# MEAN REVERSION

Jon Exley
Shyam Mehta
Andrew Smith

Presented to:
Faculty & Institute of Actuaries
Finance and Investment Conference
Brussels, June 2004

**DRAFT 21st June 2004**

ABSTRACT

Mean reversion exists in many different forms within investment markets, none of these forms is necessarily inconsistent with efficient markets. However, there is a lack of precision in what many investment practitioners mean by the term “mean reversion”. In this paper we review the various forms of this phenomena and both the evidence (if any) and implications (if any) of their existence. In doing so we propose a formal mathematical definition of what most investment practitioners seem to mean by “mean reversion”, based on the correlation of returns between disjoint intervals. We then look at some of the mathematical properties of processes that mean revert (or avert) under our definition and propose a mean reverting (or averting) model that is consistent with a simple form of market efficiency. Finally we review the actuarial implications of adopting mean revering models and highlight some important, and possibly surprising, considerations in using them in a world of market consistent valuations.

## 1. INTRODUCTION

“Is it really worse to lose £10,000 in the stock market than to lose £10,000 from a burglary? Personally I would rather lose the money in the stock market. At least it is your own fault if you lose money in the stock market, and you might make it back again"

BAJ 9 (2003) part III, p607

1.1 Much of financial theory is based on random walk models of asset prices, returns and yields. However, over the last 20 years or so the theory has been extended to look much more closely at departures from a random walk. This theory has been used for many purposes, including investment strategy, capital adequacy and the pricing and hedging of options.

1.2 The random walk is a special case of a wider class of models, which includes mean reverting or mean averting models. It has been suggested that some classes of mean reverting models reduce the capital supposedly required for many classes of insurance business, relative to results obtained from random walks, because of the degree of long-term investment risk typically retained by life offices. This has stimulated renewed actuarial interest in the mean reverting class of models.

1.3 As discussed also in our concluding remarks, an issue of cherry-picking arises: modelling mean reversion appears at first sight to allow offices to carry lower capital requirements whereas possible *causes* of apparent or actual mean reversion in asset returns, such as jumps in asset prices or mean reverting volatility are not usually modelled. These latter features would of course be associated with other attributes of returns, such as fat tails, which would probably require the setting up of *higher* capital requirement (depending on where the fat tail kicks in).

1.4 From the outset, it is important to stress that many forms of mean reversion do not imply market inefficiency, although in popular investment folklore mean reversion is regarded quite wrongly as proof of inefficiency. We will show that mean reversion may arise quite naturally in many different forms in a market where the distribution of future returns, or risk aversion, varies over time.

1.5 Our paper is set out as follows. We first consider in the first four sections what we mean by mean reversion. It may seem somewhat surprising that it takes so long to establish a definition, but it is worth noting that many investment professionals who profess to believe in mean reversion are often unable to provide a precise definition of exactly what they mean by this. As we shall see, there is no existing universal measure of "mean reversion" and the definition that we believe investment professionals often struggle towards (a form involving in fact simultaneously both mean reversion and aversion) is not the same as the standard definition of time series analysis (namely "stationarity").

1.6 Once we have defined mean reversion, and in the process considered some of the evidence for this, we go on to discuss the comparison of statistical models in terms of their mean reverting tendencies and look at some generic mathematical attributes of mean revering (or averting) asset return processes . Finally we propose a simple model that combines the intuitive investment practitioners' description of "mean reversion" with strong empirical and economic support but which could be described as either mean averting or mean reverting depending on some critical parameter settings and which could be viewed as having some quite surprising and disconcerting implications in a market consistent valuation framework. In our conclusions we discuss the implications of our results for actuarial work.

## 2. "Asset Prices Always Fall After Hitting A Maximum"

2.1 As discussed in the introduction, there are many possible definitions of mean reversion, which seems to mean different things to different people. The broadest definition is probably as follows:

**Definition 1**: *An asset model is mean reverting if asset prices tend to fall (rise) after hitting a maximum (minimum).*

The definition is popular because it is linked to a straightforward test. Look at historically extreme stock market highs, and establish whether the market subsequently fell.

2.2 Using this definition, many analysts can convince themselves that stock markets obviously mean revert. For example, (so the thinking goes), the stock market was clearly overvalued in the summer of 1987 and also in the later 1990s. This overvaluation explains the subsequent falls. Or, a common perception is that equity markets mean revert because it so happened that equities rose after the 1974 market low was reached.

2.3 The trouble with this definition is its breadth. It is a truism that markets fall after hitting a maximum – because a local maximum necessarily hits a higher value than those on nearby dates. Any process at all is mean reverting in this sense. So although the 1987 crash might satisfy this definition of mean reversion, this proof is of no use in refining the class of appropriate models. It is helpful therefore to move to a narrower definition, which will be harder to validate empirically, but also distinguishes better between classes of stochastic processes.

2.4 Before we move on, however, it is worth making the point that it can be perfectly consistent with an efficient and arbitrage free market to have some processes in an asset model where all market participants do indeed know that a high or low point has been reached.

2.5 To illustrate this, consider a trivial deterministic model. In a deterministic world stocks and cash must have identical returns and cash rates must deterministically follow the current path of forward interest rates. Here everyone can know that interest rates (bond prices) or stock returns have reached a high (low) or low (high) point without violating the trivial market efficient and arbitrage free conditions. If the current forward interest rate curve is smoothly upward or downward sloping then this could in turn be viewed as a form (possibly the simplest conceivable form) of mean reversion.

## 3. AUTOCORRELATION – NEGATIVE AND POSITIVE

3.1 Autocorrelation of returns is a well known attribute of certain discrete time asset models (described generically as VAR models, not to be confused with value at risk) that are commonly referred to as "mean reverting". This gives us our second, more mathematically precise, candidate definition of mean reversion:

**Definition 2**: *An asset model is mean reverting if returns are negatively auto-correlated.*

3.2 This is the sort of process discussed in Lee (1991):

"Under this model, which has wide intuitive appeal, a below average return in one period is likely to be followed by "compensatory" above average returns in subsequent periods. It has frequently been said for example that the fantastic returns achieved in the 1980s were really a catching up exercise to make up for the poor returns in the 1970s"

3.3 This description captures well the intuitive element of this form of mean reversion. The model is expressed in *discrete* time as follows:

$$R_t = a(R_{t-1} - \mu) + \mu + \sigma W_t \ldots (3.1)$$

where $R_t$ is the return in period $t$, $\mu$ is the mean return in a single period, $W_t$ is a standard normal variate and $a$ (<1) is the (negative) autocorrelation coefficient.

3.4 In order to assess the informal evidence for this form of mean reversion in equity markets, we looked at 100 years of equity return data in 16 countries (the decennial Dimson, Marsh & Staunton (2003) data set, split according to returns in each decade of the last century). Looking at autocorrelation over ten year periods to investigate the common intuition referred to by Lee, we found no evidence that poor (good) returns in one decade are followed by good (poor) returns in the next.

3.5 The returns in the period subsequent to a large equity market fall (and rise), for example, were broadly identical to the average return in the period as a whole (based on 16 worst and best return decades, one for each country). We also looked at the UK annual equity return data (Barclays 2003). We mined the data to find the future holding period most correlated with annual returns. This was 2 years, i.e. returns in a year are negatively correlated (-0.2) in the data set with returns in the subsequent two years. Taking the 102 years of data as a whole, we had 100 observations with 52 years exhibiting mean reversion and 48 mean aversion. The negative correlation arose primarily in respect of 6 or so specific observations.

3.6 There was evidence of mean reversion following negative return years (60% of these are followed by a period of positive return performance); in the case of positive return years 40% are followed by negative returns in the succeeding two years. Breaking this down, 8 of the 10 worst performing years were followed by positive returns over the succeeding two years and 5 of the 10 best performing years were followed by negative returns. Finally, we looked at 75 years of US equity return data. This exhibited a little less autocorrelation (-0.1 rather than -0.2) against subsequent two-year holding periods (other periods had lower correlation).

