Some observations on the random walk model
by A. D. Wilkie, United Kingdom

ABSTRACT

The random walk or logarithmic Brownian motion model for share prices is now the standard one for many purposes, particularly option pricing, for which it is very useful. In this paper, however, we draw attention to many of the practical anomalies that can be seen with this model, in the short, medium and longer terms. In the short term we know that share prices move in discrete steps and are not available at all times. In the medium term we see that it is difficult for total return indices and share prices both to be uniform random walks, because dividend yields vary; that if share prices are lognormally distributed, share indices cannot be; that the CAPM is not consistent with lognormal distributions for both share prices and portfolio returns; and that the observed distribution of log share price changes is conspicuously fat-tailed, though the effect reduces with the time step. In the longer term we observe that share prices follow a similar track to many other related monetary indices, consumer prices, wages, company dividends and company earnings, with which they can be described as being cointegrated. The object of the paper is not to draw conclusions, but to stimulate discussion.

KEYWORDS

Random walk; logarithmic Brownian motion; share prices; discrete steps; tick sizes; total returns; dividend yields; lognormal distributions; share indices; CAPM; portfolio returns; fat-tailed innovations; stochastic volatility; regime switching; dividends, earnings, consumer prices; wages; cointegration; autoregressive models.

CONTACT ADDRESS

A. D. Wilkie, Department of Actuarial Mathematics and Statistics, Heriot-Watt University, Riccarton, Edinburgh, EH14 4AS, email: A.D.Wilkie@ma.hw.ac.uk. and InQA Limited, Dennington, Ridgeway, Horsell, Woking, GU21 4QR, email: david.wilkie@inqa.com.

1 Introduction

1.1 The random walk model for share prices is widely used for many purposes, particularly for option pricing, and for these purposes it is extremely useful. In continuous time it is, in its simplest form:

$$\mathrm{dS}(t) = \mu.\mathrm{S}(t).\mathrm{d}t + \sigma.\mathrm{S}(t).\mathrm{dZ}$$

where S(t) is the price of the share or share index at time t, Z is a Brownian motion, and μ and σ are constants. The discrete time equivalent of this, with unit steps, is:

$$\ln\mathrm{S}(t) = \ln\mathrm{S}(t-1) + \mu - \sigma^2/2 + \sigma.\mathrm{Z}(t)$$

or

$$\mathrm{S}(t) = \mathrm{S}(t-1).\exp\{\mu - \sigma^2/2 + \sigma.\mathrm{Z}(t)\}$$

but we can easily write this with a different value of μ, say μ^* as

$$\mathrm{S}(t) = \mathrm{S}(t-1).\exp\{\mu^* + \sigma.\mathrm{Z}(t)\}$$

1.2 Elaborations of the model often allow "stochastic volatility" in some form, i.e. making σ into a function of t, say $\sigma(t)$. One could also make μ into a function of t, say $\mu(t)$, but since the value of μ is not relevant for many option pricing applications, this is seldom done. Either μ or σ or both could be made functions of both S and t, and therefore be clearly stochastic.

1.3 In this paper I wish to make some observations about this model, and indicate in what way it is not a perfect model for representing share prices. Many of the points I make are well known, but perhaps not all are familiar to all readers. In a fully academic paper I would include references for all my statements, but for a discussion paper such as this I shall omit all references, letting the reader fill them in from their own knowledge as they wish.

1.4 I start with the short term, and move forwards to the medium term, and then to the longer term.

2 The short term

2.1 *Discrete steps and times*

2.1.1 The continuous Brownian motion form for a share price implies that $S(t)$ exists for all t, and also that in going from some price S_1 at time t_1 to another price S_2 at time t_2 $S(t)$ has a continuous path, so that it takes every possible value between S_1 and S_2, and possibly some outside that range too.

2.1.2 Neither of these implications is fulfilled in practice. Prices of individual shares, in all markets, move in "tick sizes", in discrete steps according to the convention of the market, often pennies or cents, or sometimes half or quarter units. Some markets use or have used binary fractions down to 32nds or 64ths. However small the ticks are that may be in use, it is impracticable to price in genuinely "real" numbers, because no numerical measurement anyway can do this.

