Is There Evidence for Momentum and Mean Reversion in Asset Class Returns?

Grant Holtes

July 1, 2024

Summary

The presence of momentum and mean reversion effects in returns at the asset class level are often postulated or presumed. This note aims to provide a simple measurement of these effects from historical data. The results suggest that fixed income assets show short and long horizon momentum effects, while equities show short term momentum and long term mean reversion. However, only a small subset of the momentum effects seem significant, at very short horizons and with small absolute effect sizes. The measurement of long horizon mean reversion is greatly impaired by the small number of independent periods that are available, even in a relatively long data series.

1 Introduction

Momentum and mean reversion in asset returns can be thought of as violations of the assumption that returns are independent over time. Momentum manifests as a positive relationship between time periods - Periods of good returns are followed by more good returns and vice versa. Mean reversion provides the opposite and manifests as a negative relationship between time periods, with periods of good returns being followed by periods of poor returns.

This note simply aims to provide a measurement of momentum and mean reversion from historical returns at an asset class level. This does not address the underlying mechanics of why such an effect may be present nor does it try to explicitly decompose returns into any underlying components that contribute to the effect.

2 Methodology

2.1 A range of available approaches

A number of methods have been used to estimate momentum and mean reversion:

Variance ratio tests: If log returns are independent over time, the variance in log return will be proportional to return horizon. If this proportionality does not hold, this provides evidence for non-independence over time. The test statistic used is $VR(k_1, k_1) = \frac{VAR(r_{k_1})/k_1}{VAR(r_{k_2})/k_2}$, where r_k is the series of k period log returns. This is explored in more detail by Poterba and Summers, 1987.

Standard unit root tests: A stationary time series had a constant mean, and as such displays mean reverting characteristics. Through this, statistical tests for stationarity such as the Augmented Dicky-Fuller test can be used to test for the presence of mean reversion.

Autoregression tests: If momentum and mean reversion are defined as a relationship in returns across time periods, it is natural to measure this relationship through a regression between historical and forward looking returns. This technique is used by Fama and French, 1988, where they ran regressions to estimate $r_{t,t+T} = \beta_0 + \beta_1 r_{t-T,t}$, where r_{t_1,t_2} is the cumulative return between times t_1, t_2.

2.2 Approach taken

In its simplest form the task is to measure the correlation between returns in the previous and next periods, which is the approach taken in this note.

$$\rho = Corr(r_{t-1}, r_t) \tag{1}$$

This can be thought of as approximately equivalent to the auto-regressive technique used by Fama and French, as the correlation coefficient, ρ, should be equal to the slope coefficient in a AR(1) in a large enough sample.[1]

Of course this method is influenced by the frequency of the data used, with monthly data only answering if momentum or mean reversion is present over a one-month horizon. Allowing the look-back period, l, and forecast horizon, h, to vary allows for more varied questions to be answered.

$$\rho(l,h) = Corr(r_{t-l,t}, r_{t,t+h}) \tag{2}$$

Where r_{t_1,t_2} is the cumulative return between times t_1, t_2. All cumulative returns are normalised by using log returns and are scaled to the same frequency, monthly in this case, as in the Equation 3.

$$r_{t,t+T} = \frac{1}{T}\sum_{i=0}^{T} ln(1 + r_{t+i}) \tag{3}$$

Hypothesis tests and p-values on the correlation coefficient can be used to determine the significance of the results. The estimation of standard error and the specification of the tests naturally requires the sample size to be known and assumes that observations are independent.

Given that monthly time series data is being used, it is tempting to use rolling windows to estimate the returns for input into the estimators, utilising all the available information. However, this is violating the independence assumption and inflating the sample size. To adjust for this issue and measure its impact, three versions of the tests are run:

Naive approach: This method uses rolling time windows and the resulting sample size in the tests without adjustment.

[1] This relationship is shown in the appendix, and a version of this analysis using the autoregession model is run in parallel, which showed comparable results.

Non-overlapping: We can easily remove the issue by ensuring that no periods overlap in the estimation. This is equivalent to only using every n^{th} observation from the rolling data, where n is the larger of l, h.

Adjusted sample size: *Throwing away* observations in the non-overlapping method is sub-optimal, and can be improved by keeping all observations but adjusting the sample size to an effective sample size.

Details on the data and calculations used are provided in the appendix.

3 Results

The results are discussed in three sections. First the naive approach is used to evaluate the correlations, without adjusting for overlapping periods. Secondly the results are adjusted for sample size using the adjusted sample size method, with the non-overlapping approach included as a baseline. Finally the same analysis is done for some non-asset class series that are also of interest in the asset-class return forecasting process.

In all cases, two visualisations are used to display the results:

Correlation by horizon: These are line plots that show how the correlation changes as the look-back period, l, and forecast horizon, h are increased symmetrically, with $l = h$. Confidence intervals (0.95) are plotted with dashed lines to explore significance of the results.

Correlation grid: These plots explore the correlation for every pair of l, h. Rather than show both magnitude and significance, the plot focuses on only the significance and the sign of the correlation. Blue indicates significant positive correlation (momentum) and red indicates significant negative correlation (mean reversion). The intensity of the colour indicates the level of significance, from full saturation (p value of 0) to a white (p value of 0.2 and above).

These are shown in the example of Small-Cap US stocks, below, which appears to show both short term momentum and long term mean reversion that is robust across return horizon combination.

