Critiquing Diversification Benefits through a Simplified Conditional Correlation Estimator

Grant Holtes

March 28, 2024

Summary

This paper examines how correlations between asset classes vary across market conditions and introduces a sliding window conditional correlation estimator. Comparing with established methodologies, it assesses correlations between asset classes, focusing on their diversification potential relative to US equities. Results reveal that treasuries generally offer favorable diversification properties, while equities and corporate bonds often fall short. The study also explores how the estimator can be used to measure how correlations vary in response to macroeconomic variables such as interest rates. Overall, findings support existing literature but highlight that extreme tail observations can have an large impact on conclusions when aiming to understanding diversification properties using conditional correlation.

1 Introduction

Diversification benefit is a core outcome from portfolio design, however, during market selloffs, the diversification implied by the correlation between asset classes can seem to disappear. This is acknowledged widely enough that the statement that "all correlations go to one in a crisis" can sound like a cliche. On the other side of the return spectrum, it has also been observed that correlations can decrease with returns, reducing portfolio returns during periods of strong returns.

There exists a range of literature on this matter. Leibowitz and Bova, 2009 demonstrated that during the 2008 global financial crisis, a well-diversified portfolio comprising US stocks, US bonds, international stocks, emerging market stocks, and REITs exhibited a shift in its equity beta from 0.65 to 0.95. Further, contrary to expectations, this diversified portfolio underperformed a simpler allocation of 60% US stocks and 40% US bonds by a margin of 9 percentage points. This analysis is extended by Page and Panariello, 2018 in their comparison of the left and right tail correlations of US equities and major asset classes. With the exception of government bonds, they show that the left tail correlation is generally large and positive while the right tail is usually low to negative, suggesting that these asset classes provide less diversification benefit than their sample correlation would suggest.

It is helpful to define what properties we would want in a diversifying asset class to have. The usual definition of a good diversifier is an asset class that has a low, zero or negative correlation with the asset class being diversified. However, given the prior studies, an ideal diversifier for an asset class can be defined as having a negative correlation with the asset class when the asset class performs poorly (times of crisis), and a positive correlation when the asset class performs well. This allows the diversifier to provide downside protection when required while avoiding reducing portfolio returns when the asset class does well. As is suggested by prior studies, the opposite is often observed when diversifying US equities: The "diversifiers" have a positive correlation during times of poor equity performance and a low-to-negative correlation during times of good equity performance. This diversification benefit of a diversifying asset class can be measured through the conditional correlation: What is the correlation between asset classes when our primary asset class experience a particular return.

This paper explores this phenomenon with a new conditional correlation methodology, aiming to estimate how correlation and implied diversification properties vary across full range of returns, rather than just the left and right tails. Through this analysis we look to determine if the lack of diversification benefits exposed in previous studies also applies to periods of poor or good performance that lie between the extreme tails

2 Methodology

Let's first examine how conditional correlation has been measured in other studies. D.B. Chua and Page, 2009 use a "double conditioning" approach to calculate the correlation in the cases that that both assets have experienced a monthly return that is below or above a threshold.

$$\rho(\theta) = \begin{cases} corr(x, y|x > \theta, y > \theta) & \theta > 0 \\ corr(x, y|x > \theta, y > \theta) & \theta < 0 \end{cases} \tag{1}$$

This is iterated on by Page and Panariello, 2018, who introduce a "single conditioning" methodology, which only conditions on a single asset's returns. This better answers questions on diversification benefit, as it captures instances where the diversifier, y, has returns above the threshold, θ, while the asset being diversified, x, has returns below the threshold, and vice versa.

$$\rho(\theta) = \begin{cases} corr(x, y|x > \theta) & \theta > 0 \\ corr(x, y|x < \theta) & \theta < 0 \end{cases} \tag{2}$$