2.1.3 Many researchers are of course well aware of this feature of prices and there are many studies of the very short term movements of share prices, concentrating on ticks, their exact timing and the number of ticks between successive prices; prices do not even need to pass through every discrete tick position, but can jump across ticks. Indeed in exceptional circumstances, such as take-over bids, unusual announcements, or dramatic external events, prices may jump by relatively large amounts. I revert to this in Section 3.4.

2.1.4 The assumption that prices exist for all times t is also not fulfilled in practice. At the best one has discrete prices at a discrete set of times. Any one agent in the market cannot know prices in continuous time. His or her attention allows only a limited number of prices to be observed, even if at times these may be quite frequent. But for a theoretical Brownian motion, even a nanosecond is an eternity.

2.1.5 Further, almost all markets are open only for limited hours. National stock markets have standard trading hours, which cover only part of each day, and only some days, normally excluding weekends and holidays. There are possibilities for trading some instruments, possibly major shares and currency exchange rates round the clock by using different markets, but even if this is becoming more common it is not the standard.

2.1.6 We can summarise: although a Brownian motion might be satisfactory way of looking at the prices of shares from a suitable distance, when we look close enough, everything is discrete. A similar feature, in reverse, applies to physical objects. However smooth an object may be to our touch, when we magnify down to atomic level, it is seen to consist of molecules buzzing around in what may indeed be like Brownian motions. But the object may seem to us to be solid, continuous and smooth. The macro model can differ from the micro one. On the other hand, a solid physical object that moves say, in one dimension, from position x_1 at time t_1 to position x_2 at time t_2 does in fact have some position at all real times between t_1 and t_2, and does pass through every position between x_1 and x_2. Whether it can or does actually move with Brownian motion is a different matter.

3 The medium term

3.1 *Share prices and total returns*

3.1.1 In many applications we need to distinguish the price of a share or share index, $S(t)$, and the corresponding total return index, $T(t)$, where dividends on the share or shares in the index are reinvested in the share or share index. It is commonly assumed that both $S(t)$ and $T(t)$ share prices move with logarithmic Brownian motion, say:

$$dS(t) = \mu_S.S(t).dt + \sigma_S.S(t).dZ$$

and

$$dT(t) = \mu_T.T(t).dt + \sigma_T.T(t).dZ$$

where μ_S and σ_S apply to the share and μ_T and σ_T apply to the total return index. A mathematically easy way to connect these is to assume that dividends are paid and reinvested continuously at rate $y(t).S(t)$ at time t. From this it follows that $\mu_T = \mu_S + y(t)$ and $\sigma_T = \sigma_S$. But if μ_T and μ_S are constants then it follows that $y(t)$ is also constant. I have seen this assumption made in many financial economics papers.

3.1.2 In practice dividends on individual shares are paid at quite long discrete intervals, quarterly, half-yearly or yearly, so the assumption of continuous payment of dividends does not represent this at all well. If the share index contains a large number of shares, and dividends on these are paid at varying times throughout the year, so that, for example, some are received every week, then a continuous assumption may be acceptable within the scale of the particular situation.

3.1.3 However, the assumption of a constant value for $y(t)$ is clearly unrealistic. Dividend yields, both for individual shares and for share indices, are widely quoted, usually based on the dividends paid over the last year, modified by dividends that have been announced but have not yet been paid. It is clear that dividend yields are far from being constant, but do seems to fluctuate around some central position. In most countries a dividend yield of 1% would be considered very low and a yield of 10% would be considered very high.

3.1.4 In the United Kingdom, dividends are announced by a company usually several weeks before they are paid, and some weeks before payment the shares change from being traded “cum dividend” to “ex dividend”. The main share (and bond) indices give details each day of the “ex dividend adjustment”, which is the amount of dividend that has gone “ex” on the index for that day; in practice the accumulated “ex dividend adjustment” for the year to date is given. This allows correct total return indices to be calculated and published assuming one particular assumption about tax on dividends, and allows any investor to calculate another total return index assuming some other rate of tax on income.

3.1.5 These practical considerations make it difficult to see how the mathematical models can be reconciled. Dividend yields are observed to be not constant, and may be subject to different rates of tax in the hands of different investors, so the assumption that μ_S, $y(t)$ and μ_T are all constant with varying t is not a good representation of the facts. This may not matter in some applications, such as option pricing, where the mean return on the share can often be ignored. But there are other many other applications where the mean return is relevant.

