

Models in Finance - Class 6

Master in Actuarial Science

João Guerra

ISEG

Stochastic models of security prices

- Markets are efficient? Whatever the answer, experience shows that it is difficult to consistently outperform the market.
- It is often considered a good practice (in actuarial contexts) to construct models consistent with market efficiency.
- Lognormal model of security prices: the log of price follows a continuous-time random walk with drift, i.e. a Brownian motion with drift. The share prices follow a geometric Brownian motion.

Stochastic models of security prices

- A geometric Brownian motion is consistent with the efficiency of markets because of the property of independent increments of Bm.
- If the markets are inefficient, the lognormal model is inappropriate and we must consider other models, such as autoregressive models (one actuarial investment model of this type is the Wilkie model).
- In autoregressive model, we assume that future security price movements revert to a long-term mean (hence, are not entirely independent of past values): these models are also called mean-reverting or error correction models.

The continuous-time lognormal model

- Continuous-time lognormal model = Geometric Brownian motion.
Log-returns:

$$\log(S_u) - \log(S_t) \sim N[\mu(u - t), \sigma^2(u - t)],$$

where μ is the drift and σ is the volatility (or diffusion coefficient constant). Note that this drift μ is the drift in the log price: it is not the rate of drift of the price itself, which is $\mu + \frac{1}{2}\sigma^2$. This means that we have:

$$\begin{aligned}d(\log(S_t)) &= \mu dt + \sigma dB_t, \\dS_t &= \left(\mu + \frac{1}{2}\sigma^2\right) S_t dt + \sigma S_t dB_t.\end{aligned}$$

- The parameters μ and σ are specific to the investment considered.
- The proportional change in the price is lognormally distributed (why?) and the return on an interval $[t, u]$ do not depend on the initial value S_t .

The continuous-time lognormal model

- Equivalent way of specifying the model (considering a discrete time)

$$\log(S_u) = \log(S_t) + \mu(u - t) + \sigma \varepsilon_u \sqrt{u - t},$$

where ε_u is a series of i.i.d. standard normal r.v. (called "innovations of the stochastic process").

- Properties of the lognormal model
 - ① Mean and variance of log-returns are proportional to the length of the interval $(u - t)$ and so, volatility (standard-deviation of log-returns) is proportional to $\sqrt{u - t}$. So, if $\mu > 0$ and $\sigma \neq 0$, the expected log-return and the volatility will increase with time and tend to $+\infty$ when time $u \rightarrow +\infty$. However, the mean and volatility of log-returns over the same time period (the annual log-return for instance) remain constant (why?)
 - ② Returns over non-overlapping intervals are independent.

- 3. Value of the investment at time u :

$$S_u = S_t \exp(X_{u-t}),$$
$$X_{u-t} \sim N\left[\mu(u-t), \sigma^2(u-t)\right].$$

4. $E[S_u] = S_t \exp\left(\mu(u-t) + \frac{1}{2}\sigma^2(u-t)\right)$.
 $\text{Var}[S_u] = S_t^2 \exp\left(2\mu(u-t) + \sigma^2(u-t)\right) [\exp(\sigma^2(u-t)) - 1]$.

- Note: if X is lognormal with parameters μ and σ then

$$E[X] = \exp\left(\mu + \frac{\sigma^2}{2}\right) \text{ and } \text{Var}[X] = \exp(2\mu + \sigma^2) [\exp(\sigma^2) - 1].$$

Empirical tests of the log random walk model of security prices

- Independent returns over disjoint time intervals \implies the past history of the prices has no influence on the future returns \implies weak form of market efficiency.
- This is consistent with empirical observations that technical analysis does not lead to excess performance.
- A model weakness is that:
- Estimates of volatility σ vary widely according to what period is considered and how frequently the data samples are taken (for instance, σ increases in recessions or with financial crisis).

