3 A complex walk down Wall Street

~~HANDOUT 4 OF 7~~

3.1 Facing the stylized facts

We have discussed the random-walk model which underpins standard theoretical treatments of asset price movements. But how different is the real world, and what are the consequences of these differences? In this Chapter we set out to explore real-world market dynamics by analyzing data from two quite distinct markets. In both cases, the linear (i.e. low-order) correlation of price-changes is essentially zero implying that the associated price-series would be judged to be a 'random walk' by many industry-standard statistical packages. However we show that this data actually exhibits significant non-linear (i.e. higher-order) temporal correlations across a range of timescales, together with non-Gaussian price distributions. Our findings illustrate the general results reported in the Econophysics literature[1] that real markets tend to deviate from the standard random-walk paradigm in a fundamental way. These empirical features form part of the so-called *stylized facts* of financial markets. Any candidate market models or financial calculations concerning risk or derivative pricing, clearly ought to be made consistent with these stylized facts -- hence the motivation for the present empirical study.

Let's return for a moment to our dot-com company *risk-e.com* whose price-history is shown in Figure 1-1. Suppose we are trying to uncover what 'extra' information might be hidden in the past behaviour of the stock price $x[t]$, by analysing the statistics of $x[t]$ over recent history. In trying to specify 'recent history' we immediately face a problem. How much recent history should be included? From a statistical analysis point of view, it is usually the case of 'the more the merrier' in that the data-tests will be less susceptible to problems of finite size. However just because the data is available for the past 5 years, should we really be placing equal weight on its behaviour in 1996-7 and 2000-1? We have no guarantee that the price-process $x[t]$ is stationary on any timescale; indeed the PDF of price-changes and its associated statistics may be strongly time-dependent. Even if we assume that it is reasonably stationary, maybe we should treat the period of the 1999-2000 'bubble' as atypical, i.e. a special case. However in true complex systems style, it could be claimed that *every* case is a special case. One can always find reasons for excluding or including particular periods, particularly in hindsight having already seen the data. Herein lies one of the major problems with analyzing financial data. All such data seem to have atypical periods at some stage or another. Below are two such examples of arguably atypical behaviour, arising in the markets to be analyzed in this Chapter:

[1] See the numerous data analysis papers on www.unifr.ch/econophysics, which report studies on price-series taken from a wide range of financial markets. See also [MS] and [BP].

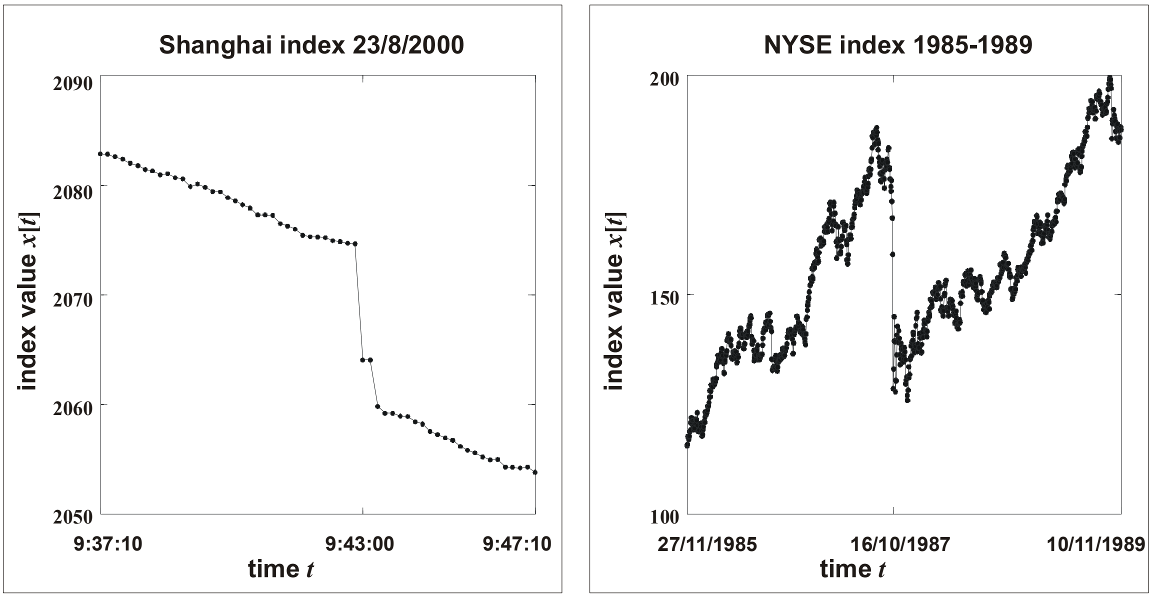

Figure 3-1 Periods of arguably atypical behaviour, arising in the Shanghai stock exchange index (left) and NYSE daily composite index (right).

So should such atypical periods be included when calculating the statistics? This is a crucially important question, given our interest in calculating risk. After all, risk is concerned with large deviations in the tails of the distribution. However the frequency of occurrence of such deviations decreases as the size of the deviation increases, hence the decision to exclude or include a few such data-points can have a profound effect on the estimation of the best-fit distribution for the tails. One could try to base this decision to exclude or include the rare events by classifying them as exogenous (i.e. generated by external effects) or endogenous (i.e. self-generated by the internal dynamics). However it is not easy to make such a distinction when considering real data, and mixtures of the two will certainly arise. In our opinion, all data points should be included when forming the statistics. Even if the stimulus was completely external, the response of the market is controlled by its own internal feedback mechanisms, hence this response carries some information about the market's internal dynamics. In short it is all part of the same complex system. However, many readers may not agree. In fact, different people looking at the same dataset could come to somewhat different conclusions as to the exact functional form of the PDF of price-changes, depending on how they decided to process the dataset.

3.2 Statistical tools and datasets

There are many statistical packages available commercially for analyzing time-series data. However such standard packages only tend to investigate the more obvious, low-order correlations (e.g. auto-correlation of price-changes). In fact it would be really surprising if such tests were to show anything significant, not least because a lot of people with access to these packages would then trade based on such correlations. Such trading would then distort or diminish the correlations themselves. Therefore it is not surprising that, as far as standard statistical tests based on these lower-order correlations are concerned, financial market data appears essentially 'random'. However as we discussed in Chapter 2, this implies nothing about the existence of higher-order temporal correlations, or deviations from Gaussian behaviour. Our goal is not to list the results of such standard tests here, but rather to exploit the points raised in Chapter 2 in order to look beyond the standard analysis. In particular we wish to see if higher-order temporal correlations can be detected over a range of different timescales, and the degree to which the resulting distributions of price-increments deviate from the Gaussian form. The appearance of such higher-order correlations and deviations from Gaussian behaviour, would sound the warning bell against automatic application of results from standard finance theory.

There are thousands of possible price-series we could look at, and hundreds of different markets. One of the significant contributions of Econophysics has been to establish that the price statistics across a wide range of supposedly different financial markets worldwide, exhibit certain universal properties[2]. It seems that despite the differences in detail between these markets in terms of how trades are registered, trading hours etc., there is something that they hold in common which is driving the market price-dynamics. Our viewpoint is that this common element is the presence of traders who are buying and selling, moving into the market or staying out, and taking actions based in part on the past behaviour of the market. In Chapters 4 and 5 we will see how this common feature can indeed give rise to the stylized facts observed across financial markets. But first, we will illustrate these common features by focusing on two market datasets which one might initially believe were quite different. These two datasets are the NYSE composite index recorded on a daily basis between 1966-2000[3] and the Shanghai stock exchange index recorded at 10-second intervals during a period of eight months in 2001-2. Apart from the fact that the two markets are in distinct geographical and temporal zones, they also differ in their history and level of liquidity. The NYSE composite index reflects the long-running, well established and highly liquid U.S. stock market. By contrast, the Shanghai stock exchange is essentially an emerging market which is relatively young in terms of global trading. Despite these structural differences, we will show that there are in fact surprising similarities between the dynamics observed in these two markets. This observation is therefore consistent with claims being made within the Econophysics community that apparently distinct financial markets can exhibit universal dynamical features[2].

Our discussion will also illustrate the common situation in practice whereby one is given either a long but low-frequency dataset such as the daily NYSE composite index[3], or a short but high-frequency dataset such as the Shanghai stock index data. Choices therefore have to be made about the

[2] Given the limited space, we cannot carry out a thorough statistical analysis over many different markets. Nor can we provide a detailed discussion of the datasets themselves. Instead the results presented here provide a flavour of the many similar ones already present in the Econophysics literature. In particular, we refer to [MS], [BP] and www.unifr.ch/econophysics.

[3] At the time of writing, this data is freely available on www.unifr.ch/econophysics

best way to analyze this data in order to minimize any bias. This type of problem is fine in academia where one can spend all day building and cleaning datasets. Practitioners on the other hand regularly face a much tougher situation: they may be handed a low-frequency dataset measured over a relatively short period and yet be expected to produce realistic calculations of derivatives prices and hedging strategies. For this reason we will return to the data analysis problem in Chapter 6, where we look at the practical implementation of a generalized derivatives theory using a much shorter version of the NYSE daily data, in order to reflect this common situation.[4] For the time being, we will stick with our large NYSE daily dataset, and our high-frequency Shanghai dataset, in order to pin down their statistical and dynamical features.

