

FULL-SCALE CURRENCY HEDGING

THIS VERSION: December 21, 2023

Megan Czasonis, Mark Kritzman, and David Turkington

Megan Czasonis is a managing director at State Street Associates in Cambridge, MA.

mczasonis@statestreet.com

140 Mt Auburn Street, Cambridge MA, 02138

617-234-9576

Mark Kritzman is the chief executive officer at Windham Capital Management in Boston, MA, and a senior lecturer at the MIT Sloan School of Management in Cambridge, MA.

kritzman@mit.edu

100 Main Street, Cambridge MA, 02142

617-901-4455

David Turkington is senior managing director and head of State Street Associates in Cambridge, MA.

dturkington@statestreet.com

140 Mt Auburn Street, Cambridge MA, 02138

617-234-9484

Key Takeaways

- The efficacy of currency hedging depends on the co-occurrence of the cumulative returns of currency forward contracts with the underlying portfolio and with each other over the investor's horizon.
- Most investors rely on beta-based hedging, which works well only if correlations estimated from a specific return sample at a specific frequency give a good approximation of the co-occurrences of cumulative returns over the investor's prospective horizon.
- Investors should consider full-scale hedging, which captures the empirical distribution of co-occurrences, as an alternative to beta-based hedging or mean-variance optimization, which relies on summary statistics.

Abstract

After years of spirited debate, most investors agree that to minimize the risk currencies add to a portfolio they should hedge its currency exposures based on its betas relative to the currencies to which it is exposed. However, this notion of hedging makes sense only if the betas reliably reflect the co-occurrences of the cumulative returns of the portfolio and currencies over the investor's horizon. And this will be true only if the correlations of currencies with the portfolio and with each other are constant across the return intervals used to estimate them and stationary through time. Neither of these conditions holds empirically. The authors propose a new currency hedging technique called full-scale hedging, which explicitly considers the full distribution of horizon co-occurrences.

FULL-SCALE CURRENCY HEDGING

It is widely accepted that to minimize the risk currencies add to a portfolio, investors should sell currency forward contracts in amounts, expressed as a percentage of total portfolio value, equal to the portfolio's betas relative to the currencies to which it is exposed. Although beta-based hedging is both intuitive and mathematically obvious, its implementation is far from straightforward. The extent to which currency hedging reduces portfolio risk depends on the co-movement of the cumulative returns of the currencies and portfolio over the investor's horizon, which can extend to several years. Investors typically estimate betas from monthly returns, assuming implicitly that correlations are invariant to the return interval used to estimate them, which is hardly true. If, instead, investors estimate betas from return intervals of equal duration to their hedging horizon, they implicitly assume that the average co-movement estimated from these longer return intervals gives a good approximation of co-movement during the investor's horizon. This too is unlikely. Given these estimation challenges, we propose an alternative approach to currency hedging called full-scale hedging.

We proceed by first describing beta-based hedging. Then we introduce the notion of co-occurrence, followed by a discussion of the challenges of using correlation to estimate co-occurrence. We then introduce full-scale hedging and discuss how it addresses the challenges of beta-based hedging. We conclude with a summary.

Beta-based hedging

To see why beta-based hedging is thought to minimize portfolio risk, consider a toy example in which there is a portfolio exposed to a single currency. We assume that the portfolio is fixed but the investor can vary exposure to a currency forward contract. The variance of this portfolio is given by Equation 1.

$$\sigma_H^2 = \sigma_U^2 + \sigma_F^2 \times W^2 + 2 \times \rho_{U,F} \times \sigma_U \times \sigma_F \times W \quad (1)$$

In Equation 1, σ_H equals the standard deviation of the hedged portfolio, σ_U equals the standard deviation of the unhedged portfolio, σ_F equals the standard deviation of the currency forward contract, W equals the weight of the currency forward contract, and $\rho_{U,F}$ equals the correlation between the unhedged portfolio and the currency forward contract. We assume that these values are based on returns denominated in the base currency of the investor as opposed to local returns.

