The Myth of
Diversification Reconsidered

**William Kinlaw, Mark Kritzman,
Sébastien Page, and David Turkington**

The Myth of Diversification Reconsidered

William Kinlaw, Mark Kritzman, Sébastien Page, and David Turkington

William Kinlaw
is head of State Street Associates, the research arm of State Street Global Markets, in Cambridge, MA.
wbkinlaw@statestreet.com

Mark Kritzman
is a founding partner of State Street Associates, chief executive officer of Windham Capital Management, and a faculty member at MIT's Sloan School of Management in Cambridge, MA.
kritzman@mit.edu

Sébastien Page
is head of global multi-asset at T. Rowe Price in Baltimore, MD.
sebastien.page@troweprice.com

David Turkington
is a senior managing director at State Street Associates, the research arm of State Street Global Markets, in Cambridge, MA.
dturkington@statestreet.com

KEY FINDINGS

- There is strong empirical evidence that asset class correlations are asymmetric, which poses complications in portfolio construction.

- Investors prefer diversification when a portfolio's main growth engine performs poorly and unification when it performs well.

- To measure correlation asymmetry caused by nonnormality, investors must adjust for changes in correlation that arise mathematically when part of a sample is excluded.

- Unlike prior research, investors should condition correlations on the performance of a single asset, not two assets.

ABSTRACT

That investors should diversify their portfolios is a core principle of modern finance. Yet there are some periods in which diversification is undesirable. When the portfolio's main growth engine performs well, investors prefer the opposite of diversification. An ideal complement to the growth engine would provide diversification when it performs poorly and unification when it performs well. Numerous studies have presented evidence of asymmetric correlations between assets. Unfortunately, this asymmetry is often of the undesirable variety: It is characterized by downside unification and upside diversification. In other words, diversification often disappears when it is most needed. In this article, the authors highlight a fundamental flaw in the way some prior studies have measured correlation asymmetry. Because they estimate downside correlations from subsamples in which both assets perform poorly, they ignore instances of successful diversification (i.e., periods in which one asset's gains offset the other's losses). The authors propose instead that investors measure what matters: the degree to which a given asset diversifies the main growth engine when it underperforms. This approach yields starkly different conclusions, particularly for asset pairs with low full-sample correlation. The authors review correlation mathematics, highlight the flaw in prior studies, motivate the correct approach, and present an empirical analysis of correlation asymmetry across major asset classes.

TOPICS

*Portfolio theory, portfolio construction, quantitative methods, statistical methods, performance measurement**

The correlation coefficient, the parameter that quantifies the degree to which two assets diversify one another, took on new significance in 1952 when Harry Markowitz published his landmark article "Portfolio Selection." Markowitz (1952) formalized the role of diversification when he showed how to construct optimal portfolios given the expected returns, standard deviations, and correlations of their component assets. Nearly 70 years after it was introduced, the mean–variance paradigm has proven surprisingly robust. However, it makes two implicit assumptions about diversification that warrant careful consideration. Because it relies on a single parameter to approximate the way each pair of assets co-vary, mean–variance optimization assumes that correlations are symmetric on the upside and downside. Moreover, the approach assumes that diversification is desirable on the upside as well as the downside. The first assumption is occasionally correct, but the second assumption never is.

Diversification is most helpful to investors when the major engine of growth in the portfolio, typically domestic equities, performs poorly. They derive benefit from assets whose returns offset this poor performance. When the growth engine is performing well, they would prefer unification, which is the opposite of diversification. The ideal complement to domestic equities would be an asset that is correlated positively when domestic equities are performing well and negatively when they are not. Put simply, investors seek diversification on the downside and unification on the upside. An adage, which has been credited both to Mark Twain and Robert Frost, defines a banker as "a fellow who lends you his umbrella when the sun is shining but wants it back the minute it begins to rain."[1] Diversification often behaves like a banker if this characterization is to be believed.