3.3 Empirical analysis

Below we show the Shanghai price-series represented by our dataset:

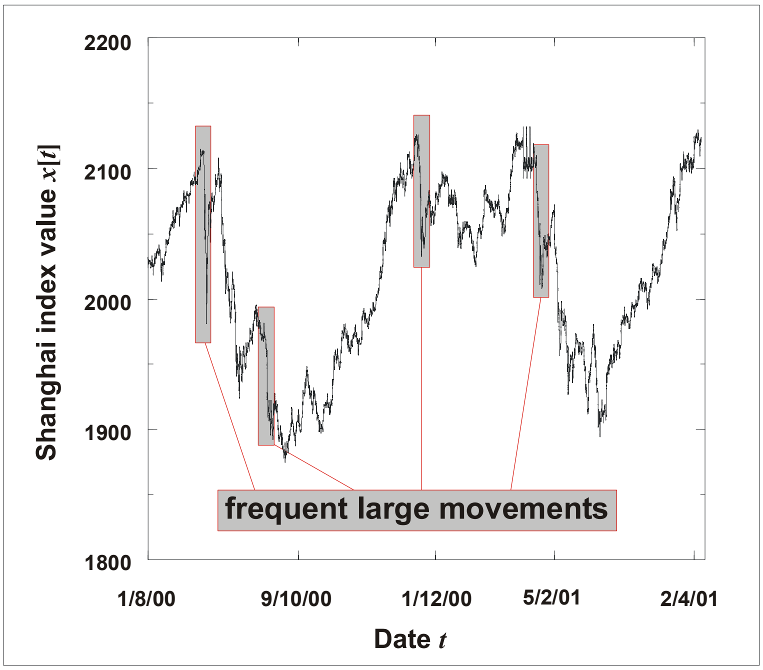

Figure 3-2 Shanghai stock market index measured at 10 second intervals over the period 1/8/00-2/4/01.

Even by eye, one can see that the index undergoes frequent large movements by contrast to the Gaussian random walk shown in Figure 2-5. A close-up of one such large change was shown in Figure 3-1. Such large changes will tend to increase the magnitude of the PDF in the tails of the distribution, hence generating fat-tailed distributions. The pioneering study of empirical data within the

[4] The conclusions reached in Chapter 6 regarding non-Gaussian behaviour and higher-order temporal correlations, are consistent with those in the present Chapter. This suggests that the stylized facts are reasonably robust to changes in size of dataset.

Econophysics field was performed by Mantegna and Stanley[5] who studied minute-by-minute data for the S&P500 index over the 6-year period 1984-1990. Our study is similar to their original analysis[5]. We form a series of log-returns (see Section 1.4) of index changes[6] for each of the two market indices:

$$z\left[t,t-\Delta t\right]\equiv z\left[t\right]=\log\frac{x\left[t\right]}{x\left[t-\Delta t\right]} \qquad (3.1)$$

where Δt is a given time-interval separating the index records, and $x\left[t\right]$ is the index value at time t. It is known from empirical studies that there exists an intra-day pattern of market activity in large financial markets[7]. A possible explanation for this intra-day pattern is the reaction to the information gathered during the hours when the market is closed, together with the fact that many liquidity traders are active near the closing hours. There is a similar intra-day pattern in the absolute changes in the Shanghai index $\left|z\left[t\right]\right|$ which we wish to remove before proceeding. The intra-day activity pattern $A\left[t_{\text{day}}\right]$, where t_{day} denotes the time during the day, can be defined as:

$$A\left[t_{\text{day}}\right]=\left\langle\left|z\left[\text{daytime}\left[t\right]=t_{day}\right]\right|\right\rangle \qquad (3.2)$$

where the operator $\text{daytime}\left[t\right]$ returns the time during the day which corresponds to an absolute time of t. The intra-day pattern $A\left[t_{\text{day}}\right]$ for the Shanghai index is shown below:

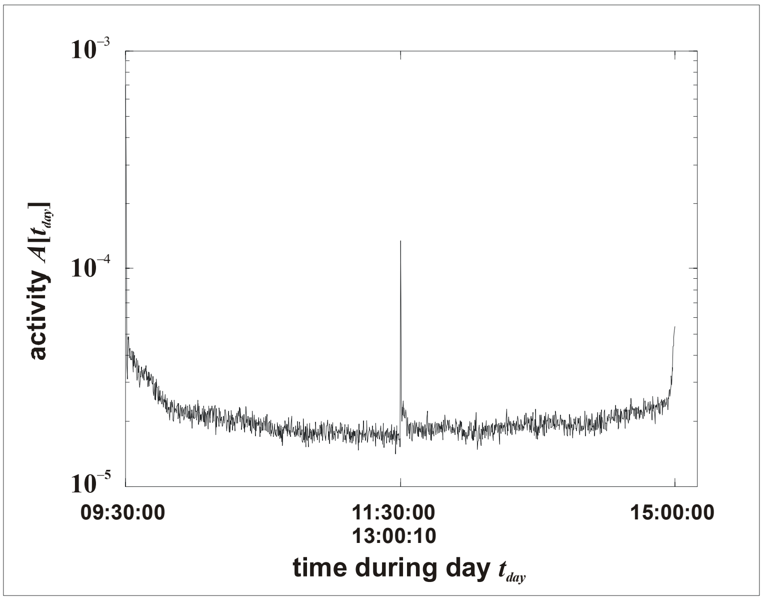

Figure 3-3 Daily (i.e. intra-day) pattern for the absolute changes of the Shanghai index.

[5] R. Mantegna and H.E. Stanley, Nature **376**, 46 (1995). See also [MS] for further discussions.

[6] We drop the explicit Δt from the log-return $z\left[t,t-\Delta t\right]$: Mantegna and Stanley's original paper is written in terms of the index-change variable Z and we want to facilitate any comparison by the reader to their original paper.

[7] P. Gopikrishnan, M. Meyer, L.A.N. Amaral and H.E. Stanley, Eur. Phys. J. B **3**, 139 (1998); P. Gopikrishnan, V. Plerou, L.A.N. Amaral, M. Meyer and H.E. Stanley, Phys. Rev. E **60**, 5305 (1999); R.A. Wood, T.H. McInish and J.K. Ord, Journal of Finance **40**, 723 (1985); L. Harris, Journal of Financial Economics **16**, 99 (1986); A. Admati and P. Pfleiderer, Review of Financial Studies **1**, 3 (1988); P.D. Ekman, The Journal of Futures Markets **12**, 365 (1992).

Similar intra-day patterns have been observed in other markets, such as the Hang Seng. In order to remove the systematic effect of this intra-day pattern, we re-define the returns for the Shanghai index as follows: $z\left[\text{daytime}[t]\right] \rightarrow z\left[\text{daytime}[t]\right] / A\left[t_{\text{day}}\right]$. Such a re-definition is obviously not applied to the NYSE composite index, since the dataset is daily and hence has no intra-day pattern. We are now ready for the analysis. We start by looking at the standard linear (i.e. low-order) correlations between returns. In particular we look at the autocorrelation of returns which is essentially just the correlation measure c_{ij} defined in Equation (2.3). Specifically we use the following normalized definition of the autocorrelation:

$$\rho\left[\{z[t]\},\{z[t-\Delta t]\}\right] = \left\langle \left(z[t]-\overline{z[t]}\right)\left(z[t-\Delta t]-\overline{z[t-\Delta t]}\right)\right\rangle \Big/ \sigma^2 \qquad (3.3)$$

where σ^2 is the variance of $z[t]$ and Δt is a given time lag. Hence if the returns $z[t]$ and $z[t-\Delta t]$ are uncorrelated, then

$$\left\langle \left(z[t]-\overline{z[t]}\right)\left(z[t-\Delta t]-\overline{z[t-\Delta t]}\right)\right\rangle = \left\langle \left(z[t]-\overline{z[t]}\right)\right\rangle \left\langle \left(z[t-\Delta t]-\overline{z[t-\Delta t]}\right)\right\rangle = 0 \qquad (3.4)$$

Below we show the autocorrelation of returns for both markets:

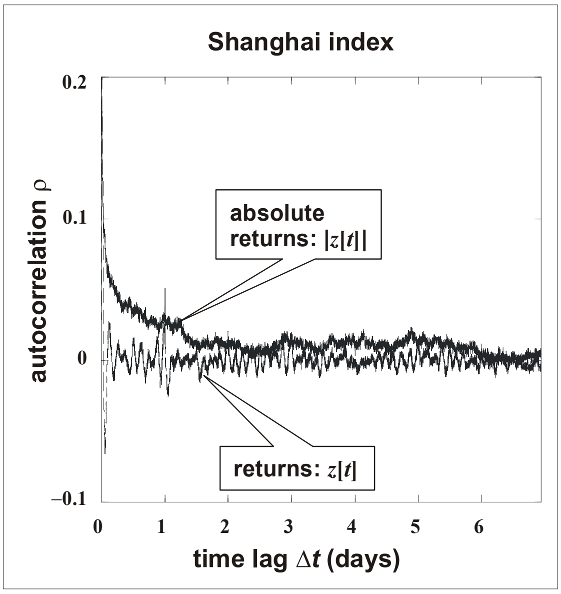

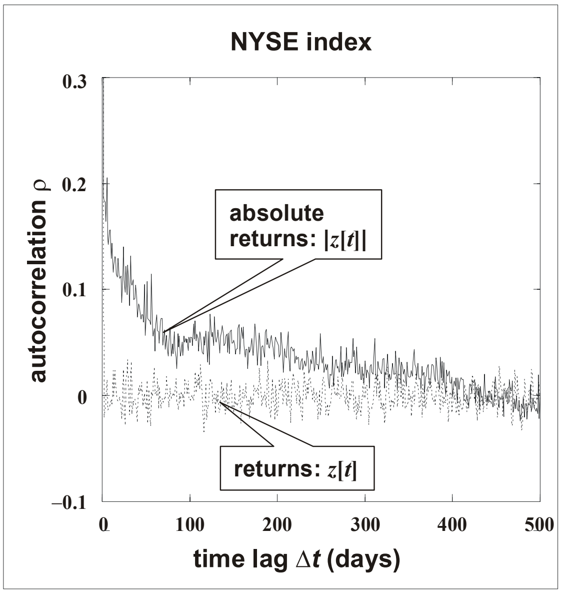

Figure 3-4 Autocorrelation of returns and absolute returns for the Shanghai index (left) and NYSE (right) as a function of the time-lag Δt

