Mean Reversion and Market Predictability

Jon Exley, Andrew Smith and Tom Wright

Abstract:

This paper examines some arguments for the predictability of share price and currency movements. We examine data from 1976-2001, and reproduce some classical regression results. On further investigation, we find biases in the classical estimation procedures. These biases are likely to overstate the extent of apparent market predictability.

The biases are also present in the standard calibration processes for many popular stochastic asset models. They lead to an under-statement of investment risk, particularly for equity investment over long horizons, and an overstatement of the extent to which investors can improve returns by dynamic sector switches.

Efficient market models provide a simpler, robust and prudent alternative approach for risk management and control purposes.

Risk Premiums – The Two Points of View

The return on a stock market or currency can be expressed as a risk free return plus a risk premium.

There are two schools of thought on risk premiums. One school argues that market movements are, to some degree predictable. There are times when markets are cheap, and in such cases there is a high prospective risk premium. At other times when markets are dear, the prospective risk premium is low.

This first school points to data apparently showing that yields can predict share price movements. Equity prices appear to rise after dividend yields have been high. They tend to fall after historically low yields. In the UK Actuarial literature, David Wilkie (1993) has popularised this view. The predictability is built into a number of popular mean-reverting investment models (Wilkie, 1995, Duval et al, 1999). For simplicity, we refer to this perspective as the "mean reverting school" in the rest of this paper.

The second school argues that prospective risk premiums are close to being constant over time. They may vary according to fluctuations in risk. However, we should always be expecting equity risk premiums of a few percent in excess of risk free rates. Risk premiums for major currencies are likely to be small.

A consequence of the second school is that price movements should be close to unpredictable. There might, however, be a small degree of predictability arising from yield differentials. For example, at times of high dividend yields we might expect lower equity capital growth in order to maintain a given expectation of total returns. We refer to this perspective as the "Efficient market school" in the rest of the paper.

These two points of view are readily testable. It should be a simple regression exercise to establish how the change in asset price over a period relates to the yield at the start of the period. In practice, of course, we may find these tests to be inconclusive.

Initial Evidence

Figure 1 shows data for five different national equity markets, based on month end data, for 1978 through 2002. We have calculated:

- The dividend yield (horizontal axis)
- The price change during the following year

These data are plotted with a best-fit straight line (one for each country). Does the positive slope of this straight line indicate that dividend yields may be able to predict subsequent market returns?

