50th

th

year anniversary

Co-Occurrence:
A New Perspective on
Portfolio Diversification

William Kinlaw, Mark Kritzman, and David Turkington

STATE STREET

William Kinlaw, CFA

Will is senior managing director and Head of Research at State Street Markets, the capital markets arm at State Street, where he is responsible for a global team of 60 data scientists, quantitative researchers and economists as well the firm's partnerships with data science startups and renowned academics at State Street Associates. The team delivers proprietary investment indicators, data-driven trade ideas, and award-winning research to more than 1,000 institutional investors around the world through its Insights research platform, mobile app, and podcast.

Will's research has focused on economic forecasting, private capital, portfolio construction, asset allocation, currency hedging, inflation, and risk. He has received six awards for peer-reviewed journal articles and the second edition of his book, "Asset Allocation: From Theory to Practice and Beyond," co-authored with Mark Kritzman and David Turkington, was published in 2021. In 2024, he and his co-authors were awarded the Harry M. Markowitz award by a panel of Nobel laureates in economics for their research on the determinants of inflation. The team was ranked #1 FX Bank for Research in the world by Euromoney magazine in 2023 and 2024. Will serves on the editorial boards of the Journal of Portfolio Management and the Journal of Alternative Investments. He holds an M.S. in finance from Boston College, a B.A. in Economics from Tufts University, and a CFA designation. He joined State Street in 2002.

Mark Kritzman, CFA

Mark Kritzman, CFA is a Founding Partner and CEO of Windham Capital Management, LLC. He is also a Founding Partner of State Street Associates, and he teaches a graduate finance course at the Massachusetts Institute of Technology. Mark served as a Founding Director of the International Securities Exchange and as a Commissioner on the Group Insurance Commission of the Commonwealth of Massachusetts. He has also served on the Advisory Board of the Government Investment Corporation of Singapore (GIC) and the boards of the Institute for Quantitative Research in Finance, The Investment Fund for Foundations, and State Street Associates. He is currently a member of the Advisory Board of the MIT Sloan Finance Group, the Board of Trustees of St. John's University, the Emerging Markets Review, the Journal of Alternative Investments, the Journal of Derivatives, the Journal of Investment Management, where he is Book Review Editor, and The Journal of Portfolio Management. He has written more than 100 articles for peer-reviewed journals and is an author of eight books including Prediction Revisited: The Importance of Observation, Asset Allocation: From Theory to Practice and Beyond, Puzzles of Finance, and The Portable Financial Analyst. Mark won Graham and Dodd scrolls in 1993, 2002, and 2010 the Research Prize from the Institute for Quantitative Investment Research in 1997, the Bernstein Fabozzi/Jacobs Levy Award 10 times, the Roger F. Murray Prize from the Q-Group in 2012, the Peter L. Bernstein Award in 2013 for Best Paper in an Institutional Investor Journal, and the Harry M. Markowitz Award in 2022 and 2023. In 2004, Mark was elected a Batten Fellow at the Darden Graduate School of Business Administration, University of Virginia. Mark has a BS in economics from St. John's University, an MBA with distinction from New York University, and a CFA designation.

David Turkington

David Turkington is Senior Managing Director and Head of State Street Associates, State Street Global Markets' decades-long partnership with renowned academics that produces innovative research on markets and investment strategy. David is a frequent presenter at industry conferences, has published more than 40 research articles in a range of journals, and serves on the editorial board of the Journal of Alternative Investments. He is the co-author of three books including "Asset Allocation: From Theory to Practice and Beyond" and "Prediction Revisited: The Importance of Observation." His published research has received the 2010 Graham and Dodd Scroll Award, five Bernstein-Fabozzi/Jacobs-Levy Outstanding Article Awards, the 2013 Peter L. Bernstein Award for best paper in an Institutional Investor journal, the 2021 and 2023 Roger F. Murray First Prize for outstanding research presented at the Q Group seminars, and the 2022 and 2023 Harry Markowitz awards for best paper in the Journal of Investment Management. David graduated summa cum laude from Tufts University with a BA in mathematics and quantitative economics, and he holds the CFA designation.

Co-Occurrence: A New Perspective on Portfolio Diversification

William Kinlaw, Mark Kritzman, and David Turkington

William Kinlaw
is a senior managing director at State Street in Cambridge, MA.
wbkinlaw@statestreet.com

Mark Kritzman
is CEO at Windham Capital Management in Boston, MA, and a senior lecturer at the MIT Sloan School of Management in Cambridge, MA.
kritzman@mit.edu

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

KEY FINDINGS

- Co-occurrence, which measures the co-movement of cumulative returns of a portfolio's assets over an investor's horizon, gives a better measure of the assets' potential to diversify a portfolio than their correlations.

- The correlations of higher-frequency returns are useful only to the extent they predict co-occurrences during an investor's horizon.

- Correlation has been a poor predictor of co-occurrence historically because it has not been stationary, nor has it been constant across the return intervals used to measure it.