This paper builds on the single conditioning approach and defines a simple sliding window estimator as: The correlation of two assets, x, y, conditional on the percentile of a condition variable, c, is defined as the correlation between x, y for all observations where c is in a range around the value of interest, as in Equation 3. There is no restriction on c, so setting $c = percentile(x)$ allows for a similar style of analysis to prior studies.

$$\rho(\theta, \epsilon) = corr(x, y|c \in (\theta - \frac{\epsilon}{2}, \theta + \frac{\epsilon}{2})) \tag{3}$$

where

- θ is the percentile of interest
- ϵ is the window width.
- $\theta \in [\frac{\epsilon}{2}, 1 - \frac{\epsilon}{2}]$

A percentile window of 40% is used throughout, as this was found to provide a good balance between preserving detail while having a sensibly large sample size of 176 monthly observations and 39 annual observations per condition given the datasets at hand.

Visually this method can be thought of as placing a fixed width window on the sorted range of c, and computing the correlation between the joint observations of x, y that are observed inside this window. To understand how the correlation varies with c, $\rho(\theta, \epsilon)$ is simply calculated for the full range of θ. An example of this is shown in Figure 1, where the grey horizontal bars indicate the "window" of large-cap stock performance that was conditioned on for each observation.

Using a conditioning bias visualisation technique borrowed from Page and Panariello, 2018, the red dashed line shows the expected result if the assets were drawn from a bivariate normal distribution with means, variances and correlations equal to their whole-of-sample values.

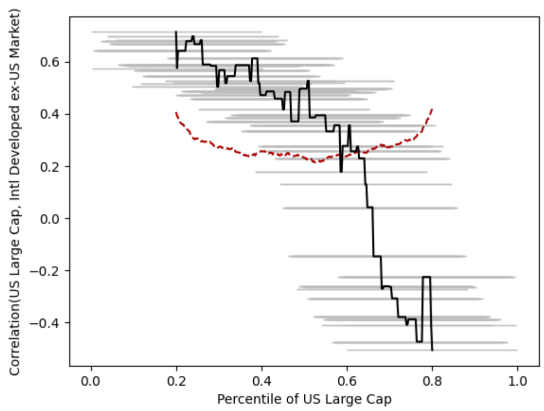

Figure 1: Sliding window of correlations between US large-cap stocks and global ex-US equities, conditional on the performance of US large-cap stocks. Annual Data, 1972-2024.

The key methodological difference in this work is the use of a constant-size sliding window rather than a one-sided percentile cutoff. This has the advantages of excluding outliers for the "middle" percentiles, using a constant sample size for all percentiles, and providing easily interoperable results, with the cost of adding noise to the results due to the smaller sample size in the "middle" percentiles.

This sliding window approach would be a good candidate to use the exponential weighting method developed by Page and Panariello, 2018, by replacing the fixed window with weights that decay as they get further from θ. This was avoided in this case to allow for the results be more easily interpreted, but in more rigorous use cases this would be a logical improvement.

3 Asset classes as Diversifiers for Large-Cap US Equities

As in prior work, the diversification benefit from a "diversifying" asset class is analysed using the correlation between asset classes, conditional on the performance of a "primary" asset class. For this analysis we assume that the primary asset class is large-cap US Equities, but this could easily be completed for any other primary asset class.

The results are broadly divided into three categories: "Good" asset classes where the assets behaved somewhat like the hypothesised "ideal" diversifier or had correlations that behaves as expected under a bivariate normal assumption; "Bad" asset classes that are poor diversifiers when analysed from a conditional correlation perspective; and "Ugly" cases where the results were too noisy to draw meaningful conclusions.

3.1 The good...

As outlined in the introduction, our definition of an ideal diversifier is one where the conditional correlation is negative when the primary asset performs poorly, and positive when the primary asset performs well. Visually, this is an upward sloping line, with the correlation increasing with increasing primary asset performance.

Very few asset classes showed this result, with only short and medium term treasuries and government bonds having convincing results, as in Figure 2 and Figure 3.