3.1.6 I have not seen any discussion of this problem in any financial economics paper, though it may well have been discussed, explained and solved in some paper that I have not seen. If so, I would have expected the results to have been more widely quoted. One approach could be to assume, first, that the market is dominated by investors subject to one rate of tax on dividends (or alternatively that the dominant investors are taxed in the same way on income and capital gains) and secondly that μ_T is constant. Variations in $y(t)$ are then exactly counterbalanced by variations in what I shall now denote as $\mu_S(t)$. Thus, if dividend yields are low, investors assume a higher mean rate of increase in share prices and vice versa if dividend yields are high.

3.1.7 This is not at all an unreasonable model for individual shares. If there is bad news about a company’s prospects, the price of its share typically falls, and the dividend yield rises, because the numerator (dividends over the past year) remains constant. But it is then quite possible that dividends next year will be cut, and the dividend yield may revert to a more average position. But prices of individual shares frequently change by much more than just to represent a one year’s change in dividend. If a share has been on a dividend yield of 3%, a fall in the share price of 30% would imply (crudely) no dividends at all for the next ten years. A rise of 30% would imply (crudely) double the dividend rate for the next ten years, or else a special once only dividend of 30% of the share price, or some other combination with the same direct financial effect.

3.1.8 But individual share prices do not seem to behave in quite this way, and share indices seem to have much larger changes in price and hence dividend yield than can be justified by immediate dividend prospects. An alternative would be to assume that “news” makes investors reassess the long-term prospects for individual share and for the aggregate of all shares, so that say a fall in dividend yield of 1% (from 5% to 4%) implies a 1% increase in the long-term rate of growth of share prices; but it is difficult to see how this can arise without a corresponding increase in share dividends, company earnings and company profits, which the news may not justify. I revert to this in Section 4.1.

3.1.9 Yet another approach would be to bring in the level of interest rates, and allow μ_T to vary with time as $\mu_T(t)$, and relate it to some relevant interest rate $r(t)$. Whether this rate should be the short-term rate or a longer-term one is another question. We could then have $\mu_T(t) = r(t) + \pi(t)$, where $\pi(t)$ is a “risk premium” on shares. It follows that $\mu_S(t) = \mu_T(t) - y(t)$ $= r(t) + \pi(t) - y(t)$, and if we assume that μ_S is constant, we can get $y(t) = r(t) + \pi(t) - \mu_S$ so that dividend yields can be seen to vary with interest rates, with adjustments. Empirically we find that dividend yields do indeed vary with interest rates to some extent, but with a quite varying differential. Indeed in some countries dividend yields have been well below nominal interest rates for quite long periods, as well as being at other times above them.

3.1.10 A further elaboration is to consider yields and interest rates on investments "in real terms". It has long been accepted that nominal interest rates adjust to reflect views about future monetary inflation (the "Fisher effect"). In a number of countries governments have issued bonds whose payments are linked to some price index, so that some measure of "real interest rates" is now available in those markets. Since company profits, earnings and dividends can be expected to be affected by monetary inflation, though not necessarily maintaining their values exactly in real terms, it may be better to compare dividend yields with "real" interest rates. We still find that dividend yields can be above and below real interest rates at different times.

3.1.11 I provide no solution to my question. But the problem of reconciling random walk models for share prices and total return indices could usefully be addressed.

3.2 *Indices and lognormal distributions*

3.2.1 If a share price is subject to the usual logarithmic Brownian motion, then, at any time t, the price, $S(t)$, is distributed lognormally. Assume that a share price index is constructed by taking shares of J companies at time t_0, indexed by $j = 1$, J, with prices $S_j(0)$, and weighted by the number of shares in issue N_j, which we shall assume remains constant during the observation period. Thus the market value of all the shares of company j at time t is $N_j.S_j(t)$. We set the value of the index at time 0, I(0), at some convenient value and calculate a constant K from the relation: $I(0) = K.\Sigma_j\, N_j.S_j(0)$. Then the index at time t is proportional to the market value of all the shares at time t and is calculated by the formula: $I(t) = K.\Sigma_j\, N_j.S_j(t)$. There are other ways of calculating share index values, but this is the most common one method for major indices.