To solve for the risk minimizing exposure to the currency forward contract, we take the partial derivative of portfolio variance with respect to the weight of the currency forward contract and set this quantity equal to zero, as shown.

$$\partial \sigma_H^2 / \partial W = 2 \times \sigma_F^2 \times W + 2 \times \rho_{U,F} \times \sigma_U \times \sigma_F = 0 \quad (2)$$

$$W = -\rho_{U,F} \times \sigma_U / \sigma_F = -\beta \quad (3)$$

Equations 2 and 3 reveal that the risk minimizing hedge ratio is equal to the negative of a portfolio's beta with respect to the currency forward contract. We determine the risk minimizing hedge ratios for a portfolio that is exposed to more than one currency by regressing

its returns on the returns of the respective currency forward contracts. The negative of the beta coefficients from this regression are the risk minimizing hedge ratios.

This beta-based solution implicitly assumes that investors are willing to sell short currency forward contracts in amounts greater than the fraction of the portfolio allocated to the corresponding foreign assets or to purchase currency forward contracts and thereby increase exposure to a currency. In practice, however, investors typically sell currency forward contracts only up to the amount allocated to the corresponding foreign asset, and they do not purchase currency forward contracts to increase exposure to a currency.

Also, investors sometimes consider a currency forward contract's expected return, including hedging costs which can be thought of as a positive expected return (assuming the forward contract is sold), when deciding what fraction of currency exposure to hedge. In these situations, investors seek to maximize expected utility rather than minimize risk. Equations 4, 5, and 6 derive the exposure to a currency forward contract that maximizes expected utility, considering both risk reduction and expected return for a portfolio that is exposed to a single currency.

$$E(U) = \mu_U + \mu_F \times W - \lambda \times \sigma_U^2 - \lambda \times \sigma_F^2 \times W^2 - \lambda \times 2 \times \rho_{U,F} \times \sigma_U \times \sigma_F \times W \quad (4)$$

$$\partial \sigma_H^2 / \partial W = \mu_F - 2 \times \lambda \times \sigma_F^2 \times W - 2 \times \lambda \times \rho_{U,F} \times \sigma_U \times \sigma_F = 0 \quad (5)$$

$$W = \mu_F / (2 \times \lambda \times \sigma_F^2) - \rho_{U,F} \times \sigma_U / \sigma_F \quad (6)$$

In Equations 4, 5, and 6, $E(U)$ equals expected utility, μ_U equals the expected return of the unhedged portfolio, μ_F equals the expected return of the currency forward contract, and λ equals the investor's risk aversion coefficient, which quantifies the amount of expected return the investor is willing to sacrifice in exchange for one unit of risk reduction.

For a portfolio that is exposed to more than a single currency, we employ mean-variance optimization to determine the optimal hedge ratios if the currency forward contracts have non-zero expected returns or if the exposures to the currency forward contracts are constrained, as is the custom.

It is also important to note that it would be more efficient to solve for the asset weights and currency forward contract positions simultaneously, though this is seldom done. Investors typically choose the underlying assets as a first step to constructing a portfolio and then choose what fraction of the currency exposure to hedge. Although this two-step approach is mathematically inefficient, it has practical merit to the extent investors have greater confidence forecasting asset class performance than they do forecasting the performance of currency forward contracts. Investors can tie asset class valuations to relatively predictable cash flows whereas currency forward contract returns often depend on less predictable factors such as geopolitical influences.

Both beta-based hedging and mean-variance optimization depend crucially on the correlation of the portfolio with the currency forward contracts and the currency forward contracts with each other, but these correlations are useful only if they give a good approximation of how the portfolio and currency forward contracts co-move cumulatively over the investor's horizon. This cumulative co-movement is called co-occurrence.