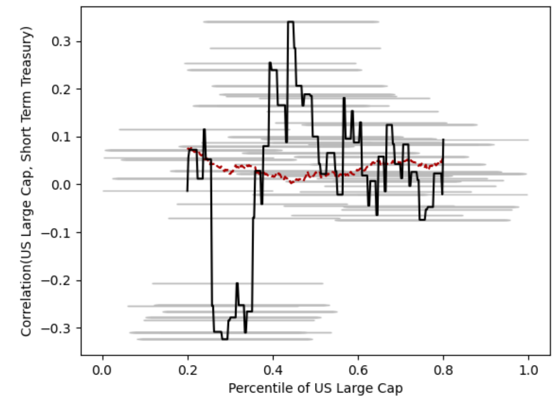

Figure 2: Sliding window of correlations between US large-cap stocks and short term treasuries, conditional on the performance of US large-cap stocks. Annual Data, 1972-2024.

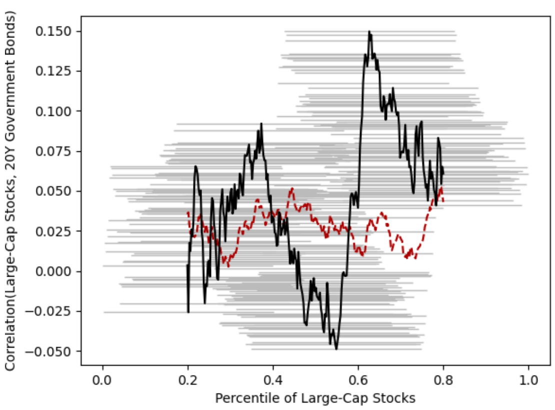

Figure 3: Sliding window of correlations between US large-cap stocks and 20Y government bonds, conditional on the performance of US large-cap stocks. Monthly Data, 1926-2023.

There are also a few assets that follow the behavior that their sample correlation would imply, with the observed conditional correlations seeming to match the results from the simulated bivariate normal distribution. These include small caps, when measured at a monthly frequency, as in Figure 4.

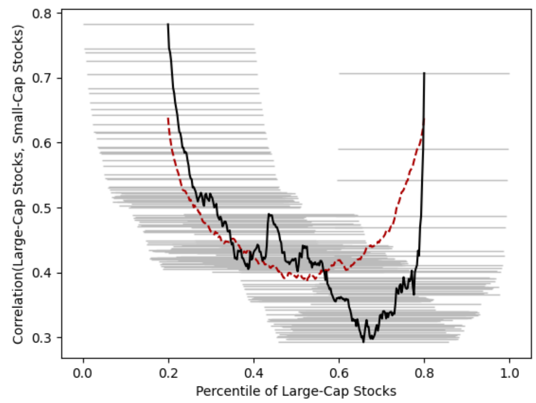

Figure 4: Sliding window of correlations between US large and small cap stocks, conditional on the performance of US large-cap stocks. Monthly Data, 1926-2023.

3.2 The bad...

Many more asset classes showed a downward sloping curve, which by our definition makes them a bad diversifier, being positively correlated with losses and negatively correlated with gains. Some notable examples include:

International equities, both ex-US equities and emerging markets, as shown in Figure 1 and Figure 5.

Corporate Bonds, High Yield Credit as in Figure 6

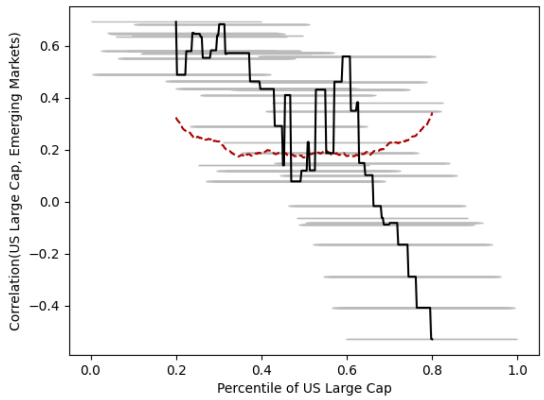

Figure 5: Sliding window of correlations between US large-cap stocks and emerging markets equities, conditional on the performance of US large-cap stocks. Annual Data, 1972-2024.