As can be seen, the autocorrelation function of returns is essentially zero for all time-lags $\Delta t \neq 0$. This behaviour is consistent with the standard finance theory model of uncorrelated price-increments. We therefore turn to look at non-linear (i.e. higher-order) temporal correlations. As mentioned in Section 2.2.1, one such measure is given by the autocorrelation of the *absolute* value of returns $\left|z[t]\right|$. Since this removes the sign from the returns, it is a measure of the temporal correlations in the *size* of the return fluctuations. This correlation measure can be obtained from Equation (3.3) simply by inserting moduli around every z -dependent term, i.e. $z[t] \to |z[t]|$ and so on. Figure 3-4 confirms that this autocorrelation measure is non-zero over a wide range of time lags Δt. We can conclude that *the non-linear (i.e. higher-order) temporal correlations in the returns* $z[t]$ *survive over a surprisingly long period*. In fact for the NYSE, these correlations last for several months. To understand better why there is such a long decay of the autocorrelation of absolute returns, we examine a 'time-dependent' or 'local' volatility from the returns' series. This quantity can be determined by calculating the standard deviation of the fluctuation in returns over a small time-window taken in the vicinity of a given time t. Figure 3-5 below shows how this local volatility then varies as a function of time t for the NYSE:

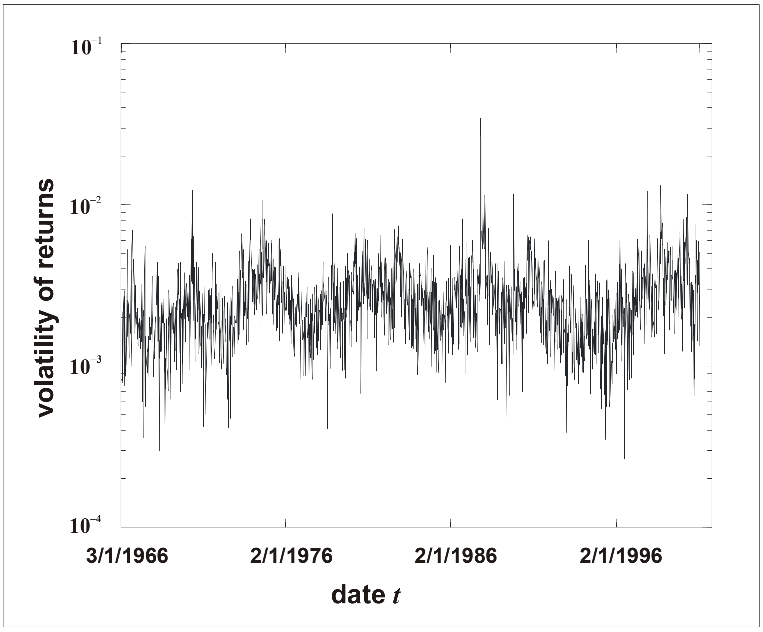

Figure 3-5 Local volatility as a function of time for the NYSE index.

This local volatility is far from constant[8]. It exhibits frequent large deviations and a bursty structure associated with clustering. This clustering reflects the tendency for volatile periods in the market to generate further volatility in the immediate future: volatility breeds volatility. This finding is therefore consistent with the long decay in the autocorrelation of absolute returns in Figure 3-4. On a microscopic level, one could relate this volatility clustering to market-activity as reflected by the number of trades etc. Such market activity therefore appears to be a highly non-stationary process.

Next we look at the full distribution of returns using the methodology of Mantegna and Stanley. Mantegna and Stanley proposed that a so-called truncated Levy distribution represents a good model for the PDF of price-changes, in contrast to the Gaussian paradigm. The truncated Levy distribution comprises a Levy distribution over the main central portion of the PDF, with an

[8] We note that even the intra-day volatility can vary significantly.

approximately exponential truncation in the outer tails. A Levy distribution is the name given to a general class of fat-tailed distributions of which the Lorentzian is a particular example. The characteristic feature of a stable symmetric Levy distribution is that it has power-law tails for large deviations:

$$p[z] \equiv p_{L,\alpha}[z] \sim |z|^{-(1+\alpha)} \quad \text{for} \quad |z| \to \infty \qquad (3.5)$$

where the parameter α satisfies $0 < \alpha < 2$ in order to guarantee that the PDF is stable under convolution. It follows that the second moment, and hence the variance and standard deviation, are *infinite*. There is no simple analytic expression for $p_{L,\alpha}[z]$, except when $\alpha = 1$ which corresponds to the Lorentzian distribution shown in Figure 2-3:

$$p_{L,\alpha=1}[z] = \frac{C}{z^2 + C^2\pi^2} \xrightarrow{|z|\to\infty} C|z|^{-2} \qquad (3.6)$$

Although it is not so easy to see, the case $\alpha = 2$ yields the Gaussian. In this limit, the tails are no longer power-law, but are instead exponential -- equivalently one can say that the region over which power-law behaviour appears goes to zero for $\alpha \to 2$. Like the Gaussian and the Lorentzian, the Levy distribution is also stable under convolution for $0 < \alpha < 2$. The consequence of this stability is the following. Suppose that the PDF for price-changes over a given time-increment is truly Levy, and that these price-changes are i.i.d.. When we then produce the price-change PDF over any larger time-increment by convoluting the PDF with itself as discussed in Chapter 2, the PDF of price-changes will remain Levy on *all* timescales. Hence we *never* achieve convergence to a Gaussian. Looking back at the conditions for the Central Limit Theorem to hold as discussed in Section 2.2.3.4, we can see why this arises: although the price-changes might be i.i.d. and we put ourselves in the limit of $n \to \infty$, the variance of single-step price changes is *infinite* hence the conditions of the CLT for convergence to a Gaussian are not met. This is simply a consequence of the fact that the power law in the tail of $p_{L,\alpha}[z]$ corresponds to $\alpha < 2$, yielding an infinite variance. Hence although the Levy distribution is appealing as a model PDF for price-changes because of the empirical evidence for power-law tails, this feature of non-convergence to a Gaussian is less attractive. For this reason, the *truncated* Levy distribution was introduced: it retains power-law tails out to large z, but then becomes exponential[9] hence guaranteeing a *finite* variance and an *eventual* convergence to Gaussian under convolution.

We will first summarize our conclusions concerning the PDF for price-change returns, before showing the numerical results themselves. For *both* the NYSE and the Shanghai market, the following statements hold: the distribution of returns shows apparent scaling behavior, which cannot be modelled

[9] There is no unique definition of a truncated Levy distribution regarding the exact value of z at which the power-law behaviour stops. Nor is there a rule concerning the precise form of the PDF once it has stopped, other than saying that it should then have a faster decay. The defining feature is just that the variance is now finite and eventual convergence to the CLT result of Gaussian is therefore guaranteed. In practice, additional parameters are introduced to parametrize the truncated Levy distribution. We refer to [MS] and [BP] for a more detailed discussion of these distributions.

by a Gaussian distribution. The non-Gaussian dynamics of the stochastic process underlying the time series of returns, can be modelled quite well using a truncated Levy distribution. A power-law behaviour is observed for the probability of zero returns. The power-law behaviour in the tails (i.e. large z) drops slower than Gaussian, but faster than a Levy stable distribution. This ensures the existence of a large but finite standard deviation (i.e. volatility) for these markets.