Co-occurrence

Co-occurrence measures the co-movement of the cumulative returns of a pair of assets over an investor's horizon, which is what determines the efficacy of currency hedging. It is important to keep in mind, that unlike combining two assets that both have positive weights, in which we case we would seek low co-occurrence, we are selling short a currency forward contract to hedge currency risk. Therefore, the currency forward contract's hedging potential is stronger to the extent it co-occurs positively with the underlying portfolio.

Consider, for example, the co-movement of the returns of a portfolio and a currency forward contract, which we call variables X_P and X_C , which both have observations for time periods $i = 1, 2, \dots, N$. We wish to measure co-occurrence for a single observation of i which represents one period. Using the sample means, \bar{x}_P and \bar{x}_C , and the sample standard deviations, σ_P and σ_C , we convert each observation into a standardized z-score as follows:

$$Z_{i,P} = \frac{x_{i,P} - \bar{x}_P}{\sigma_P} \quad (7)$$

$$Z_{i,C} = \frac{x_{i,C} - \bar{x}_C}{\sigma_C} \quad (8)$$

The co-occurrence of X_P and X_C for observation i is defined as:

$$c_i(P, C) = \frac{Z_{i,P} Z_{i,C}}{\frac{1}{2}(Z_{i,P}^2 + Z_{i,C}^2)} \quad (9)$$

This measure of co-occurrence has the following desirable properties, which allows us to interpret it as a pure measure of the point-in-time alignment of an observation of two variables.

- The highest value is +1, which occurs when both assets move by the same extent in the same direction.
- The lowest value is -1, which occurs when both assets move by the same extent in opposite directions.
- The value is zero if either asset has a z-score of zero.¹
- The value may equal any number between -1 and +1, indicating the extent of alignment.
- The value indicates direction and not extent: any points that lie along a line through the origin (a scalar multiple of $z_{i,P}$ and $z_{i,C}$) have the same co-occurrence.

We must also define the joint informativeness of an observation of X_P and X_C as shown by Equation 10:

$$info_i(P, C) = \frac{1}{2}(z_{i,P}^2 + z_{i,C}^2) \quad (10)$$

This perspective enables us to view the traditional full-sample Pearson correlation coefficient as a weighted average of the co-occurrence of each observation, in which each observation's weight equals its informativeness as a fraction of the total informativeness of the sample.

$$\rho(P, C) = \sum_{i=1}^N w_i c_i(P, C) \quad (11)$$

$$w_i = \frac{info_i(P, C)}{\sum_{k=1}^N info_k(P, C)} \quad (12)$$

We place greater weight on observations of co-occurrence that come from large magnitude returns because these returns convey more information than observations of small magnitude returns, which mostly reflect noise.

To summarize, co-occurrence measures the alignment of the cumulative returns of a pair of assets in each investment period. When applied to the alignment between a portfolio and a currency forward contract, it determines the efficacy of hedging that occurs in that period. Correlation is equal to a weighted average of co-occurrence across all such periods in a data sample, where each observation's weight is equal to the informativeness (joint magnitude) of its returns.

Our foregoing discussion of the connection of co-occurrence to correlation assumes implicitly that we estimate correlations from return intervals that correspond to the length of the investor's hedging horizon. However, investors with multi-year horizons do not typically estimate correlations this way. Typically, investors estimate correlations from shorter return intervals and assume they are invariant to the return interval used to estimate them. Or they estimate correlations from longer horizon returns that match the duration of their hedging horizon. However, neither approach gives a reliable estimate of co-occurrence during an investor's prospective hedging horizon, because correlations estimated from one return interval diverge from correlations estimated from a different return interval, even for the same period, and correlations estimated from return intervals that match the duration of the investor's hedging horizon will likely differ substantially from the co-occurrence that obtains during the investor's hedging horizon owing to sampling error. We next discuss these two estimation challenges.²