ABSTRACT

Investors seek to diversify a portfolio by combining assets that have low correlations with each other, but a correlation coefficient is a useful guide to diversification only under special conditions that rarely hold in practice. If returns are not serially independent and if they do not conform to a stable multivariate normal distribution, correlation will not provide a reliable guide to the asset co-occurrences that a portfolio is likely to experience. The authors provide evidence that assets frequently violate these assumptions of serial independence, stationarity, and normality. These empirical realities cast serious doubt on the usefulness of full-sample correlations to measure an asset's potential to diversify a portfolio. The authors introduce a measure of co-occurrence as a fundamental building block of diversification. An informativeness weighted average of co-occurrence across a full sample equals correlation, but co-occurrence offers an intuitive and flexible way to assess complexities of co-movement that lie beyond the grasp of any summary statistic. The authors present an alternative technique for diversifying a portfolio that explicitly considers the empirical prevalence of co-occurrences and thus the nonnormality of returns.

arry Markowitz, in his landmark article "Portfolio Selection" (1952), taught investors to diversify a portfolio by combining assets that have low correlations with each other. Unfortunately, a correlation coefficient is a useful guide to diversification only under special conditions that rarely hold in practice. First, returns must be independently distributed over time for correlation to offer a good estimate of how asset returns co-occur cumulatively over an investor's prospective horizon. Second, returns must conform to a normal distribution for a correlation coefficient to characterize the range of possible asset co-occurrences accurately. Both assumptions are at odds with reality. We offer evidence that correlations differ depending on the return interval used to estimate them, which belies the notion that returns are serially independent. Additionally, we show that many assets exhibit significantly different

patterns of co-movement across high- and low-interest-rate regimes and between turbulent and quiescent market conditions and that they vary asymmetrically around return thresholds. This evidence strongly contradicts the assumption that returns are identically distributed according to a normal distribution. Correlation is a summary statistic that averages across many instances of co-occurrence. The empirical realities of returns demand that we view co-occurrence with greater precision.

We proceed by first defining co-occurrence and describing its role as the fundamental building block of correlation. Next, we show how serial dependence in returns undermines the usefulness of short-term correlations as a guide to multi-period co-occurrence. We then provide evidence that co-occurrence varies across regimes and is asymmetric around return thresholds. Finally, we propose a portfolio construction technique that accounts for all these complexities by considering the full spectrum of co-occurrences directly, rather than relying on correlation, which is an oversimplified summary of the diversification properties of assets.

CO-OCCURRENCE

Co-occurrence measures the cumulative co-movement of pairs of assets during a specified horizon, which is what determines diversification. If the cumulative return of a potentially diversifying asset converges to the cumulative returns of the other portfolio assets during the investor's horizon, it provides no diversification regardless of how it comoved over shorter intervals within the investor's horizon or how it comoved on average during prior periods of equal length to the investor's horizon. Therefore, correlation is a useful guide to diversification only if it gives a good estimate of co-occurrence during an investor's prospective horizon.

To be more specific, consider the co-movement of the returns of two different assets, which we call variables X_A and X_B, which both have observations for time periods $i = 1, 2, \ldots, N$. We wish to measure co-occurrence for a single observation of i that represents one time period. Using the sample means, $\overline{x}_A$ and $\overline{x}_B$, and the sample standard deviations, σ_A and σ_B, we convert each observation into a standardized z-score as follows:

$$z_{i,A} = \frac{x_{i,A} - \overline{x}_A}{\sigma_A} \tag{1}$$

$$z_{i,B} = \frac{x_{i,B} - \overline{x}_B}{\sigma_B} \tag{2}$$

The co-occurrence of X_A and X_B for observation i is defined as

$$c_i(A,B) = \frac{z_{i,A} z_{i,B}}{\frac{1}{2}(z_{i,A}^2 + z_{i,B}^2)} \tag{3}$$

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 $z_{i,A}$ and $z_{i,B}$) have the same co-occurrence.

We must also define the joint informativeness of an observation of X_A and X_B as shown by Equation (4):

$$info_i(A,B) = \frac{1}{2}(z_{i,A}^2 + z_{i,B}^2) \tag{4}$$

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.

$$\rho(A,B) = \sum_{i=1}^{N} w_i c_i(A,B) \tag{5}$$

$$w_i = \frac{info_i(A,B)}{\sum_{k=1}^{N} info_k(A,B)} \tag{6}$$

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 random noise.

The equivalence of this definition of correlation with the traditional Pearson formula occurs because $\sum_{k=1}^{N} info_k(A,B) = N - 1$, therefore

$$\rho(A,B) = \sum_{i=1}^{N} w_i c_i(A,B) \tag{7}$$

$$\rho(A,B) = \frac{1}{N-1} \sum_{i=1}^{N} info_i(A,B) c_i(A,B) \tag{8}$$

$$\rho(A,B) = \frac{1}{N-1} \sum_{i=1}^{N} z_{i,A} z_{i,B} \tag{9}$$

$$\rho(A,B) = \frac{1}{\sigma_A \sigma_B} \frac{1}{N-1} \sum_{i=1}^{N} (x_{i,A} - \overline{x}_A)(x_{i,B} - \overline{x}_B) \tag{10}$$

$$\rho(A,B) = \frac{Cov(A,B)}{\sigma_A \sigma_B} \tag{11}$$

---

[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.