Now we turn to the numerical results themselves. For each market, we use the original time-series to generate a new time-series of returns, where each return is calculated for a fixed time-increment Δt . Since we are interested in how the markets behave over different time-scales, we then repeat this procedure for various Δt . We then plot the PDF of these returns for each timescale Δt . These figures are shown below:

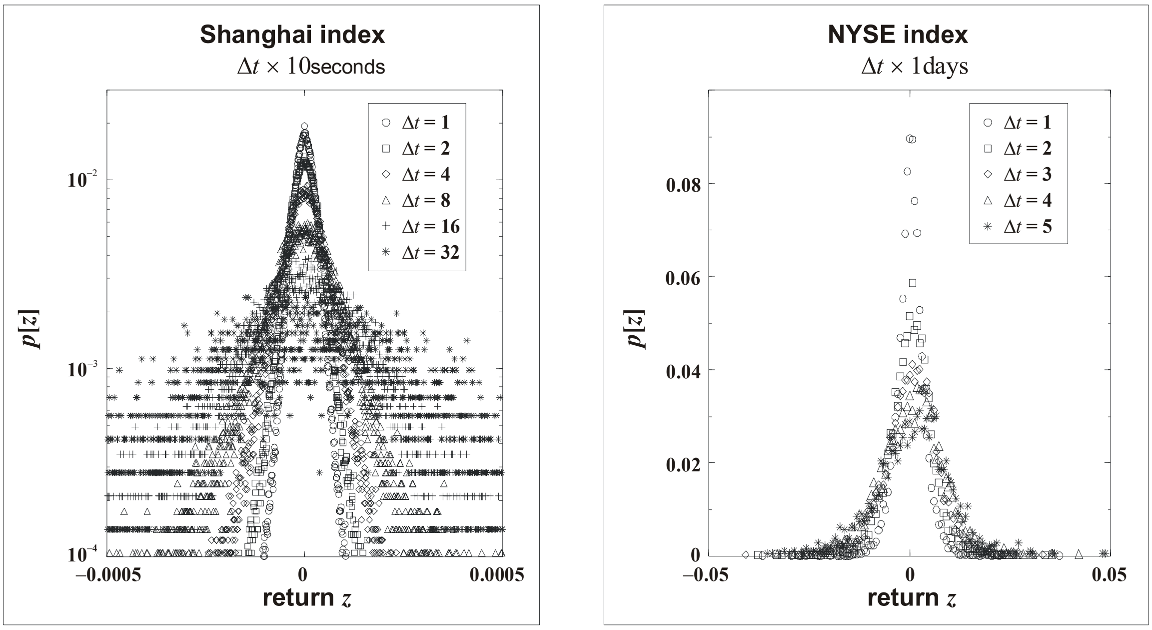

Figure 3-6 PDF of returns over different time-increments Δt , for the Shanghai market index (left) and NYSE composite index (right). For the Shanghai market, Δt is in units of 10 second intervals, while for NYSE Δt is in units of days.

The sets of PDFs look quite different for each market, and for different timescales Δt within each market. As expected, the distributions are roughly symmetrical with a wider spread for increasing Δt . However, the positive and negative tails of the distributions are larger than that of a Gaussian process[10]. Since larger Δt implies less data points, it is difficult to determine the parameters characterizing the distributions just by investigating the spreads. Hence we will study the peak values at the centre of the distribution $p[z=0]$, i.e. the probability that the return is zero, as a function of Δt .

[10] In addition to the visual appearance, this statement can be checked by trying to fit Gaussians to these PDFs. Despite the different possible criteria that one could use to fit such a Gaussian, the central peak in the data always tends to be narrower and higher while the tails in the data are always fatter. This is shown for example in Figure 3-9. In Chapter 6, we report values of the excess kurtosis for the NYSE data which further confirms its non-Gaussian nature.

In this way, we can investigate the point in each probability distribution with the largest amount of associated data. Graphs of these peak heights as a function of Δt are shown below:

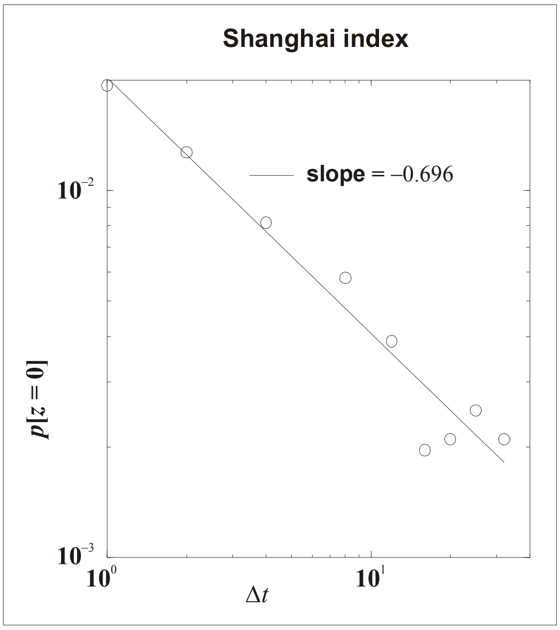

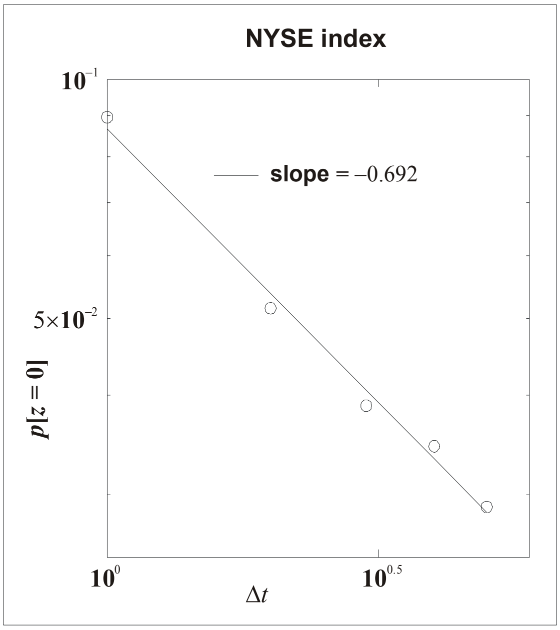

Figure 3-7 A log-log plot showing the central peak of the PDF $p[z=0]$ as a function of the time-increment Δt for the Shanghai market index (left) and the NYSE composite index (right).

The data are fit well by a straight line on the log-log plot in each case, and hence exhibit a power-law behaviour. The slopes are -0.6957 (Shanghai) and -0.6922 (NYSE) which are very close in value[11]. In particular, both are different from the value expected for a Gaussian PDF, as we now show. From Equation (2.25) we know that the PDF for a Gaussian would have the form:

$$p_G[z] = \left[\frac{1}{2\pi\Delta t\sigma^2}\right]^{\frac{1}{2}} e^{-\frac{z^2}{2\Delta t\sigma^2}} \tag{3.7}$$

where σ^2 is the variance for a single timestep and Δt is the number of timesteps per time-increment. Hence $\log p_G[z=0] = -0.5\log \Delta t + \text{constant}$, yielding a slope of value -0.5 . Hence *both* markets give slopes which are similar to each other, but *neither* gives a value close to the Gaussian result. The scaling behavior of the associated non-Gaussian process in both markets can be seen from Figure 3-7 to survive over a wide range of Δt . Let's assume for the moment that the main central portion of the distribution of returns can be described by a Levy stable distribution. The Levy distribution, while not known in closed form, can be expressed as the following exact integral, in terms of the index α and an extra parameter γ :

$$p_{L,\alpha}[z] \equiv \frac{1}{\pi}\int_0^\infty e^{-\gamma\Delta t|q|^\alpha} \cos(qz)\,dq \tag{3.8}$$

[11] A very similar value of -0.712 ± 0.025 was observed by Mantegna and Stanley for the S&P500.

where $e^{-\gamma\Delta t|q|^{\alpha}}$ is the so-called characteristic function of a symmetrical Levy stable process (see [MS] for more details). The probability of zero-return is hence given by

$$p[z=0]=p_{L,\alpha}[0]\equiv\frac{\Gamma\left[\frac{1}{\alpha}\right]}{\pi\alpha(\gamma\Delta t)^{\frac{1}{\alpha}}} \qquad (3.9)$$

where Γ is the Gamma function. Hence $\log p[z=0]=-(1/\alpha)\log\Delta t+\text{constant}$ which has a slope $-(1/\alpha)$. Taking the reciprocal of the empirical slope values found above, we obtain $\alpha=1.44$ for Shanghai and $\alpha=1.44$ for NYSE (both to three significant figures). This is to be compared to $\alpha=2$ for a Gaussian. We note that in the original scaling study of the S&P 500, Mantegna and Stanley found essentially the same value of $\alpha=1.40\pm0.05$.

We now investigate whether this Levy scaling can be extended across the entire PDF of returns for both markets and for all timescales Δt. To do this, we perform the transformations:

$$z_s\equiv z\Big/\left[(\Delta t)^{\frac{1}{\alpha}}\right] \qquad (3.10)$$

$$p_s[z_s]\equiv(\Delta t)^{\frac{1}{\alpha}}\,p_{L,\alpha}[z]=(\Delta t)^{\frac{1}{\alpha}}\,p_{L,\alpha}\left[(\Delta t)^{\frac{1}{\alpha}}z_s\right] \qquad (3.11)$$

If the PDFs were Gaussian and hence $\alpha=2$, it is easy to see that Equations (3.10) and (3.11) become:

$$z_s\equiv z\Big/\left[(\Delta t)^{\frac{1}{2}}\right] \qquad (3.12)$$

$$p_s[z_s]\equiv(\Delta t)^{\frac{1}{2}}\,p_G\left[(\Delta t)^{\frac{1}{2}}z_s\right]=(\Delta t)^{\frac{1}{2}}\left[\frac{1}{2\pi\Delta t\sigma^2}\right]^{\frac{1}{2}}e^{-\frac{\left[(\Delta t)^{\frac{1}{2}}z_s\right]^2}{2\Delta t\sigma^2}}=\left[\frac{1}{2\pi\sigma^2}\right]^{\frac{1}{2}}e^{-\frac{z_s^2}{2\sigma^2}} \qquad (3.13)$$

which is *independent* of the time-increment Δt. Hence all curves would then appear identical when plotted in the transformed variable z_s. This is exactly the property of self-similarity that we mentioned in Section 2.2.3.4 of Chapter 2. This Gaussian scaling is a special case of the Levy scaling result in Equations (3.10) and (3.11), and suggests that *if* the PDFs were approximately Gaussian, then they should appear to collapse onto a single curve under the transformations of Equations (3.10) and (3.11) with $\alpha\approx2$. Since we have already found evidence that $\alpha=1.44$ for both markets[12], we will instead use this value to perform the scaling in Equations (3.10) and (3.11). The rescaled PDFs are shown below, i.e. $p_s[z_s]$ vs. z_s for different Δt. We emphasize that they correspond to the same data as in Figure 3-6.