Divergence

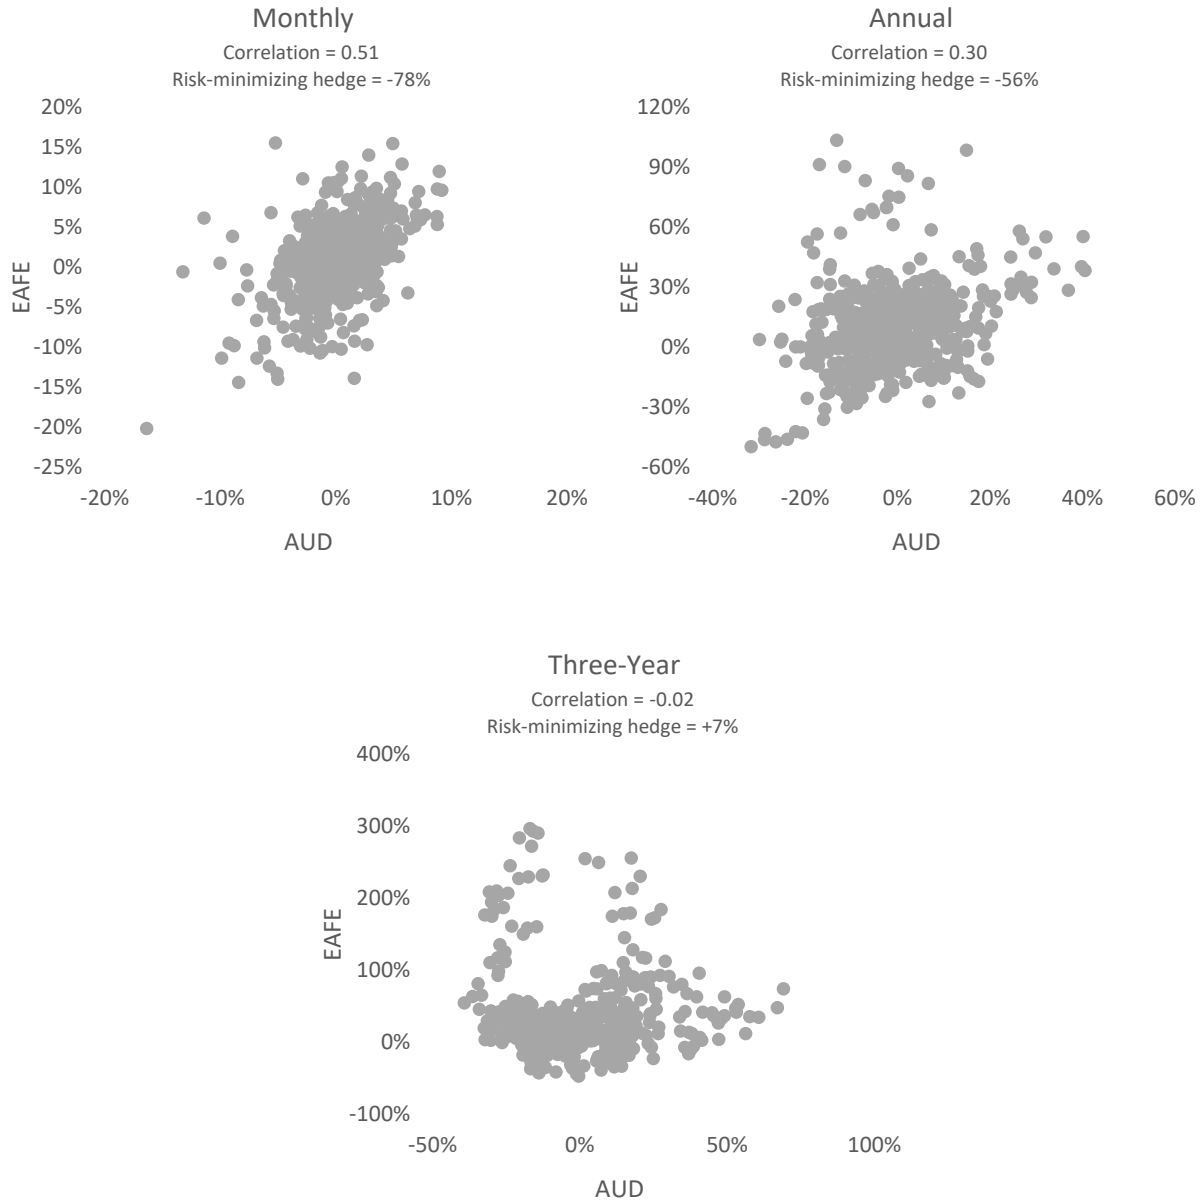
Divergence refers to the notion that correlations estimated from a given return interval will differ from correlations estimated from a different return interval, even for the same measurement period, if either asset's autocorrelations or their lagged cross-correlations are non-zero at any lag.³