This approach to calculating correlation is also useful because it allows us to compute a helpful measure of co-movement for a subsample, Φ. We define a subsample measure of co-movement as a weighted average of the same quantities over a subsample of the observations:

$$\rho_\Phi(A,B) = \sum_{i \in \Phi} w_i c_i(A,B) \tag{12}$$

$$w_i = \frac{info_i(A,B)}{\sum_{k \in \Phi} info_k(A,B)} \tag{13}$$

As a measure of subsample co-movement, $\rho_\Phi(A,B)$ has the following desirable properties:

- The value ranges between −1 and +1, as it does over the full sample.
- The value is defined for any size subsample, even for a single observation, in which case it converges to the observation's co-occurrence.
- If co-occurrence is symmetric around zero, the expected subsample estimate based on positive values will equal that of negative values, and both will equal the full sample correlation.
- The measure is centered around the full-sample mean, as opposed to the subsample mean, which facilitates interpretation of subsample properties in the context of the broader distribution. This is analogous to the way that semivariance is computed around the full sample mean, rather than the mean of downside returns. If it were not, we might erroneously conclude that down-side returns are less volatile than the full sample since they are clustered around a "local" negative mean.

To summarize, co-occurrence measures the alignment of the cumulative returns of a pair of assets in each investment period. It determines the amount of diversification 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. We can compute subsample averages of co-occurrence to study the regime-specific properties of diversification, similar to how we can compute subsample averages of returns to study how they differ from a broader average. Thus, co-occurrence offers an intuitive way to study subsample co-movement, which always lies between −1 and +1 and which converges to co-occurrence when evaluated for a single observation and converges to the cor-relation coefficient when evaluated across every observation in a sample.

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 our investment horizon. We framed our discussion this way for expository convenience. However, investors with multiyear horizons do not typically estimate correlations this way. Typically, they estimate correlations from shorter interval returns and assume they are invariant to the return interval used to estimate them. Therefore, to evaluate the efficacy of correlation as a guide to cumulative co-occurrence, we must consider two broad features of the data: 1) the autocor-relations and lagged cross-correlations of short interval returns, which determines the divergence of co-movement across return intervals; and (2) the nonstationar-ity of co-movement across different regimes and around return thresholds, which determines the nonnormality and asymmetry of co-occurrence.

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 nonzero at any lag.[2]

Consider the following example, which is contrived from 20 return observations of two hypothetical assets.

|	
Monthly correlation	0.50
Autocorrelation of first asset	–0.50
Autocorrelation of second asset	–0.50
Correlation of first asset with lag of second asset	–0.20
Correlation of second asset with lag of first asset	–1.00
Bimonthly correlation	–0.20

Exhibit 1 shows the cumulative returns of these hypothetical assets. It reveals that even though these assets' monthly returns are 50% correlated, they diverge significantly over the full period, resulting in a –0.20 bimonthly correlation.

Exhibit 2 confirms, in dramatic fashion, that divergence is not merely a hypothetical phenomenon. It shows scatter plots of US equity and emerging market equity returns for monthly, annual, and 5-year return intervals over the same 25-year period. Whereas the correlation of monthly returns was 0.66, the correlation of annual returns was 0.49, and the correlation of 5-year returns was –0.08.

Equations (14) and (15) show explicitly how nonzero lagged correlations affect the relationship between high- and low-frequency standard deviations and correlations,

EXHIBIT 1
Cumulative Returns of Hypothetical Assets with Positive Correlation
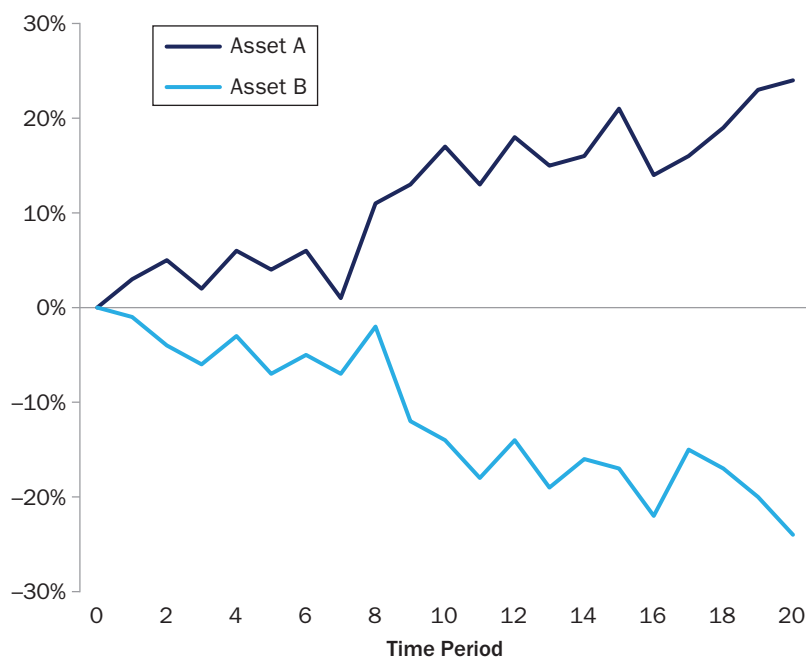
[2] For a detailed discussion of divergence, see Kinlaw, Kritzman, and Turkington (2014).