[12] Of course, all markets will not necessarily have this same value of α. However research studies to date suggest that all markets show similar scaling properties, with values of α deviating from the Gaussian value of 2.

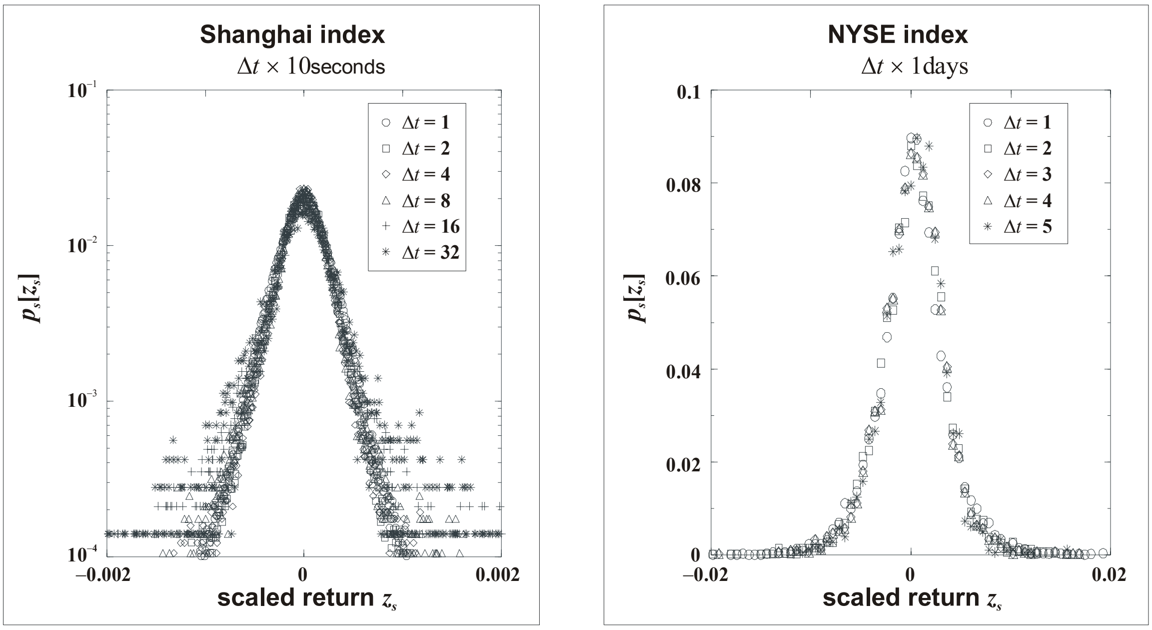

Figure 3-8 Re-scaled plot of Figure 3-6 using the transformations in Equations (3.10) and (3.11). For both the Shanghai market index (left) and the NYSE composite index (right), we use $\alpha = 1.44$ to 3 significant figures. This value of α comes from the slopes in Figure 3-7.

The data-collapse for $\alpha = 1.44$ is quite impressive, and reasonably complete[13] apart from some data points in the tails of the Shanghai dataset. The closer one looks to the central point $z_s = 0$, the stronger is the extent of this data-collapse. These observations imply that the Levy distribution is a better description than the Gaussian to describe the dynamics of the price-process for small-to-medium sized returns. Below we show the Shanghai data, together with the Levy distribution corresponding to $\alpha = 1.44$, and a 'best-fit' Gaussian chosen to have the same standard deviation as the experimental data. Also shown is a narrower Gaussian whose standard deviation was chosen to be significantly smaller than that of the experimental data in order to fit the central portion of the data curve.

[13] In order to be consistent, we have performed this re-scaling using the same value of α as that deduced from Figure 3-7. We note however that the data-collapse can be improved by choosing α close, but not exactly equal, to 1.44. This makes sense, because of the natural error in estimating the slope in Figure 3-7.

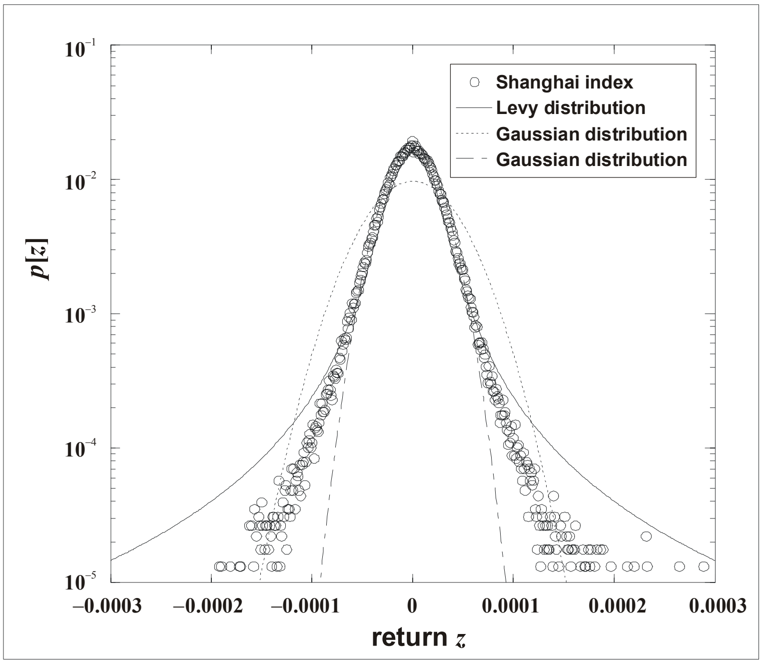

Figure 3-9: PDF of returns for the Shanghai market data with $\Delta t = 1$. This plot is compared to a stable symmetric Levy distribution using the value $\alpha = 1.44$ determined from the slope in Figure 3-7. The agreement is very good over the main central portion, with deviations for large z. Two attempts to fit a Gaussian are also shown. The wider Gaussian is chosen to have the same standard deviation as the empirical data. However, the peak in the data is much narrower and higher than this Gaussian, and the tails are fatter. The narrower Gaussian is chosen to fit the central portion, however the standard deviation is now too small. It can be seen that the data has tails which are much fatter and furthermore have a non-Gaussian functional dependence.

One can extend this investigation in order to determine if a truncated Levy distribution can be used to describe the stochastic process for larger returns which are outside the Levy stable region. Since one would in this case like an accurate representation of the tails, it is better to work with the cumulative distribution $p[z < Z]$ for each timescale Δt, where $p[z < Z]$ is the probability that the return z is less than a given value Z. While we won't go through the specific analysis here, we make the comments that the tails generally decay to zero with an exponent $\alpha > 2$ which is therefore outside the Levy stable regime[14]. Hence the earlier claim that the variance will be large but finite.

Finally we comment on one further simple test. We saw earlier that the autocorrelation function of returns is zero for anything but the smallest time-scales in both markets. This implies that the returns are essentially uncorrelated. If we assume they are also identically distributed, then the results of Chapter 2 suggest that the standard deviation $\sigma_{\Delta t}$ should increase as $\sigma_{\Delta t} \propto (\Delta t)^{\frac{1}{2}}$. However, non-zero residual correlations or non-stationarity of the price-process could alter this relationship. In Chapter 6 we investigate this further for the NYSE dataset. The numerical results show that the $(\Delta t)^{\frac{1}{2}}$ scaling is not followed for all Δt. This again provides evidence for a price-process which is beyond standard finance theory. In order to further explore the presence of higher-order correlations using

[14] See for example, the original paper of Mantegna and Stanley for the S&P500. For similar analysis applied to the Hang Seng index, see B.H. Wang and P.M. Hui, Eur. Phys. J. B **20**, 573 (2001).

such scaling ideas, Bouchaud and Potters [BP] suggested investigating the scaling relationship $\left\langle {z_{t,t-\Delta t}}^{q} \right\rangle \sim (\Delta t)^{\varsigma_q}$, where $z_{t,t-\Delta t}$ are the returns for a given time-increment Δt and q is an arbitrary power. The case $q = 2$ represents the scaling of the variance. The use of a function ς_q allows the Δt-dependence to change exponent over different timescales. This enables an investigation of so-called multiscaling by mapping out ς_q as a function of q.

3.4 Challenging the standard theory

We have uncovered various empirical facts associated with real markets. Our particular examples concerned two structurally and historically different markets, and yet similar statistical results were obtained for each of these markets. In particular these results lie *outside* the paradigm of standard finance theory. We stress that these results are just an illustration of the large number of such studies that have appeared in the Econophysics field. Taking these statistical studies as a whole, the evidence strongly suggests that despite the nuances that different markets seem to have, there is a basic set of qualitative *stylized facts* which markets across the world seem to exhibit. Although there is as yet no complete list of such facts, the following features which we illustrated in Section 3.3 do appear:

- fat-tailed PDF for price-changes, with non-trivial scaling properties
- slow decay of the autocorrelation of absolute value of price-changes
- volatility clustering
- fast decay of the autocorrelation of price-changes

The ability to reproduce these features represents an important test for any candidate market model. However the Gaussian random-walk paradigm does not include the first three of these features. The standard model is therefore an approximation to the truth. It may not be a bad approximation in terms of giving the general form of the distribution of a given financial variable, but the ‘devil is in the details’.