3.7 The empirical support for this intuitive form of mean reversion is therefore far from convincing. Furthermore, the difficulty of deciding over which period to look for this effect highlights a particular problem. Lee's definition refers to the 1980s "catching up" with the 1970s – so the returns in these decades were negatively auto-correlated if we looked at discrete decades– but presumably if we had just looked at annual periods we would have found that the returns in the 1980s were generally above average and all returns in the 1970s generally below average. This would seem to maybe suggest an element of *positive* autocorrelation if we change the time step. A hint of this potential for confusion in the intuitive explanation is indeed revealed in Lee's definition of positive autocorrelation (our emphasis in italics):

"Under this model, if past returns have been better (or worse) than average, then the return in the next period is likely to be a little worse/(better) (*i.e. there is a tendency to revert to the mean*), but still better (worse) than average."

Positive auto correlation is in discrete time the same process as above, but with $a$>1

3.8 So it seems that there is some uncertainty as to whether mean reversion is a positive or negative auto correlation phenomenon – or possibly both? This is quite an important issue to nail. A shown by Lee, the asymptotic volatility of the continuously compounded return is:

$$\left(\frac{\sigma\sqrt{t}}{(1-a)}\right)$$

Clearly the sign of $a$ can make a not inconsiderable difference, even for relatively small absolute values.

3.9 One way of reconciling this positive versus negative autocorrelation conundrum is to look at an unambiguous example of "mean reversion", namely pull to parity on a bond. The process whereby a bond price always converges on its nominal amount as it approaches maturity is, after all, a pretty clear example of a statistical process that always knows where it is heading, eventually.

3.10 "Pull to parity" reversion implies that *any* rise or fall in a bond price must result in a corresponding fall or rise in the returns over subsequent periods to maturity, as a consequence of the fact that the bond matures at £100, and applies whether or not interest rates mean revert (which we discuss further below). The form of reversion implies that the initial shock to the bond price will be negatively correlated with the subsequent returns over the period to maturity period. On the other hand, these subsequent returns will all tend to be above average as a consequence of this initial shock (positive autocorrelation). This pull to parity effect seems to be at the heart of the intuition in the above definition – returns are supposed to catch up after the initial shock and there is a mix of both positive and negative autocorrelation.

3.11 An interesting observation (of wider application) is that even this – apparently unambiguous – form of mean reversion does depend on our frame of reference. Instead of looking at the absolute bond return after the shock, we might prefer to look for positive or negative autocorrelation in the *excess* bond returns over and above the cash rate. In many ways, particularly when looking at equity returns, for example, this is a more sensible way to look at the question of mean reversion. However, since the initial shock comes from a rise or fall in interest rates, we might find that the higher bond returns subsequent to the shock are offset by the higher or lower cash rate – so the *excess* return in subsequent periods is unaffected by the shock. This analysis would then reveal no autocorrelation at all, positive or negative in the excess return.

3.12 In summary there seems to be some arbitrariness in the choice of time period of returns under the autocorrelation definition of mean reversion – we have a combination of positive and negative autocorrelation that has many analogies with the effects of "pull to parity" in a bond. We would prefer a definition of mean reversion that is invariant under choice of time interval. Observing that mean reversion involves some negative autocorrelation followed by some positive autocorrelation is not particularly helpful. We have also noted in passing the important question of the frame of reference - it is necessary to determine whether we are talking about excess returns (over cash) or total returns. In the case of classical pull to parity, mean reversion effects disappear if we work with excess return over cash, we shall see this again in our subsequent discussion of equity mean reversion.

## 4. STATIONARITY

*Definition*

4.1 The equation for the return process $R_t$ in equation (3.1) in the last section is in fact an example of a so called stationary process. If we wander outside the field of asset models into the world of econometric or other time series modelling then we would find stationarity as the unambiguous and widely understood definition of "mean reversion". Thus our difficulties in applying this definition can't be in the formula itself. Perhaps the problem is in our attempt to define mean reversion in terms of stationary *returns*. Lets therefore look at a wider definition.

**Definition 3**: *An asset model is mean reverting if interest rates (and volatilities), yields or growth rates are stationary*

4.2 A stationary process has the same distribution at every point in time, unconditional on the immediate past. Under suitable conditions, the sample distribution of observations over a very long time period will converge to the stationary distribution.

4.3 If an observation falls high up in the tail of the stationary distribution, it is likely that the following observation will be nearer to the long term average. This can give the appearance of a force driving observations over time towards a long term mean. This mean reverting force is countered by the influence of random noise which pushes the process away from its current value.

4.4 The simplest form of stationary process is a first order autoregressive process. Assuming normal distributions, equation (3.1) above can be generalised as follows (for $-1 < A < 1$)

$$X_0 \sim N\left(\mu, \frac{\sigma^2}{1-A^2}\right) \qquad \ldots(4.1)$$
$$X_{t+1} = A(X_t - \mu) + N(\mu, \sigma^2)$$

4.5 As mentioned above, stationary series have proved fruitful for analysing economic quantities such as interest rates (and volatilities), or dividend yields because at first sight it is plausible that these have a natural long term mean level. However, taking interest rates as an example, the underlying process is essentially unknowable, being determined for example, at the whim of Government intervention in the markets. For example, although the Government may have a current inflation rate target of 2% or 2.5%, the next Government might choose 5%. Similarly, there is no underlying economic theory which suggests that dividend yields should be stationary. This means that one needs to exercise a good deal of judgment in interpreting the results of a model which assumes that these processes are in fact stationary.

*Interest rates*

4.6 Evidence for mean reversion (albeit relative to the current yield curve rather than a fixed constant) in interest rates is not hard to come by (nor as we repeatedly stress does it violate any concept of arbitrage free or efficient markets). Given the implied exponential decay in the volatility of forward rates, as discussed below, the simplest measure of this phenomenon is the average level of decline in implied volatility of forward interest rates (as measured by the implied volatilities of caplets and floorlets, for example). When volatility is measured as the standard deviation of forward rates divided by the square root of time, one observes a decay pattern, particularly for long tenor (20 yr+) interest rates, ie. strong mean reversion. For example, at the current time the annual rate of decay in volatility is around 6% pa.

4.7 Some of this apparent mean reversion may be an artefact arising because nominal interest rates cannot decline to below zero, and therefore the pattern of future volatility is bounded from below by a positive interest constraint. When volatility is measured as the standard deviation of the log of forward rates divided by the square root of time, one observes only a modest decay pattern, i.e. low mean reversion. For example, at the current time the annual rate of decay in volatility measured in this way is approximately 2% pa. This is another important observation with wider implications – strong evidence for mean reversion on the basis of one distributional *assumption* for a model may be weak evidence under another assumption. A test for stationarity may turn out to be joint test for stationarity *and* a particular distributional assumption.

4.8 The classical example of a mean reverting interest rate model is the Hull & White (1990) extended Vasicek model – and we will use this model to illustrate some important aspects of stationarity. Strictly the Hull & White model replaces the constant mean in the above definition with a time dependent $\mu_t$ but otherwise this is an example of a stationary process). The time dependent mean ensures that the drift of the short rate follows the slope of the initial forward rate curve and when the short rate deviates from the initial curve, it is pulled back to it at rate $a$.

4.9 As suggested above, it can be shown that this model implies that the innovations to the short rate must also perturb the forward rate curve with perturbations that decay exponentially at rate $a$, reflecting the anticipated reversion of the short rate back to the mean. So in this model exponentially decaying forward rate volatilities and mean reversion in the short rate are equivalent.

4.10 When expressed as exponentially decaying forward rate volatility, mean reversion in interest rates does not of course sound at all inconsistent with efficient or arbitrage free markets. However, at first sight mean reverting short rates might appear to give rise to free lunches to those who know where rates are heading. This illusion is easily dispelled though. Knowledge that interest rates will mean revert to the original forward rate curve is of course factored into this decaying response of longer bonds to changes in interest rates - prices of longer dated bonds already reflect the expected mean reversion of the short rate – which is why the volatility of forward rates decays exponentially.