3.2.2 Now if each of the share prices $S_j(t)$ is distributed lognormally, then $I(t)$ is distributed as the weighted sum of J lognormally distributed variables (which may or may not be correlated), and it is *not* distributed lognormally (except in the exceptional and improbable case where the prices of all J shares are perfectly correlated). So $I(t)$ cannot be described as following geometric Brownian motion. If the number of shares J is sufficiently large, the distribution of $I(t)$ will approach normality. Alternatively, if $I(t)$ is in fact lognormally distributed, then it cannot be derived from the sum of J variables which are also lognormally distributed; but it is difficult to see what distribution the individual shares might have.

3.2.3 It is possible to do a little algebra to see what the distribution of an index might look like. We know that the skewness (the scaled third moment) of a normally distributed random variable is zero and the kurtosis (the scaled fourth moment) is 3. We also know that the skewness and kurtosis of a lognormally distributed random variable depend only on the value of σ (the standard deviation of the logarithms of the variable) and not on the value of μ (the mean of the logarithms). They are given by:

$$\text{skewness} = (\delta^6 - 3\,\delta^2 + 2)/(\delta^2 - 1)^{3/2}$$

and

$$\text{kurtosis} = (\delta^{12} - 4\,\delta^{6} + 6\,\delta^{2} - 3)/(\delta^{2} - 1)^{2}$$

where $\delta = \exp(\sigma^2/2)$.

3.2.4 We now assume that J *independent* lognormally distributed random variables are added together and the result is divided by J. This could represent the return on an index or a portfolio of independent and initially equally weighted shares where the share returns are lognormal. The mean for the portfolio is the same as the mean of the individual share returns, and the standard deviation of the index reduces as J increases. So also do the skewness and kurtosis. If we put $\sigma = 0.3$, a highish but not unreasonable value for the annual standard deviation for individual shares, we get the following values of the skewness and kurtosis of the index return, depending on the number of shares, J.

J	Skewness	Kurtosis
1	0.95	4.64
2	0.67	3.82
3	0.55	3.55
4	0.47	3.41
5	0.42	3.32
10	0.30	3.16
15	0.25	3.11
20	0.21	3.08
30	0.17	3.05
40	0.15	3.04
50	0.13	3.03
100	0.09	3.02

3.2.5 We see that the excess kurtosis (in excess of 3) reduces quite quickly as the number of shares is increased. By 10 shares one would hardly notice the excess unless the dataset were very long. The skewness seems to decline to zero a little more slowly, but, again unless there is very long dataset the standard error of any estimate would not indicate that the index returns were other than normal. Of course, I have assumed independence, and share movements within one market are noticeably correlated. I discuss this problem further in Section 3.3.

3.2.6 An alternative method for calculating indices, used in the past, but less popular nowadays, is the geometrical equally weighted index. This usually uses only a few leading shares in the relevant market and ignores the weights. It assumes that the investor puts equal investments into each of the J shares, and maintains equal weightings at all times. Thus we start with $I(0) = K.\exp(\Sigma_j \ln S_j(0))$ and then each day calculate $I(t) = K.\exp(\Sigma_j \ln S_j(t))$. If we follow through the algebra, we find that, if each of the share prices, $S_j(t)$, is lognormally distributed, and if rebalancing is done continuously, $I(t)$ is also lognormally distributed.

3.2.7 The old *Financial Times* "30 share" index was constructed in this geometrical way, and I believe it still is, though it has been superseded in importance by the FTSE 100 index. The latter type of index, arithmetically weighted by the number of shares in issue, is much better for representing an institutional portfolio. The geometric type is better at representing the investments of an individual investor who might well put a small and equal amount into each of a small number of leading shares; but such an investor would hardly be able to maintain equal amounts by market value in each share.

3.3 *Portfolio selection and the CAPM*

3.3.1 The Markowitz portfolio selection model assumes a single time period, say (0, 1), and a number of securities for each of which the return over the time period is a random variable, of which the mean, variance and correlation with the returns of other securities are known. Portfolios are formed at time 0 and the mean and variance of the return on each portfolio at time 1 is calculated. (By a *return* of R I mean that an investment of 1 at time 0 returns R at time 1, a *rate of return* of 100(R – 1)%.) The distribution of returns is not explicitly stated, though it is assumed that returns with higher means and lower variances are preferred, which is consistent with these returns being normally distributed, and second order stochastic dominance applying.