The conceptual argument against any non-random walk behaviour of markets is related to arbitrage, i.e. the idea that there should be ‘no-free lunch’ in the financial markets. In short the argument goes like this: ‘Surely if there was an opportunity for predicting the future behaviour with any certainty at all, then someone would have found it and traded on it – hence they would have traded away this opportunity.’ The same argument is also applied to propose that no such opportunities can exist between markets, e.g. it should not be possible to buy futures contracts at the same time as assets, and then reverse the trade in some way in order to make money. If that opportunity existed, someone would find it – and the opportunity would gradually disappear as the person traded on it. However there are various practical arguments that one can make against this theoretical hypothesis:

(i) In order to exploit an opportunity, you have to find it in the first place. In other words, in order that no such opportunity exists, somebody somewhere would have to find it, and then trade on it. After all, you can only pick up a 'free' $20 bill off the ground if you happen to stumble across it. In order to guarantee exploitation, such opportunities should therefore be somewhat 'obvious'. However, higher-order temporal correlations are, almost by definition, extremely well hidden and hence non-obvious.

(ii) Even assuming such an opportunity has been found, and someone has started trading on it, the timescale over which this opportunity then dies away may not be that small. This is particularly true if the person trades in small amounts (which he/she is likely to do if they don't want anyone else to notice). Alternatively, the transaction costs may be too large to take full advantage. For example in emerging markets, it may need an insider within the country to make the trade in order to avoid various tax penalties and legal barriers.

(iii) Suppose that the markets *are* effectively random at some instant in time. The fact that there are chartists who *think* they see patterns in this randomness, and seem to have a set toolbox for interpreting the next movements, means that patterns could be introduced into the subsequent market dynamics merely by the collective actions of these traders. This issue is the subject of Chapters 4 and 5.

In short, we feel that one could justifiably turn the arbitrage question around and ask: 'why *shouldn't* there be some form of arbitrage opportunity, albeit limited?'

3.5 Toward a general stochastic process framework

In Chapter 2 we discussed i.i.d. variables using probability theory. However we saw in Section 3.3 that real price-changes do not seem to have this i.i.d. property. For completeness, we will therefore give a brief discussion of how probability theory can be extended to discuss non-i.i.d.variables. Suppose that y_1 and y_2, measured at times t_1 and t_2 respectively, correspond to two market variables. Some possibilities include:

- y_1 and y_2 are prices for the same asset, or increments in that asset price
- y_1 and y_2 are volumes of trade, or increments in volume of trade
- y_1 is a price or increment in the price, while y_2 is a volume or increment in the volume
- y_1 and y_2 are measurements of the daily volatility, or increments in the daily volatility.

In addition, y_1 and y_2 could represent an infinite number of functions (both linear, and non-linear) of these prices, volumes, volatilities, etc. Equally, y_1 and y_2 could refer to different assets or markets. Given the global nature of today's markets and common news sources, there could be any number of financial variables which are non-i.i.d., or which have become non-i.i.d. over time.

The joint probability distribution for having an outcome (e.g. a price) y_1 at time t_1, and an outcome (e.g. a price) y_2 at time t_2, is given by $p[y_2,t_2;\,y_1,t_1]$. If we consider a string of variables (e.g. prices) obtained at a sequence of times, we have

$$p[y_N,t_N;y_{N-1},t_{N-1};\ldots;y_1,t_1] \equiv p[\{y,t\}] \tag{3.14}$$

A priori, we don't know whether there are any hidden correlations, nor do we know the probability distribution of the individual variables (e.g. prices) at each time. We can rewrite this expression exactly in terms of conditional probabilities as:

$$\begin{aligned} &p[y_N,t_N;y_{N-1},t_{N-1};\ldots;y_1,t_1] \\ &= p[y_N,t_N \mid y_{N-1},t_{N-1};\ldots;y_1,t_1]\,p[y_{N-1},t_{N-1};\ldots;y_1,t_1] \\ &= p[y_N,t_N \mid y_{N-1},t_{N-1};\ldots]\,p[y_{N-1},t_{N-1} \mid y_{N-2},t_{N-2};\ldots]\,p[y_{N-2},t_{N-2};y_{N-3},t_{N-3};\ldots] \\ &= \prod_{i=2}^{N} p[y_i,t_i \mid y_{i-1},t_{i-1};y_{i-2},t_{i-2};\ldots;y_1,t_1]\,p[y_1,t_1] \end{aligned} \tag{3.15}$$

where $p[y_i,t_i \mid y_{i-1},t_{i-1};\ldots;y_1,t_1]$ is the conditional probability to see an outcome y_i at time t_i given a history of earlier values $y_{i-1},t_{i-1};y_{i-2},t_{i-2};\ldots;y_1,t_1$. If we know the conditional probabilities, and we know the past history, then we can determine the probability distribution of future values. In the simplifying case that the distribution in Equation (3.14) is invariant under an arbitrary translation of the origin of time for all N, the process is said to be stationary. Note that if only the means and covariances are independent of time, then the process is called weak stationary or wide-sense stationary. We can then rewrite each term in Equation (3.15) as follows:

$$p[y_i,t_i \mid y_{i-1},t_{i-1};y_{i-2},t_{i-2};\ldots;y_1,t_1] = p[y_i,t \mid y_{i-1},t-T_1;y_{i-2},t-T_2;\ldots;y_1,t-T_{i-1}] \tag{3.16}$$

If this conditional probability can be limited to a finite history for all t, we can write

$$p[y_i,t \mid y_{i-1},t-T_1;y_{i-2},t-T_2;\ldots;y_1,t-T_{i-1}] = p\left[y_i,t \mid \underbrace{y_{i-1},t-T_1;y_{i-2},t-T_2;\ldots;y_{i-n},t-T_n}_{n \text{ terms}}\right] \tag{3.17}$$

We will now simplify the notation by writing $y_i,t \to y_t$, hence the above equation becomes:

$$p[y_t \mid y_{t-T_1},y_{t-T_2},\ldots] = p\left[y_t \mid \underbrace{y_{t-T_1},y_{t-T_2},\ldots;y_{t-T_n}}_{n \text{ terms}}\right] \tag{3.18}$$

which describes an n'th order Markov process. In the simplifying case that $n=1$, i.e. it only depends on the previous value, then $p[y_t \mid y_{t-T_1},y_{t-T_2},\ldots]$ is equivalent to $p[y_t \mid y_{t-T}]$ which is usually just referred to as a Markov process. This is the case of the coin-toss price model of Chapter 2, and hence the random-walk price model in standard finance theory, with y_{t-T} representing the price at time $t-T$ and y_t representing the price at time t. If y and t are both discrete variables, then we can simply replace $p\left[y_t \mid y_{t-T}\right] \to p\left[y_t \mid y_{t-1}\right]$ and we hence have a Markov chain (see [F] and [G]). Note that we can always convert an n'th order Markov process for a scalar variable y into a first-order Markov process in an n-dimensional variable, e.g. by converting

$$y_t, y_{t-1} \to \underline{w}_t = \begin{pmatrix} y_t \\ y_{t-1} \end{pmatrix} \tag{3.19}$$

For a simple Markov chain with M possible values of y_t (i.e. M possible 'states' $y_{t,\alpha}$ where $\alpha = 1, 2, \ldots M$) we have

$$p\left[y_{t+1,\beta}\right] = \sum_{\alpha=1}^{M} p\left[y_{t+1,\beta} \mid y_{t,\alpha}\right] p\left[y_{t,\alpha}\right] \tag{3.20}$$

Defining an M-component vector of the state probabilities $\underline{p}_t = \left\{p\left[y_{t,1}\right], p\left[y_{t,2}\right] \ldots p\left[y_{t,M}\right]\right\}$ and a matrix of transition probabilities $\underline{\underline{P}}_{\alpha\beta} = p\left[y_{t+1,\alpha} \mid y_{t,\beta}\right]$ then the update for all states can be written:

$$\underline{p}_{t+1} = \underline{\underline{P}} \cdot \underline{p}_t \tag{3.21}$$

Applying this repeatedly with $\underline{\underline{P}}$ being time-independent, gives

$$\underline{p}_{t+n} = \left[\underline{\underline{P}}\right]^n \cdot \underline{p}_t \tag{3.22}$$

The powers of the matrix $\underline{\underline{P}}$ therefore determine the temporal evolution of the system. If it is possible to get from every state to every other state of the system, then the system is said to be ergodic. In a steady state whereby $\underline{p}_{t+1} = \underline{p}_t = \underline{p}$, then we have that $\underline{p} = \left[\underline{\underline{P}}\right] \cdot \underline{p}$. In this way we can find the stationary probabilities for occupying each of the M states. The continuous-time analog of this Markov chain equation follows a similar derivation, but with the summation being replaced by an integral.

We have presented a brief methodology for describing the stochastic properties of price-models which go beyond the coin-toss price model of Chapter 2. The use of $n > 1$ Markov models will allow a degree of temporal correlation and dependence to be incorporated, in addition to the possibility of non-identical distributions at different times. In the next Section, we will explicitly use an $n = 2$ (i.e. second-order) Markov chain calculation to look at the effect of temporal correlations on earnings in a model market setting. Although the market scenario itself is quite artificial, it illustrates very well how counter-intuitive the investment process can be in the presence of temporal correlations.