EXHIBIT 2

Scatter Plot of US Equity and Emerging Market Equity Returns at Monthly, Annual, and Five-Year Intervals (1988–2022)

Monthly Intervals

Annual Intervals

Five-Year Intervals

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, $x_t + \cdots + x_{t+q-1}$, is given by Equation (14).

$$\sigma(x_t + \cdots + x_{t+q-1}) = \sigma_x \sqrt{q + 2\sum_{k=1}^{q-1}(q-k)\rho_{x_t,x_{t+k}}} \qquad (14)$$

In Equation (14), 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 (15).

$$\rho(x_t + \cdots + x_{t+q-1}, \; y_t + \cdots + y_{t+q-1}) =$$

$$\frac{q\rho_{x_t,y_t} + \sum_{k=1}^{q-1}(q-k)(\rho_{x_{t+k},y_t} + \rho_{x_t,y_{t+k}})}{\sqrt{q + 2\sum_{k=1}^{q-1}(q-k)\rho_{x_t,x_{t+k}}} \sqrt{q + 2\sum_{k=1}^{q-1}(q-k)\rho_{y_t,y_{t+k}}}} \qquad (15)$$

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 (14). 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. Of course, it would be quite challenging to estimate all these autocorrelations and lagged cross-correlations, which, as we show later, is unnecessary given our proposed method for constructing portfolios.

We next show that correlations differ across regimes even when they are estimated from the same return intervals.

REGIME DEPENDENCE

Interest Rate Regimes

Perhaps the most important co-occurrence is that of bonds with equities, which history has shown is not stationary, but rather dependent on the level of interest rates. Moreover, this dependency makes sense when we consider how bonds contribute to diversification. When interest rates are high, bonds offer a competitive return to equities, especially on a risk-adjusted basis. During high-interest-rate regimes, bonds offer only modest diversification, but investors are willing to include them in their portfolios for their expected return. When interest rates are low, bonds are not competitive with equities from the standpoint of expected return, whether risk adjusted or not. Their purpose in a low-interest-rate regime is to offset potential equity losses. Therefore, when interest rates are low investors own bonds only if they are significantly negatively correlated with equities. This notion is borne out empirically. When 10-year Treasury yields were above their median of 5.71%, the co-movement of stocks and bonds was 0.47, whereas it was −0.18 when 10-year Treasury yields were below their median.[3]

---

[3] Co-movement is calculated as the monthly co-movement of annual returns and is based on the S&P 500 Total Return Index (from Robert Shiller's website) and the Bloomberg Long US Government Bond Index (from Datastream) from 1974 to 2022. We partition the data into two interest rate regimes using a threshold equal to the median 10-year Treasury yield during this period, which was 5.71%.

Stability Regimes

Co-occurrence also differs depending on whether market conditions are turbulent or quiescent. Whereas it is straightforward to observe the level of interest rates, turbulence is more difficult to measure. Turbulence occurs when disruptive events cause asset returns to depart significantly from their typical pattern of behavior. For example, one or more assets may produce a return that is substantially above or below its average return or a pair of assets that usually move together may drift apart. We use a statistic called the Mahalanobis distance to measure unusual return behavior, as shown in Equation (16).

$$MD = (x_i - \overline{x})\Omega^{-1}(x_i - \overline{x})' \tag{16}$$

In Equation (16), *MD* equals the Mahalanobis distance, x_i is a vector of asset returns, $\overline{x}$ is the average of the asset returns, and Ω^{-1} is the inverse covariance matrix of the asset returns.

The term $(x_i - \overline{x})$ captures the extent to which an asset's return is independently unusual. By multiplying this measure of independent unusualness by Ω^{-1}, the inverse of the covariance matrix, we capture unusual interactions across the pairs of assets, and at the same time, we standardize the values by dividing by variance. We then post-multiply by the transpose of $(x_i - \overline{x})$, to convert the result of all these operations into a single number, which represents the unusualness of a set of returns for a given period such as a day, a month, or a year.