4.11 At risk of labouring this simple point, the important issue to highlight is therefore that aside from simple pull to parity "mean reversion" (discussed previously), under this model of stationary interest rates, bond *prices* do not display any mean reversion – *mean reverting interest rates do not lead to mean reverting bond prices*.

*Volatilities*

4.12 It is appropriate to mention here also the possibility that *volatility* (of interest rates, stock prices, commodity prices etc) can also display similar mean reversion or stationarity again without introducing mean reversion in the corresponding asset price (i.e. an option on a bond, stock or commodity). This is of course another leg up in terms of sophistication in our hierarchy of modelling development, but the analogy with interest rates is direct to the extent that existing option prices can already embed a mean revering volatility term structure assumption – to the extent that volatility follows this anticipated process, option prices will not themselves display mean reversion. Option traders cannot necessarily earn free lunches from the knowledge that volatilities mean revert.

4.13 The evidence for mean reversion in equity market return volatility is also very strong. UBS kindly supplied us with data for 3-month and 5-year implied volatilities on the FTSE 100 (approximately the last 10 years) and DJ Euro-Stoxx Indices (approximately the last 5 years). The annualised rate of mean reversion for these series measured against a model where log volatility reverts to a constant level was 76%pa, 39%pa, 93%pa and 88%pa for the four series. Clearly, these specific numbers just relate to the average over the observation period and may not necessarily relate to current levels of mean reversion.

4.14 Deterministic mean reversion of volatility – directly analogous to an upward or downward term structure of interest rates is feasible (and indeed often a convenient assumption for pricing options), although it implies time inhomogeneity. In simple terms the world would be becoming a more or less volatile place over time in a deterministic manner. In other words we are saying that by time we reach 2010 the world will be a more or less volatile place than we see today.

4.15 A more realistic model would perhaps suggest a mean reverting stochastic process for volatility, such as the GARCH model (which incidentally also introduces fat tailed distributional properties). The GARCH model again basically follows the form of our second definition, with the asset variance being the random variable ($X_t$). We will return to this model later.

4.16 So we have two good examples of stationarity in asset models of the form described above where the evidence for its existence in data is fairly unambiguous and the implications for the efficient market hypothesis of this form of "mean reversion" are basically nil – bonds or options (or both) can already reflect either deterministic or stochastic mean reversion in interest rates or volatility so that we do not observe any associated mean reversion in *prices* – asset prices can continue to describe a random walk despite the stationarity of these associated processes. Market efficiency is not violated either in the physical or the options market.

*Equity Dividend Yields*

4.17 Another commonly cited example of mean reversion in the actuarial literature is of course the equity dividend yield. Over the 100 years to 2001, analysis of the data suggests, on the face of it, a rate of equity dividend yield mean reversion of around 23% pa. The data is similar to the evidence for equity return mean reversion discussed in the previous section: of 100 observations, 41 exhibited mean aversion and 59 mean reversion. Most of the mean reversion was associated with reversion following an extreme market decline, also as discussed above.

4.18 Tests for equity yield mean reversion are however prone to a number of biases and the association between the apparent mean reversion and extreme market decline is an immediate cause for concern - tests for mean reversion often assume normality of returns and yet events (such as 1974) that have a very low probability in this distribution (and suggest non normality) often make a large contribution to the quoted confidence of the test. More importantly, however, the evidence for mean reversion in equity dividend yields has been substantially over stated in the past, due to biases in the regressions used, as discussed in Exley, Smith & Wright (2002). After correcting for these biases it can be argued that the evidence for stationarity of equity dividend yields is far weaker than the evidence for stationarity in interest rates and volatilities.

4.19 However, even if equity dividend yields *were* stationary we can use some of our insights gained so far to see why equity price movements could still follow a time invariant random walk or at least the behaviour could be perfectly consistent with efficient markets (the risk premium may simply vary consistently over time).

4.20 A good place to start is to look at the analogy between an equity and an inflation linked bond (we can consider an equity to behave like an index linked bond with the retail prices index replaced by a dividend index). If the real yield on this bond rises (by analogy, the dividend yield rises) the price of the asset falls. There are two reasons why the yield could rise (in a situation where the nominal interest rate remains constant) either the risk premium on the asset has risen unexpectedly or expected growth of the income (inflation or dividends) has unexpectedly fallen. We are careful to add the words "unexpected" here as we have already explained (at some length) expected rises or falls can be already discounted in the term structure of interest rates.

4.21 The easy one to start with is of course the case where an unexpected rise in equity yields is associated with an unexpected fall in dividend growth expectations. If these expectations are in the form of a decaying perturbation to some trend growth (i.e. dividend growth expectations are also stationary) – then it would be perfectly natural to see mean reverting dividend yields but the subsequent equity returns would be the same as they were before the unexpected news arrived – i.e a random walk with unchanged mean. The initial price fall that triggered the rise in yields would simply ensure that the future price appreciation compensated for the lower (although mean reverting back to trend) dividend growth. We would not observe any predictability in the subsequent asset price changes – *a high dividend yield would not predict higher returns*.

4.22 That was the simple case. Now let's consider the more complicated example of an unexpected rise in the equity risk premium. The effect of this unexpected change would be a fall in the price of equities, followed by a higher risk premium in subsequent periods. This is really a form of "pull to parity" effect. As in our simple conventional interest rate example subsequent returns would be viewed as neither negatively nor positively autocorrelated if they were measured relative to the new risk premium.

4.23 Of course a mean reverting risk premium is not necessarily inconsistent with efficient markets if either risk or risk aversion are also mean reverting. We will return to this later. The fact that the apparent mean reversion effect disappears when we consider returns in excess of the market price of risk will also turn out to be a very important observation when we consider the risk neutral version – and hence option valuation implications – of this form of mean reversion.

*Stationary Drift in Asset Models*

4.24 So although we remain extremely sceptical as to whether equity dividend yields are stationary, we have in the process of discussing explanations for this phenomenon arrived at the possibility of a stationary process for the equity risk premium that could generate the sort of intuitive but not yet properly defined "mean reverting" behaviour in equity prices described in our second definition – call it pull to par or call it a "catching up period" after a fall.

4.25 The difference between this process and our autocorrelation process is subtle, but the key is that we have introduced a second process determining the risk premium that is distinct from the process for equity returns. The subtle comparison is thus:

**Positive autocorrelation**:

$$R_t = a(R_{t-1} - \mu) + \mu + \sigma W_t \ldots (3.1) \text{ (again)}$$

**A stationary equity risk premium** ($X_t$) might also look something like:

$$X_t = aX_{t-1} + \sigma_X W_{Xt} \ldots (4.2a)$$

$$R_t = X_t + \mu \ + \sigma_R W_{Rt} \quad ..(4.2b)$$

4.26 In the second case the *return on the asset* $R_t$ is a function of the risk premium (defined as $X_t + \mu$) *and* the effect of shocks to the asset price itself, which may be positively (or more likely negatively, under the "pull to parity" analogy) correlated with shocks to the risk premium. Rather than assuming that the *total return* is positively autocorrelated or stationary (as $aR_{t-1}$ appears on the R.H.S of equation (3.1)), we now have a stationary *drift* process.

4.27 This twin process approach gets us very close to a workable definition of mean reversion is assets such as equities.

## 5. Increments over Disjoint Intervals Negatively Correlated

5.1 Many analysts think of share prices in two portions – a true underlying value, plus or minus a short term fluctuation (associated with the fluctuating equity risk premium above). This suggests that the market value fluctuates around a fundamental or true value, but that the true value is also subject to random movements.

5.2 One way of approaching this is to describe the mean revering process by $X_t+W_t$ . This gives us a model similar to above of the form:

$$\begin{pmatrix} X_0 \\ W_0 \end{pmatrix} \sim N\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \frac{v_{XX}}{1-A^2}\begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}\right]$$
$$\begin{pmatrix} X_{t+1} \\ W_{t+1} \end{pmatrix} \sim N\left[\begin{pmatrix} AX_t \\ W_t+\mu \end{pmatrix}, \begin{pmatrix} v_{XX} & v_{XW} \\ v_{XW} & v_{WW} \end{pmatrix}\right] \quad \ldots.(5.1)$$

Note we have rigged the starting point so that $X_0+W_0 = 0$.

