

The Journal of **Alternative** *Investments*

Hedging and Constructing Portfolios of Active Strategies: *Strategy Time Horizon and Estimation Errors*

Alexander Rudin and William M. Marr

JAI 2018, 21 (1) 39-47

doi: <https://doi.org/10.3905/jai.2018.21.1.039>

<http://jai.iijournals.com/content/21/1/39>

This information is current as of July 5, 2018.

Email Alerts Receive free email-alerts when new articles cite this article. Sign up at:
<http://jai.iijournals.com/alerts>

Institutional Investor Journals

1120 Avenue of the Americas, 6th floor,
New York, NY 10036, Phone: +1 212-224-3589

© 2017 Institutional Investor LLC. All Rights Reserved

IIJ Institutional
Investor
Journals

Downloaded from <http://jai.iijournals.com/> by guest on July 5, 2018

Hedging and Constructing Portfolios of Active Strategies: *Strategy Time Horizon and Estimation Errors*

ALEXANDER RUDIN AND WILLIAM M. MARR

ALEXANDER RUDIN is a managing director at State Street Global Advisors in Stamford, CT. alexander_rudin@ssga.com

WILLIAM M. MARR is a senior managing director and global head of the Alternative Investments Strategies Group at State Street Global Advisors in Stamford, CT. william_marr@ssga.com

There are many portfolio management tasks where forming views on future correlation of assets is essential. Hedging unwanted systematic exposures is one such task; portfolio construction is another. It is well established (see, for example, Lehmann and Casella (1998) and Evans, Hastings, and Peacock (1993)) that to the extent our views on future correlations are obtained through analyzing finite historical samples, they are associated with estimation errors. Judging the magnitude of such estimation errors is therefore a key component of the investment process. If those errors happen to exceed the estimated values themselves, we may be compelled to disregard historical estimates as unreliable for all practical purposes and resort to postulated (assumed) values of correlations between assets (for example, by assuming that correlations are zero). Acting on “noise” may lead to erroneous hedging decisions or sub-optimal portfolio performance out of sample; both are outcomes that a prudent investor would prefer to avoid. Conversely, if estimation errors are deemed to be low, we can view historical correlation estimates as being “significant beyond reasonable doubt” and can use them to guide our actions.

A recent article by Engle, Focardi, and Fabozzi (2016) reviewed the closely related topic of factor-based modeling in investment management and addressed how the

number of factor variables interplays with sample lengths, including statistical methods for forecasting covariance matrices. The topic is complex and deeply technical, especially for dynamic or “active” strategies whose risk exposures evolve over time. In the present article, we discuss specific challenges associated with forecasting correlations between active strategies and suggest a simple heuristic for estimating errors associated with such forecasts. We also conduct a few numerical experiments to illustrate our intuition and provide a practical framework for its implementation.

The main difference between active strategies—hedge funds, risk premia, etc. and more traditional assets is the dynamic nature of their risk exposure. For most of these strategies, this dynamic nature is associated with some sort of Average Time Horizon, whether it is a moving momentum calculation window for a trend following strategy, or the frequency of reported earnings for an equity value premium. Our main contribution lies in observing that for active strategies with Average Time Horizon \mathcal{L} extending over multiple data points within the historical sample ($\mathcal{L} \gg 1$) (but smaller than the available sample size, L), the number of independent observations in that sample is closer to L/\mathcal{L} than to L . This leads to the correlation estimation error scaling down with sample size as $1/\sqrt{L/\mathcal{L}}$, which is much

slower than the $1/\sqrt{L}$ scaling that traditional thinking may suggest. This fact leads to significant modifications in how we hedge and construct portfolios of active strategies based on historical correlation estimates.

We will start with a hedging problem when an investor wants to remove systematic exposure to a “market” from a given strategy, we evaluate an optimal hedge ratio for such procedure, and decide whether hedging is warranted. A t -statistic reflecting our intuition is suggested.

We will then embark on a broader discussion of a portfolio construction problem. A new “litmus test” is suggested for determining whether a correlation matrix across multiple active strategies may be effectively estimated from a historical sample.