Exhibit 3 shows how patterns of co-occurrence differ between turbulent and quiescent market regimes. In this case, we estimate the Mahalanobis distance, given by Equation (16), for each year from 1988 through 2022 across five asset classes: US equities, foreign equities, emerging market equities, Treasuries, US corporate bonds, and commodities. Then we compute subsample co-movement between US

EXHIBIT 3

Subsample Co-Movement for US Equities vs. Other Assets during Turbulent and Quiescent Periods (1988–2022)
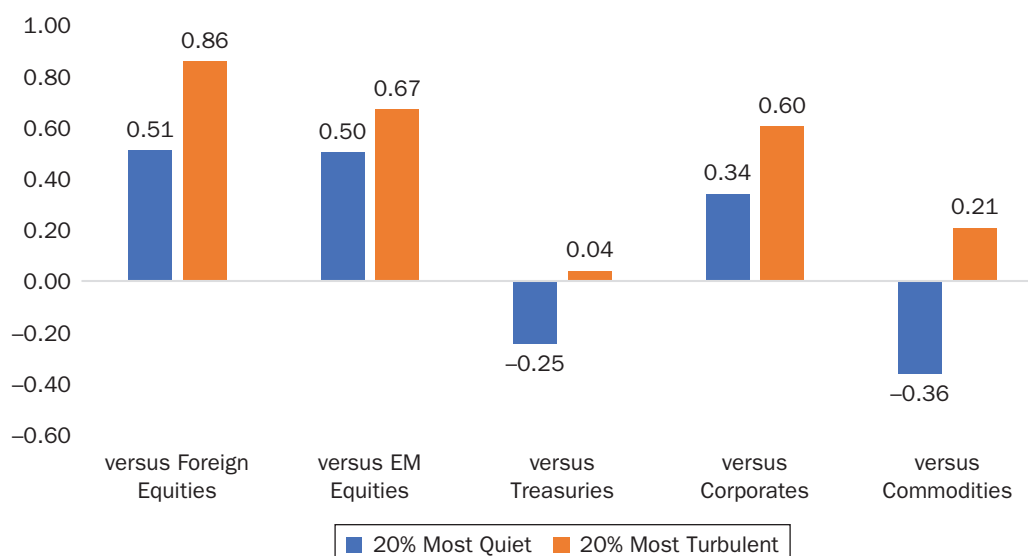
NOTES: This exhibit shows the monthly return co-movement for the 20% turbulent years as given by Equation (16) and the 20% quiet years between 1998 (when data for emerging markets become available) and 2022. We use the S&P 500 Total Return Index (from Robert Shiller's website) for US equities, the Bloomberg US Treasury Bond Index (from Datastream) for Treasuries, the MSCI Emerging Markets Total Return Index (from Datastream) for emerging market equities, the Bloomberg US Corporate Bond Total Return Index (from Datastream) for US corporate bonds, and the GSCI Commodities Index for commodities.

equities and each other asset class during the 20% most turbulent years and the 20% quietest years, using Equation (12).

Exhibit 3 reveals that the co-movement of US equities with emerging market and foreign equities during turbulent periods is significantly higher than it is during quiescent periods. Similarly, the correlation between US equities and US corporate bonds is also significantly higher during turbulent periods when equities tend to decline and credit spreads widen.

Turbulence can coincide with unusual positive returns as well as negative returns. We expect, and shall observe in Exhibit 5, that Treasuries act as a diversifier when equities perform poorly. However, there are also turbulent periods where Treasuries and equities rise in unison, which is why we do not observe a more negative correlation during turbulent environments in Exhibit 3.
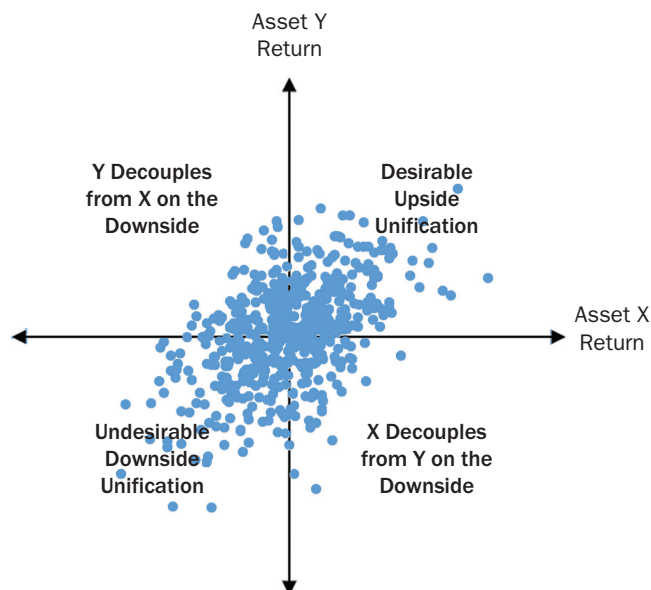
Return Thresholds

Now, let us consider how correlations differ around return thresholds, which is complicated for several reasons. First, we must consider when diversification is beneficial to a portfolio and when it is not. The simulated scatter plot in Exhibit 4 illustrates the difference between beneficial diversification and diversification that harms a portfolio. Consider two hypothetical assets, X and Y. Imagine that asset X is the main growth driver in the portfolio and that we have selected asset Y to diversify it. When asset X is performing well, we would prefer that asset Y also perform well. These outcomes reflect desirable upside unification and are associated with the upper-right quadrant of Exhibit 4. Alternatively, when asset X is experiencing losses, we would prefer that asset Y decouple from asset X to offset those losses. Therefore, we should consider only the left-hand side of the scatter plot to measure the potential of asset Y to diversify unfavorable performance of asset X. And if we wish to measure an asset's potential to reinforce favorable performance of asset X, we should consider just the right-hand side of the scatter plot.