In this article, we review correlation mathematics and show how to distinguish true correlation asymmetry from the illusory correlation shifts that arise as an artifact of how the data are partitioned. We then highlight a fundamental flaw in the way several prior studies have measured correlation asymmetry. Because they estimate downside correlations from subsamples in which both assets perform poorly, they ignore instances of successful diversification (i.e., periods in which one asset's gains offset the other's losses). We propose instead that investors measure what matters: the degree to which a given asset diversifies the main growth engine when it underperforms (see, e.g., Page and Panariello 2018). This approach yields different conclusions, particularly for asset pairs with low correlation. To compare the two approaches, we present an empirical study of correlation asymmetry across six major asset classes. Finally, we show how investors can employ full-scale optimization to construct portfolios that exploit correlation asymmetry by increasing average correlation on the upside when it is beneficial and reducing average correlation on the downside when it is not.

LITERATURE REVIEW

Page and Panariello (2018), Page (2020), and many others have argued that despite the wide body of published research, many investors still do not fully appreciate the impact of correlation asymmetries on portfolio efficiency or, perhaps more importantly, exposure to loss. During left-tail events, diversified portfolios may have greater exposure to loss than more concentrated portfolios. Leibowitz and Bova (2009) showed that during the 2008 global financial crisis, a portfolio diversified across US equities, US bonds, foreign developed equities, emerging market equities, and real estate investment trusts (REITs) saw its equity beta rise from 0.65

to 0.95, and the portfolio unexpectedly underperformed a simple 60/40 US stock/bond portfolio by 9%.

Studies on tail dependence (how crashes tend to happen at the same time across markets) corroborate these findings. For example, Garcia-Feijóo, Jensen, and Johnson (2012) showed that when US equity returns are in their bottom 5%, non-US equities, commodities, and REITs also experience significantly negative returns, beyond what would be expected from full-sample correlations. Hartmann, Straetmans, and de Vries (2010) showed that currencies co-crash more often than would be predicted by a bivariate normal distribution. Similarly, Hartmann, Straetmans, and de Vries (2004) estimated that stock markets in G-5 countries are two times more likely to co-crash than bond markets. Van Oordt and Zhou (2012) extended pairwise analysis to joint tail dependence across multiple markets and reached similar conclusions. These studies ignore asymmetries, however, between the left and right tails. They either focus on the left tail or use symmetrical measures of tail dependence, such as the joint *t*-distributions.

Prior research suggests that correlation asymmetries are closely related to the concept of risk regimes. Financial markets tend to fluctuate between a low-volatility state and a panic-driven, high-volatility state (see, e.g., Kritzman, Page, and Turkington 2012). In fact, Ang and Bekaert (2015) directly linked the concept of regime shifts to rising left-tail correlations. But what causes regime shifts? A partial answer is that macroeconomic fundamentals themselves exhibit regime shifts, as documented for inflation and growth data.

In normal markets, differences in fundamentals drive diversification across risk assets. During panics, however, investors often sell risk irrespective of differences in fundamentals. Huang, Rossi, and Wang (2015), for example, showed that sentiment is a common factor that drives both equity and credit-spread returns—beyond the effects of default risk, liquidity, and macro variables—and suggested that sentiment often spills over from equities to the credit markets.

Related studies in the field of psychology suggest that to react more strongly to bad news than good news is human nature. Fear is more contagious than optimism. In an article titled "Bad Is Stronger than Good," Baumeister et al. (2001) explained that "Bad information is processed more thoroughly than good.... From our perspective, it is evolutionarily adaptive for bad to be stronger than good."

The literature is divided, however, on the correct approach to measuring correlation asymmetry. As we will explain further, some studies (e.g., Longin and Solnik 2001; Ang and Chen 2002; Chua, Kritzman, and Page 2009) estimate downside correlations by conditioning on the returns of both assets simultaneously. Others, such as Page and Panariello (2018), Gulko (2002), and Garcia-Feijóo, Jensen, and Johnson (2012), conditioned on the returns of a single asset. In this article, we show why the former approach is flawed and argue that the latter approach results in a more useful measure of correlation asymmetry.

CORRELATION MATHEMATICS

Market participants often remark that "correlations go to one" when the markets are in turmoil, but such differences do not necessarily prove that the bivariate return distribution is nonnormal or that returns emanate from more than one regime. Sub-sample correlations change naturally as an artifact of how we partition the sample, even if the underlying distribution is normal. We must therefore account for these effects to detect correlation asymmetry properly.

