financial modelling with jump processes. CRC Press (2004)
10. Bollerslev, T., Engle, R.F., Nelson, D.B.: ARCH models. In: Engle, R.F., McFadden, D. (eds.) Handbook of Econometrics, vol. 4. Elsevier Science, Amsterdam (1994)
11. Muzy, J.-F., Delour, J., Bacry, E.: Modelling fluctuations of financial time series: from cascade process to stochastic volatility model. Eur. Phys. J. B **17**, 537–548 (2000)
12. Borland, L.: Option Pricing Formulas based on a non-Gaussian Stock Price Model. Phys. Rev. Lett. **89**(N9), 098701 (2002)
13. Borland, L.: A Theory of non-Gaussian Option Pricing. Quantitative Finance **2**, 415–431 (2002)
14. Borland, L., Bouchaud, J.-P.: A Non-Gaussian Option Pricing Model with Skew. Quantitative Finance **4**, 499–514 (2004)
15. Tsallis, C., Bukman, D.J.: Anomalous diffusion in the presence of external forces: Exact time-dependent solutions and their thermostatistical basis. Phys. Rev. E **54**, R2197 (1996)
16. Borland, L.: Microscopic dynamics of the nonlinear Fokker-Planck equation: a phenomenological model. Phys. Rev. E **57**, 6634 (1998)
17. Vellekoop, M.H., Nieuwenhuis, J.W.: On option pricing models in the presence of heavy tails. Quant. Finance **7**, 563–573 (2007)
18. Borland, L., Bouchaud, J.-P.: On a multi timescale statistical feedback model for volatility fluctuations. The Journal of Investment Strategies, 1–40 (2012)
19. Preis, T., Kenee, D., Stanley, H.E., Helbing, D., Ben-Jacob, W.E.: Quantifying the Behavior of Stock Correlations Under Market Stress. Nature, October 2012
20. Lillo, F., Mantegna, R.: Variety and volatility in financial markets. Phys. Rev. E **62**, 6126–6134 (2000)
21. Sornette, D.: Why stock markets crash: Critical Events in Complex Financial Systems. Princeton University Press (2002)
22. Kaizoji, T.: Power laws and market crashes. Prog. Theor. Phys. Suppl. **162**, 165–172 (2006)
23. Raffaelli, G., Marsili, M.: Dynamic instability in a phenomenological mode of correlated assets. J. Stat. Mech., 8001 (2006)
24. Borland, L.: Statistical Signatures in Times of Panic: Markets as a Self-Organizing System. Quant. Finance **12**(9), 1367–1379 (2012)
25. Borland, L., Hassid, Y.: Market Panic on Different Time-Scales. ArXiv e-prints 1010.4917 (2010)
26. Haken, H.: Synergetics: an introduction. Springer (1977)