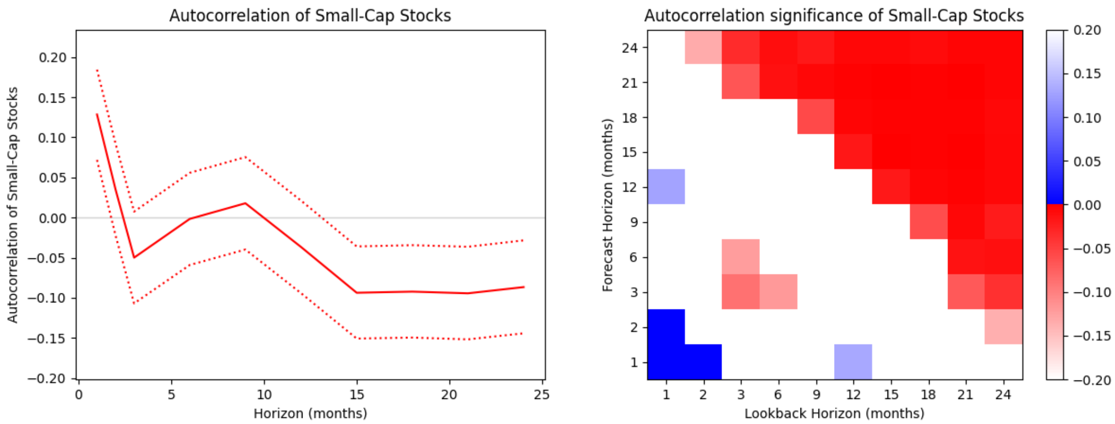

Figure 1: Autocorrelation of Small-Cap US stock returns, naive standard errors.

3.1 Naive approach

The use of un-adjusted standard errors provides encouraging results, with some asset classes showing significant momentum (equities, treasuries, short term government bonds, corporate bonds), some show significant mean reversion (large and small-cap stocks). These results seem to be robust to changing the look-back and forecast horizons.

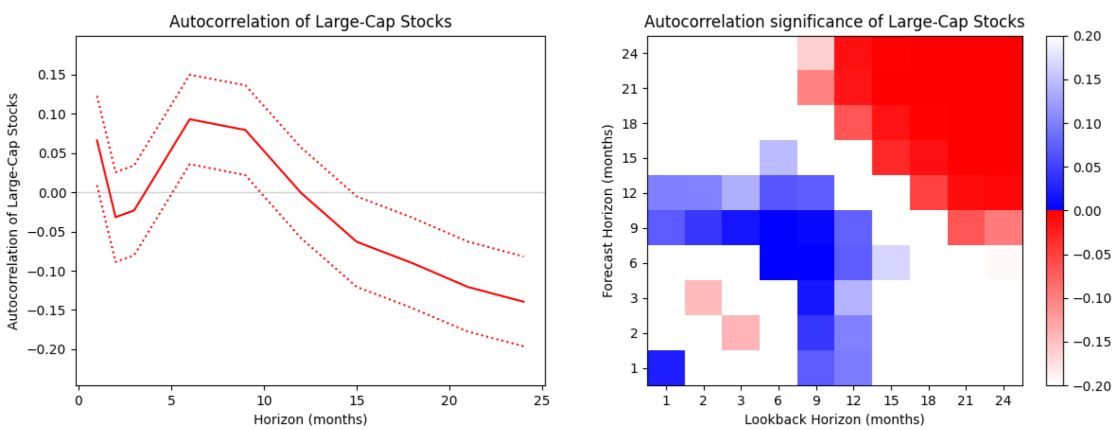

Figure 2: Autocorrelation of Large-Cap US stock returns, naive standard errors.

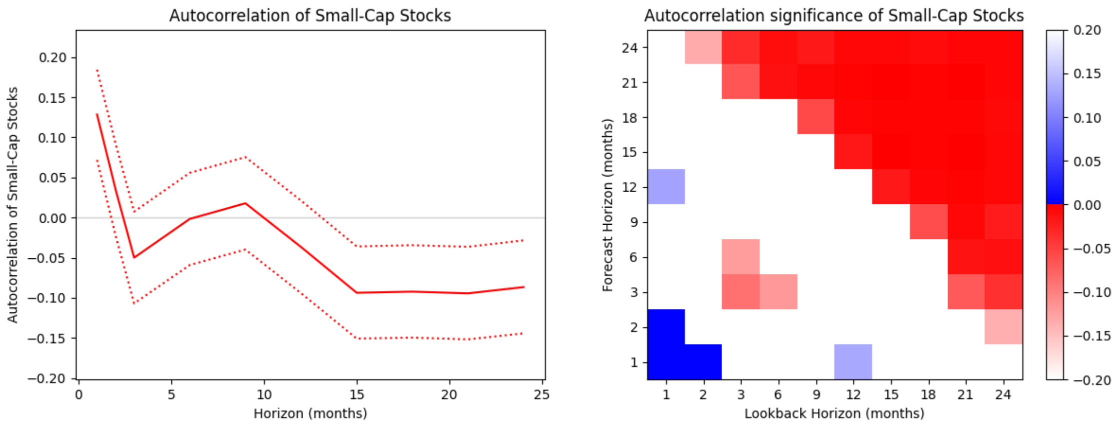

Figure 3: Autocorrelation of Small-Cap US stock returns, naive standard errors.

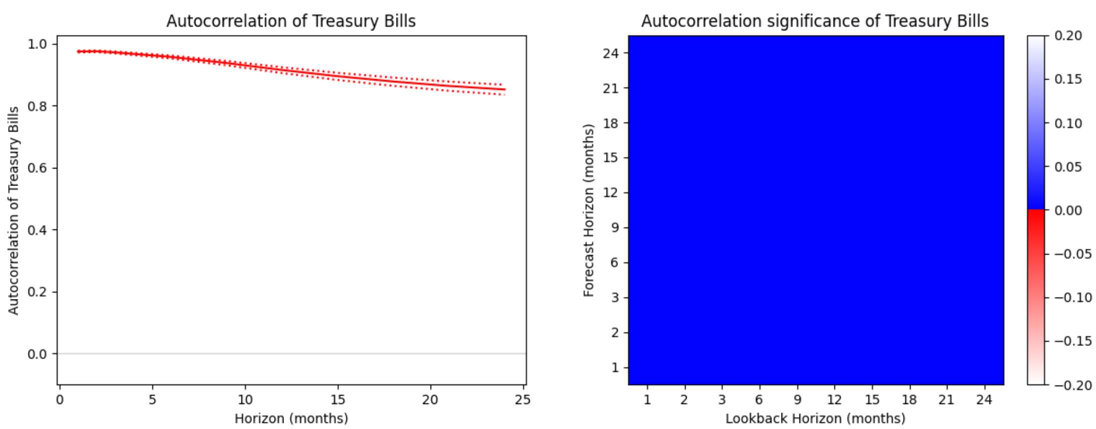

Figure 4: Autocorrelation of Treasury returns, naive standard errors.