3.3.2 This dominance criterion does not, however, apply to lognormally distributed returns, for which the right criterion is that a return with mean M and variance V dominates a return with mean less than M and variance greater than that along the line from (0, 0) to (M, V). (0, 0) is the point where the return is zero with certainty, giving a *rate of return* of –100%. This only means that the "efficient frontier" stops at the point where the tangent from (0, 0) meets the curve of minimum variance portfolios, rather than at the minimum variance point. If a risk-free security exists, then there is no change in the efficient frontier at all.

3.3.3 However, this does not help us out entirely. If the distributions of the individual shares are lognormal, then the distribution of a portfolio is not lognormal, just as for a market index. We could redefine the way of holding the portfolio from "buy and hold" to "continuously rebalance". If the individual share prices are subject to geometric Brownian motion, and if we continuously rebalance the portfolio so that the original proportions, by market value, are preserved at all times, then the portfolio return is indeed lognormal. But the simple additivity of means and variances (and covariances) is destroyed and needs to be reconsidered.

3.3.4 Within the Capital Asset Pricing Model (CAPM) it is then typically assumed that the return on share j over (0, 1), R_j, is related to the return on "the market", denoted R_M, by:

$$R_j = \alpha_j + \beta_j.R_M + e_j$$

where α may be zero, or may depend on the means of R_M and R_j and e_j is distributed in some unspecified way, with mean zero and variance σ_j^2. Again, the distribution is typically unspecified, but the implicit assumption of normal distributions throughout is often made. However, if R_M and e_j are both lognormally distributed, then R_j is not. It is not easy to see how this can be reconciled.

3.3.5 A further problem with the CAPM is that the market portfolio is often taken to be the portfolio that contains all the relevant shares, or indeed all risky assets. It is further assumed that the e_j terms above are all independent. But if the market portfolio is indeed the sum of all the shares, then not only is the weighted sum of the α_j terms zero and the β_j terms unity (which is not difficult to arrange), but also the weighted sum of the e_j terms is always identically zero, and does not just have mean zero. Thus the e_j terms cannot all be independent; if all but one is known, then the value of the last is determined and must be that value that makes the sum zero. There must be some negative correlation between some of the e_j terms.

3.4 *The lognormal assumption*

3.4.1 Geometric Brownian motion implies that the return on shares or on a share portfolio over any period is lognormally distributed. Thus the distribution of the logarithm of the return should be normal. A normal distribution has skewness of zero and kurtosis of 3. We can examine this from actual data. I have available returns on U.K. share indices at monthly intervals for a long period, and also monthly returns for many other countries over shorter periods. For this purpose I look only at changes in the logarithm of the share price index, not any total return index, for which the results would be similar.

3.4.2 We find that, for the U.K., for 980 monthly changes (differences of the logarithms of a series of chain-linked share price indices at the end of each month) from end-December 1923 to end-August 2005, the skewness is –0.12 and the kurtosis is 11.53. These, if the logarithmic Brownian motion hypothesis were true, would be normally distributed. The skewness is nothing exceptional, well within the confidence interval for a normal distribution; the standard error, assuming normality, is 0.08. However the kurtosis is large, well outside the standard error of 0.16. So the monthly differences are reasonably symmetrical, but quite strongly fat-tailed.

3.4.3 It is interesting to look also at the differences, taken in pairs, in threes, … yearly, etc. A problem here is that for steps of j months we have j subseries, that taking steps from 0 to j, 1 to j+1, 2 to j+2, etc, until we reach j to 2j, which is the second term of the first subseries. Looking at the U.K. series in this way the results are illuminating. Taking j = 2, i.e. two-monthly steps, we have two subseries, the first starting with the step from December 1923 to February 1923, and so, and the second starting with the step from January 1924 to March 1924 and so on. For the first subseries we get skewness = 0.56, kurtosis = 15.84, and for the second subseries skewness = –0.42 and kurtosis = 7.93. Inspection of the data shows that the first subseries includes January and February 1975, when successive log changes were 0.42 and 0.21, whereas the second subseries splits these into two, but combines October and November 1987, with log changes of –0.31 and –0.11.