Empirical tests

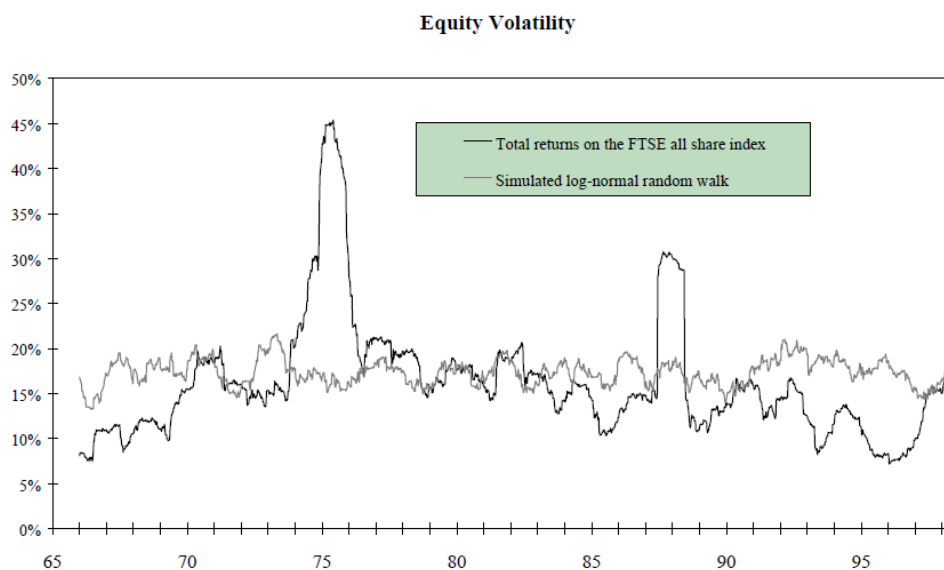


Figura:

Rolling annual volatility of total returns on FTSE All-Share index versus the rolling volatility of a simulated lognormal rand. walk with same mean and stand.

deviation

Empirical tests

- The model fails to capture the differing volatilities observed and the simulated log random walk is much less volatile than the actual returns graph, which has volatility peaks around 1974 and 1987.
- Evidence from option prices: the Black-Scholes formula gives option prices in terms of anticipated values of volatility over the term of the option. From the market observed option prices, one can deduce (working backwards) the "implied volatility": the value of σ consistent with observed option prices.
- Prices of options tell us implicitly what the market "believes" the volatility of the security price will be.

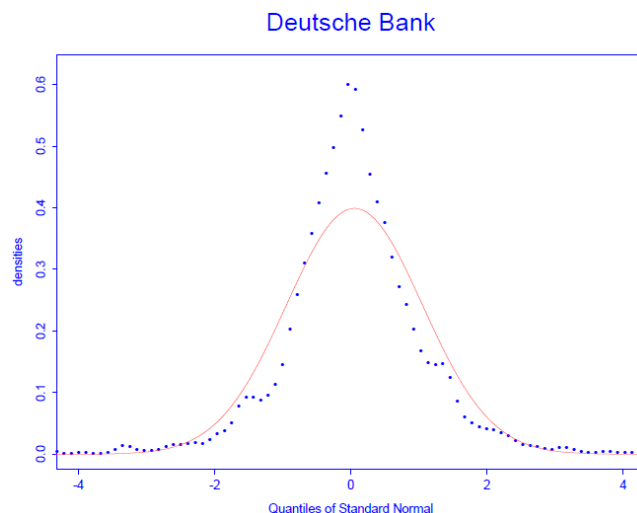
- Analysis of historical option prices suggests that volatility expectations fluctuate widely over time.
- Conclusion: empirical evidence \implies assumption of constant volatility is inappropriate.
- In order to model the fluctuations of volatility, one can use a process for the volatility. One class of models with this feature is the ARCH class: autoregressive conditional heteroscedascity models.

Empirical tests

- What about the drift parameter μ ? Is it constant over time?
- Theoretical reason that could imply that μ is not constant: it is reasonable to suppose that investors require a risk premium on equities relative to bonds. Therefore, if interest rates are high, we might expect that μ should be high as well.
- Mean reversion: Are markets mean reverting? If they are mean reverting, after a fall, rises are more likely, and market falls are more likely following a rise.
- There appears to be some empirical evidence of mean reversion but this evidence depends heavily on the market evolution after a small number of dramatic crashes.
- On the other hand, there is also some empirical evidence of "momentum effects": a rise one day is more likely to be followed by another rise the next day.

Empirical tests

- Normality assumption: market crashes appear more often than one would expect from a normal distribution of the log-returns (the empirical distribution has "fat tails" when compared to the Normal). Moreover, days with very small changes also happen more often than the normal distribution suggests (more peaked distribution).



Empirical tests

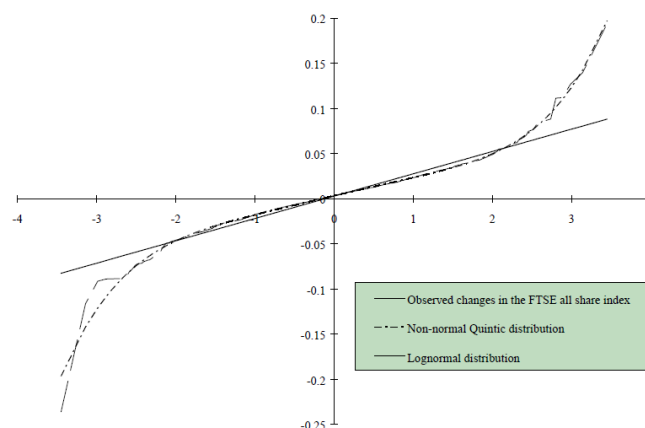
- The geometric Brownian motion model produces continuous price paths, but jumps and discontinuities are an important feature of real market prices.
- The fat tails and jumps justify the consideration of Lévy processes for modelling security prices.
- How to compare two distributions, if we are interested in their "tail" behavior? One way is using Q-Q (quantile-quantile) plots.