3.6 Effects of temporal correlations in a market

3.6.1 Winning by losing

We have seen empirically that subtle temporal correlations can lie hidden in financial time-series. There is no quick recipe to say how one might exploit such correlations. However we would like to demonstrate, using a simple example, how the presence of correlations can produce counter-intuitive effects in terms of trading and investment.

Suppose you are a day-trader who focuses on two particular markets, call them A and B. At the beginning of each day, you have to decide whether to invest in market A or market B. Having decided, you invest a fixed amount to open up a position in your chosen market. At the end of the day you close out that position. You therefore either win or lose on that day according to whether you made a profit or a loss. To keep things simple, we imagine that you receive one unit of reward if you win, and are fined one unit of reward if you lose. We will also assume that the subtleties of your position can be subsumed into a simple probability of winning. Let's start with market A. In the absence of transaction costs and the market-maker's spread, we will take[15] your probability of winning in market A to be $p_{\text{win, A}} = 0.5$. You would therefore neither win or lose on average. If we include the effect of transaction costs and the market-maker's spread, your probability of winning is effectively lowered to $p_{\text{win, A}} < 0.5$. Now consider market B. Again your probability of winning in the presence of transaction costs and the market-maker's spread, is $p_{\text{win, B}} < 0.5$. Hence in both markets you will lose on average.

Now comes the remarkable result which we will subsequently prove. In the presence of certain types of temporal correlation in either market A or B or both, you can actually *win* on average by switching between A and B, even though you would lose if you played repeatedly in just one of the two markets. Hence the phrase: *winning by losing*. Moreover you don't have to have a complicated investment plan for switching – random switching between A and B will work fine, as will periodic switching[16]. This remarkable effect, which is known as the Parrondo effect[17], is a completely alien notion within standard finance theory. The Parrondo effect sounds wonderful for the investor, however there is a catch. Just as you can systematically win on average by switching between two particular types of losing markets, you can also systematically lose on average if you switch between two particular types of winning markets. It all depends on the temporal correlations and associated dynamics in these two markets, as we will now show.

For simplicity we will assume the particular scenario in which the probability $p_{\text{win, A}}$ of winning in market A is independent of past outcomes, just like the coin-toss of standard finance theory. However we assume that in market B, the trader's investment strategy is such that the probability $p_{\text{win, B}}$ of him winning now depends on his past success. We will denote $X[t]$ as the trader's capital at the beginning of day t, or equivalently at the end of day $t-1$. If he wins on day t, his capital becomes $X[t+1] = X[t]+1$. If he loses, it becomes $X[t+1] = X[t]-1$. Since his success in

[15] This is consistent with the standard finance idea of random-walk markets.
[16] An example of periodic switching would be: Monday and Tuesday, market A; Wednesday and Thursday, market B; Friday and Monday, market A etc.
[17] J.M.R. Parrondo, G.P. Harmer, and D. Abbott, Phys. Rev. Lett. **85,** 5226 (2000).

market B depends on his previous success, his capital $X[t]$ is no longer an $n=1$ Markov process. We will consider the specific case where his success in market B depends on his success in the previous two timesteps. Following Section 3.5, we can form a Markov chain in terms of the state $\{X[t-1]-X[t-2],\, X[t]-X[t-1]\}$ whose components are the change in capital during day $t-2$ and the change in capital during day $t-1$ respectively. We summarize the probabilities for winning and losing in the following diagram:

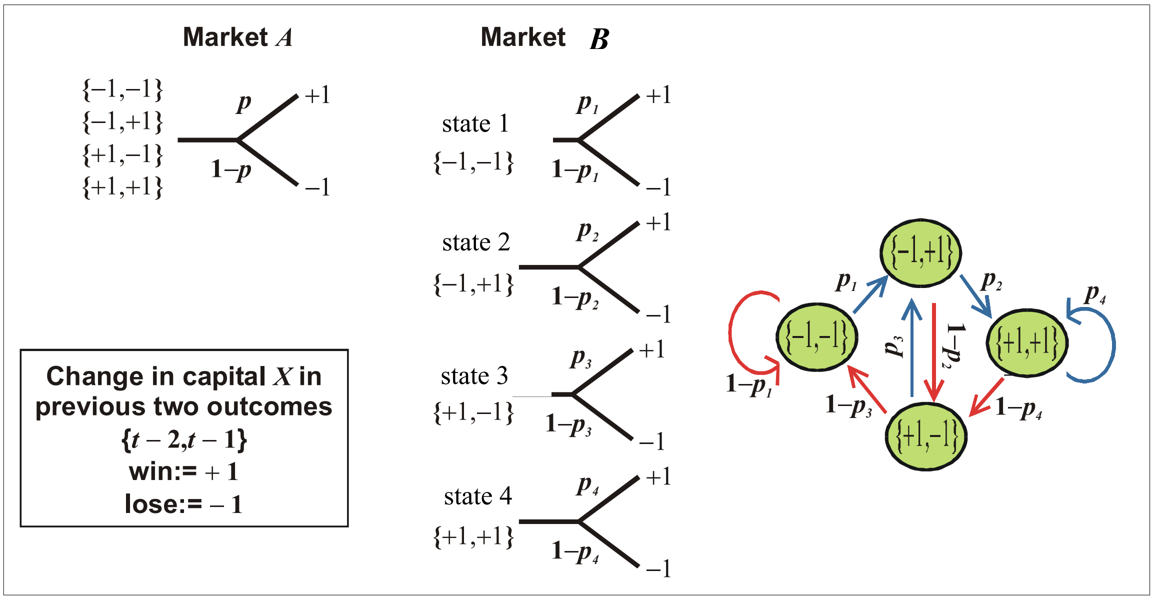

Figure 3-10 Markets A and B and the respective probabilities of winning (+1) and losing (-1) based on the present state of the vector describing the change in capital in two previous outcomes. Right-hand diagram gives a diagram showing the four possible states, and the transition probabilities between these states for market B.

We start the analysis by considering Market B which contains the non-trivial temporal correlations. For the moment we assume that the trader is only playing in market B. Hence we can define the following vector state which is written in terms of his capital X_B at timesteps $t-2$, $t-1$ and t from playing market B:

$$Y_B[t]=\begin{pmatrix} X_B[t]-X_B[t-1] \\ X_B[t-1]-X_B[t-2] \end{pmatrix} \qquad (3.23)$$

Clearly this vector can only take four distinct states:

$$Y_B[t]: \begin{matrix} \begin{pmatrix}-1\\-1\end{pmatrix} & \begin{pmatrix}+1\\-1\end{pmatrix} & \begin{pmatrix}-1\\+1\end{pmatrix} & \begin{pmatrix}+1\\+1\end{pmatrix} \\ \begin{pmatrix}\text{lose}\\\text{lose}\end{pmatrix} & \begin{pmatrix}\text{win}\\\text{lose}\end{pmatrix} & \begin{pmatrix}\text{lose}\\\text{win}\end{pmatrix} & \begin{pmatrix}\text{win}\\\text{win}\end{pmatrix} \\ \text{state 1} & \text{state 2} & \text{state 3} & \text{state 4} \end{matrix} \qquad (3.24)$$

We now define $\pi_i[t]$ *as the probability that the agent is in state* i at time *t*. The probabilities $\pi_i[t]$ form a vector $\underline{\pi}[t]$ with four components, since there are four possible values of i

$$\underline{\pi}[t] = \begin{pmatrix} \pi_1[t] \\ \pi_2[t] \\ \pi_3[t] \\ \pi_4[t] \end{pmatrix} \tag{3.25}$$

Referring back to the possible probabilities for winning in market B given a particular state, we can write a dynamical matrix equation for the evolution of this probability vector:

$$\underline{\pi}[t+1] = \underline{\underline{A}}\ \underline{\pi}[t] \tag{3.26}$$

where

$$\underline{\underline{A}} = \begin{pmatrix} 1-p_1 & 0 & 1-p_3 & 0 \\ p_1 & 0 & p_3 & 0 \\ 0 & 1-p_2 & 0 & 1-p_4 \\ 0 & p_2 & 0 & p_4 \end{pmatrix} \tag{3.27}$$