3.4.4 This irregularity continues as the number of months increases. The value of the kurtosis generally reduces a bit, but still irregularly, and the skewness continues to show wide contrasts. For example, at five-month intervals the skewness ranges from –0.92 to 0.54, and the kurtosis from 4.42 to 10.77. At yearly intervals we get the results as shown in the following table:

Year ending	Skewness	Kurtosis
January	–0.20	2.98
February	–0.21	2.72
March	–0.50	3.45
April	–0.43	2.79
May	–0.76	3.63
June	–0.91	4.16
July	–0.80	4.10
August	–0.76	4.60
September	–0.85	7.00
October	–0.74	7.77
November	–0.54	8.89
December	–0.32	7.57

3.4.5 This provides a salutary warning. It is tempting to analyse data at yearly intervals; some series are available only yearly. Yet one can see that one's conclusions might depend strongly on which month one chose as the year-end. If one chose any month from end-January to end-May one would not consider this data very fat-tailed; but any month from end-September to end-December would lead to the opposite conclusion. Months from December to February do not look seriously skew; months from May to October do.

3.4.6 However, it is clear that this data belies the assumption of logarithmic Brownian motion. This feature is well known. The observed data have been modelled in a variety of ways:

a) fat-tailed distributions for the changes;
b) regime-switching;
c) conditional heteroscedasticity (stochastic volatility)

3.4.7 These have some fundamentally different features. The fat-tailed distribution model assumes that the distribution of innovations is constant, but the exceptional values observed come from some fat-tailed distribution, and their occurrence does not affect subsequent innovations. Among the fat-tailed distributions that have been suggested are the *t*-distribution and the difference between two positive distributions, e.g. lognormal, gamma, etc. An extreme fat-tailed distribution is the α-stable or stable Paretian one; however this has infinite variance, which is inconvenient, and it does not converge to normality as the number of steps is increased, which does not fit the observations so well.

3.4.8 Regime-switching models assume that one set of parameters is in force at one time and another set at another time; switches between the two or more states happen at random, usually as a Markov chain. It suits the approach that extreme events, such as the "9/11" attack on New York, the outbreak of a war, a natural catastrophe, or some other major external event causes revised attitudes within the stock market. Yet it is not clear whether any external political events caused the stock market crashes of 1929, 1987 or 2000, and a more gradual shift in attitude might be more plausible.

3.4.9 Stochastic volatility models assume a slow change in the underlying volatility of what is still normally taken as a lognormal model for the innovations conditional on that volatility. Typically they focus on the σ parameter, which is what interests those engaged in option pricing, rather than the mean, the μ parameter. Regime switching models generally assume that the values of both parameters change between regimes.

3.4.10 There is, of course, nothing to stop combinations of these models. Either the regime switching model or the stochastic volatility model could have fat-tailed innovations. The stochastic volatility model could be subject to switches between different regimes, and also have fat-tailed innovations, possibly with different distributions in each. But the number of parameters required to describe such complex models may be too great for the data to stand. In any case they are a long way from simple logarithmic Brownian motion, even if this is the original inspiration and a tolerable first approximation.

4 The longer term

4.1 *Dividends and earnings*

4.1.1 In Figure 1 I show a graph of five UK indices: consumer prices, wages, share prices, share dividends and share earnings, at monthly intervals from December 1918 to August 2005. Not all start at the same point. In order to demonstrate their similarity I have plotted them all with black lines and omitted a key. The point is not that you should look at the differences, but at the similarities. All have risen in the same sort of way, by a large factor. Share prices oscillate most; share dividends and earnings next; consumer prices and wages least. But all have a very similar overall pattern.