Method:

- 1 Rank the n observations of the actual log-returns in ascending order and denote them by $\log y_i$, $i = 1, \dots, n$.
- 2 Find the log returns, $\log x_i$ such that the normal distribution function has the value i/n , i.e. $\Phi(\log x_i) = \frac{i}{n}$.
- 3 Plot the pairs $(\log x_i, \log y_i)$.

Empirical tests

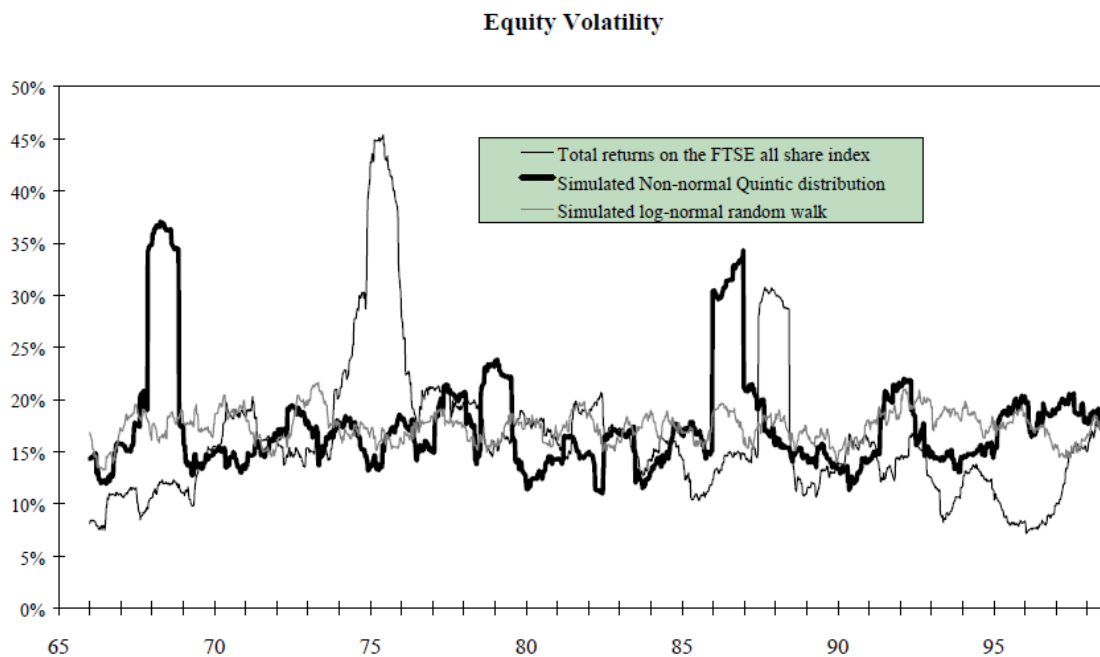
- The closer the distribution is to a lognormal, the straighter will be the resulting plot.



- Figure: Q-Q plot for the FTSE index: the actual returns have many more extreme events ("fat tails"). Fitted quintic polynomial distribution. x -axis is measured in stand. dev. from mean (of $\log x_i$) and the y -axis is simply $\log(y_i)$.

Empirical tests

- What about if we consider the volatility of a simulation based on a non-normal distribution (like the quintic polynomial)?



Empirical tests

- Despite the non-normal quintic distribution having a constant volatility, this process gives rise to volatility which has the same characteristics as the observed volatility from the equity market.
- Modelling capital markets requires the use of distributions which more accurately reflect the returns observed. These distributions provide an improved description of the varying volatility without requiring volatility to be modelled as a stochastic process.
- One measure of these non-normal features is the Hausdorff fractal dimension of the price process.
- Brownian motion has a fractal dimension of $\frac{3}{2}$.
- Empirical investigations of market returns often reveal a fractal dimension around 1.4.

Market efficiency

- Many of the empirical deviations from the random walk do not imply market inefficiency.
- Example: periods of high and low volatility could arise as a consequence of new information arriving in large measure or in small measure. Market jumps are consistent with the idea of arrival of information in discrete packets rather than continuously.
- The hypothesis of market efficiency can be difficult to disprove.
- Many statistical tests are based on assumptions of normality. If we reject normality, the other hypothesis should be re-tested and, in particular, the evidence for time varying mean and volatility is greatly weakened: these apparent effects could result from the properties of an underlying fractal process.