IS AN “OPTIMAL HEDGE RATIO” ALWAYS OPTIMAL?

Hedging of unwanted systematic exposure is often an essential part of an investment strategy. The process typically involves determining the appropriate hedge ratio β for the strategy and then, for each \$1 of exposure to a strategy, also establishing $-\beta$ times \$1 position in whatever asset an investor is attempting to hedge with. Broad equity market exposure is often hedged, as is exposure to interest rates, for example.

A traditional approach to determining an “optimal” hedge ratio is an attempt to minimize hedged portfolio variance. If r_i are strategy returns and x_i are “market” returns, an exposure that we are trying to hedge, the goal of a traditional hedging strategy is to find such hedge ratio β that minimizes variance of the hedged portfolio $r_i - \beta x_i$. Calculation is straightforward and yields “optimal” hedge ratio of

$$\beta^* = \rho_{xr} \frac{\sigma_r}{\sigma_x} \quad (1)$$

where ρ_{xr} is a correlation coefficient between the strategy returns and market returns and σ_r and σ_x are strategy and market volatilities respectively.

But is such a hedge ratio always optimal? In most cases, investors understand “optimality” in the sense of maximizing risk-adjusted returns, for example, as expressed through a Sharpe ratio. That may or may not yield the same answer as minimizing volatility.

The Sharpe ratio of the hedged strategy where the hedge ratio is given by equation (1) is

$$SR_H = SR_r \frac{1 - \rho_{xr} \left(\frac{SR_x}{SR_r} \right)}{\sqrt{1 - \rho_{xr}^2}} \quad (2)$$

Here $SR_r = \mu_r / \sigma_r$ is the Sharpe ratio of the unhedged strategy and SR_x is the same ratio for the “market.” Elementary analysis shows that SR_H is not always higher than the Sharpe ratio of the unhedged strategy SR_r . Instead, assuming that all Sharpe ratios are positive and correlation ρ_{xr} is positive as well (which provides the original inclination to hedge), hedging according to equation (1) improves the risk-adjusted returns of the unhedged strategy ($SR_H > SR_r$) only if

$$\rho_{xr} > \frac{2S}{1 + S^2}, \quad (3)$$

where $S = SR_x / SR_r$. The result is easy to interpret. When hedging, you replace a portfolio where you are long the strategy with a portfolio where you are long the strategy and short the hedge. Improvement of a Sharpe ratio associated with such replacement depends on both the correlation between the strategy and the hedge and on their standalone risk-adjusted returns. Equation (3) effectively tells us that hedging according to equation (1) improves risk-adjusted returns only if the correlation coefficient is large enough; otherwise it leads to a Sharpe ratio reduction and may be unwarranted.

HEDGING AMID ESTIMATION UNCERTAINTY

Equations (1) and (3) would provide a very straightforward guide to hedging if not for the elephant in the room—estimation uncertainty. We never really know precisely the values of volatilities, correlations, or Sharpe ratios. We could either infer them from some exogenous considerations, or estimate them from the historical data available to us.

Estimating errors associated with the process of calculating statistics from *finite* samples is a well-established area of mathematics. While the details depend on what specific measure we are trying to analyze, it has been shown that for the normally distributed random

variables, estimation errors generally decrease inversely proportionally to the square root of the sample size L (see, e.g., Lehmann and Casella 1998 and Evans, Hastings, and Peacock 1993). For mean and standard deviation, estimation errors are approximated by:

$$Error(\hat{\mu}) = \sqrt{Var\{\hat{\mu}\}} = \frac{\hat{\sigma}}{\sqrt{L}} \quad (4a)$$

$$Error(\hat{\sigma}) = \sqrt{Var\{\hat{\sigma}\}} \approx \frac{\hat{\sigma}}{\sqrt{2 * (L - 1)}} \quad (4b)$$

For the correlation coefficient, estimation error (for relatively large sample with $L \gg 1$) is also inversely proportional to the square root of the sample size after the Fisher transformation is applied:

$$Error(\hat{z}) = \sqrt{Var\{\hat{z}\}} \approx \frac{1}{\sqrt{L}},$$

$$\text{where } \hat{z} = \frac{1}{2} \ln \left(\frac{1 + \hat{\rho}}{1 - \hat{\rho}} \right). \quad (5a)$$