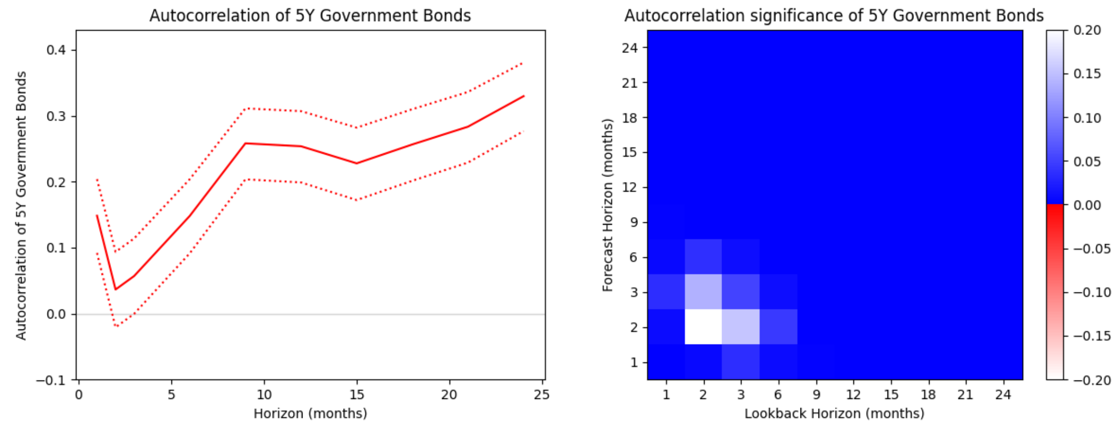

Figure 5: Autocorrelation of 5Y Government Bond returns, naive standard errors.

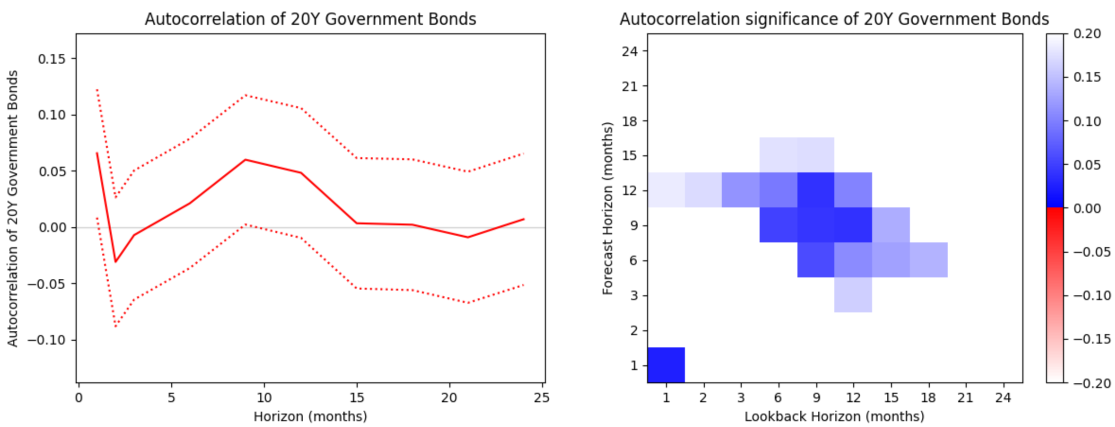

Figure 6: Autocorrelation of 20Y Government Bond returns, naive standard errors.

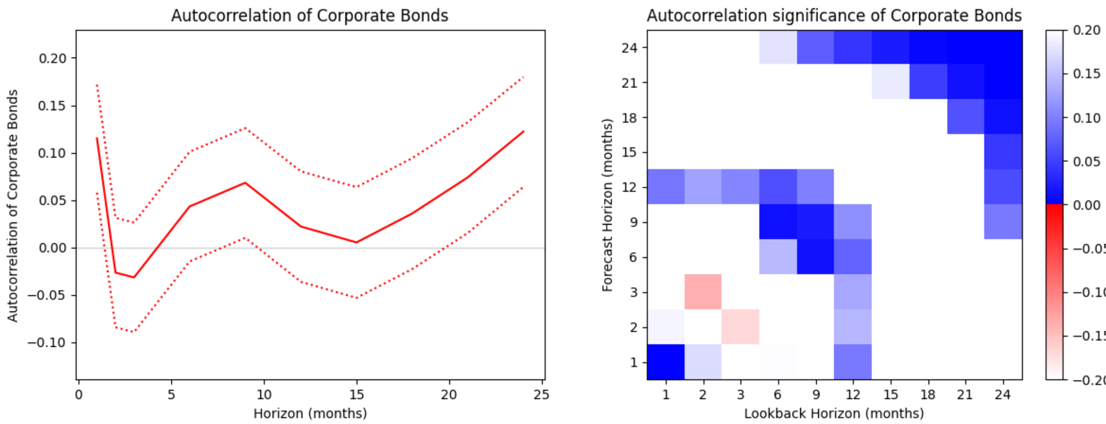

Figure 7: Autocorrelation of Corporate Bond returns, naive standard errors.

3.2 Adjusted for effective sample size

Adjusting standard errors for overlapping periods provides results that are less convincing. As the longer return horizons generate more overlapping periods, the reduction in significance is more pronounced, with the promising mean reversion results being removed in large-caps and small-caps. Some evidence for short term momentum remains for both equities and bonds.

The correlation by horizon line plot shows the adjusted sample size results (black lines) and the non-overlapping results (grey lines), while the correlation grid shows only the adjusted sample size results.