Longin and Solnik (2001) introduced the notion of exceedance correlation, which they defined as the correlation between two assets when the returns of both assets are

EXHIBIT 1
Return Observations for Two Assets
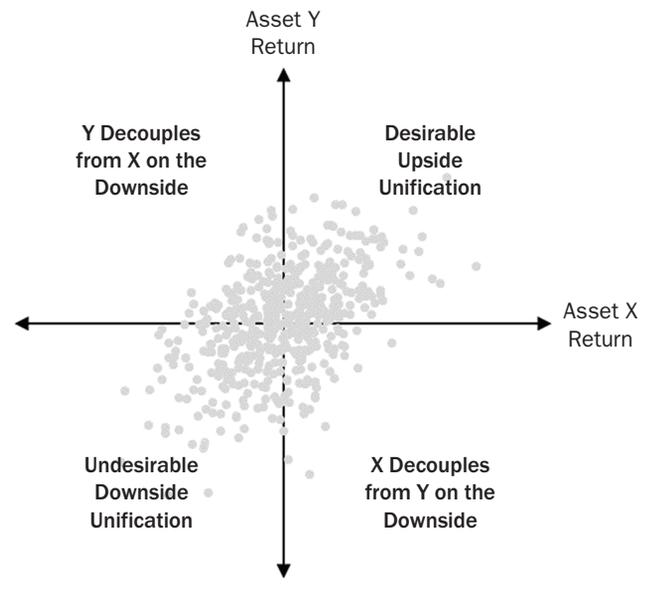
either above or below a given threshold. For example, they might estimate the exceedance correlation for US and foreign equities from the subsample of returns in which both asset classes suffer losses of 10% or more. Chua, Kritzman, and Page (2009) applied the same approach to a variety of asset classes, country equity markets, hedge fund styles, and fixed-income segments. They found pervasive evidence of correlation asymmetry that is unfavorable to investors. However, this specification of exceedance correlation suffers from a fundamental flaw: It misses important instances of diversification because it ignores outcomes in which, in the previous example, US equities perform poorly but foreign equities perform well.

Exhibit 1 presents a simple illustration of this point. We use Monte Carlo simulation to generate 500 returns for assets X and Y that conform to a bivariate normal distribution with identical means equal to zero, standard deviations equal to 20%, and a correlation of 0.50. 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 follow suit. These outcomes reflect desirable upside unification and are associated with the upper right quadrant of Exhibit 1. On the other hand, when asset X is suffering losses, we would prefer that asset Y decouple from asset X to offset those losses. These outcomes reflect desirable downside decoupling of asset Y and are associated with the upper-left quadrant of Exhibit 1. These are the periods in which asset Y successfully diversifies asset X. Finally, the lower left quadrant of Exhibit 1 is associated with very unpleasant outcomes in which asset X underperforms and asset Y fails to diversify it.

It is evident from this illustration that we must consider the entire left side of the distribution to measure properly the diversification potential of asset Y with respect to asset X. It would not be informative to focus only on the lower left quadrant because these are the instances in which diversification has already failed. Yet this is the way several other studies have measured downside correlations.

Exhibit 2 presents a comparison of these two approaches for the bivariate distribution presented in Exhibit 1 with a threshold value of 0%. The left panel shows the subsample in which both assets' returns are below the threshold; this is the approach taken by Longin and Solnik (2001), Chua, Kritzman, and Page (2009), Ang and Chen (2002), and others. We submit that investors should instead estimate downside correlations from the subsample of returns shown in the right panel, in which the return of asset X is below a particular threshold, regardless of the return of asset Y. This approach introduces additional complexity because we are now able to estimate two downside correlation coefficients for each pair of assets, one conditioned on the returns of each asset. This doubles the number of correlation coefficients we must potentially consider. It is, however, reasonable for investors to focus their attention on a few of the portfolio's main growth engines, the assets that contribute the largest share of portfolio risk, as the conditioning assets.

The upside and downside correlations between assets X and Y, conditioned on the returns of asset X, are given by Equation 1.