We could translate equation (5a) into the estimate for $Error(\hat{\rho})$ observing that for the relatively small errors, $Error(\hat{\rho}) \approx \frac{\partial \hat{\rho}}{\partial \hat{z}} Error(\hat{z})$. Expressing $\hat{\rho}$ through \hat{z} and taking the derivative, we arrive at the following estimate for the correlation sampling error:

$$Error(\hat{\rho}) = \sqrt{Var\{\hat{\rho}\}} \approx \left(\frac{1}{\cosh(\hat{z})} \right)^2 \frac{1}{\sqrt{L}}. \quad (5b)$$

Here $\hat{\mu}$, $\hat{\sigma}$, $\hat{\rho}$ are sample estimates obtained by analyzing the historical sample.

What are the practical implications of this for hedging in light of equation (3)? The first observation we would like to make is that precise estimation of Sharpe ratios from historical samples is extremely difficult. Expected returns for the financial assets (and strategies) tend to be small as compared to their volatilities, at least over moderate time horizons. At the same time, according to equation (4a), the *uncertainty* in estimating those expected returns is determined by these comparably large volatilities. As a result, most investors rely on exogenous (a.k.a. “assumed”) values

for Sharpe ratios for strategies (and markets) based on their economic validity or empirical evidence of similar strategies (just like the expected Sharpe ratio of a trend following strategy with a given time horizon on a single market is often estimated by simultaneously analyzing such strategies across many time horizons and hundreds of markets). The good news is that here we are trying to estimate the *ratio* of Sharpe ratios of the market and the strategy and for most realistic cases we can assume such a ratio to be below one.

That leaves the correlation coefficient as the only variable worthy of estimating from historical data for the purposes of including it in equation (2). The traditional way of making decisions based on levels of “noisy” variables is to build a t -statistic reflecting the extent to which the value of the noisy variable—in this case correlation—rises above the noise levels. Equation (3) leads to the following form of a t -statistic appropriate for our hedging problem:

$$t\text{-stat} \stackrel{\text{def}}{=} \frac{\hat{\rho} - \frac{2S}{1+S^2}}{Error(\hat{\rho})}. \quad (6)$$

Here S is our view on the ratio of risk-adjusted returns between the “market” and the “strategy” and $Error(\hat{\rho})$ is the estimation error for correlation. In the absence of other considerations, we would expect this error to be well estimated by equation (5b), with L being the number of data points in the historical sample.

Will this method work equally well for both traditional assets and *active* strategies? As we will show in the next section, it does work reasonably well for traditional assets. At the same time, using unmodified equations (5b)–(6) for active strategies can *significantly underestimate estimation errors*. As it turns out, determining the appropriate “size” for a given historical sample is not as straightforward in the case of active strategies as it may seem to be at first.

DEFINING THE “SAMPLE SIZE” WHEN THE STRATEGY IS ACTIVE

Let us begin by taking a traditional asset—the DAX Index (“DAX”)—and attempt to measure its correlation to the “market” defined as S&P 500 Index (“SPX”). In our numerical experiment we took 10 years of historical data between 2007 and 2016, in the form of both a daily

time series and a weekly one. We then created a number of “samples” by calculating the correlations between the DAX and the SPX over moving windows of size L ranging from 10 to 50 and we studied the average correlation across the samples and the dispersion of that correlation, as an empirical way to gauge estimation error. Notably, for the daily time series L ranging “between 10 and 50” meant 10 to 50 *days*, while for the weekly time series the same L meant 10 to 50 *weeks*. Exhibit 1 presents results of this experiment.

What can we infer from Exhibit 1? Both weekly and daily data-based estimation errors indeed reduce as sample size increases, in accordance with the statistical theory. Also, albeit imperfect, equation (5b) provides an important base line in our error estimation process. This is quite helpful because it implies that whenever studies of the type described in Exhibit 1 are not feasible (e.g., due to lack of data) it is sufficient to count number of data points in the sample and equation (5b) would yield a reasonable estimation error for the correlation coefficient and help determine whether hedging is indeed warranted.