5.3 This has the same structure as the form of mean reversion that we proposed at the end of the previous section and is close to a model that could plausibly describe share prices, or at least log share prices.

5.4 However, we still have a major challenge to decide whether the actual process is, or is not, of this form. This test is particularly difficult as we do not in practice observe $X$ and $W$ separately, but only the log share price, that is $X_t$ and $W_t$. We could conceivably test stationarity given observations of $X$, but the need to unscramble $X$ from the random walk $W$ adds a whole new layer of complexity. This would of course be easy if we could identify $X$ (the "equity risk premium") with an observable such as the equity dividend yield.

5.5 However, as discussed earlier, the old faithful actuarial indicator lets us down badly – we believe (see Exley, Smith & Wright 2002) that there is no conclusive evidence that equity dividend yields mean revert nor, even if they did, that this mean reversion can predict the equity risk premium ($X$ in the above formula).

5.6 This type of model also suffers from identifiability problems. In particular, if

$$v_{XX} + (1+A)v_{XW} = 0 \ldots(5.2)$$

then the process is indistinguishable from a random walk (for proof of this result see Appendix 1) . So once again, to establish that a process is a sum of a random walk and a stationary process, is not to eliminate the possibility that the original process might also be a random walk.

5.7 This is a surprising result – we have successfully generalised the intuitive association between mean reversion and autocorrelation but it turns out that it is still not sufficient to define mean reversion in terms of this two stage process. Negative correlation between the shocks to the asset price and shocks to the risk premium does not necessarily generate what most practitioners would regard as "mean reverting" behaviour. This leads us to our final definition.

**Definition 4:** *A process is mean reverting if increments over disjoint intervals are negatively correlated*.

5.8 We will see this also appropriately generalises the notion of a stationary process, and captures the intuitive notion that a fall is more likely after a rise (namely the concept of autocorrelation discussed in definition 2). However, this definition does truly exclude some possible processes that satisfy these notions of stationarity without being mean reverting in a practitioner's sense of the word. Therefore, we can separate models into mean reverting and not mean reverting, with each *a priori* being plausible.

5.9 For example, let us consider the processes previously discussed, and evaluate the covariance of increments. Let us take $r < s < t < u$. Then

$$\mathbf{Cov}\left(X_s + W_s - X_r - W_r, X_u + W_u - X_t - W_t\right)$$
$$= \frac{v_{XX} + (1+A)v_{XW}}{1-A^2}\left\{\begin{array}{l}(1-A^s)(1+A^{u-s}) \\ -(1-A^r)(1+A^{u-r}) \\ -(1-A^s)(1+A^{t-s}) \\ +(1-A^r)(1+A^{t-r})\end{array}\right\}$$

We can rearrange the curly brackets as follows:

$$\mathbf{Cov}\left(X_s + W_s - X_r - W_r, X_u + W_u - X_t - W_t\right)$$
$$= -\frac{v_{XX} + (1+A)v_{XW}}{1-A^2} A^{t-s}(1-A^{u-t})(1-A^{s-r})$$

We can see that if A > 0 then mean reversion is equivalent to vXX + (1+A)vXW > 0.

## 6. MEASUREMENT OF MEAN REVERSION

*Statistical Model Comparisons*

6.1 As observed in the Introduction to this paper, unlike means, standard deviations and correlations, there is no single accepted numerical measure of mean reversion. To say that a model has mean reversion of 5%, or of 500, is meaningless. Several comparisons exist of model statistics. See for example Smith (1996), Lee and Wilkie (2001) or PriceWaterhouseCoopers (2003). None of these comparisons extend to measures of mean reversion.

6.2 It does not help that mean reversion means different things to different people. In the parlance of econometrics a mean reverting process should be stationary, but as we have seen, this is not a useful definition given that the main focus of attention is likely to be the "mean reversion" or otherwise of stock prices. A broader usage of the term has therefore become common in the asset modelling community, but this usage conflicts with classical time series terminology. This confusion over terminology has hitherto complicated the creation of standard parameter definitions.

6.3 However, the development in the previous section provides us with a useful way forward in defining a standard diagnostic tool to establish the extent to which a particular model displays mean reversion. Such a diagnostic tool would clearly have value to regulators, end users and others seeking to understand key differences between various models.

*Short and long term volatility*

6.4 What then would be the generic parameters such a benchmark model should show? Let us suppose that the mean reversion acts like a time-varying drift as discussed above. Over very short time periods, the drift is difficult to observe, because it is a drift per unit time and you do not have many units of time. Or put another way, as you reduce the time interval the volatility of the process (typically of order $\sqrt{dt}$) increasingly dominates the drift (or order $dt$) and nuances of the drift process will be dwarfed by the crude volatility term.

6.5 So, sampled at frequent intervals over short time spans, a mean reverting process could look a lot like a random walk, with some volatility $\sigma_S$. This is our first generic parameter. As short-term volatility is easy to compute – and we have few constraints on data – the data may become rather granular if we look on a second by second basis but within reasonable limits we can get quite a good handle on short term volatility by using small time steps. We might expect different models to show some consensus on this parameter, at least for two alternative models calibrated to the same data set.

6.6 Let us now look at very long holding periods. A mean reverting model shows dependence of returns from one period to the next, but that influence has to decay if the two periods are separated by a large time interval. Over periods separated by a wide interval, we might expect something close to independence, because in a market efficient environment information that impacts prices on one day is not likely to be hugely correlated with information that happens many years later. Therefore, over very long time scales, most models should look something like a random walk. The volatility of this long term random walk will not necessarily be equal to the short term volatility $\sigma_S$ but instead would take some different value $\sigma_L$, which forms our second parameter.

6.7 As this second parameter describes volatilities over long holding periods, it is intrinsically difficult to measure and will depend heavily on calibration. Whereas we have (almost) unlimited access to non overlapping short time periods (e.g. hourly data) there are not many independent ten, twenty or thirty year time periods to work with. We can expect $\sigma_L$ to vary from one model to the next, according to the selected model-building methodology.

6.8 The difference between long and short-term volatilities already gives one way to describe mean reversion. Strong mean reversion is characterised by long-term volatility well below the short-term volatility. For a random walk, the long and short term volatilities are equal.

*Measuring Volatilities: History and Simulations*

6.9 There are two situations requiring volatility measurement. If we are *calibrating* a particular model we need to measure volatility in actual data. If we are *comparing* two or more different models (lets assume that they are black boxes) by looking at simulated data generated by the models.

6.10 The first situation relates to either historic volatility, based on one observed time series or implied volatility derived from derivative prices. Essentially the major problem in the case of taking measurements from historic data is sampling error. The observed history is only one possible outcome of thousands that could have occurred, and there is no guarantee the observed history is representative in any way. Lack of data is a major problem, resulting in subtle biases we will later consider.

6.11 The problems may be different, and generally less acute in the case of calibrating against derivative price data, but they still exist. If a model is calibrated against derivative prices then the choice of derivatives may be important – the model may accurately price the chosen calibration instrument but be hopeless at pricing other instruments. The issue of time homogeneity is also important – a time decaying volatility structure may be indistinguishable from the effects of mean reversion in a time inhomogeneous model. This distinction is unimportant for derivative pricing on day one, but it matters greatly if a model is marched forward in time.

6.12 The second situation we need to consider involves the estimation of volatility from simulation output (generated from a black box model, say). Here, lack of data is not a problem. If sampling error is significant one can simply run more simulations from the black box. Instead, the difficulty that arises is the proliferation of different ways of calculating volatilities.

6.13 For example, consider a 1-year volatility. We could measure this as the volatility of the return in year 1, or alternatively as the forward start volatility of return between t=20 and t=21. For most models, the forward start volatility will be higher. This is because for the forthcoming year, we already know current levels of interest rates and the returns in immediately preceding years. To some extent this may help us predict the next year's returns. Some of the volatility in the first year returns will be reduced by the fact that all simulations start from common initial conditions. The forward start volatility may be higher because each simulation starts from a different point at t=20.