EXHIBIT 2
### Subsamples of Returns for Assets *X* and *Y* in Which One or Both Assets Underperform

Subsample of Observations Where X < θ and Y < θ, θ = 0%

Subsample of Observations Where X < θ, θ = 0%
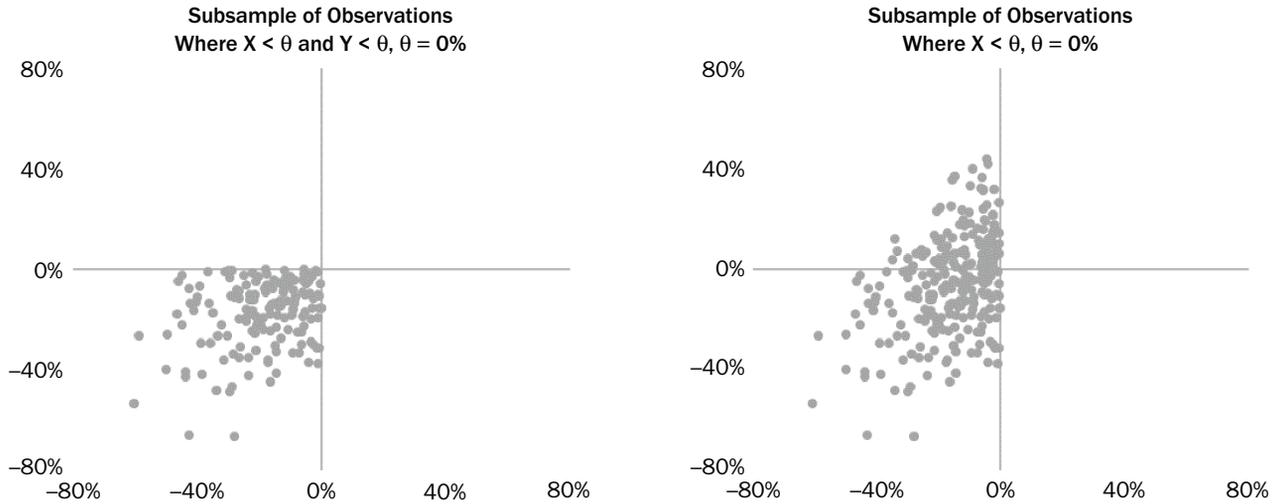
$$\rho(\theta) = \begin{cases} corr(x, y \,|\, x > \theta) \;\; if \;\; \theta > 0 \\ corr(x, y \,|\, x < \theta) \;\; if \;\; \theta < 0 \end{cases} \tag{1}$$

where ρ is the upside or downside correlation, *x* and *y* are observed returns for each asset, and θ is the threshold applied to the returns of asset *X*. In practice, we express θ in units of standard deviation above or below the mean of asset *X*. Thus, if we set θ equal to +1.0, we would evaluate the top portion of Equation 1 and esti- mate correlation for all observations in which the return of asset *X* is one standard deviation or more above its mean. If we set θ equal to −1.5 we would evaluate the bottom portion of Equation 1, which focuses on the subsample of returns in which asset *X* is 1.5 standard deviations or more below its mean. In the unique case in which θ equals zero, we evaluate both equations to estimate an upside correlation (in which returns are above the mean) and a downside correlation (in which returns are below the mean). Because the bivariate normal distribution is symmetric, the expected upside and downside correlations will be identical in this instance and for any instance in which the thresholds have the same absolute value.

Given Equation 1 and setting θ equal to zero, the upside and downside correla- tion associated with the bivariate distribution defined earlier is 0.33, as opposed to 0.50 for the full sample. If we modify Equation 1 to apply the threshold to both assets' returns, the upside and downside correlation is 0.27. Were these estimates derived from real data rather than simulated data, we might be tempted to conclude that diversification increases in the extremes for this pair of assets. However, this interpretation would be incorrect. These differences are an artifact of conditional correlation math and do not indicate any change in the relationship between the two assets in the tails. Exhibit 3 shows how the expected upside and downside correla- tions change as a function of threshold value (θ).