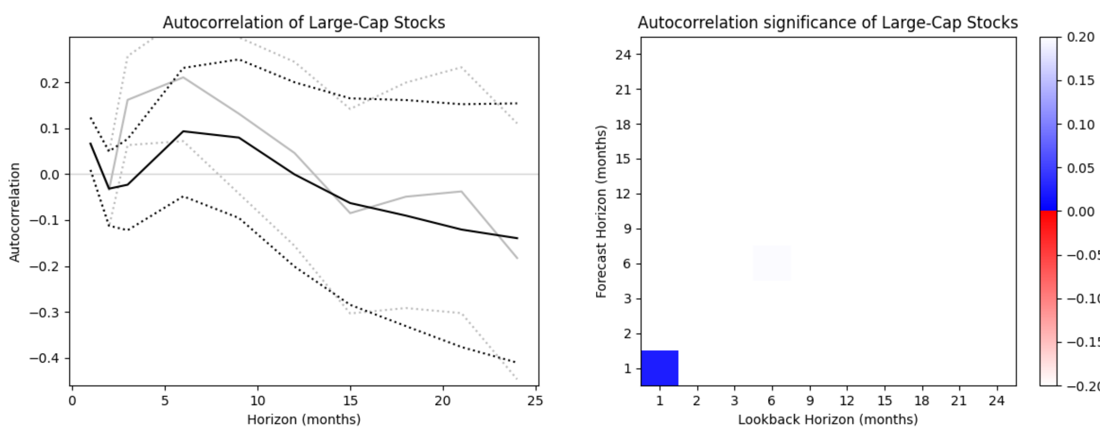

Figure 8: Autocorrelation of Large-Cap US stock returns, adjusted standard errors.

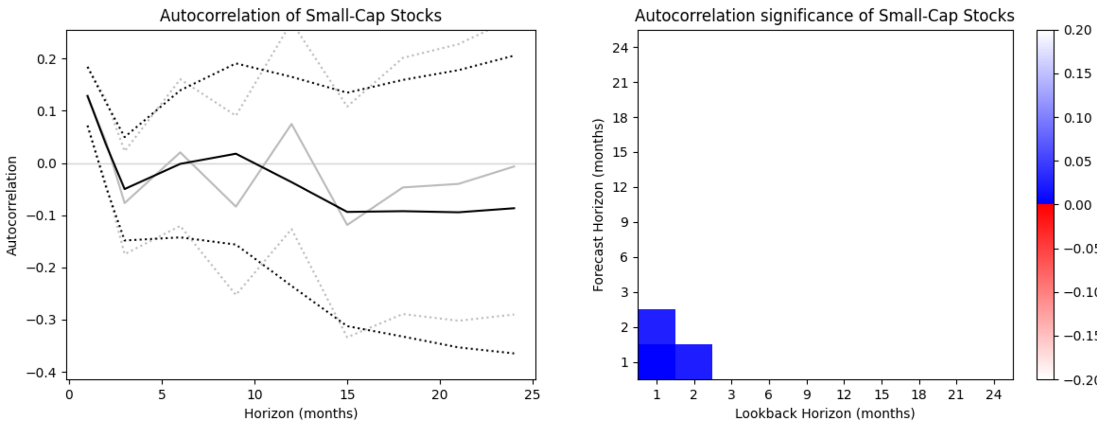

Figure 9: Autocorrelation of Small-Cap US stock returns, adjusted standard errors.

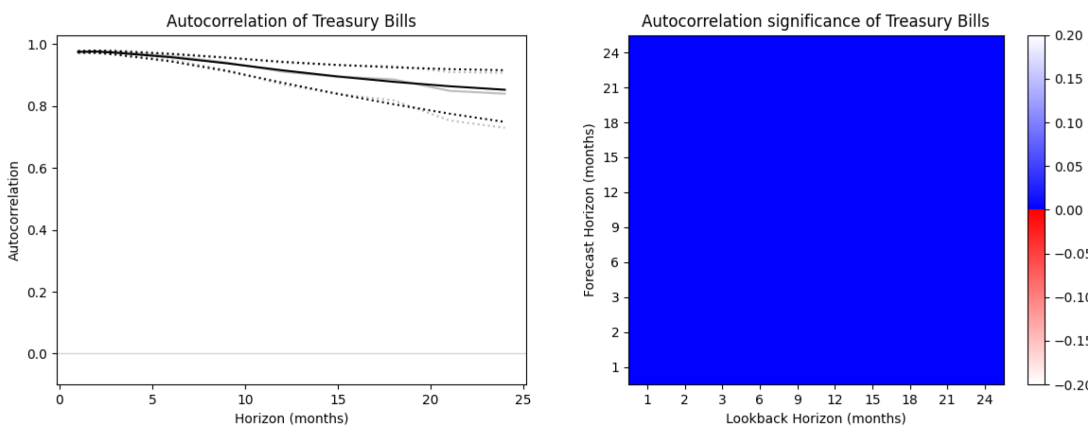

Figure 10: Autocorrelation of Treasury returns, adjusted standard errors.

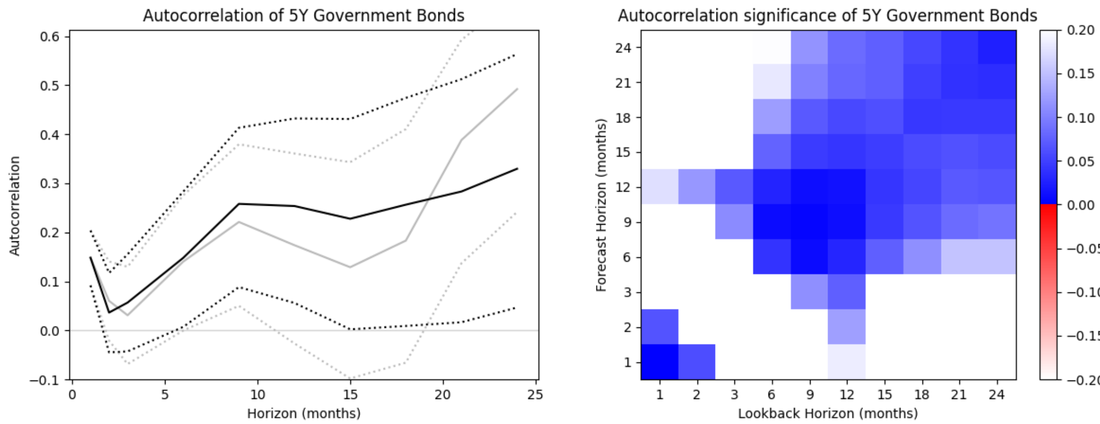

Figure 11: Autocorrelation of 5Y Government Bond returns, adjusted standard errors.