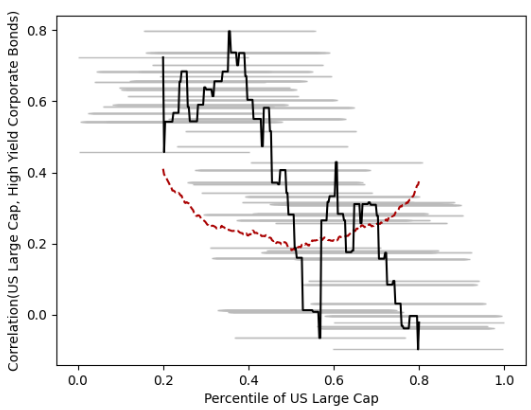

Figure 6: Sliding window of correlations between US large-cap stocks and high yield corporate bonds, conditional on the performance of US large-cap stocks. Annual Data, 1972-2024.

3.3 ...And the ugly

The remaining asset classes don't show any convincing pattern. These include REITs, cash, short-term investment-grade corporate bonds. It is entirely possible that the result from these asset classes are pure sampling noise[1], with the simulated bivariate normal distribution results showing similar patterns when the simulation sample size is dropped to the observed sample size.

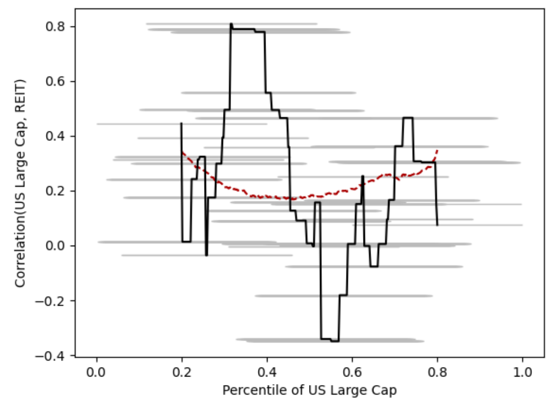

Figure 7: Sliding window of correlations between US large-cap stocks and REITs, conditional on the performance of US large-cap stocks. Annual Data, 1972-2024.

4 Comparison with previous work

The results found with the sliding window estimator are less decisive than the results found in previous work, which deserves some investigation to answer if this is a function of the data or methodology used. The approach from Page and Panariello, 2018 was implemented as in Equation 2[2].

A good example of this is long-term treasuries - The previous methodology suggests that this as-

[1] Arguably the same could be said for many of these results.

[2] This implementation is without the exponential weighting of observations that is used in the original paper, which allows for closer comparison to the sliding window method but does not significantly alter results.

set class is an ideal diversifier, with results in Figure 8 showing a convincing upwards slope from -1 to 1. This same result is not clear in the sling window methodology, as in Figure 9. It is worth noting that the same correlations are reported on both graphs in two points: The ends of the sliding window graph include the tails, and the corresponding samples are measured at the 40^{th} and 60^{th} percentiles in the tail correlation graph.

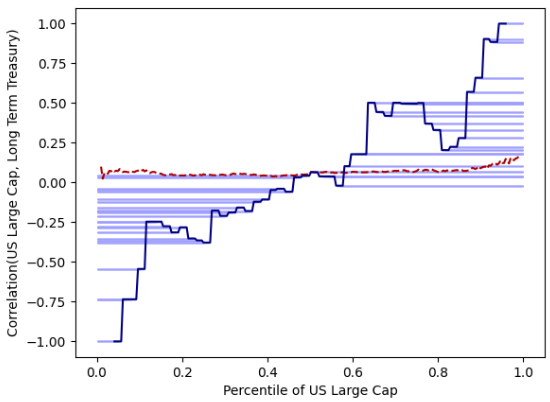

Figure 8: Tail correlations between US large-cap stocks and long-term treasuries, conditional on the performance of US large-cap stocks, with method from Equation 2. Annual Data, 1972-2024.