EXHIBIT 4

**Scatter Plot of Two Hypothetical Asset Returns**

Exhibit 4 also reveals why we cannot simply calculate conventional correlation coefficients from a truncated sample to measure an asset's diversification properties. The correlation of Y with X when the returns of X are below a threshold will be lower than the full sample correlation because we have excised the opposing returns in the upper-right quadrant, which otherwise would pull up the correlation. In this example, the correlation of the thus-truncated return sample in relation to the respective subsample means is 0.33 as opposed to 0.50 for the full sample. We might be tempted to conclude that diversification increases given this lower correlation for the truncated sample, but this interpretation would be incorrect. These differences are an artifact of conditional correlation math with subsample means and do not indicate any change in the relationship between the two assets.[4] By contrast, our proposed measure of co-occurrence and its summary for a subsample as given by Equation (12) retain the perspective of the full-sample means, which makes them straightforward to interpret.

Exhibit 5 shows subsample co-movement for US equities versus each of the other five asset classes

---

[4] For more on truncated sample correlations, see Kinlaw et al. (2021) and Czasonis et al. (2022).

EXHIBIT 5
Subsample Co-Movement for US Equities vs. Other Assets during High- and Low-Return Periods (1988–2022)

introduced in Exhibit 3. In this case, we report co-movement during the 20% best years for US equities as well as for the 20% worst years, using Equation (12). We again observe that co-movement among the equity asset classes is significantly higher during periods when US equities perform poorly. In fact, US and emerging market equities hardly co-move at all during periods when US equities perform well. In other words, they exhibit little desirable upside unification but an excess of undesirable downside unification. Treasuries present the opposite profile: during periods when US equities perform well, Treasuries co-move positively. When US equities experience drawdowns, investors flee to the safety of Treasury bonds, causing them to rally. This is an example of desirable downside decoupling.

We have thus far argued that diversification depends on co-occurrence, the degree to which cumulative returns co-move synchronously or drift apart during an investor's horizon. We have also argued that correlation is a useful guide to co-occurrence only if returns conform to a multivariate normal distribution, which requires them to be independent and identically distributed. We then provided stark evidence that returns are neither independent nor identically distributed. Now we show how to construct a portfolio based on an empirical sample that reflects co-occurrence across multiple regimes and return intervals. Therefore, the approach accommodates nonzero autocorrelations and lagged cross-correlations, multiple regimes, asymmetry of co-occurrence around return thresholds, and asymmetric preferences for diversification, without any consideration of conventional correlation statistics.

FULL-SCALE OPTIMIZATION

Mean–variance optimization is a heuristic that yields an approximation of the optimal portfolio based on statistics, such as correlation, that summarize a return distribution. Investors sometimes refer to it as parametric optimization because it depends on parameters to approximate the true distribution. We propose as an alternative full-scale optimization, which considers every multivariate return observation for every asset in

a return sample. It therefore evaluates co-occurrences explicitly and addresses all the complexities of empirical co-occurrence that we previously discussed.[5]

Before we describe full-scale optimization, it is worth discussing the divergence of correlations as a function of the return interval used to estimate them. If an investor's only concern is to diversify cumulative returns over her prospective investment horizon, it would be of no importance if the correlation of multiyear returns diverged from the correlation of monthly returns. The investor would only need to consider whether the correlation of multiyear returns gives a good estimate of the co-occurrence of returns over her prospective multiyear horizon. But it is natural and appropriate to worry about within-horizon outcomes as well. If a portfolio suffers significant losses along the way to the conclusion of the investment horizon, an investor may alter the portfolio's risk profile prematurely or by necessity if external constraints demand a more cautious investment posture. All of this is to say that investors may have both a short-term horizon as well as a long-term horizon.

Full-scale optimization addresses this multihorizon sensitivity by using a return sample that comprises both short-term returns and multiyear returns for the same overall period.[6]

The key to incorporating both high- and low-frequency returns in a single return sample is to balance these returns properly, which means that every multiyear return must include all the short-term returns that go into it. And if the multiyear returns overlap, the short-term returns must be repeated accordingly. Thus, if an investor has a monthly horizon as well as a five-year horizon, there should be 60 times as many monthly returns as five-year returns. If there were fewer than 60 monthly returns for every five-year return, the five-year returns would have an outsized effect on the result.

After we build our multihorizon sample, we calculate the portfolio return of a candidate asset mix, expressed as a vector of portfolio weights, for each short-term period and each long-term period. We then convert these short-term and long-term portfolio returns into utility values given a chosen utility function. Next, we sum all the utilities for the short-term and long-term returns and store this value. We then repeat this process for another asset mix and carry on in this fashion until we have evaluated enough portfolios to ensure that one of the portfolios offers the highest possible utility.

We illustrate this process using annual and nonoverlapping cumulative five-year returns from January 1988 through December 2022 for six asset classes. For the annual horizon, we assume a kinked utility function that assumes power utility with a curvature (risk aversion) parameter of 1.5 for returns above the kink and a slope penalty of 1 below the kink. We locate the kink at −5%. For the five-year horizon, we assume values of 2.0, 5.0, and 0%, respectively. This formulation is consistent with the notion that investors are more averse to long-term losses than short-term losses.