6.14 In exceptional circumstances, however, the forward start volatility may be lower than the one-year volatility. This can happen in a model with stochastic volatility whose starting conditions specify a starting volatility well above the long term mean.

6.15 From a time series analysis point of view, we are most interested in unconditional volatilities that have minimal dependence on initial conditions. In other words, our ideal definition of a one-year volatility should be forward start volatility, starting some point in the long future. In practice, simulation data may well be supplied with a finite horizon, and so some extrapolation may be required to develop the limiting forward start volatility.

*Maximum predictability*

6.16 We have now defined and clarified notions of long and short-term volatility. The next important property is required to describe where the transition occurs between short-term behaviour and long-term behaviour. We capture this using the notion of maximum predictability.

6.17 We concluded section 5 with a useful and quite broad definition of mean reverting processes as one where returns in one period are negatively correlated with those in other periods. We have argued that this correlation is likely to be close to zero for very small periods, and also for very long periods. It follows that there is some finite period length, which has a largest absolute correlation with the return in some subsequent period. We call this the most predictive period.

6.18 It seems reasonable to assume that if one interval is most predictive of another, then these two intervals will be adjacent, because every intervening interval results in a loss of information. For similar reasons, a symmetry argument suggests that the optimally predictive interval pair will be the same length. This optimally predictive period term is an important attribute of the time series, as is the value of the optimum correlation achieved. This gives two more attributes of a mean reverting series, which we can define in generic terms.

6.19 These attributes are shown below for the Wilkie model

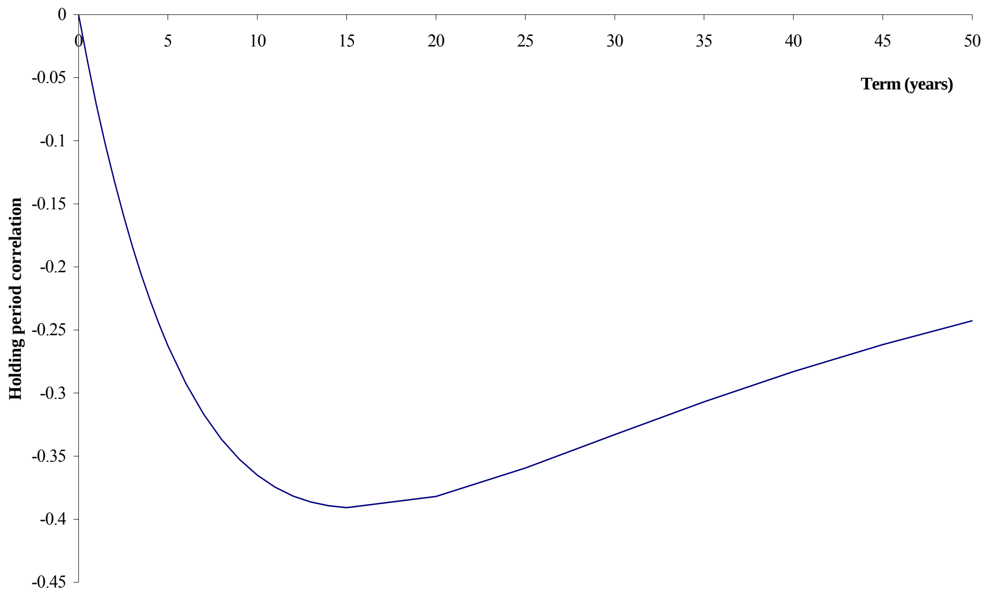


6.20 The optimal predictive term may be close to 1 year, or even less. For annually projected series it is therefore helpful to extrapolate the observed volatilities to shorter time frames, so that a non-integer optimal predictive length can be estimated.

6.21 Thus far, we have defined four numerical attributes of a mean reverting process:

- Short-term volatility
- Long-term volatility
- Optimal predictive term
- Optimal predictive correlation

These are mutually constrained to some extent. For example, if the long term volatility is close to the short term volatility, then the process is much like a random walk, and so we would be surprised to find a large optimal predictive correlation. We are not here suggesting that optimal predictive term and correlation are perfect measures. Correlation and its structure varies over time. One needs to be pragmatic and analyse a data set to see whether correlation is so different from 0 that the concept of optimal predictive correlation adds to the analysis.

*Prediction based on Return History*

6.22 We have examined the distribution of return in one period given the return over a previous period. The observed correlation gives us a possible predictor of future returns. Given a predictor of a future price, we can decompose the price variance as:

- The variance of the predictor (“$\sigma_{XX}$”, in our earlier model)
- Plus residual variance (“$\sigma_{WW}$”)

This decomposition will vary according to the information basis of the predictor. The predictor could be based on:

- A ratio such as the dividend or earnings yield
- The historic return over a particular past period
- The entire history of historic returns

6.23 We have described the first of these using the optimal predictive correlation. The second form of predictor was discussed, in terms of past data, in section 2. Economically, it is also relevant to establish the extent to which future movements can be predicted given *all* past movements. This reconciles to the economic notion of weak-form market efficiency. Therefore, we have another potential statistic:

- Predictive correlation based on total history

As this is the result of a multivariate conditional expectation, the calculation is delicate, both from the point of view of inverting large matrices and also the tendency for sampling error to generate spurious evidence of predictability. The regression actually gives us an $r^2$ statistic from which we extract a square root to express as a correlation.

6.24 We could choose the time horizon over which to measure the predictive correlation; the most natural horizon for comparison purposes is the previously defined optimal predictive term.

*Information structures*

6.25 Time series properties are usually defined in terms of forward start volatilities. However, many actuarial investigations start from known initial conditions and therefore the initial start volatilities are more relevant. An information structure is a way of relating initial start volatilities to the forward start volatilities. Information structures are discussed in Appendix 2.

## 7. MEAN REVERSION AND RANDOM WALKS – MATHEMATICAL FORMULATION

*Random Walks*

7.1 We consider models of log asset prices. Let $S_t$ be a total return index of prices, and write

$$S_t = \exp(Z_t)$$

Our processes are defined for all real $t$, both in the past ($t < 0$) and the future ($t > 0$).

7.2 Under a classical random walk, the change in $Z_t$ over a time interval $h$ is normal, with mean $\mu h$ and variance $\sigma^2 h$. The parameter $\mu$ is called the *drift* and $\sigma$ the *volatility*.

*Difference Stationary Processes*

7.3 A random walk is an example of a *difference stationary* process, because the distribution of the difference $Z_{t+h} - Z_t$ depends on the holding period $h$ but not on the start of the period $t$. There are many other possible difference stationary processes besides the random walk, including mean-reverting processes.

7.4 In each of these definitions, we are concerned only with unconditional distributions, which know nothing of the history of the model. In a projection model, we are more usually concerned with conditional distributions. For example, one might start a projection model at time 0, and assume knowledge of interest rates, past returns and so on up to time 0. A stochastic model might then generate distributions conditional on what was known at time 0. The return distribution in year 1 might well be different from that of year 5, simply because at time 0, more is known about year 1.

7.5 However, this would not necessarily violate stationarity, as these distributions are conditional ones. To get unconditional distributions, we would need to start projecting from a long time back, say t = -1000 years.

*Volatility Structure*

7.6 For a general difference stationary process, we can define the *variance function*, that is the variance $v(h)$ of the log return over a holding period $h$. In symbols.

$$v(h) = \mathbf{Var}(Z_{t+h} - Z_t)$$

From this, we can define the volatility function (also known as volatility term structure):

$$\sigma(h) = \sqrt{\frac{v(h)}{h}} = \sqrt{\frac{\mathbf{Var}(Z_{t+h} - Z_t)}{h}}$$

For a random walk, the volatility function has a constant value σ.