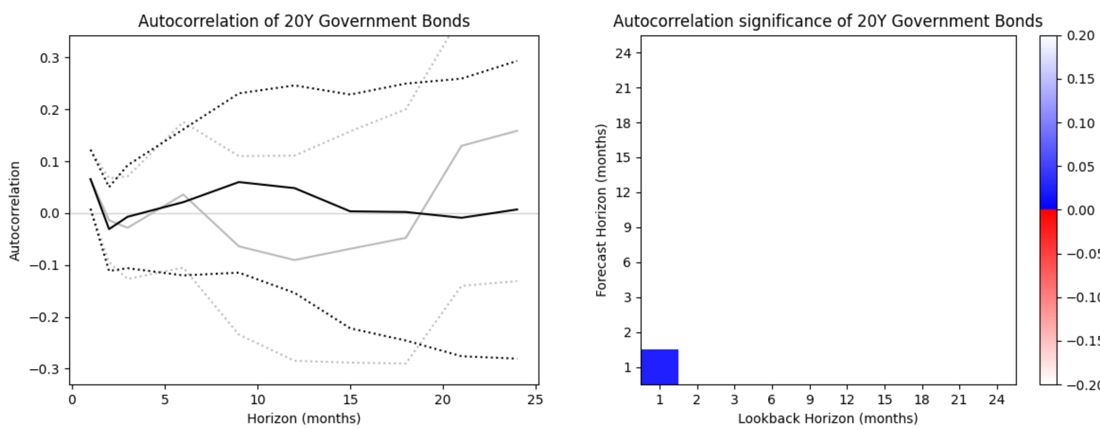

Figure 12: Autocorrelation of 20Y Government Bond returns, adjusted standard errors.

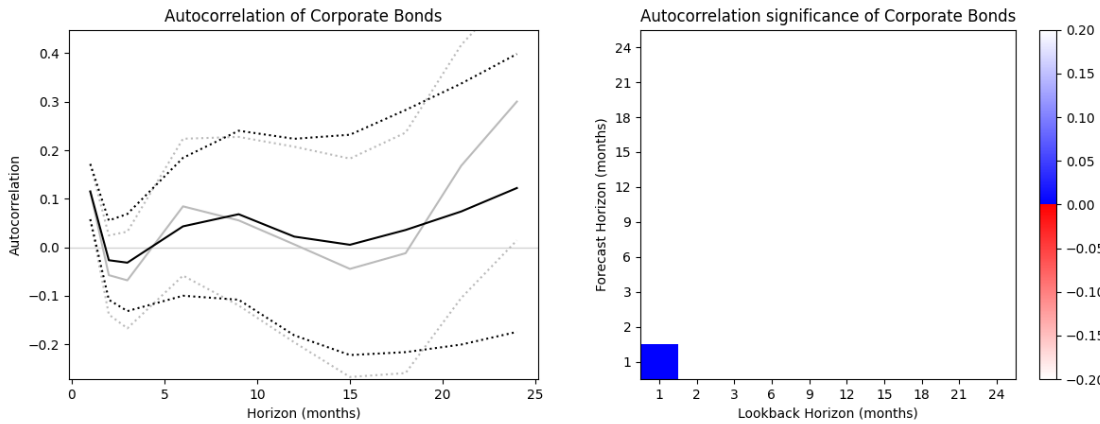

Figure 13: Autocorrelation of Corporate Bond returns, adjusted standard errors.

3.3 Other variables of interest

It is interesting to apply this same analysis to some non-asset class variables that are often discussed as key determinants to asset class returns:

1. Inflation and Interest Rates
2. Earnings and Dividend growth rates

All of these variables show momentum to various horizons. Inflation and interest rates shows a high degree of autocorrelation all the way out to a 24 month horizon, the longest measured. This is an unsurprising result given that inflation tends to be a slow moving variable. Given that inflation is linked to interest rates via monetary policy, it follows that interest rates show similar properties.

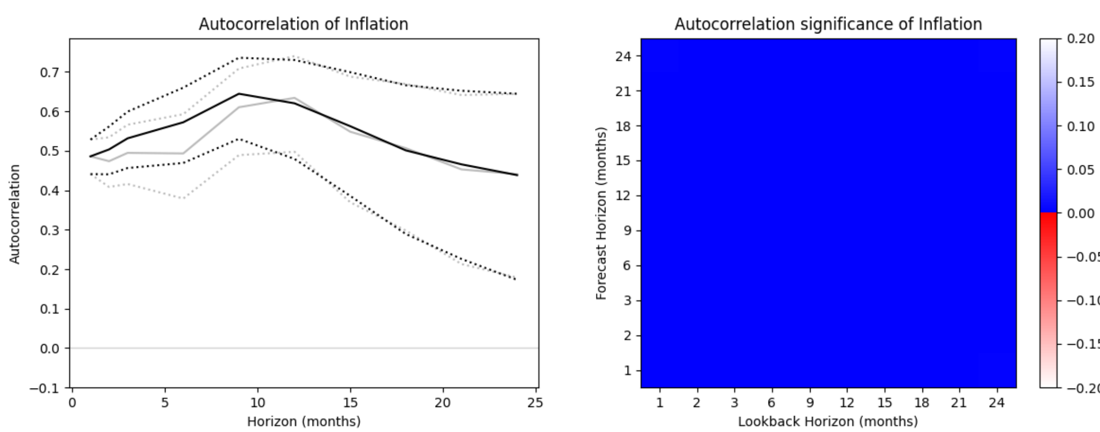

Figure 14: Autocorrelation of inflation, adjusted standard errors.