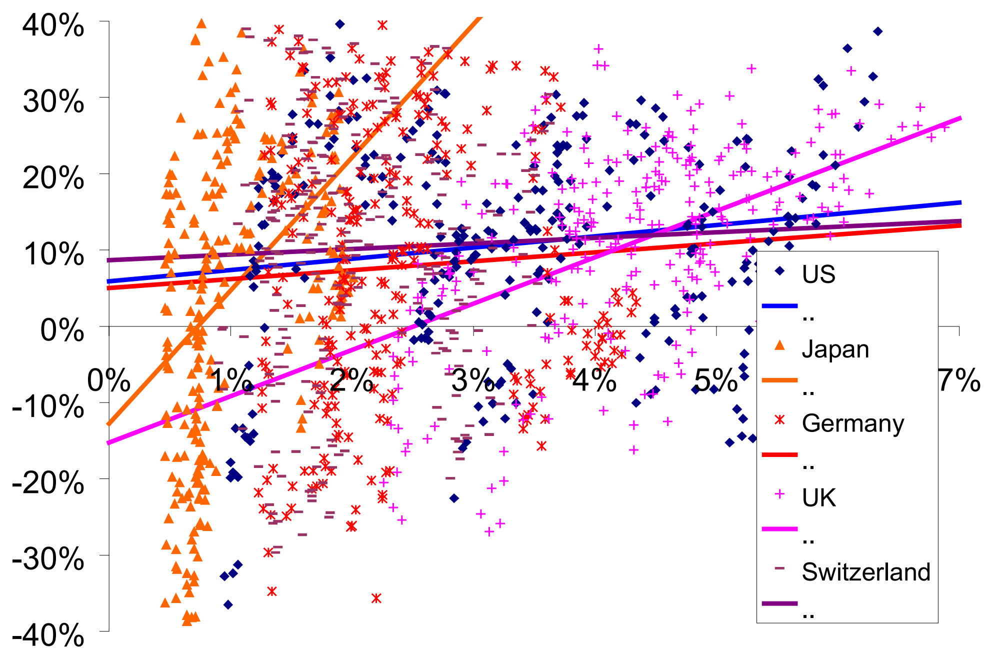

The positive slope parameter (about 6 for the UK) suggests that investing when dividend yields are 1% higher then "normal", results in an expected return around 6% higher over the next year. We will see that broadly similar results hold for many different indices over many time periods and across many major world stock markets. This conclusion seems to support the mean reverting school.

The potential practical application of this idea would be to examine dividend yields relative to some normal level, and to use this as a signal for buy and sell decisions in the market. To be confident in this tool, we would need to be confident not only in the regression slope, but also in confident in our estimate of the normal level of dividend yields. In the UK, many have regarded actuarial efforts to set "long term" dividend yield assumptions, for example in the statutory pension funding basis, as embarrassingly unsuccessful, as dividend yields have repeatedly moved out of their recent historic trading ranges.

Eliminating the Predictability

We can take these observations one step further, and evaluate the long term effect on returns, as dividend yields move from whatever their current level is, towards their long term norm.

Our idea is to express:

log share price = pure random walk + η * (dividend yield - mean dividend yield)

We could regard our pure random walk as some measure of underlying value, while the second term represents departures from this fundamental value, represented as temporary cheapness or dearness. Even without trying to guess short-term value, long-term investors may be concerned mostly with the risk of changes in fundamental value. They can ride out the short-term effect of fluctuations in market sentiment. If our estimated pure random walk is less volatile in the short term than the share price itself (as turns out to be the case) then this implies long term investors get some sort of time diversification benefit, which might lead them to see a special attraction in the equity market.

Our characterisation of the random walks is that its changes are not predicted by dividend yields. Our process for estimating η is then to start with a guess, then to calculate:

guess pure random walk
= log share price – (guess η) * (dividend yield - mean dividend yield)

We then regress changes in our guess random walk against the initial dividend yield. When the slope is zero, we have guessed the right value of η.

We can estimate η for various holding periods. We can regress the changes in our guessed random walk over a period of *k* months, against the dividend yield at the start of the k month period. The chart below shows estimates of η for five different equity markets, for increments *k* varying from 1 month through 60 months:

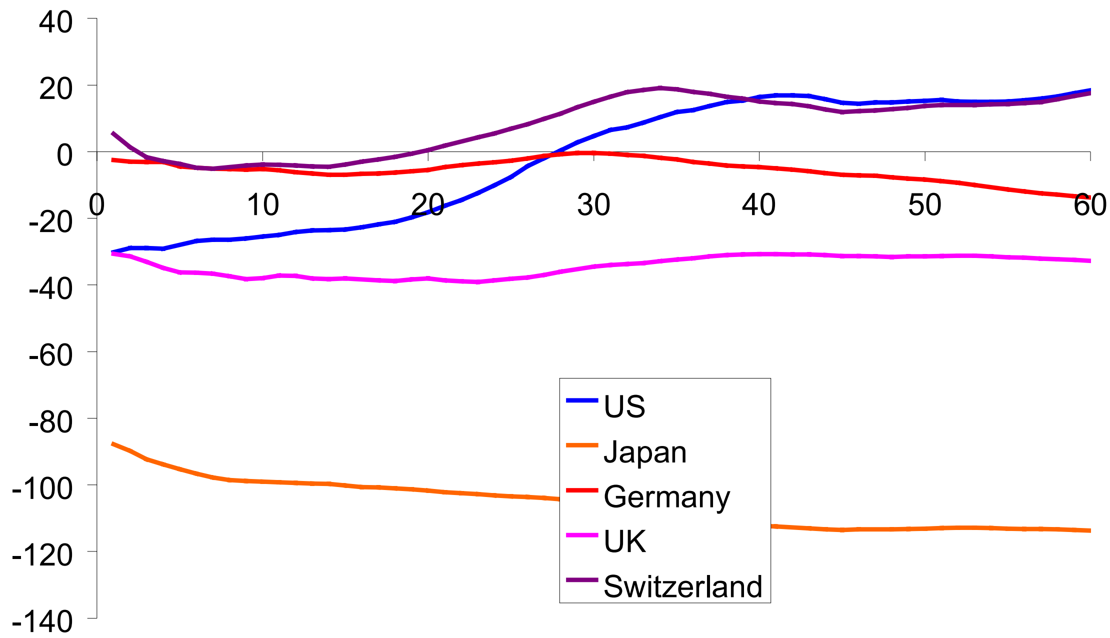

These parameters were estimated using rolling k-moth periods, that is, using overlapping data. We can verify that for many of these markets, and at least for holding periods of 20 months or fewer, we estimate a negative value of η. This is consistent with the results of Fama and French (1988) and Wilkie (1993). Interestingly, we can verify Wilkie's observation that the effect persists over longer holding periods, but only in the UK and Japan.

The parameter η gives us a picture of the long-term effect of abnormal dividend yields. For example, suppose η = -20 and that dividend yields are 1% above their normal value. This would imply that current prices are 20% below their normal value. We would forecast an upward correction of 20% in prices to occur at some future point. In a trading context, this would probably trigger a buy signal.

Cause for Concern?

We now present a simplified example, which could give some explanation for the apparent predictability in share price moves, even when markets are efficient.

We consider a simple model of a random walk, over four periods. A share price starts at 10, and can increase or decrease by 1 each period. This gives the following 16 possible price sequences:

Case	t=0	t=1	t=2	t=3	t=4
1	10	9	8	7	6
2	10	9	8	7	8
3	10	9	8	9	8
4	10	9	8	9	10
5	10	9	10	9	8
6	10	9	10	9	10
7	10	9	10	11	10
8	10	9	10	11	12
9	10	11	10	9	8
10	10	11	10	9	10
11	10	11	10	11	10
12	10	11	10	11	12
13	10	11	12	11	10
14	10	11	12	11	12
15	10	11	12	13	12
16	10	11	12	13	14

We now consider regression to relate the price increase over a year to the price at the start of the year.

A regression analysis of historic data carried out after time t=4 would be based on just one of the 16 possible price sequences, namely, the one that had actually happened. The plot would consist of four points. However, if you plot the price increase against the price at the start of each period for any one of these 16 possible paths, you find that the slope of the best-fit line is never positive: in 12 of the 16 cases it is negative, in the other 4 it is zero (cases 1, 8, 9 and 16). So, for example, if each of the 16 possibilities were equally likely, the historical evidence after time t=4, would be likely to suggest (if analysed by regression in this way) that a high share price now suggest a fall over the next year, and vice versa. In other words: that there is mean reversion and that markets are inefficient. However, it is clear from the simple random walk construction of this model that this would be a false conclusion.

The high chance of a negative regression slope simply reflects the fact that what (with the benefit of hindsight) was a relatively high price must usually have been followed by a price fall. Otherwise, it would not appear, with hindsight, to have been a particularly high price. In particular, any local maximum must have been followed by a fall, and any local minimum by a rise. This is of no use whatsoever for predictive purposes: future price movements may still be completely independent of the current price level.

Another way of looking at this is that if we have a random walk, the increments are independent and identically distributed. However, by plotting them against price at the start of each period, we are not plotting them in a random order: we are partially sorting them so that the positive increments tend to fall to the left and the negative increments to the right. Not surprisingly this tends to give a negative slope.

We now return to our original observation: we did not regress returns against prices, but against initial dividend yields. However, the issues are broadly the same.