Figure 1. UK indices of consumer prices, wages,
share prices, dividends and share earnings 1918-2005

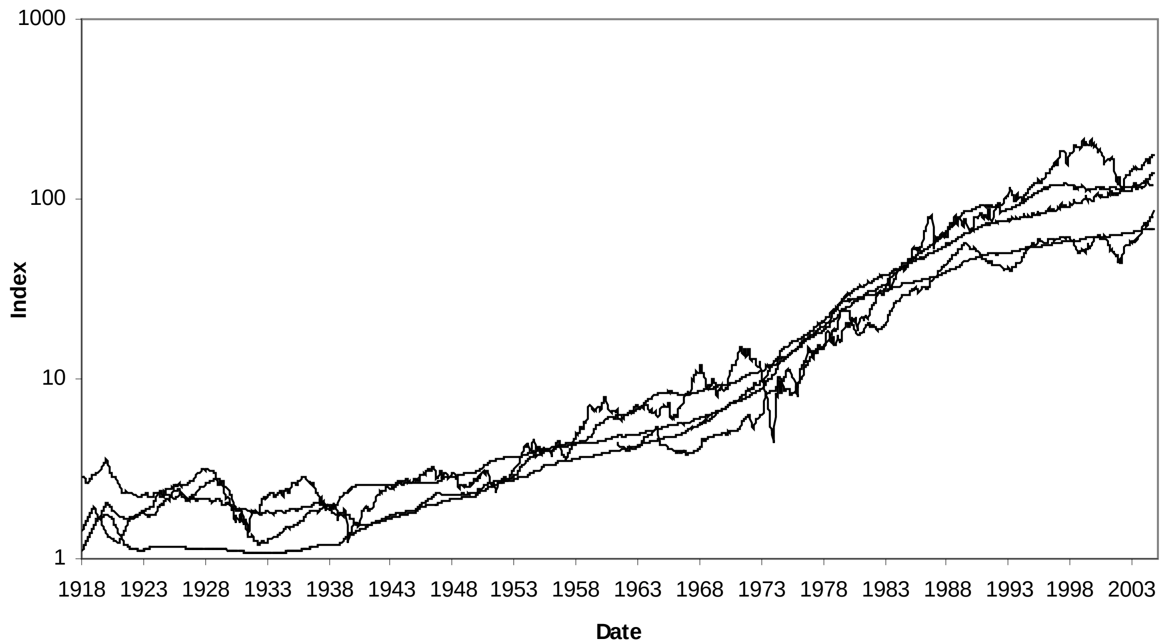

4.1.2 It is perhaps not surprising that all these indices move in a similar way. All represent some values measured in the same currency units (U.K. £). All are connected with the economy of the United Kingdom, though in different ways. Share prices, dividends and earnings are likely to more closely connected with one another than any would be with consumer prices or wages. Series with this sort of graph are likely to be at least correlated, so that changes in one (or the changes of the logarithm) are correlated with simultaneous or lagged changes in another (in time series analysis described as a "vector autoregressive model"); or possibly also "cointegrated", a time series model in which some function of two (or more) series, such as the difference between the logarithms, or the log ratio (e.g. dividend yield for dividends and share prices, or pay-out ratio for dividends and earnings, or P-E ratio for share prices and earnings) is a stationary, possibly autoregressive, time series. Thus if the two series get "too far apart" in some sense, the cointegration feature tends to pull them together.

4.1.3 The weak form of the efficient market hypothesis states that one cannot forecast future changes in share prices by studying the past history of share prices. This is quite plausibly true in the short term. Even if share prices are found to be related to dividends or earnings, these change quite slowly in the short term. For individual companies dividends and earnings change at discrete intervals, possibly quarterly, possibly half-yearly or even annually. Changes in them cannot directly affect share prices in the short run. But over longer periods the situation may be different.

4.1.4 The price of an individual share, or of an index, S, can be decomposed into a dividend, D, and a dividend yield, Y:

$$S = D / Y$$

Dividends in turn can be decomposed into earnings, E, and a pay-out ratio, P:

$$D = E \times P$$

Earnings can be decomposed into a consumer price index value, Q, and a real earnings figure, R:

$$E = Q \times R$$

Combining these we can put:

$$S = Q \times R \times P / Y$$

So changes in log S can be decomposed into:

changes in consumer prices;
changes in real earnings;
changes in pay-out ratio;
changes in dividend yield, or "rating"

4.1.5 I do not go into details, but I now summarise what would be much too long a statistical investigation to reproduce in this paper. Consumer prices change as an integrated autoregressive time series. Thus they are like a random walk, but with autocorrelated changes. They have no fixed mean, but their rate of change often behaves like a stationary time series (in some countries there have been periods of "hyperinflation" when the behaviour even of the rate of inflation has seemed to non-stationary). Real company earnings per share, at least in the U.K., seem to be stationary; they have had no obvious drift, either up or down, but fluctuate around a constant level, or perhaps a very slowly moving level. Pay-out ratios must in general be stationary; companies cannot let them go too high, and generally do not let them go too low (there are conspicuous exceptions where individual companies have found it better to retain all their earnings to finance growth; but no company does this for ever).