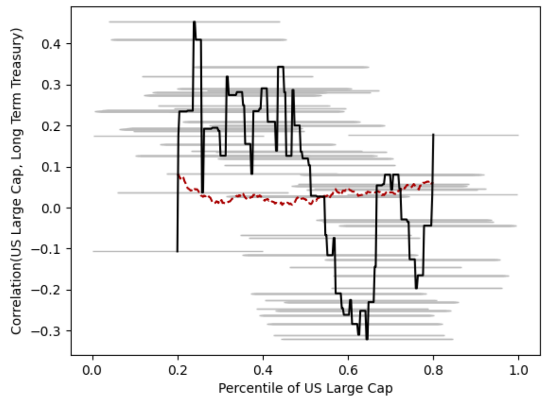

Figure 9: Sliding window of correlations between US large-cap stocks and long-term treasuries, conditional on the performance of US large-cap stocks. Annual Data, 1972-2024.

The main differences between the sliding window approach and the methods used in previous papers are:

Inclusion of tails: The previous methods always calculate correlations from the cutoff to the end of the tail, resulting in samples that always include far tail observations. In the method introduced in this paper, these tail observations are excluded from windows that are not at each end of the distribution.

Variable sample size: The previous methods use a sample size that decreases as the cutoff is pushed further into the tails, while the window approach ensures a constant sample size is used.

This is not to say that excluding the tails is the "right" approach. After all, tail events are the most impactful to a portfolio and are the types of events that investors seek diversification from. However, using extremely small sample sizes to infer information on future correlations is likely to be influenced by a small handful of observations.

5 Correlations Conditional on Macroeconomic State

By conditioning on a third variable, the moving window estimator can be used to understand how the correlation between asset classes varies based on the macroeconomic environment.

5.1 Stocks, Bonds and Interest Rates

Does the stock-bond correlation change with interest rates? This can be explored in Figure 10, which suggests that the correlation between short-term treasuries and stocks increases with increasing interest rates. This same pattern is also seen for intermediate and long term treasuries. Interestingly this relationship is strongest in the middle of the ranges of windows, when the tails are excluded from the correlation calculations.

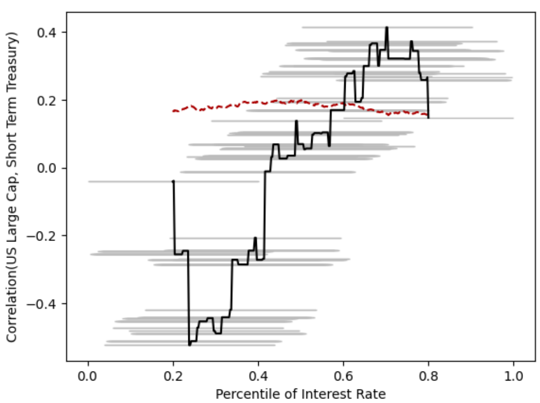

Figure 10: Sliding window of correlations between US large-cap stocks and short term treasuries, conditional on interest rates. Annual Data, 1972-2024.

The opposite pattern is observed for global bonds, where their correlation with stocks appears to decrease as US interest rate increases. Given the persistence of interest rates and their general downward trend in the observed data, its entirely possible that these relationships are structural changes to the correlations over time, rather than a relationship with interest rates specifically[3].