Absolute dividend levels (monetary amounts) have been moderately stable, at least compared to the volatility of share prices. So if there is a high total return in a particular year, most of this comes from a price rise, therefore the dividend yield falls. The dividend yield would have been higher at the start of the year, than at the end. So dividend yields that are high, with the benefit of hindsight, tend to be followed by high total returns: hence the positive slope in the regression of total real returns against initial dividend yield. Once again, this is useless as a way of making money, since it is only with hindsight that we know whether dividend yields at a point in time were high or not.

Regression and Chartist Arguments

Statistical tests used in this way are really just a formalisation of the naive chartist argument:

"Every time the price reached a high in the past it subsequently fell. The price has reached a high now, so we can expect its next move to be downward."

The first sentence is vacuous of course: by definition a high is followed by a fall. In the second sentence the second part is not a deduction from the first part (as implied by the word "so"), it is just another way of saying "the price has reached a high now". In efficient markets it is impossible to know this.

Quantifying the Biases

These biases have been identified in the finance literature, but have not received great prominence. Kendall (1954) gives an asymptotic estimate of the bias in estimating the mean reversion coefficient. Stambaugh (1986) and Mankiw & Shapiro (1986) consider bivariate systems equivalent to a discrete time version of our model. Stambaugh (1999) gives more analytical formulas for the bias, under restrictive conditions.

Share price analyses are often more complex than the models whose statistical bias is easily analysed, for example because of the use of overlapping intervals. We therefore turn to Monte Carlo tools. We make a simple model where the log share price (actually, a total return index) is a pure random walk. The annual return has a lognormal distribution, with a mean of 10% and a standard deviation of 20%. The dividend amount is 6% of a 5 year moving average price.

For 100,000 simulations, and for data periods of 5 through 20 years, we estimate η using ordinary least squares and annual holding periods. Remember that if we had an efficient market, we might hope to estimate η close to zero.

The chart below shows the 5th, 25th, 50th, 75th and 95th percentiles of estimated η.

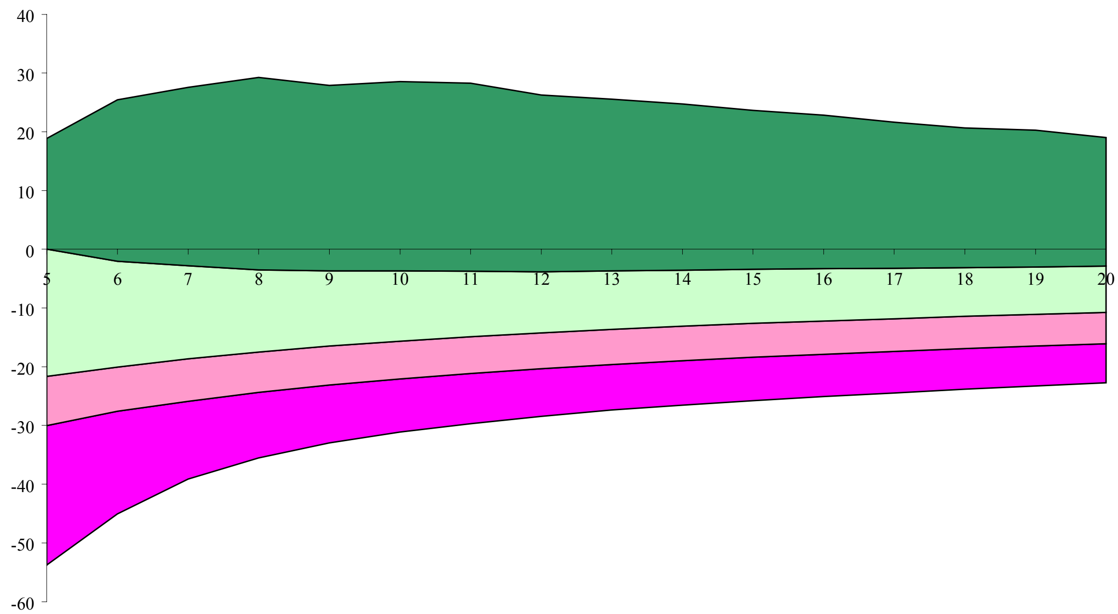

We can see a strong bias towards negative values of η. In comparison to our previous empirical work, we see that only Japan stands out as not being readily explained by this simple model. In particular, the negative values of η estimated in the UK are only around the median statistical bias. These do not constitute evidence for auto-regression in share prices.