7.7 We now consider a more general class of processes, where $\sigma(h)$ varies by $h$. We assume that the long and short limits exist and are given by $\sigma_L$, $\sigma_S$ respectively.

$$\sigma_S = \sigma(0) = \lim_{h \downarrow 0} \sqrt{\frac{v(h)}{h}} = \lim_{h \downarrow 0} \sqrt{\frac{\mathbf{Var}(Z_{t+h} - Z_t)}{h}}$$

$$\sigma_L = \sigma(\infty) = \lim_{h \uparrow \infty} \sqrt{\frac{v(h)}{h}} = \lim_{h \uparrow \infty} \sqrt{\frac{\mathbf{Var}(Z_{t+h} - Z_t)}{h}}$$

These describe processes which behave over short time scales like a random walk with volatility $\sigma_S$, and also like a random walk over long time scales, but with a different volatility $\sigma_L$.

7.8 There are several possible choices for volatility functions with the desired limits $\sigma_S$ and $\sigma_L$. For calibration purposes, it is convenient for $v(h)$ to be linear in both $\sigma_S^2$ and $\sigma_L^2$. We consider the following cases:

| Possible volatility term structures given $\sigma_S$, $\sigma_L$. | | |
|---|---|---|
| Case | volatility function $\sigma(h)$ | variance function $v(h)$ |
| I | $\sigma(h) = \sqrt{\dfrac{\sigma_S^2 + \alpha h \sigma_L^2}{1 + \alpha h}}$ | $v(h) = \dfrac{\sigma_S^2 h + \alpha \sigma_L^2 h^2}{1 + \alpha h}$ |
| II | $\sigma(h) = \sigma_L \sqrt{1 + \left( \dfrac{\sigma_S^2}{\sigma_L^2} - 1 \right) \dfrac{1 - e^{-\alpha h}}{\alpha h}}$ | $v(h) = h\sigma_L^2 + (\sigma_S^2 - \sigma_L^2) \dfrac{1 - e^{-\alpha h}}{\alpha}$ |
| III | $\sigma(h) = \sqrt{\sigma_L^2 + \dfrac{\sigma_S^2 - \sigma_L^2}{\alpha h} \left[ 1 - \left( \dfrac{k}{k + \alpha h} \right)^k \right]}$ | $v(h) = h\sigma_L^2 + \dfrac{\sigma_S^2 - \sigma_L^2}{\alpha} \left[ 1 - \left( \dfrac{k}{k + \alpha h} \right)^k \right]$ |

7.9 These are all continuous monotone volatility functions. In each case, the parameter $\alpha > 0$ determines the speed with which the volatility term structure moves from the short limit to the long limit, a larger α implying faster convergence. The value of $k$ determines long-range dependency, with large values of $k$ denoting low memory between returns on distant intervals. We can verify that these expressions do indeed have limiting volatility $\sigma_S$ for small $h$ and $\sigma_L$ for large $h$.

*Mean Reversion and Aversion*

7.10 We now investigate the concepts of mean reversion and mean aversion. A difference stationary process is *reverting* if changes are negatively autocorrelated, that is, if a price rise is more likely to be followed by a price fall. Conversely, if a rise is more likely to be followed by another rise then the process is *averting*.

7.11 We can measure this by examining the covariance of changes over two disjoint periods. In symbols, we look for the covariance $\mathbf{Cov}(Z_s - Z_r, Z_u - Z_t)$, for $r<s<t<u$. We then classify difference stationary processes as follows:

| | |
|---|---|
| Reverting | $\mathbf{Cov}(Z_s - Z_r, Z_u - Z_t) \le 0$ for all $r<s<t<u$. |
| Averting | $\mathbf{Cov}(Z_s - Z_r, Z_u - Z_t) \ge 0$ for all $r<s<t<u$. |
| Indeterminate | $\mathbf{Cov}(Z_s - Z_r, Z_u - Z_t) > 0$ for some $r$, $s$, $t$, $u$<br>but $\mathbf{Cov}(Z_s - Z_r, Z_u - Z_t) < 0$ for other $r$, $s$, $t$, $u$ |

It so happens that we can calculate the covariances given only the variance function v($h$). This is because of the following identities:

| | | |
|---|---|---|
| $v(t\text{-}s)$ | $\mathbf{Var}(Z_t - Z_s)$ | |
| $v(t\text{-}r)$ | $\mathbf{Var}(Z_t - Z_r)$ | $\mathbf{Var}(Z_s - Z_r) + \mathbf{Var}(Z_t - Z_s)$<br>$+ 2\mathbf{Cov}(Z_s - Z_r, Z_t - Z_s)$ |
| $v(u\text{-}s)$ | $\mathbf{Var}(Z_u - Z_s)$ | $\mathbf{Var}(Z_t - Z_s) + \mathbf{Var}(Z_u - Z_t)$<br>$+ 2\,\mathbf{Cov}(Z_t - Z_s, Z_u - Z_t)$ |
| $v(u\text{-}r)$ | $\mathbf{Var}(Z_u - Z_r)$ | $\mathbf{Var}(Z_s - Z_r) + \mathbf{Var}(Z_t - Z_s)$<br>$+ \mathbf{Var}(Z_u - Z_t)$<br>$+ 2\mathbf{Cov}(Z_s - Z_r, Z_t - Z_s)$<br>$+ 2\mathbf{Cov}(Z_s - Z_r, Z_u - Z_t)$<br>$+ 2\mathbf{Cov}(Z_t - Z_s, Z_u - Z_t)$ |
| $v(u\text{-}r)$ + v($t\text{-}s$)<br>- $v(t\text{-}r)$ - $v(u\text{-}s)$ | | $2\mathbf{Cov}(Z_s - Z_r, Z_u - Z_t)$ |

Concave and convex functions $f$ are defined as follows:

| | | |
|---|---|---|
| $f(x)$ convex | $f[\lambda x + (1\text{-}\lambda)y] \le \lambda f(x) + (1\text{-}\lambda)f(y)$ | all $x$, $y$ and $0 \le \lambda \le 1$ |
| $f(x)$ concave | $f[\lambda x + (1\text{-}\lambda)y] \ge \lambda f(x) + (1\text{-}\lambda)f(y)$ | |

We now claim that:

An integrated process is *averting* if and only if the function $v(h)$ is convex
An integrated process is *reverting* if and only if the function $v(h)$ is concave

7.12 We demonstrate the mean reverting case; the mean averting case is similar. Firstly, let us demonstrate that a concave $v$ implies a mean reverting $Z$. The convex condition implies that

$$v(t-r) \ge \frac{u-t}{u-t+s-r} v(t-s) + \frac{s-r}{u-t+s-r} v(u-r)$$

$$v(u-s) \ge \frac{s-r}{u-t+s-r} v(t-s) + \frac{u-t}{u-t+s-r} v(u-r)$$

Adding these two equations gives the negative covariance as claimed.

7.13 The converse is harder. By taking $s = t = (u+r)/2$ we can demonstrate the concave condition for $\lambda = ½$. Other values follow by a more complicated continuity argument.

## 8. AN EXAMPLE OF A MEAN REVERTING MODEL

*Background and Definition of Model*

8.1 In Exley, Mehta & Smth (1996) we set out a mean reverting equity model that was consistent with efficient markets, based on a model for risk aversion that resulted in risk premia rising sharply after markets have fallen. In this section we propose what is arguably an even simpler model, to the extent that it is based on a simple linear market price of risk, similar to CAPM.

8.2 Our approach follows Duan (1996) by defining the following GARCH process for stock prices:

$$Ln\left(\frac{S_{t+1}}{S_t}\right) = r + \lambda\sigma_{t+1} - \frac{1}{2}\sigma_{t+1}^2 + \sigma_{t+1}Z_{t+1}$$
$$\sigma_{t+1}^2 = \beta_0 + \beta_1\sigma_t^2 + \beta_2\sigma_t^2(Z_{t+1} - \theta)^2$$

where $Z_t$ is a standard normal variate in the real world measure and $\beta_0 > 0$, $\beta_1, \beta_2 \geq 0$. Here $\lambda$ is of course the standard "market price of risk" so this model can be seen to be entirely consistent with a very simple efficient market model.