Exhibit 1 confirms, in dramatic fashion, that divergence is not merely a hypothetical phenomenon. It shows scatter plots of MSCI EAFE and AUD returns (both denominated in USD) for monthly, annual, and three-year return intervals over the same period beginning in 1980. The correlation of monthly returns was 0.51, and the standard deviations of EAFE and AUD returns were 17.1% and 11.2% respectively, implying that an investor would minimize the risk that exposure to the AUD adds to the portfolio by selling an AUD forward contract equal to 78% of the portfolio's value. If, instead, an investor uses annual returns to estimate correlation and standard deviations, the risk minimizing hedge ratio would equal 56% of the portfolio's value. And based on three-year returns, AUD exposure actually reduced the volatility of the EAFE portfolio, implying the minimum risk hedge ratio would be a long forward position in AUD in an amount equal to 7% of the portfolio's value.

Exhibit 1: Scatter Plot of MSCI EAFE and AUD Returns (USD Base)

Monthly, Annual, and Three-Year Intervals

(January 1980-November 2023)



As we mentioned, the divergence of estimates of standard deviations and correlations across different return intervals occurs either because returns have non-zero autocorrelations

or lagged cross correlations at one or more lags. Equations 13 and 14 show explicitly how non-zero lagged correlations affect the relationship between high and low frequency standard deviations and correlations, respectively. These calculations assume that the instantaneous rates of return of the assets are normally distributed with stationary means and variances.

The standard deviation of the cumulative continuous returns of x over q periods, $x_t + \dots + x_{t+q-1}$, is given by Equation 13.

$$\sigma(x_t + \dots + x_{t+q-1}) = \sigma_x \sqrt{q + 2 \sum_{k=1}^{q-1} (q-k) \rho_{x_t, x_{t+k}}} \quad (13)$$

In Equation 13, σ_x is the standard deviation of x measured over single-period intervals. Note that if the lagged autocorrelations of x all equal zero, the standard deviation of x will scale with the square root of the horizon, q .

Now we introduce a second asset, y , whose continuously compounded rate of return over the period $t - 1$ to t is denoted y_t . The correlation between the cumulative returns of x and the cumulative returns of y over q periods, is given by Equation 14.

$$\rho(x_t + \dots + x_{t+q-1}, y_t + \dots + y_{t+q-1}) = \frac{q \rho_{x_t, y_t} + \sum_{k=1}^{q-1} (q-k) (\rho_{x_{t+k}, y_t} + \rho_{x_t, y_{t+k}})}{\sqrt{q + 2 \sum_{k=1}^{q-1} (q-k) \rho_{x_t, x_{t+k}}} \sqrt{q + 2 \sum_{k=1}^{q-1} (q-k) \rho_{y_t, y_{t+k}}}} \quad (14)$$

The numerator equals the covariance of the assets taking lagged cross-correlations into account, whereas the denominator equals the product of the assets' standard deviations as described by Equation 13. This equation allows us to assume values for the autocorrelations of x and y , as well as the lagged cross-correlations between x and y , to compute the correlations and standard deviations that these parameters imply over longer horizons. It would be quite


challenging to estimate all these autocorrelations and lagged cross-correlations but, as we show later, is unnecessary given our proposed method for hedging currency exposure.