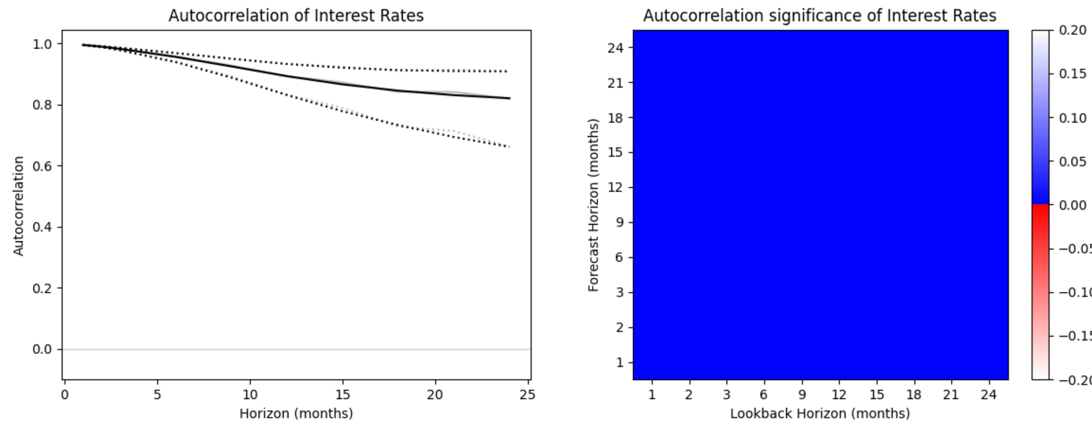

Figure 15: Autocorrelation of interest rates, adjusted standard errors.

Both inflation and interest rates have been shown to be highly persistent, with Ólan Henry, 2004 suggesting that shocks to US inflation show infinite persistence, while the UK and Japan have multiple inflation regimes, with inflation only showing persistence in certain regimes. This high degree of persistence is likely contributing to the measured momentum.

Given the strength of these results, it is interesting to consider the differences versions of these series - how is the change in inflation linked to previous changes in inflation? In the results below change in inflation shows strong mean reversion across horizons while the change in interest rates shows short term momentum only.

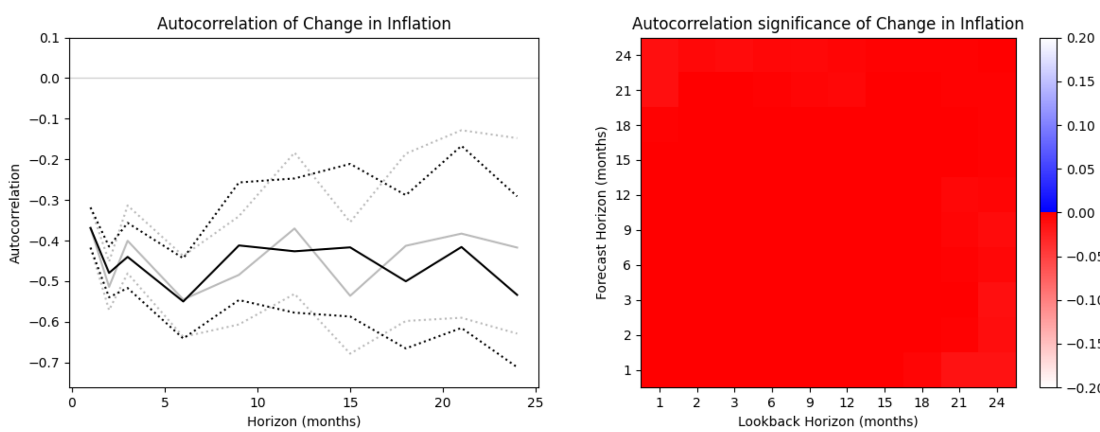

Figure 16: Autocorrelation of change in inflation, adjusted standard errors.

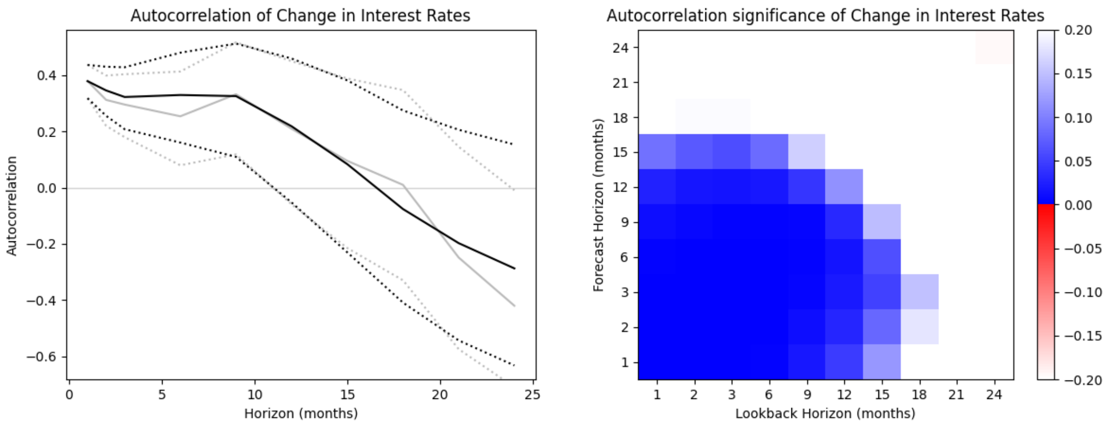

Figure 17: Autocorrelation of change in interest rates, adjusted standard errors.

Both dividend growth and earnings growth show significant momentum, with the horizons being longer for dividend growth than earnings growth, with the difference in horizon potentially being an artifact of dividend smoothing. Earnings growth also shows significant mean reversion over longer horizons of one to two years.