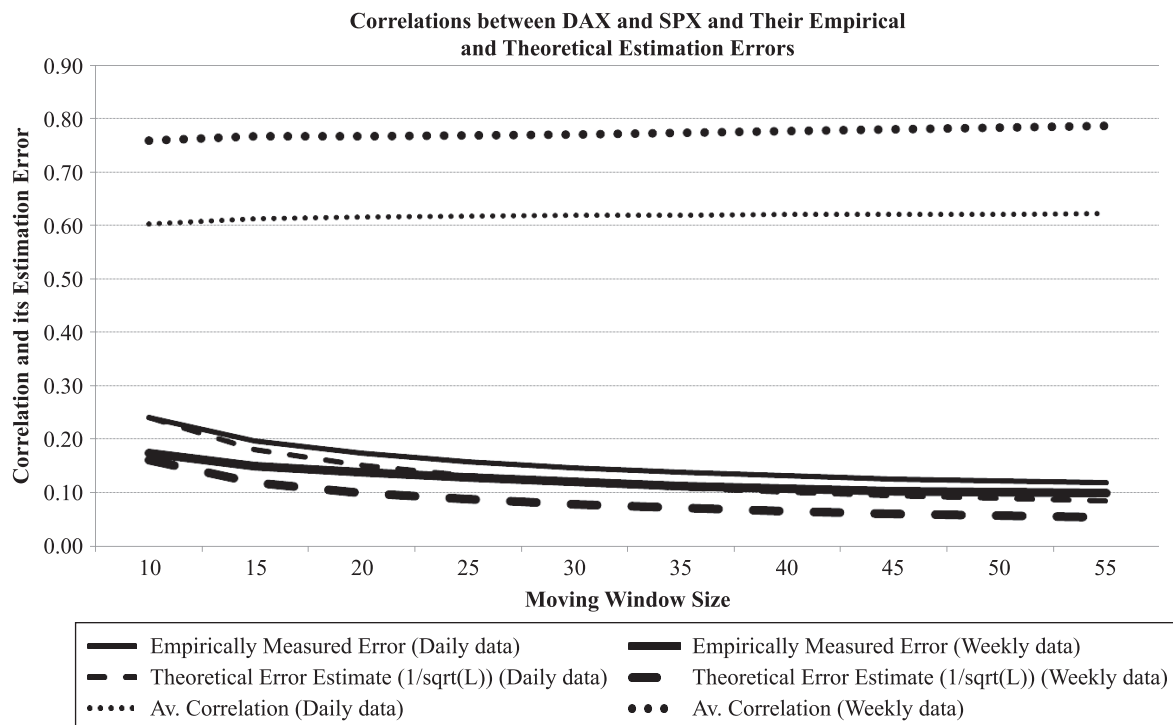
Our final observation from Exhibit 1 is that both empirically observed and theoretically estimated errors for correlations are well below average sample correlations for daily and weekly data. In other words, the correlation between the DAX and the SPX seems to indeed exist “beyond reasonable doubt.” If we were deciding whether to hedge DAX exposure with SPX, we could use our views on the relative Sharpe ratios of DAX and SPX and equation (3) as a guide in helping us to make a decision.

Let us now replace our buy and hold position in the DAX index with a simple *active* strategy applied to that index. We define “active” strategies as those where the portfolio changes its composition (and hence risk characteristics) over time. Let’s use a very simple trend following strategy formulated as follows:

- At the end of every time period t we measure the average DAX index return for the moving window that starts at $t - Z + 1$ and ends at t .
- If that average return is positive, then for the $t + 1$ period we hold a long \$1 position in the DAX

EXHIBIT 1

The DAX and the SPX



index; if it is negative—we hold a short \$1 position in the DAX index.

We will name such strategy “DAX Trend (Z)”. Let us choose value of $Z = 30$ (results are similar for a broad range of Z s), apply it to the same data set as used above, and recreate our earlier study of estimation error. Results are shown in Exhibit 2.