Sampling error

Given the divergence of standard deviations and correlations estimated from different return intervals, we might consider estimating these parameters from the return interval that matches the duration of the investor's hedging horizon. This approach is problematic, however, because it assumes implicitly that the full sample parameters give the best estimates of the minimum risk hedge ratio for a specific prospective horizon. It fails to consider the dispersion of longer horizon co-occurrences throughout the estimation sample or any asymmetry or other features of the distribution of co-occurrences throughout the sample. It assumes that there is only one co-occurrence that matters - the average co-occurrence, which may have never occurred historically.

Exhibit 2 illustrates how co-occurrences can vary dramatically over time. It shows the time series of three-year co-occurrences of MSCI EAFE and AUD (denominated in USD). Though their Pearson correlation was close to zero over this period (-0.02), there were some three-year periods, such as the early 2000s, when their cumulative returns moved in tandem, and other three-year periods, such as the mid-1980s, when their returns were highly divergent.

Exhibit 2: Three-Year Co-occurrences Overlapping Monthly
(January 1980-November 2023)



We have described how the efficacy of currency hedging depends on the co-occurrence of the underlying portfolio with the currency forward contracts over the investor's hedging horizon. We have also explained and offered evidence of why it is unwise to use correlations estimated either from short interval returns or return intervals that match the duration of the hedging horizon. For these reasons, beta-based hedging or its extension, mean-variance optimization, will likely fail to identify the appropriate hedge ratios. We therefore propose an alternative approach to currency hedging called full-scale hedging, which effectively addresses these estimation challenges.

Full-scale hedging

Beta-based hedging and mean-variance optimization are heuristics that yield approximations of the optimal in sample hedge ratios, even before one considers sampling error. These approximations rely on parameters such as correlation, that summarize a return distribution. By contrast, full-scale hedging explicitly considers the returns for every period in the sample for the portfolio and the currency forward contracts. It therefore considers the full distribution of co-occurrences of the portfolio with the currency forward contracts and with each other explicitly as well as other features of the portfolio and currency forward contract returns. We implement full-scale hedging as follows.⁴

1. We construct a sample of overlapping returns, calculated over the same interval as our hedging horizon, for the portfolio and the currency forward contracts used to hedge the portfolio's currency risk.
2. We re-mean the currency forward contract returns to have average returns of 0.0%.⁵
3. We specify a utility function that exhibits aversion to loss.
4. While holding fixed the portfolio's weight at 100%, we compute the total return of a hedged portfolio given weights for the currency forward contracts that range from -100% of the portfolio's exposure to the corresponding currency to 0%, for every period in our sample.
5. We convert the hedged portfolio return into utility for every period.
6. We add the utilities across all the periods and store this value.
7. We repeat this process for as many combinations of currency forward contract weights needed to identify the weights that minimize risk for the hedged portfolio.

Some explanation of this process may be helpful. If our hedging horizon is three years, for example, we could use cumulative three-year returns that overlap monthly. By overlapping the returns, our result is less sensitive to the start date of our sample. We re-mean the currency forward contract returns to equal 0.0%, so that full-scale hedging minimizes risk without regard to return. If we were to optimize exposure to currency forward contracts based on their returns as well as risk, we would need to consider more carefully the choice of the utility function. We fix the portfolio's weight to equal 100% so that the solution depends only on the amount of currency exposure hedged and not on the composition of the portfolio. By restricting the currency forward contract weights to range from -100% of the currency exposure to 0, we preclude hedging more than the portfolio's exposures to the currencies or adding currency exposure beyond the portfolio's embedded exposure.

Exhibit 3 shows a stylized full-scale hedging example in which we consider a portfolio that has a 30% exposure to the Euro and a 20% exposure to the Yen. Our sample comprises 20 years of cumulative three-year returns that overlap monthly. We have 35 fewer periods than 240 because beyond the 205th three-year period we can no longer observe full three-year periods. The first trial considers a fully hedged portfolio. The sum of its utilities equals 30.1266. This number assumes that all the three-year returns are considered though we only show four of them. The second trial hedges less of the portfolio's Euro exposure, which produces a slightly higher utility. We proceed in this fashion iteratively considering different exposures to the Euro and the Yen until we have maximized utility.