We estimate expected returns as equilibrium returns derived from a full sample regression of each asset class's returns on those of a market portfolio, assuming a 2% risk-free rate, a 4% market portfolio premium, and a market portfolio allocated as shown in Exhibit 6. For full-scale optimization, we also re-mean the monthly and five-year returns in the

EXHIBIT 6

Market Portfolio and Equilibrium Returns

	Market Portfolio	Equilibrium Returns
US Equities	30%	8.3%
Foreign Equities	20%	8.7%
EM Equities	10%	11.8%
Treasuries	20%	3.2%
Corporates	15%	4.4%
Commodities	5%	5.8%

[5] The term "full-scale optimization" was coined by Paul A. Samuelson in a written correspondence to Mark Kritzman on July 16, 2003.

[6] For a more thorough discussion of a multihorizon return sample, see Kritzman and Turkington (2022).

EXHIBIT 7
Optimal Portfolios

	Full-Scale Multihorizon	Mean–Variance Yearly	Mean–Variance Five-Year
Weights			
US Equities	19.6%	32.1%	47.6%
Foreign Equities	21.3%	19.5%	19.5%
EM Equities	14.3%	10.9%	0.0%
Treasuries	0.0%	17.3%	0.0%
Corporates	42.8%	15.2%	32.9%
Commodities	2.0%	5.0%	0.0%
Annual Statistics			
Standard Deviation	12.4%	12.3%	12.6%
Frequency above Kink (–5%)	85.7%	82.9%	80.0%
Best	29.3%	24.5%	25.3%
Worst	–28.0%	–28.4%	–29.2%
5-Year Statistics			
Standard Deviation	42.1%	41.2%	33.4%
Frequency above Kink (0%)	100.0%	71.4%	85.7%
Best	110.5%	102.4%	83.6%
Worst	3.5%	–0.4%	–0.9%

mixed-frequency return sample to have the same equilibrium returns that we used in the mean–variance analyses.

The optimal full-scale and mean–variance portfolios are shown in Exhibit 7. We constrain all three portfolios to have the same expected returns. We report the following performance statistics for both one- and five-year horizons: standard deviation, frequency of returns above the kink, best return, and worst return.

The yearly and five-year mean–variance portfolios have the lowest standard deviation over their respective horizons. The other performance statistics reveal, however, that standard deviation does not tell the whole story. The full-scale, multihorizon optimal portfolio performs best in terms of avoiding losses below the 0% and 5% threshold values. Its best return and worst return are also higher than that of the mean–variance portfolios, for both annual and five-year horizons. Given its more favorable range of returns, its higher volatility may reflect desirable upside volatility. Or put differently, it tilts toward assets that unify with US equities when they perform well. This example illustrates how investors can use full-scale optimization to balance multihorizon performance objectives by blending higher- and lower-frequency returns. The method accounts for the actual co-occurrence of asset returns at every measurement interval.

SUMMARY

We introduced a new measure of portfolio diversification called co-occurrence, which captures the co-movement of cumulative returns over a given investment period. We showed that an informativeness weighted average of a sample's co-occurrences equals the traditional correlation coefficient, and we showed how co-occurrence offers an intuitive and flexible tool with which to analyze the complex realities of diversification that defy characterization by any summary statistic. We then argued that the full-sample correlation is a useful guide to diversification only if it gives a good estimate of co-occurrence over an investor's prospective horizon, and that

this would be true only if returns are serially independent. We went on to show how nonzero autocorrelations or lagged cross-correlations cause correlations between assets to differ depending on the return interval used to estimate them, and we presented evidence that returns were neither individually nor jointly serially independent, which belies the notion of multivariate normality. We also provided evidence that co-occurrence, irrespective of the intervals used to estimate it, is not stationary across regimes and that it is not symmetric around return thresholds, which contradicts the notion of multivariate normality. And given asymmetric patterns of co-occurrence, we pointed out that although investors prefer downside diversification, they are averse to upside diversification. We then introduced a portfolio construction technique called full-scale optimization, which inherently accounts for all these complexities of diversification without consideration of conventional correlation, including the possibility that investors have more than a single investment horizon.

APPENDIX A

ALTERNATIVE MEASURES OF SUBSAMPLE CO-MOVEMENT

We derive the measure of subsample co-movement in Equation (12) from the co-occurrence and informativeness of subsample observations, each of which reflects the mean and variance of the full sample. Alternatively, one might compute the correlation of the observations in a subsample Φ based on re-estimated means and variances for the subsample. Both approaches can be used to detect correlation asymmetry by comparing statistics for subsamples that are above or below symmetric thresholds for the return of one of the assets.[7] There are some advantages to using co-occurrence, as in Equation (12). First, it is easier to interpret because the means and variances remain constant. Second, we may rely on the fact that for a normal distribution, the subsample statistic above and below the average of one of the assets equals the full sample statistic. If we re-compute means and variances, this property does not hold (the subsamples have less correlation around their subsample means). For more extreme thresholds, the subsample correlation for a normal distribution becomes increasingly small, and a proper empirical analysis must account for this mechanical fact.