*Continuous Time*

8.3 It can be shown, by expanding the quadratic term in the second equation, noting that $Cov(Z_t, Z_t^2) = 0$, and applying the continuous limit results of Foster & Nelson (1994), that this process reduces to the continuous time system:

$$dLnS_t = (r + \lambda\sigma_t - \frac{1}{2}\sigma_t^2)dt + \sigma_t dZ_{1t}$$
$$d\sigma_t^2 = (\beta_0 - \theta\sigma_t^2)dt - 2\beta_2\gamma\sigma_t^2 dZ_{1t} + \beta_2\sigma_t^2 dZ_{2t}$$

where $\theta = 1 - \beta_1 - \beta_2(1 + \gamma^2)$ and $dZ_{1t}, dZ_{2t}$ are *independent* standard Brownian motions.

8.4 This GARCH process therefore turns out, in the continuous time limit, to be remarkably similar to the standard process introduced in section 4. In particular note that

$$Cov(d\sigma_t^2, d\ln(S_t)) = -2\beta_2\gamma\sigma_t^{3/2}$$

However, the processes are not identical. Importantly, the "risk premium" is $\lambda\sigma_t = \lambda\sqrt{\sigma_t^2}$ and not $\lambda\sigma_t^2$. Nevertheless, although this complicates the distributional properties, it does not affect the basic similarity in the structure. In particular, note that if the innovation to the process is positive, so $\sigma_t^2$ increases, then so does its square root $\sigma_t$ and if the innovation is negative then both $\sigma_t^2$ and its square root $\sigma_t$ decrease. We are however limited to positive is premiums in this model, which is actually an enhancement to the earlier process.

8.5 There is quite strong empirical evidence to support this form of model for stock prices. For example, Duan (1996) successfully fits the process to actual stock price data and then uses the parameters in a risk neutral version of the model (see below) to give a reasonable (but not exact) fit to option prices.

*Implications from Earlier Results*

8.6 However, the similarity of the model with our earlier work gives us the immediate insight that despite the apparent "pull to parity" effect of the negative covariance between the risk premium and the stock price process, it does not follow that the model necessarily describes a mean reverting process as we defined in Section 5.

8.7 Thus we have a model with some empirical support that for some parameter settings will be mean reverting, and for others it will actually be mean averting. This is particularly relevant in actuarial applications that tend to push the limits of models in terms of asymptotic behaviour (actuaries may be interested in the volatility of equities thirty or even fifty years hence). Quite small parameter changes can often flip the asymptotic behaviour of these models from the well behaved to the pathological.

8.8 In many ways this encapsulates our views on mean reversion. As with our 1996 paper, we stress that it is perfectly possible to construct a model that is consistent with efficient markets and displays mean reversion (or mean aversion). The problem is not with the principle of building such models but with the strength of the evidence actually in favour of mean reversion (rather than mean aversion or random walks) as an intrinsic, natural feature of equity market returns. Mean reversion does not appear to be "hard coded" into the structure of plausible market models.

*Market Consistent Applications – Costs of Guarantees*

8.9 However, even if equity markets did display strong mean reversion under this model, there may be a rather awkward implication associated with adopting this model in a market consistent valuation framework. Not only does the model *not* display mean reversion in the risk neutral world, but the model tends to suggest "fat tailed" distributions of equity returns that will typically increase the costs of guarantees measured on a market consistent basis. This is because the risk neutral version of the model is as follows;

$$dLnS_t = (r - \frac{1}{2}\sigma_t^2)dt + \sigma_t dZ_{1t}$$
$$d\sigma_t^2 = (\beta_0 - \theta^* \sigma_t^2)dt - 2\beta_2\gamma^*\sigma_t^2 dZ_{1t} + \beta_2\sigma_t^2 dZ_{2t}$$

where $\gamma^* = \gamma + \lambda$ and $\theta = 1 - \beta_1 - \beta_2(1 + \gamma^{*2})$.

8.10 The negative correlation between the risk premium and the stock process has disappeared in the risk neutral world – and all we are left with is the possibility of fat tails which will probably actually increase the cost and value of guarantees on a market basis. This is another important general observation – a process that is mean reverting in the real world may not mean revert in the risk neutral world and vice versa.

## 9. CONCLUSIONS AND IMPLICATIONS

*Conclusions*

9.1 There is strong evidence for mean reversion in interest rates and volatilities but rather weak evidence for mean reversion (and possibly mean aversion) in equity markets. The supposed mean reversion in equity dividend yields is in our opinion very likely to be a red herring, created by various biases in the regressions used.

9.2 The existence of various types of mean reverting features in asset yields or returns does not in any way contradict the assumption that markets are efficient. We discussed this in our earlier paper on market efficiency with reference to varying risk aversion. In this paper we have discussed various other forms of mean reversion, and proposed a mean reverting (or possibly mean averting) equity model that is again perfectly consistent with a very simple "market price of risk" approach.

*Pension Funds*

9.3 In the case of pension fund applications of these models, the issue of mean reversion is of course largely irrelevant in any event. This is because standard practice involves steering Trustees towards an asset allocation that is strongly conditioned on a consensus allocation to equities (60% or 70% for example). Changing a model to introduce mean reversion or aversion in this environment would be likely to be offset by focusing on more or less risk averse objectives – who is to say for example whether Trustees should focus on a 10% probability tail or the 0.1% tail.

9.4 However, even in pensions it is relevant in as much as there is an obvious need to show the amount of risk being borne and using a mean reverting model will often appear to show less risk.

*Insurance Companies*

9.5 Insurance companies are managed much more prudently than pension funds in the UK and need to set aside capital to meet risk, not just reserves to meet expected costs.

9.6 Typically a mean reverting model is used to suggest that an unhedged long equity position needs less capital than implied in a non-mean reverting model. This result is communicated to the regulator and so the difference can be as stark as the company being allowed to trade or not.

9.7 In the UK particularly, the 'statutory' valuation rate of interest and resilience tests used by with-profits life offices when establishing reserves depend in a complex way on the levels of dividend and earnings yields. In turn, unless the alternative 'realistic' valuation bites, the process used for projecting dividend and earnings yields make a large difference to any assessment of fair value of liabilities and of risk, because changes in these parameters can drive substantial shifts in investment policy.

9.8 One of the most important attributes of an investment model used in a UK with-profits context is therefore to properly project dividend and earnings yields. A constant dividend yield assumption (i.e. 100%pa mean reversion) will likely have very different results from using a model with say 5%pa mean reversion in dividend yields towards a market forward dividend yield. Our conclusion that there is something distinctly fishy about the alleged mean reversion in dividend (and similarly earnings) yields is therefore particularly relevant here. However, this is old ground.

9.9 Of more relevance today in the context of life insurance is the issue is the issue of cherry picking of models to minimise capital requirements trailed in our introduction to this paper. Our example model is particularly interesting in this regard, for two reasons:

(1) The basic structure of the model has strong empirical and theoretical support, but mean reversion is not hard-coded into the model structure. It might show mean reversion or aversion. Nailing one's colours to a model that displays mean reversion under current parameter settings could lead to unintended consequences if a year down the road a re-parameterisation suggests strong mean aversion.

(2) Mean reversion (or aversion) in the real world does not translate into mean reversion (or aversion) in the risk neutral world. Thus there is not compression (or expansion) of long term volatility in the pricing of guarantees relative to a random walk model arsing from this particular effect. However , the *quid pro quo* for introducing the possibility of mean reversion (aversion) into a plausible economic model is the possible introduction of fat tailed returns and these will typically increase the cost of guarantees, often quite substantially, if the same capital adequacy model is used for market consistent valuation.

So you pay your money for a model and you make your choice. It would appear however, that if regulators insisted on economic consistency in mean reverting models – in the form of a plausible economic rationale for mean reversion, then such models may become somewhat less popular.