The results have changed dramatically when compared with a buy and hold strategy. “Naïve” $1/\sqrt{L}$ approach (with L being the number of historical data points) to estimation error leads us to the conclusion that correlation, at least for the weekly data sample, is significant and hence hedging may be warranted. At the same time, such a conclusion is not supported by empirically observed errors, which are much higher and imply that hedging is unwarranted.

The contradiction can be resolved if we recall that equations (5a–b) were derived under the assumption that each observation in the sample is *independent* from the others (this is why the L used in those equations is often called “number of degrees of freedom” as opposed to the “sample size”). By simultaneously

using data over a certain time horizon (Z) to make decisions at a given point in time, an active strategy like DAX Trend (Z) reduces the effective number of independent data points in the sample from the sample size L to the sample size divided by the average time horizon of the strategy, which we will denote as \mathcal{L} . The longer the time horizon of the strategy is, the longer the average trade takes, the smaller number of independent observations we actually have in our sample, and hence the higher estimation error we should be prepared to face.


We can now combine our new understanding of equations (5a–b) and (6) into the first result of this article:

$$t\text{-stat} \stackrel{\text{def}}{=} \frac{\hat{\rho} - \frac{2S}{1+S^2}}{\text{Modified Error}(\hat{\rho})} \quad (7a)$$

$$\text{Modified Error}(\hat{\rho}) = \left(\frac{1}{\cosh(\hat{z})} \right)^2 \frac{1}{\sqrt{L/\mathcal{L}}}, \quad (7b)$$

EXHIBIT 2

DAX Trend (30) and SPX



where \mathcal{L} is the Average Time Horizon of the active strategy. We need to replace sample length L with the effective sample length $\tilde{L} = L/\mathcal{L}$ before using the traditional error estimation formulae in order to incorporate the active nature of the investment strategy.

CALCULATING MODIFIED ERRORS USING THE VOLATILITY-BASED TURNOVER CONCEPT

For the trend following strategy described above, it is easy to link Average Time Horizon \mathcal{L} to Z . How about a more general case? Sometimes \mathcal{L} may be estimated by looking at the strategy formulation, time horizon of signals, and the history of strategy turnover. If we are fortunate to have a full history of portfolio weights for all strategies (or full description of those strategies allowing us to reproduce those weights), we should be able to estimate the Average Time Horizon consistently across strategies. One way to do that¹ is to utilize the Volatility-Based Portfolio Turnover concept introduced recently by Gnedenko and Yelnik (2016).

The Volatility-Based Portfolio Turnover was introduced to provide a consistent framework for determining how actively we are changing the risk profile of the portfolio, while executing an active strategy. It ensures that certain intuitive preferences regarding turnover estimation hold true. For example, we would often prefer trading activity in more risky assets within a portfolio to contribute more to the turnover than trading in its less volatile components. Or we would expect that replacing an asset with a similar one (a roll of futures contracts near the expiration is a good example of such a situation) would only minimally contribute to the turnover.

For an active strategy involving a single asset (we kept notation consistent with earlier sections of our paper), Gnedenko and Yelnik's definition of the Volatility-Based Portfolio Turnover is

$$volPT \stackrel{\text{def}}{=} \frac{1}{L} \sum_{i=1}^L \sigma_i |w_i - w_{i-1}|. \quad (8)$$

¹The authors are grateful to Igor Yelnik for pointing this out and for bringing the Gnedenko and Yelnik (2016) article to their attention.

Here $\{w_i\}$ are historical weights and σ_i is volatility for the asset at t_i .

In other words, Volatility-Based Portfolio Turnover is defined as the averaged sum of volatilities corresponding to the *difference portfolios* (this measure is easy to annualize or use for any other “standard” time interval). In the general case of a multi-asset portfolio, expression transforms into

$$volPT \stackrel{\text{def}}{=} \frac{1}{L} \sum_{i=1}^L \sqrt{\Delta_i^T \Sigma_i \Delta_i} \quad (9)$$

where $\Delta_i = w_i - w_{i-1}$ is the difference portfolio at time t_i and Σ_i is the covariance matrix.