4.1.6 That leaves changes in dividend yield, or "rating". I find it difficult to imagine that this is not stationary. It has gone up and down over the years. In "bubble" periods (e.g. 1929 or the late 1990s) dividend yields may go to very low levels (and P-E ratios to very high levels). In what seem times of gloom (e.g. the mid 1970s in the U.K.) they may do the opposite. But if the economy survives (and sometime it might not), a reversion to the mean seems to occur.

4.1.7 If share prices are decomposed in this way, one can see that the expected rate of growth in share prices can also be so decomposed. I revert to the problem I discussed in Section 2.1. The mean rate of growth of a share index and the mean rate of growth of a total return index cannot both be constant if dividend yields vary. But if at times the expected rate of growth of one of the features I mention above might be different from normal, this might justify a higher or lower share price, hence a higher or lower dividend yield (or PE ratio), so that the mean total return might remain constant with a varying dividend yield and a varying mean return on the share price.

4.1.8 What would, in such a model, justify above average share prices could be: above average real earnings growth, or particularly low pay-out ratios that might rise, or the belief that already high ratings will increase further. But this is the fuel of bubbles, if investors buy shares only because they believe that their prices will rise, but they can rise only if other investors can be persuaded in the same way. Bubbles of this sort normally burst.

4.1.9 In addition to these points we should consider the premium that shares have over "risk-free" investments. Unfortunately "risk free" has multiple meanings. Some would consider fixed cash as risk-free. That gives as many risk-free assets as there are different currencies in the world. Others would consider long-term bonds linked to a country's consumer price index as roughly risk-free. That gives as many risk-free assets as there are different price indices to which such bonds are linked. The price indices may be heavily correlated, but not perfectly so. Yet others would be happy with long-term conventional bonds. It may depend on the investor's liabilities or preferences. The "risk-free" rates of interest on all these are not all the same. Some of the discrepancies reflect different views about particular currencies and possible changes in the exchange rates, so a fuller model would need to take these into account. One can describe exchange rates as an autoregressive series wandering around a "purchasing power parity" moving central rate.

4.1.10 All this indicates to me that the simple, perhaps even naïve, random walk model has its place in financial economics and in actuarial thought, but that a comprehensive model for share price changes in the long term needs to be much more complex and needs to considered much more fully than it often is.

4.1.11 It can be countered that the medium or strong form of the efficient market hypothesis suggest that investors do take all these factors into account, and that share prices do therefore move rationally, in the light of all these available facts. I am unconvinced. Different investors, even rational ones, have very different time horizons, very different liabilities, very different cash flows, and very different tax positions. It would be interesting to see or to develop a model that looked at the equilibrium position in this case.

4.1.12 Different individual investors also may have different attitudes, which may be naïve or be well-considered. Some (the trend followers) may buy shares because the price has gone up and they believe that others are buying shares; others may buy them because they always buy some shares and (if anything) believe in "pound cost averaging" (e.g. those making regular investments into a share-linked investment vehicle); yet others may buy them because the price has gone down, and they see good buying opportunities (if they have cash available – one can never assume that an investor can readily borrow); these form the counterbalance. But a market may include different proportions of these attitudes at different times, and individuals may change their views, perhaps slowly, as events unfold. I have not seen, nor developed, a satisfactory model to describe such a variegated market.

5 Conclusion

This paper leads to no conclusions, but I hope opens the way forward to others to consider some of the points expressed in it, and to seek better models that can consistently fit facts and theory, without some of the awkwardnesses that I have pointed out. But I start where I began, that the Brownian motion model is extremely useful for option pricing, whereas more realistic models quickly run into mathematical difficulties. Of course since the real world does not behave exactly like the option pricing model, hedging in accordance with the theoretical proportions can at best be an approximation, even if a good one, and one then has to investigate the possible hedging errors. But that is another story.