## REFERENCES

**Barclays (2003) "**Equity Gilt Study 2003" Published by Barclays Capital equitygiltstudy.com

**Dimson, E, Marsh, P & Staunton, M** "Global Investment Returns Yearbook 2003" Published by ABN AMRO ISBN 0-9537906-3-0

**Duan, J.C** (1996) "Cracking the Smile" RISK **9** No.12 (December)

**Exley, C.J, Smith, A.D & Wright, T** "Mean Reversion and Market Predictability" available at http://www.gemstudy.com/investment.htm

Exley, C.J, Mehta, S.J.B & Smith, A.D "Market Efficiency" available at http://www.gemstudy.com/investment.htm

**Foster, D. P. & Nelson D. B.** (1994) "Continuous Record Asymptotics for Rolling Sample Variance Estimators" Working Paper, The Wharton School.

**Lee, P. (1991)** "Just How Risky Are Equities Over the long Term?" Staple Inn Actuarial Society paper (24th September).

**Wilkie, A.D (1995) "**More on a Stochastic Investment Model for Actuarial Use" British Actuarial Journal **1** 777-945.

## APPENDIX 1 PROOF OF RESULT (EQUATION 5.2)

We are seeking the condition for the process described in equation (5.1) to be indistinguishable from a random walk.

Write

$$\begin{pmatrix} X_{t+1} \\ W_{t+1} \end{pmatrix} = \begin{pmatrix} AX_t \\ W_t + \mu \end{pmatrix} + Z_{t+1}$$

then, by induction,

$$\begin{pmatrix} X_t \\ W_t \end{pmatrix} = \begin{pmatrix} 0 \\ \mu t \end{pmatrix} + \begin{pmatrix} A^t & 0 \\ 0 & 1 \end{pmatrix}\begin{pmatrix} X_0 \\ W_0 \end{pmatrix} + \begin{pmatrix} A^{t-1} & 0 \\ 0 & 1 \end{pmatrix} Z_1 + \ldots + \begin{pmatrix} A^2 & 0 \\ 0 & 1 \end{pmatrix} Z_{t-2} + \begin{pmatrix} A & 0 \\ 0 & 1 \end{pmatrix} Z_{t-1} + Z_t$$

Now we can work out the covariance:

$$\begin{aligned}
\mathbf{Cov}\left[\begin{pmatrix} X_s \\ W_s \end{pmatrix}, \begin{pmatrix} X_t \\ W_t \end{pmatrix}\right] &= \frac{v_{XX}}{1-A^2}\begin{pmatrix} A^s & 0 \\ 0 & 1 \end{pmatrix}\begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}\begin{pmatrix} A^t & 0 \\ 0 & 1 \end{pmatrix} \\
&+ \begin{pmatrix} A^{s-1} & 0 \\ 0 & 1 \end{pmatrix}\begin{pmatrix} v_{XX} & v_{XW} \\ v_{XW} & v_{WW} \end{pmatrix}\begin{pmatrix} A^{t-1} & 0 \\ 0 & 1 \end{pmatrix} \\
&+ \begin{pmatrix} A^{s-2} & 0 \\ 0 & 1 \end{pmatrix}\begin{pmatrix} v_{XX} & v_{XW} \\ v_{XW} & v_{WW} \end{pmatrix}\begin{pmatrix} A^{t-2} & 0 \\ 0 & 1 \end{pmatrix} \\
&+ \ldots \\
&+ \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}\begin{pmatrix} v_{XX} & v_{XW} \\ v_{XW} & v_{WW} \end{pmatrix}\begin{pmatrix} A^{t-s} & 0 \\ 0 & 1 \end{pmatrix}
\end{aligned}$$

This is:

$$\begin{aligned}
\mathbf{Cov}\left[\begin{pmatrix} X_s \\ W_s \end{pmatrix}, \begin{pmatrix} X_t \\ W_t \end{pmatrix}\right] &= \frac{v_{XX}}{1-A^2}\begin{pmatrix} A^{s+t} & -A^s \\ -A^t & 1 \end{pmatrix} \\
&+ \begin{pmatrix} A^{s+t-2} v_{XX} & A^{s-1} v_{XW} \\ A^{t-1} v_{XW} & v_{WW} \end{pmatrix} \\
&+ \begin{pmatrix} A^{s+t-4} v_{XX} & A^{s-2} v_{XW} \\ A^{t-2} v_{XW} & v_{WW} \end{pmatrix} \\
&+ \ldots \\
&+ \begin{pmatrix} A^{t-s} v_{XX} & v_{XW} \\ A^{t-s} v_{XW} & v_{WW} \end{pmatrix}
\end{aligned}$$

We can sum the geometric progressions analytically:

$$\mathbf{Cov}\left[\begin{pmatrix} X_s \\ W_s \end{pmatrix}, \begin{pmatrix} X_t \\ W_t \end{pmatrix}\right] = \frac{v_{XX}}{1-A^2}\begin{pmatrix} A^{s+t} & -A^s \\ -A^t & 1 \end{pmatrix} + \begin{pmatrix} A^{t-s}\dfrac{1-A^{2s}}{1-A^2}v_{XX} & \dfrac{1-A^s}{1-A}v_{XW} \\ A^{t-s}\dfrac{1-A^s}{1-A}v_{XW} & sv_{WW} \end{pmatrix}$$

Putting these together, we find:

$$\mathbf{Cov}\left(X_s + W_s, X_t + W_t\right) = \frac{(1-A^s)(1+A^{t-s})}{1-A^2}\left\{v_{XX} + (1+A)v_{XW}\right\} + sv_{WW}$$

Hence the condition for the process to be indistinguishable from a random walk:

$$v_{XX} + (1+A)v_{XW} = 0$$

## APPENDIX 2: INFORMATION STRUCTURES

With a rich information structure, initial start volatilities will be low, because information common to all simulations will play a large part in explaining the variability of the first year's returns. With a sparse information structure, initial start volatilities will be close to their forward-start cousins, indicating that the first year returns have little additional predictability.

Information structures are of a fundamentally different nature to the parameters already discussed. Our defined parameters can be estimated, in principle, either from simulations or from a sufficiently long data series.

Information structures define how much information is known, at a point in time, regarding future outcomes, based not only on a series itself but other possible underlying drivers in a model. We could have two models with different information structures but for which the time series properties of a particular series, viewed on its own, are indistinguishable. We will develop a methodology to extract these statistics from simulated data.

Extraction from historic data is much more problematic, to say the least. This is because historic data on a time series itself does not necessarily encapsulate all that is known about future prices. Extra information might also come from other series. For example, an interest rate series might be predictive of movements in equity or currency markets. A specified finite list of possible predictors can give an estimate of how much the base series can be predicted, but this estimate will be a lower bound because we might have omitted a particularly predictive series from our analysis. Unfortunately, we cannot overcome this problem by throwing in more and more potential predictors, because sampling error takes over and any list of possible predictors will appear strongly predictive as the number of series approaches the number of time points in the data set.

Lee and Wilkie (2001) draw an analogy between information structures and the select period in a mortality table. Population mortality alone can tell us nothing about the effect of selection. The select period effect depends critically on the underwriting criteria adopted by a life office, and how those correlate to an individual's mortality.

Thus, we regard information structures as a property of a simulation data set, but not a readily measurable property of historic data. This measurement difficulty should not surprise us. Information structure is about the extent to which future share price can be predicted, not only from its own history, but also from other available data. In other words, it is a test of semi-strong market efficiency (see Fama and French, 1971).

The historical efficiency of particular markets is hotly debated – and much of that debate depends on the unanswerable question of whether, at a market extreme, analysts could have known whether a reversal was imminent. On the other hand, there is no dispute about whether particular models describe an efficient market. For example, it is accepted that Wilkie's (1986) model describes an inefficient market while Smith's (1996) Jump Equilibrium model describes an efficient market.

We can thus add a fourth possible basis for the predictor:

- All information in the model's history

The more information that is included, the greater is the variance of the predictor and the lower is the residual variance. The first three of these predictors could all in principle be extracted from historic time series, while the last requires simulation data.