Gnedenko and Yelnik (2016) also introduced a concept of Effective Number of Trades (“ENT”) per unit of time as $volPT$ normalized by an average portfolio risk $\bar{\gamma} = 1/L \sqrt{\sum_{i=1}^L w_i^T \Sigma_i w_i}$.

A possible analytical framework for evaluating Average Time Horizon \mathcal{L} for a generic active strategy is to calculate it as an inverse of the Effective Number of Trades over a “unit” time period of our historical data sample (1 day or 1 week):

$$\mathcal{L} = \frac{\sqrt{\sum_{i=1}^L w_i^T \Sigma_i w_i}}{\sum_{i=1}^L \sqrt{\Delta_i^T \Sigma_i \Delta_i}} \quad (10)$$

Let us use this definition—equation (10)—to calculate \mathcal{L} for the DAX Trend (Z) strategy. Having accomplished that exercise, we would put ourselves in a position to evaluate correlation estimation error using equation (7b). Exhibit 3 summarizes the result (we only included results where $\tilde{L} = L/\mathcal{L} \gg 1$).

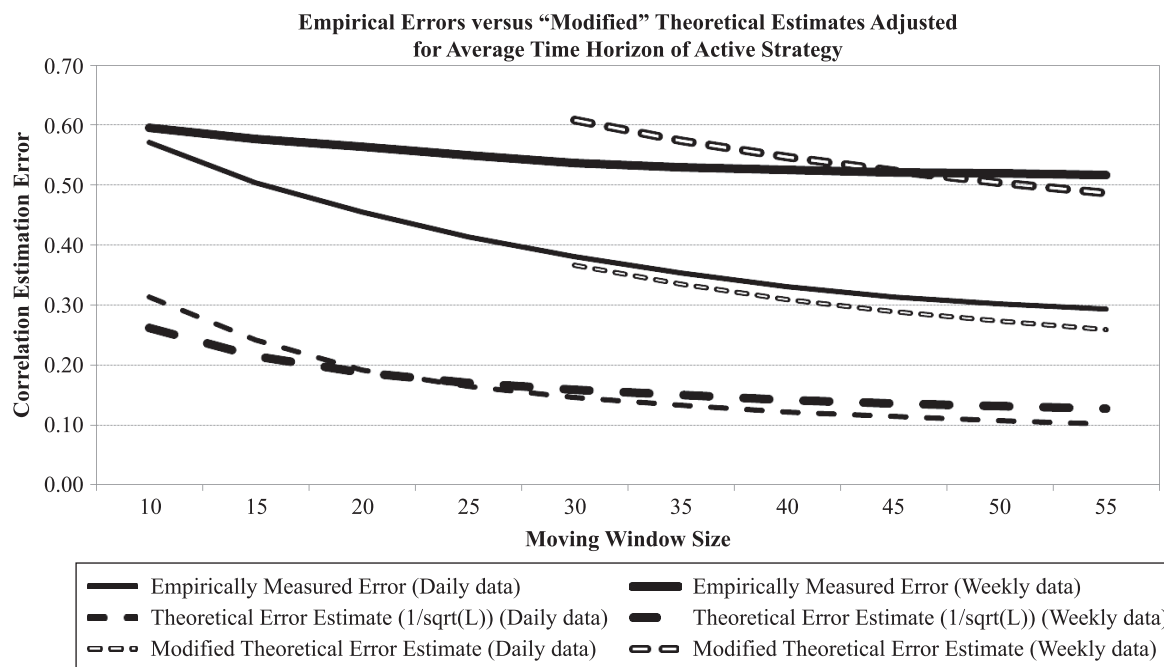
We can observe that using equation (7b) in combination with equation (10) yields estimation errors forecasts that are in good accordance with the empirical results validating our intuition.

THE ROLE OF ESTIMATION ERRORS IN PORTFOLIO CONSTRUCTION

So far we have looked at correlation estimation errors from the angle of hedging unwanted “market” exposure for a single strategy. Let us now comment

EXHIBIT 3

Modified Estimates



on a more general problem of portfolio construction, as determining correlations or other “interactions” between strategies is an essential step of that problem. Building reliable estimators for a correlation matrix based on historical, finite sample data has been studied by a wide variety of researchers and practitioners, who have provided many insights over time on how to “clean” or “shrink” those matrices in the attempt to reduce data fitting in portfolio construction and improve out of sample results. Providing a thorough review of these methods is beyond the scope of this article; Haft (1980) introduced correlation matrix shrinkage for the first time, while Ledoit and Wolf (2004), Bun, Bouchaud, and Potters (2016), and Engle, Focardi, and Fabozzi (2016) provide excellent early and recent reviews of further research advances in this subject.