Fixing the Tests

The remaining question relates to whether there is any fix we can apply to the statistical tests. If slopes and correlation coefficients are not the right way to proceed, what tests are up to the job?

In general, removing the statistical bias is a hard issue. We've given an example based on a four-step model, where the bias is evident. As the number of time periods increases, the bias reduces, ultimately to zero. Unfortunately, for some crucial economic time series, this asymptotic convergence is slow. Even with samples of hundreds of data points, classical tests may give spurious evidence of market inefficiencies.

How do we live with this? The first step is to ensure robust statistical inference. The simplest way to get a handle on small sample properties is by numerical methods such as Monte Carlo simulation. Without this, we don't know whether the asymptotic result is useful or not.

Conclusions

Let us return to our opening theme – that of the empirical evidence of inefficiencies in equity pricing. Empirical evidence is the acid test of any model. Such testing is a daunting task, as however carefully a model is crafted, there always seem to be some facts, which the model poorly explains. In the case of efficient market models, the historically observed correlations between dividend yields and subsequent market performance may have seemed an insurmountable hurdle to acceptance. It is a salutary lesson that, after all, efficient market models comfortably account for such correlations.

The biases we have seen also apply to calibration of mean reverting models. Statistically, at least using ordinary least squares, we are doomed to overstate the extent of market predictability. Our autoregressive models are therefore likely to understate the risk of equity investment over longer terms. We will also be overconfident in our estimates of long-term means for mean-reverting processes.

If such models are used for risk management purposes, we have a further degree of circularity. Autoregressive models may give a signal that a market is cheap or dear. But to code this value judgement into a risk management model, is to benchmark our trading view against itself. A favourable view of our own skill is likely to be the result – which is flattering for the model developer but poor risk management practice.

Efficient market models, where prices are inherently unpredictable, are more robust for risk management processes. They are also capable of generating many results that were previously thought conclusive proof of market inefficiency.

Acknowledgements

We are grateful to Elliot Varnell for extracting our data sample, and to Shyam Mehta for useful discussions, in particular drawing our attention to Nelson & Myung (1993). Any remaining errors are those of the authors only.

References

Duval, Donald B; Teeger, Mark H; Yakoubov, Yakoub H; (1999) The TY model. A stochastic investment model for asset and liability management. Staple Inn Actuarial Society. Also available at http://www.sias.org.uk/papers/model.pdf

Fama, Eugene F and French, Kenneth R (1988). Dividend yields and expected stock returns, *Journal of Financial Economics* 22, 3-25.

Kendall, M G (1954). Note on bias in the estimation of autocorrelation, *Biometrika* 41, 403-404.

Mankiw, N Gregory and Shapiro, Matthew D (1986). Do we reject too often? *Economics Letters* 20, 139-145.

Nelson, Charles R and Kim, Myung J (1993). Predictable Stock Returns: The Role of Small Sample Bias. Journal of Finance, Vol XLVIII, No 2 (June 1993), pp 641-661

Stambaugh, Robert F (1986) Bias in regression with lagged stochastic regressors, CRSP Working Paqpewr No 156, University of Chicago

Stambaugh, Robert F (1999). Predictive regressions. Journal of Financial Economics, 54, 375-421.

Wilkie, A David (1993) Can Dividend Yields Predict Share Price Changes? *Proceedings of the 3rd AFIR Colloquium*, vol 1, pp335-347.

Wilkie, A David (1995). More on a Stochastic Asset Model for Actuarial Use. *British Actuarial Journal* Vol 1 p 777-946