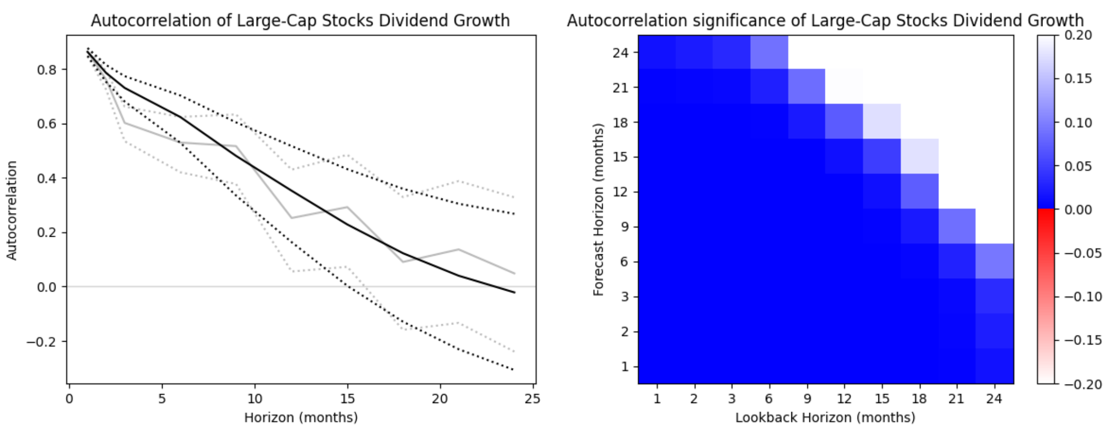

Figure 18: Autocorrelation of dividend growth rates, adjusted standard errors.

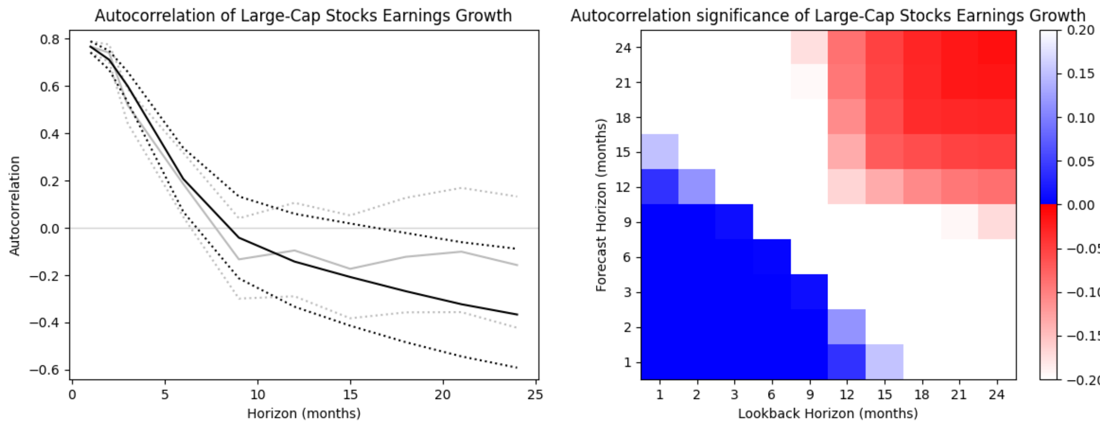

Figure 19: Autocorrelation of earnings growth rates, adjusted standard errors.

4 Conclusion

The empirical analysis presented in this note provides some evidence for momentum in asset-class returns, although the effect appears to be limited to very short horizons. Correlations are generally small and only significant over a one-month period.

Evidence for mean reversion is less convincing. While equities and long-term government bonds exhibit negative correlations at longer horizons, the small sample size makes these results statistically insignificant. Mean reversion is often discussed over longer periods than those measured here, but the decreasing effective sample size at extended horizons further diminishes the likelihood of finding significant results.

Future analysis could benefit from decomposing returns and examining the properties of these components to establish a stronger empirical or theoretical basis for these effects. The results for earnings and dividend growth showed promising significance and would be interesting to investigate further.

Note:

Views expressed are the author's, and may differ from those of JANA investments. This material does not constitute investment advice and should not be relied upon as such. Investors should seek independent advice before making investment decisions. Past performance cannot guarantee future results. The charts and tables are shown for illustrative purposes only.

References

CFA-Institute. (2024).

Fama, E., & French, K. (1988). Permanent and temporary components of stock prices. *Journal of Political Economy*.
Ólan Henry, K. S. (2004). Is there a unit root in inflation? *Journal of Macroeconomics*.
Poterba, J., & Summers, L. (1987). Mean reversion in stock prices: Evidence and implications. *Massachusetts Institute of Technology*.
Robert Shiller, Y. S. o. M. (2024).

5 Appendix

5.1 Correlation - Regression equivalence

In large enough samples sizes the correlation between a variable and a lag of itself is equivalent to the slope in a linear autogregressive model with one lag (AR(1)). This correlation - regression equivalence can be shown by examining the correlation formula.

The correlation coefficient between a variable X_t and its lag X_{t-1} is given by:

$$p = \frac{\sum_{t=1}^{n-1}(X_t - \bar{X})(X_{t-1} - \bar{X})}{\sqrt{\sum_{t=1}^{n-1}(X_t - \bar{X})^2 \times \sum_{t=1}^{n-1}(X_{t-1} - \bar{X})^2}} \tag{4}$$

For large samples the below holds

$$\sum_{t=1}^{n}(X_t - \bar{X})^2 \approx \sum_{t=1}^{n}(X_{t-1} - \bar{X})^2 \tag{5}$$

This simplifies the denominator of the correlation equation

$$p = \frac{\sum_{t=1}^{n}(X_t - \bar{X})(X_{t-1} - \bar{X})}{\sum_{t=1}^{n}(X_t - \bar{X})^2} \tag{6}$$

This is the same as the equation for the slope, β_1 in an autoregressive model with one lag:

$$\beta_1 = \frac{\sum_{t=1}^{n}(X_t - \bar{X})(X_{t-1} - \bar{X})}{\sum_{t=1}^{n}(X_t - \bar{X})^2} \tag{7}$$

This methodology can still apply when using different horizons for the lookback and forecast horizon, but the mean of the lookback and forecast returns series need to be the same. Intuitively the solution to this is to annualise the returns of the two series to remove the scaling difference due to time window length. This ensures the approximation in Equation 5 holds.