Statistical researchers may vary in ways they “clean” correlation matrices, but they agree on two principles:

- (1) It is generally a bad idea for the portfolio construction process to use correlation matrices directly calculated using historical data, as such an approach introduces systematic errors that may

lead to portfolios with grossly underestimated out-of-sample risk.

- (2) While for a single correlation coefficient, the quality of a historical estimate is driven by the criteria $L \gg 1$, where L is the number of data points in a historical sample, for the portfolio optimization problem this criteria changes to $L/N \gg 1$, with N being the number of assets in the portfolio. If L/N ratio is not sufficiently high the, “signal” contained in the historical sample becomes effectively indistinguishable from “noise” from the estimation error rendering historical correlation matrices useless.

Earlier in this article we illustrated that for active strategies with Average Time Horizon \mathcal{L} we need to replace sample length L with the effective sample length $\tilde{L} = L/\mathcal{L}$ before placing it into the error estimation formulae. The same intuition drives us to the conclusion that the criteria we should use in determining whether historical correlation matrices can be indeed fruitfully used for portfolio optimization purposes should be modified in a similar fashion and become

$$\frac{\tilde{L}}{N} = \frac{L}{\mathcal{L}N} \gg 1 \quad (11)$$

What are the practical implications of this conclusion for the construction of portfolios of active strategies? For construction schemes that incorporate the realized correlation matrix with a weight depending on how confident the investor is in such matrix, a t -statistic incorporating equation (11) could be designed to serve as such a weight. For correlation matrix cleaning techniques based on random matrix analysis—a.k.a. eigenvalue “clipping”—the modified ratio introduced in equation (11) could be used for calculating the eigenvalue cut-off value.

How material are modifications implied by equation (11) for portfolios of hedge funds or risk premia? Let us start with the case of hedge funds. A typical history available to a hedge fund allocator is 5–10 years of data, while the typical number of managers in a diversified hedge fund portfolio is between 10 and 30. Taking a middle of each range, we arrive at $L/N \sim 90/20 = 4.5$, not much higher than 1 even before adjusting for the Average Time Horizon of the underlying hedge fund strategies. Taking into account that many hedge funds execute active strategies with time horizons longer than a single month (long-term trend following, systematic macro, and many of the credit strategies, for example), equation (11) would argue strongly against using portfolio construction schemes for hedge fund portfolios that rely heavily on historical correlation matrices.

What about “alternative risk premia” portfolios that have received so much attention in recent years? Such portfolios tend to be much more transparent to investors, with more historical data often being available, including daily and weekly performance going back for many years. That would increase the sample size from 60–120 data points (for 5–10 years) to thousands of data points, seemingly providing enough data for the robust historical correlation analysis and subsequent portfolio optimization. Unfortunately, the intuition incorporated into equation (11) drives us to a much more conservative conclusion. Most risk premia strategies are characterized by low risk turnover and long time horizons, making $\tilde{L} = L/\mathcal{L} \ll L$. A risk premia strategy like PPP-based currency value counts its typical trade length in years no matter whether its performance is supplied on a

monthly, weekly, or daily basis, or even every minute. The *effective* number of independent data points in the 10-year sample will still be just a few points, making the historical correlation matrix involving this strategy and similar strategies highly noisy and ineffective for the portfolio optimization purposes.

CONCLUDING REMARKS

An important comment worth making is that the situation when the condition of equation (11) is not met or the t -statistic (7a–b) is small does not necessarily mean a complete lack of correlations. It merely means that such correlations cannot be discovered through the statistical analysis of historical samples. We can still *postulate* non-zero correlations from structural considerations when we have a reason to believe they exist. An obvious but still key example of this situation is when an active strategy is put on top of a stable, non-cash benchmark. In this case, considerations above only apply to the excess return of the strategy, while correlation of the benchmark return stream to other assets in the portfolio should, of course, be fully incorporated into the analysis.