The trend is reversed for co-occurrence. As we condition on increasingly extreme thresholds for the return of one asset, the weighted average co-occurrence of the subsample rises in the case of a normal distribution. This effect is not artificial. It occurs because more-extreme observations are more informative, which means they are less polluted by noise. To the extent there is a relationship between two assets, it is likely to be most apparent among extreme-magnitude observations, resulting in higher co-occurrence.

APPENDIX B

CO-OCCURRENCE AND CORRELATION SURPRISE

For some applications, it is interesting to compare the co-occurrence of an observation with the weighted average co-occurrence of all observations in the sample (which equals the correlation). In other words, we can ask the question: how surprising is a given observation of co-occurrence? Kinlaw and Turkington (2013) define a measure of correlation

---

[7] It is preferable to condition on the returns of one asset rather than jointly conditioning on both assets performing well versus both assets performing poorly. The most impactful diversification occurs when one asset is down and the other is up. Double conditioning eliminates these most important observations. For more discussion on this topic, see Kinlaw et al. (2021).

surprise for any number of assets as the ratio of two Mahalanobis distances: one that accounts for the actual correlations among the assets and one that assumes all the correlations are zero. The Mahalanobis distances can be expressed in terms of normalized z-scores, in which case the denominator equals the Euclidean distance of the normalized z-scores:

$$CS_i = \frac{z_i P^{-1} z_i'}{z_i z_i'} \tag{B1}$$

In Equation (B1), z_i is a row vector for a chosen observation, P^{-1} is the inverse of the asset correlation matrix, and $'$ denotes matrix transpose. Intuitively, the numerator represents joint unusualness in magnitude and orientation, while the denominator represents unusualness in magnitude only. The ratio represents unusualness in orientation only. For two assets, the standard matrix inversion formula gives

$$P^{-1} = \frac{1}{1 - \rho^2} \begin{pmatrix} 1 & -\rho \\ -\rho & 1 \end{pmatrix} \tag{B2}$$

Next, we multiply out the matrices for correlation surprise:

$$CS_i(A,B) = \frac{1}{1 - \rho^2} \frac{z_{i,A}^2 + z_{i,B}^2 - 2\rho z_{i,A} z_{i,B}}{z_{i,A}^2 + z_{i,B}^2} \tag{B3}$$

$$CS_i(A,B) = \frac{1}{1 - \rho^2} \left(1 - \frac{\rho z_{i,A} z_{i,B}}{\frac{1}{2}(z_{i,A}^2 + z_{i,B}^2)} \right) \tag{B4}$$

$$CS_i(A,B) = \frac{1 - \rho c_i(A,B)}{1 - \rho^2} \tag{B5}$$

Equation (B5) shows how correlation surprise for a pair of assets depends on both co-occurrence and the full-sample correlation. Correlation surprise equals 1 when co-occurrence equals correlation. The informativeness-weighted average of correlation surprise also equals 1. The least surprising co-occurrence for positively correlated assets is 1, and the most surprising co-occurrence for positively correlated assets is −1. The reverse is true for negatively correlated assets. For assets that are strongly correlated, the maximum possible surprise from a divergence is very large. For uncorrelated assets, every observation is equally surprising and has a correlation surprise equal to 1.

Correlation surprise measures the degree to which an observation diverges from the typical direction of co-movement. We can substitute correlation surprise for co-occurrence in Equation (12) to quantify the degree of correlation surprise that is present in a sample, relative to the full sample correlation as a baseline. For positively correlated assets, correlation surprise events may offer welcome diversification during bad times. But correlation breakdowns are not always good. We can use correlation surprise to stress test the performance of any portfolio that is formed by assuming average correlations. In addition, Kinlaw and Turkington (2013) show that correlation surprise often serves as a leading indicator of volatility spikes.

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

Czasonis, M., M. Kritzman, B. Pamir, and D. Turkington. 2022. "The Role of Cryptocurrencies in Investor Portfolios." *The Journal of Alternative Investments* 24 (4): 18–32.

Kinlaw, W., M. Kritzman, S. Page, and D. Turkington. 2021. "The Myth of Diversification Reconsidered." *The Journal of Portfolio Management* 47 (8): 124–137.

Kinlaw, W., M. Kritzman, and D. Turkington. 2014. "The Divergence of High- and Low-Frequency Estimation: Causes and Consequences." *The Journal of Portfolio Management* 40 (5): 156–168.

Kinlaw, W., and D. Turkington. 2013. "Correlation Surprise." *Journal of Asset Management* 14 (December): 385–399.

Kritzman, M., and D. Turkington. 2022. "History, Shocks, and Drifts: A New Approach to Portfolio Formation." *The Journal of Portfolio Management* 48 (3): 142–152.

Markowitz, H. 1952. "Portfolio Selection." *The Journal of Finance* 7 (1): 77–91.