Hence we have formed a Markov chain, as discussed in Section 3.5. We are interested in the steady-state behaviour, hence we are interested in the solutions in the long time limit $t \to \infty$. In particular, we would like to know if there are any solutions of $\underline{\pi}[t+1] = \underline{\underline{A}}\,\underline{\pi}[t] \equiv \underline{\pi}$, hence:

$$(\underline{\underline{A}} - \underline{\underline{I}})\,\underline{\pi} = 0 \tag{3.28}$$

which represents four coupled equations for the individual components. Solving these equations yields:

$$\underline{\pi} = \begin{pmatrix} \pi_1 \\ \pi_2 \\ \pi_3 \\ \pi_4 \end{pmatrix} = \frac{1}{N} \begin{pmatrix} (1-p_3)(1-p_4) \\ p_1(1-p_4) \\ p_1(1-p_4) \\ p_1 p_2 \end{pmatrix} \tag{3.29}$$

where $N = (1-p_3)(1-p_4) + 2p_1(1-p_4) + p_1 p_2$. Notice that while $\pi_1 + \pi_2 + \pi_3 + \pi_4 = 1$, $p_1 + p_2 + p_3 + p_4 \neq 1$ in general. The probability of winning in a generic run in market B in this stationary regime is therefore given by:

$$p_{\text{win, B}} = \sum_{i=1}^{4} \pi_i\ p_i = \frac{p_1(p_2 + 1 - p_4)}{(1-p_4)(1-p_3+2p_1) + p_1 p_2} = \frac{1}{2 + \frac{c}{s}} \tag{3.30}$$

where $c = (1-p_4)(1-p_3) - p_1 p_2$ and $s = p_1(p_2 + 1 - p_4)$. Since *s* is always positive, it is *c* which controls whether the trader is winning on average in market B, i.e. $c < 0$ and hence $p_{\text{win, B}} > \frac{1}{2}$, or losing, i.e. $c > 0$ and hence $p_{\text{win, B}} < \frac{1}{2}$. The criterion that the trader is losing in market B is therefore $c > 0$ and hence

$$(1 - p_4)(1 - p_3) > p_1 p_2 \tag{3.31}$$

The criterion that the trader is losing when playing in market A is

$$(1 - p) > p \tag{3.32}$$

Now we consider the situation in which the trader randomly switches between markets A and B. Since he chooses markets A and B at random by flipping an unbiased coin, then the same analysis follows as for market B, as long as we make the replacement

$$p_i \to p_i' = \frac{p_i + p}{2} \tag{3.33}$$

Hence the criterion that he loses in the random-switching scenario is given by

$$(1 - p_4')(1 - p_3') > p_1' p_2'$$
$$\left(1 - \frac{p_4 + p}{2}\right)\left(1 - \frac{p_3 + p}{2}\right) > \left(\frac{p_1 + p}{2}\right)\left(\frac{p_2 + p}{2}\right) \tag{3.34}$$

Hence the criteria that he would lose in either market A or market B, but wins by randomly switching between them, are given by the following inequalities:

$$(1 - p) > p$$
$$(1 - p_4)(1 - p_3) > p_1 p_2 \tag{3.35}$$
$$(2 - p_4 - p)(2 - p_3 - p) < (p + p_1)(p + p_2)$$

If we make the following choice of probabilities:

$$p = \frac{1}{2}, \quad p_1 = \frac{9}{10}, \quad p_2 = p_3 = \frac{1}{4}, \quad p_4 = \frac{7}{10}, \tag{3.36}$$

then markets A and B are both fair, i.e. $p_{\text{win,A}} = p_{\text{win,B}} = \frac{1}{2}$. In order to make markets A and B losing, we will now choose

$$p = \frac{1}{2} - \varepsilon, \quad p_1 = \frac{9}{10} - \varepsilon, \quad p_2 = p_3 = \frac{1}{4} - \varepsilon, \quad p_4 = \frac{7}{10} - \varepsilon, \tag{3.37}$$

and ask: for what range of ε are the three criteria in Equation (3.35) satisfied? Setting these probabilities into the inequalities, yields the answer

$$0 < \varepsilon < \frac{1}{168} \tag{3.38}$$

This is the range of ε such that markets A and B are each losing on average, while switching randomly between them results in winning on average. This effect is counter-intuitive, and originates entirely from the nature of the temporal correlations. It has no analog in the random-walk world assumed by standard finance theory.

3.6.2 Drawdowns and crashes

In this section we explore another consequence of the existence of higher-order temporal correlations, related to the properties of potential 'drawdowns' or crashes[18]. Such drawdowns represent moments when the market undergoes a set of downward moves over consecutive timesteps. There is no fixed timescale over which a drawdown will last – both this, and the magnitude of the drawdown, will depend on the precise nature of the temporal correlations. Here we provide a toy model of a price-process which exhibits such higher-order temporal correlations while appearing 'random' from the point of view of lower-order correlations. The model was discussed by Anders Johansen and Didier Sornette[19] and serves as an illustration of the practical importance of higher-order temporal correlations.

Suppose we have a price-process $x[t]$ such that the price-change at timestep t can be written:

$$\Delta x[t,t-1] \equiv \Delta x[t] = x[t] - x[t-1] = \varepsilon[t] + \varepsilon[t-1]\varepsilon[t-2] \qquad (3.39)$$

where $\varepsilon[t]$ is a white-noise process with zero mean and unit variance, and we have used the shorthand notation $\Delta x[t] \equiv \Delta x[t,t-1]$. In particular, we will take the following simple form:

$$\varepsilon[t] = \begin{cases} +1 \\ -1 \end{cases} \quad \text{with probability } 0.5 \qquad (3.40)$$

The mean value of the price-change $\Delta x[t]$ becomes:

$$\langle \Delta x[t] \rangle = \langle \varepsilon[t] \rangle + \langle \varepsilon[t-1]\varepsilon[t-2] \rangle \qquad (3.41)$$

The term $\langle \varepsilon[t-1]\varepsilon[t-2] \rangle$ can be either positive or negative with equal probability, and hence the mean value of the price-change is zero. Next consider the mean of the product of price-changes at different times, i.e. the autocorrelation function $\langle \Delta x[t]\Delta x[t'] \rangle$. Again by writing out the various cases, it is straightforward to show that $\langle \Delta x[t]\Delta x[t'] \rangle = 0$. Hence the price-series produced looks 'random' since its mean and autocorrelation are both zero. In other words, passing the resulting price-series $x[t]$ through standard statistical software packages that just look at low-order correlations such as the autocorrelation, would conclude that the price-series is a random walk.

However this is not the case. If we look at higher-order correlation functions, we begin to find that non-zero correlations appear which should be absent in a strict random walk. In other words, there are temporal correlations in the resulting price-series $x[t]$ which are buried from sight of any standard linear statistical analysis tools. Consider for example the three-point correlation function $\langle \Delta x[t-2]\Delta x[t-1]\Delta x[t] \rangle$. In a random-walk model, this should be zero – but in the present case it is

[18] 'Drawdown' is the term typically used to describe gentler crashes.
[19] A. Johansen and D. Sornette, preprint xxx.lanl.gov/cond-mat/0010050. See also J. of Risk **4**, No. 2 (2001).

non-zero. Furthermore the conditional mean $\langle \Delta x[t] | \Delta x[t-2], \Delta x[t-1] \rangle$, i.e. the mean of $\Delta x[t]$ given the values of $\Delta x[t-2]$ and $\Delta x[t-1]$, is not always zero. In particular,

$$\begin{aligned}\langle \Delta x[t] | \Delta x[t-2], \Delta x[t-1] \rangle &\propto \Delta x[t-2] \Delta x[t-1] \\ &\neq \langle \Delta x[t] \rangle = 0\end{aligned} \qquad (3.42)$$

Since this conditional quantity is non-zero in general, the price-series $x[t]$ has some level of predictability. This is a consequence of the market having memory: the expected value of its next movement depends on the product of the past two movements.

We now illustrate the dramatic effect that such higher-order temporal correlations can have on the size and duration of a drawdown, defined as a drop in the price $x[t]$ from a local maximum in $x[t]$ to a local minimum in $x[t]$. Using Equation (3.39), we can generate the possible price-changes at each timestep. These are shown in the following table:

$\varepsilon[t-2]$	$\varepsilon[t-1]$	$\varepsilon[t]$	$\Delta x[t]$
+1	+1	+1	+2
+1	+1	-1	0
+1	-1	+1	0
+1	-1	-1	-2
-1	+1	+1	0
-1	+1	-1	-2
-1	-1	+1	+2
-1	-1	-1	0

The largest price-change per timestep is +2 or –2. More importantly, we can also say something about the largest drawdown and its duration. Looking at the table, it can be seen that the sequence of coin-tosses $\varepsilon[t]$ giving the largest drawdown is $\pm 1, -1, +1, -1, -1, \pm 1$ which leads to a series of price-changes $0/+2, -2, -2, +2/0$. Hence the longest drawdown has duration of 2 timesteps and magnitude of 4 units. The temporal correlations, which were buried in the price-series and which showed up in an indirect way within higher-order correlation functions, have now surfaced very clearly in the drawdowns. To see this explicitly, imagine what the corresponding answer would be for the usual random coin-toss price model where

$$\Delta x[t] \equiv x[t] - x[t-1] = \varepsilon[t] \qquad (3.43)$$

Again the mean and autocorrelation are both zero, so Equation (3.43) appears indistinguishable from the price-model of Equation (3.39). However the truth is quite different: in particular, the coin-toss price-series in Equation (3.43) can generate drawdowns of unlimited length and size, corresponding to an infinite string of tails (T) which generates a string of -1's. In particular, a drawdown of length at least n (i.e. a sequence of at least n outcomes -1) has a probability $(0.5)^n$ of occurring. Hence there is a small but finite probability of obtaining drawdowns of any length and hence any size. This is in sharp contrast to the model with temporal correlations given by Equation (3.39) which can never generate drawdowns larger than 4 or longer than 2 timesteps. Hence we have seen that the presence of higher-order temporal correlations in a price-series can have dramatic effects on the price dynamics, yielding results which are quite different from the standard finance random-walk model. We note that while the effect of the higher-order correlations in the present model was to *limit* the size of the drawdowns, the opposite can of course occur. A priori it is impossible to say what will happen without knowing the correlation details, hence the importance of accounting for such higher-order temporal correlations in any theoretical finance model.