A somewhat related situation is one where we *know* that an otherwise active strategy maintains a definitive and persistent bias. For example, if an Equity Long/Short hedge fund maintains a long bias of 30%, it is completely legitimate to use $\beta_0 = 0.3$ as a base line value for the hedge ratio and *then* use the considerations outlined above to gauge if we can confidently affirm (using t -statistic (7)) that this hedge fund has an equity market long bias above or below β_0 .

Another useful example is a well-known FX carry strategy where an investor builds a relative value portfolio of currencies, going long those with higher prevailing interest rates and short those with lower prevailing rates. FX carry trades are likely to have a left-tail correlation with equity indexes irrespective of the specific composition of the carry portfolio at a given point in time. Investors in low-yielding countries are taking a risk (going against their home country bias) to invest in a higher yielding country, and will likely pull that capital back when equities collapse and they get nervous. One can build a model for such behavior that implies non-zero correlations, even if purely statistical analysis does not demonstrate a sufficient level of confidence. It is the universal behavioral bias that creates the correlation

structure in this case rather than anything intrinsic to a specific set of assets or a given historical data sample.

In many situations involving active strategies, it is difficult to either build reliable structural models leading to definitive expected correlations, or to infer those correlations from historical data. Alternative investments practitioners have long recognized this, which could help explain popularity of non-optimization portfolio construction methods where correlation takes a distinctly secondary role, including various forms of risk parity (see Da Silva, Lee, and Pornrojnangkool 2008; Lee 2011; Asness, Frazzini, and Pedersen 2012; Rudin and Marr 2016; for reviews of the approach and additional literature). Our present findings endorse this choice when it comes to active strategies characterized by either low frequency of available historical data, or long time horizons for trading, or both.

The authors are grateful to John Dolfen, Geoffrey Kelley, James Lewis, and Igor Yelnik for valuable discussions.

REFERENCES

- Asness, C., A. Frazzini, and L. Pedersen. 2012. "Leverage Aversion and Risk Parity." *Financial Analysts Journal* 68 (1): 47–59.
- Bun, J., J.P. Bouchaud, and M. Potters. 2016. "Cleaning Correlation Matrices." *Risk.net*, April.
- Da Silva, A., W. Lee, and B. Pornrojnangkool. 2008. "Risk Budgeting." In *Handbook of Finance, Vol. 2, Investment Management and Financial Management*, edited by F. Fabozzi, 195–220. Hoboken, NJ: John Wiley & Sons.
- Engle, R. F., S. M. Focardi, and F. J. Fabozzi. 2016. "Issues in Applying Financial Econometrics to Factor-Based Modeling in Investment Management." *The Journal of Portfolio Management* 42 (5): 94–106.
- Evans, M., N. Hastings, and B. Peacock. 1993. *Statistical Distributions*, NY: Wiley.
- Gnedenko, B., and I. Yelnik. 2016. "Speed and Dimensions of Trading." SSRN Archive 2745278, March.
- Haft, L. 1980. "Empirical Bayes Estimation of the Multivariate Normal Covariance Matrix." *Annals of Statistics* 8:586–597.
- Ledoit, O., and M. Wolf. 2004. "Honey, I Shrunk the Sample Covariance Matrix." *The Journal of Portfolio Management* 31, no. 1 (Fall): 110–119.
- Lee, W. 2011. "Risk-Based Asset Allocation: A New Answer to an Old Question?" *The Journal of Portfolio Management* 37 (4): 11–28.
- Lehmann, E. L., and G. Casella. 1998. *Theory of Point Estimation*, New York, NY: Springer Verlag.
- Rudin, A., and W. M. Marr. 2016. "Investor Views, Drawdown-Based Risk Parity, and Hedge Fund Portfolio Construction." *The Journal of Alternative Investments* 19 (2): 63–69.

To order reprints of this article, please contact David Rowe at drowe@ijjournals.com or 212-224-3045.