Exhibit 3: Full-Scale Hedging

First Trial									
Periods	Asset Class Returns			Portfolio Weights			Portfolio Return	Utility Function	Utility
	Portfolio	Euro	Yen	Portfolio	Euro	Yen			
1st Three Years	18.15%	-10.31%	-5.25%	100%	-30%	-20%	22.29%	U = ln(W)	0.2012
2nd Three Years	25.03%	2.04%	-3.33%	100%	-30%	-20%	25.08%	U = ln(W)	0.2238
3rd Three Years	-12.67%	7.59%	8.14%	100%	-30%	-20%	-16.58%	U = ln(W)	-0.1812
↓									
205th Three Years	45.20%	5.90%	11.88%	100%	-30%	-20%	41.06%	U = ln(W)	0.3440
Sum of 205 three-year utilities									30.1266
Second Trial									
Periods	Asset Class Returns			Portfolio Weights			Portfolio Return	Utility Function	Utility
	Portfolio	Euro	Yen	Portfolio	Euro	Yen			
1st Three Years	18.15%	-10.31%	-5.25%	100%	-25%	-20%	21.78%	U = ln(W)	0.1970
2nd Three Years	25.03%	2.04%	-3.33%	100%	-25%	-20%	25.19%	U = ln(W)	0.2246
3rd Three Years	-12.67%	7.59%	8.14%	100%	-25%	-20%	-16.20%	U = ln(W)	-0.1767
↓									
205th Three Years	45.20%	5.90%	11.88%	100%	-25%	-20%	41.35%	U = ln(W)	0.3461
Sum of 205 three-year utilities									30.2916

We now illustrate full-scale hedging for a 60/40 foreign equity/bond portfolio from a USD base.⁶ We seek to minimize risk⁷ over a three-year horizon, relying on a sample of rolling, three-year returns over the period January 1995 through November 2023. The right column of Exhibit 4 reports the full-scale hedge ratios for the portfolio's five largest currency exposures. For comparison, the left column shows risk minimizing hedge ratios based on the conventional practice of estimating betas from monthly returns over the same period.⁸ In the bottom panel, we show the volatility and downside risk of three-year returns for the hedged portfolio assuming the respective beta-based and full-scale hedge ratios.

Exhibit 4: Risk Minimizing Hedge Ratios for 60/40 Foreign Equity/Bond Portfolio (USD Base)

Beta-Based Hedging versus Full-Scale Hedging

(January 1995-November 2023)

	Beta Monthly	Full-Scale 3-Year
Hedge Ratios		
EUR	100%	100%
JPY	34%	0%
GBP	100%	100%
CHF	100%	0%
AUD	100%	100%
3-Year Statistics of Hedged Portfolio		
Standard Deviation	19.5%	18.9%
10% Worst	-8.4%	-5.5%

Compared to the beta-based approach, full-scale hedging produces different hedge ratios that provide greater risk reduction, both in terms of volatility and left-tail risk. This is because the beta-based hedge ratios implicitly assume that the full sample correlation of monthly returns gives a reliable estimate of the unknown co-occurrences that will prevail during the investor's hedging horizon. However, this would be true only if the autocorrelations and lagged cross-correlations of the higher frequency returns from which the betas are estimated were equal to zero at all lags, or if the betas were estimated from cumulative multi-horizon returns that matched the investor's hedging horizon and these returns were normally distributed. Full-scale hedging considers all the features of the data including the dispersion of co-occurrences as well as non-normalities and outliers.