5.2 Standard Errors, Significance and Rolling Data

Analysis with time series data often requires adjustments to significance statistics to account for the autocorrelation inherent in the data. This is especially important when using rolling window data, where observations overlap. The two adjustment approaches explored are outlined below.

Effective sample size method: The effective sample size adjustment method adjusts degrees of freedom based on the number of rolling windows used in the analysis. It assumes each observation within the rolling window is independent and calculates the effective sample size by dividing the total sample size by the length of the rolling window. T-values and p-values are then calculated using the adjusted standard errors and the effective sample size.

Non-overlapping method: The simplest fix is to only keep every m^{th} observation of the m period rolling data, then proceed with the analysis as per usual.

See calculations section below for details.

5.3 Data

Data was sourced from the CFA SBBI dataset (CFA-Institute, 2024), which provides asset class returns at a monthly frequency from 1926. This was extended to include earning and dividend growth rates, sourced from Robert Shiller, 2024.

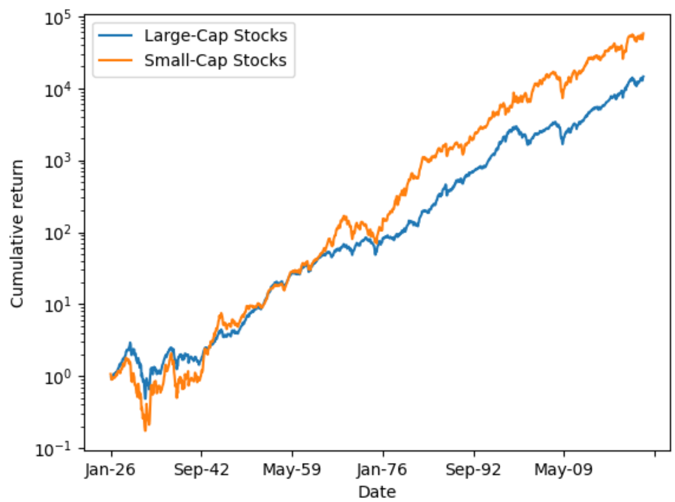

Figure 20: Data used for equities - Expressed as cumulative returns

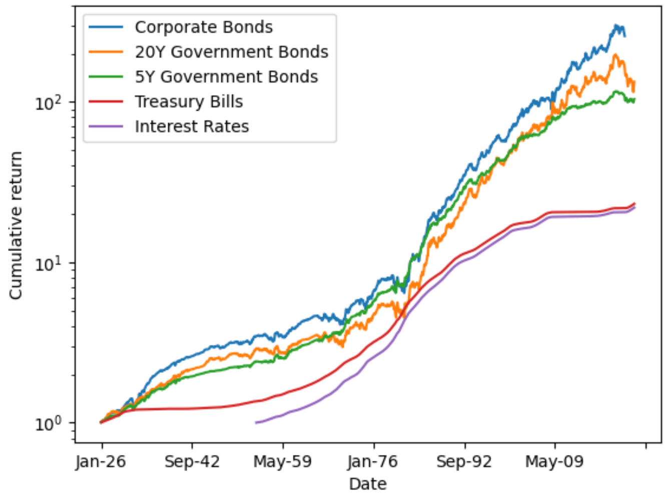

Figure 21: Data used for bonds - Expressed as cumulative returns

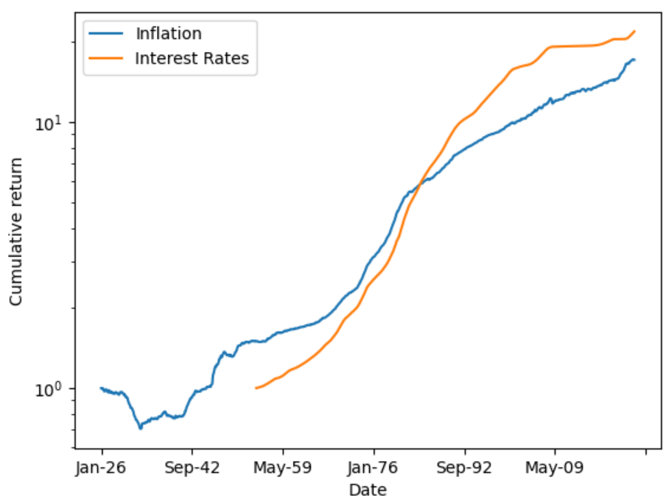

Figure 22: Data used for inflation and interest rates - Expressed as cumulative returns

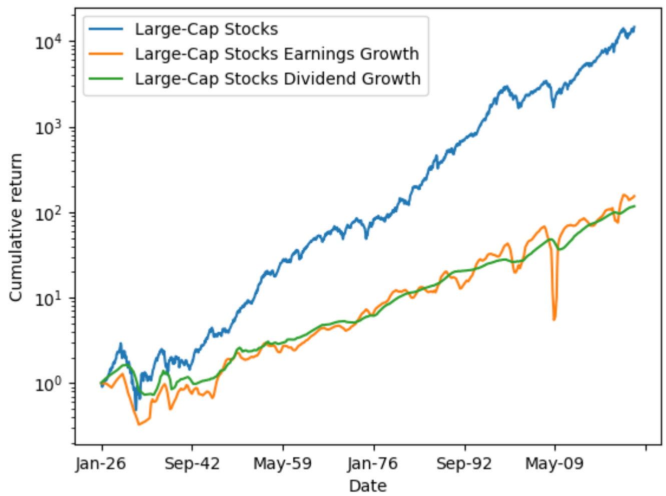

Figure 23: Data used for earnings and dividend growth rates - Expressed as cumulative returns

5.4 Calculations

The t-statistic on the correlation coefficient is given by:

$$t = \frac{\rho\sqrt{n-2}}{\sqrt{1-\rho^2}} \tag{8}$$

The adjusted sample size, n_{eff} is estimated as

$$n_{eff} = \frac{n}{max(l,h)} \tag{9}$$

The adjusted sample size can then be used in the equations for t-statistic, critical t-statistic, as well as p-value and confidence intervals.

$$t = \frac{\rho\sqrt{n_{eff}-2}}{\sqrt{1-\rho^2}} \tag{10}$$