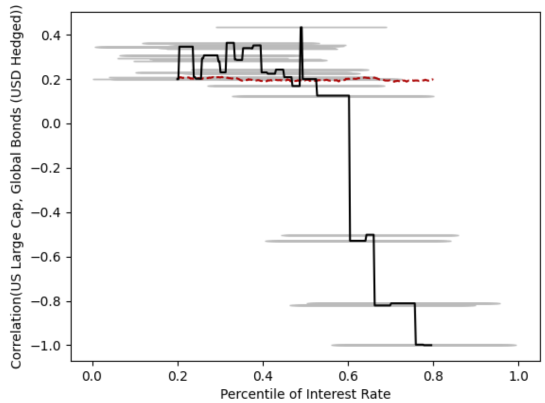

Figure 11: Sliding window of correlations between US large-cap stocks and global bonds, conditional on interest rates. Annual Data, 1972-2024.

5.2 Stocks, Bonds and Change in Interest Rates

Often the direction of interest rate movement is of more interest than the absolute level of rates. This shows a similar pattern to the pre-differencing analysis, as in Figure 12, where treasuries are more correlated with stocks when rates are increasing, however, the magnitude of the increase is smaller and visually appears to be noisier.

[3]See Figure 13 for the stock-bond correlation ”conditional on date”, which is just a time series.

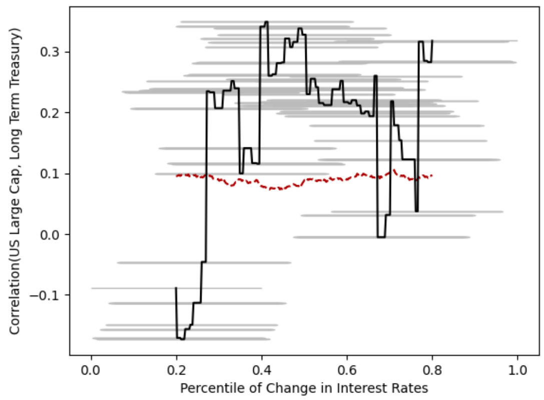

Figure 12: Sliding window of correlations between US large-cap stocks and long term treasuries, conditional on change in interest rates. Annual Data, 1972-2024.

6 Conclusions

This investigation aimed to explore the well-trodden path of conditional correlations between asset classes through the perspective of a simplified estimator to confirm the existing findings and explore how macro variables may also impact correlations. The results largely align with existing literature, including replicating the findings that treasuries are good diversifiers and have desirable conditional correlation properties, while equity asset classes and corporate bonds show poor conditional correlation properties. However, a number of other asset classes that were also previously seen as having poor conditional correlations, such as REITs, show less decisive results when the tail observations are cut out in the window estimator. This suggests that far tail observations may be driving large proportion of the results when using the other estimators.

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

D.B. Chua, M. K., & Page, S. (2009). The myth of diversification. *Journal of Portfolio Management*.

Leibowitz, M., & Bova, A. (2009). Diversification performance and stress-betas. *The Journal of Portfolio Management*.

Page, S., & Panariello, R. (2018). When diversification fails. *Financial Analysts Journal*.

7 Appendix

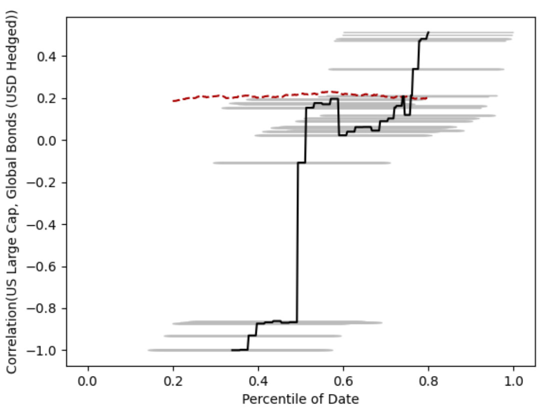

Figure 13: Sliding window of correlations between US large-cap stocks and global bonds, conditional on date. Annual Data, 1972-2024.