Summary

We first showed conceptually why beta-based hedging is thought to minimize the risk currencies add to a portfolio. We then introduced the notion of co-occurrence which captures the co-movement of cumulative returns over an investor's hedging horizon and therefore determines the efficacy of hedging. We described the linkage of co-occurrence and correlation and discussed the conditions by which correlation, and by extension, beta give a good approximation of horizon co-occurrence. We then offered evidence that these required conditions do not hold empirically. We proposed, as an alternative to beta-based hedging, full-scale hedging which explicitly considers the entire distribution of co-occurrences, including features such as dispersion and asymmetry. We then computed hedge ratios based on beta-based hedging and full-scale hedging given the same historical period. Our analysis revealed that full-scale hedging produces significantly different solutions than beta-based hedging by accounting for nuances that summary statistics fail to consider.

Notes

This material is for informational purposes only. The views expressed in this material are the views of the authors, are provided “as-is” at the time of first publication, are not intended for distribution to any person or entity in any jurisdiction where such distribution or use would be contrary to applicable law and are not an offer or solicitation to buy or sell securities or any product. The views expressed do not necessarily represent the views of Windham Capital Management, State Street Global Markets®, or State Street Corporation® and its affiliates.

References

Aruda, N., A. Bergeron, and M. Kritzman. “Optimal Currency Hedging: Horizon Matters.” *The Journal of Alternative Investments*, Vol. 23, Issue 4, March 2021.

Kinlaw, W., M. Kritzman, and D. Turkington. “The Divergence of High- and Low-Frequency Estimation: Causes and Consequences.” *The Journal of Portfolio Management*, Vol. 40, Issue 5, Special 40th Anniversary Issue, 2014.

Kinlaw, W., M. Kritzman, and D. Turkington. “Co-occurrence: A New Perspective on Diversification” forthcoming in *The Journal of Portfolio Management*.

Kritzman, M. and D. Turkington. “History, Shocks, and Drifts: A New Approach to Portfolio Formation.” *The Journal of Portfolio Management*, Vol. 48, Issue 3, February 2022.

¹ If both assets have z-scores of zero, the equation is technically undefined, as zero divided by zero. However, in this rare instance we should define co-occurrence to be equal to zero. Its value will not have any influence on further assessments of sample (or subsample) correlation, because as we will soon argue, co-occurrence must be scaled by the informativeness of an observation, which itself equals zero in the case of two zero z-scores.

² See Aruda, Bergeron, and Kritzman (2021) for a comprehensive of the effect of horizon on optimal currency hedging.

³ For a detailed discussion of divergence, see Kinlaw, Kritzman, and Turkington (2014).

⁴ Investors may care about short-term outcomes as well as long term outcomes. In this case, one could augment the return sample to include returns measured over more than one interval. It is critical, though, the sample include all the shorter interval returns that go into the longer interval returns. If the sample, for example, includes three-year returns and monthly returns, there must be 36 times as many monthly returns as three-year returns, even if the three-year returns overlap. For more detail about this multi-horizon approach, see Kritzman and Turkington, 2022.

⁵ If we had convictions about currency returns, we could re-mean the returns to align with these convictions.

⁶ We proxy foreign equities using market capitalization weighted MSCI country equity indices for the EAFE universe excluding Israel (due to data limitations). We proxy foreign bonds using market capitalization weighted FTSE

country government bond indices for the same universe of countries excluding Israel, Hong Kong, and Singapore (due to data limitations). Portfolio weights and currency exposures reflect market values as of November 30, 2023.

⁷ We assume a kinked utility function that assumes power utility with a curvature (risk aversion) parameter of 2 for returns above the kink and a slope penalty of 5 below the kink. However, because we re-mean the currency forward contract returns to equal 0.0%, full-scale hedging minimizes risk without regard to return. If we were to optimize exposure to currency forward contracts based on their returns as well as risk, the choice of utility function would have a greater impact on optimal hedge ratios.

⁸ For simplicity, we describe the approach as beta-based hedging because it minimizes risk. However, in order to constrain the hedge ratios to fall within the range of -100% to 0, we employ mean-variance optimization and assume zero expected returns for the currency forwards used to hedge.