Exhibit 3 reveals that the conditional correlations decrease as the absolute value of the threshold increases. At a threshold value of positive or negative 20%, corre- sponding to one standard deviation, the upside and downside correlation is 0.25 if we condition on asset *X*. It is 0.18 if we condition on both assets. Exhibit 3 also reveals that, all else equal, we should expect a higher correlation by construction when we

EXHIBIT 3

Conditional Correlations for the Bivariate Normal Distribution as a Function of θ
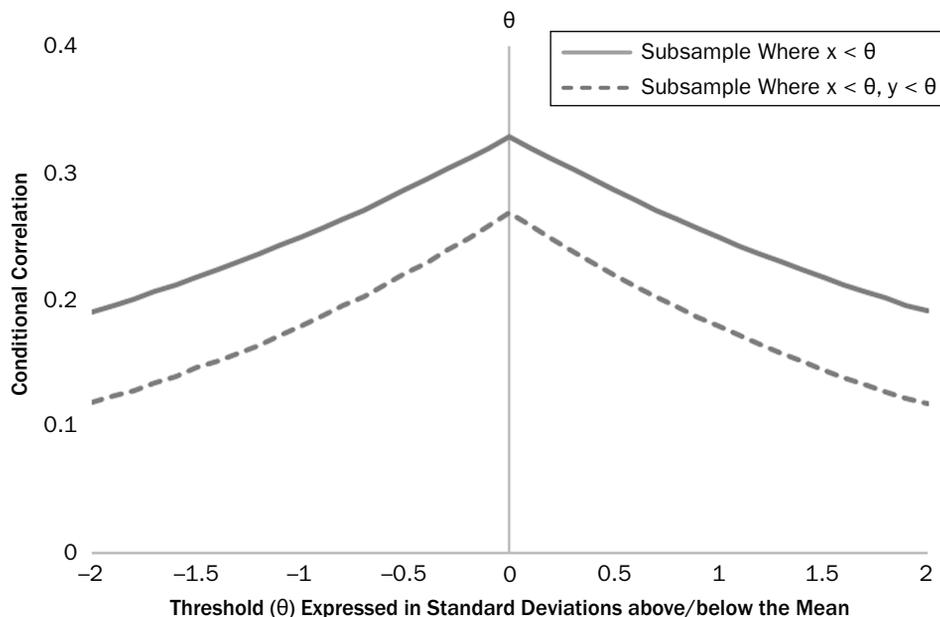
NOTES: The exhibit shows conditional correlations estimated from subsamples of a bivariate normal distribution with means equal to zero, standard deviations equal to 20%, and a correlation of 0.50. We employ Monte Carlo simulation to create a sample of 100 million observations. We then estimate conditional correlations for each threshold, θ, using Equation 1 which conditions on asset X only. We then estimate conditional correlations using an analogous equation in which the returns of both assets are above or below the threshold, θ.

condition on one asset than when we condition on both. When we estimate upside and downside correlations from empirical data, we must compare them to these expected values. Only if we observe material differences between the empirical and expected correlation profiles can we conclude that the observed correlation asymmetry is a symptom either of a nonnormal bivariate distribution or multiple distributions. In these instances, investors should consider adjusting explicitly for correlation asymmetry when they construct portfolios or estimate downside risk exposure.

Exhibit 4 shows the empirical correlation profile for US equities and foreign developed equities based on monthly returns starting in January 1976 and ending in December 2019.[2] It also shows the expected correlation profile for the corresponding bivariate normal distribution. These two asset classes have a full sample correlation of 0.66. Because the two asset classes have different volatilities, we standardize each return series by subtracting the mean from each monthly observation and dividing this quantity by the standard deviation. Exhibit 4, Panel A shows the correlation profile in which the threshold is applied to both assets. Exhibit 4, Panel B shows results conditioned on US equity only, as given by Equation 1.

We can draw two conclusions from Exhibit 4. First, in both cases, the empirical downside correlation is higher than implied by the bivariate normal distribution. In the right tail, we observe the opposite. This correlation profile represents downside unification and upside diversification and is therefore undesirable to investors. Second, we observe that the two correlation profiles are quite similar: It does not appear to make much difference whether we condition on the returns of one asset or both assets. This is because the two assets have a relatively high correlation to start. As a result, most observations are in the lower-left and upper-right quadrants anyway, and we therefore retain most of them when we impose the double condition.

EXHIBIT 4
Empirical Correlation Profile for US Equities and Foreign Developed Equities

Panel A: Correlation Profile for US and Foreign Equity, Conditioned on Both Assets

Threshold (θ) Expressed in Standard Deviations above/below the Mean

Panel B: Correlation Profile for US and Foreign Equity, Conditioned on US Equity Only

Threshold (θ) Expressed in Standard Deviations above/below the Mean
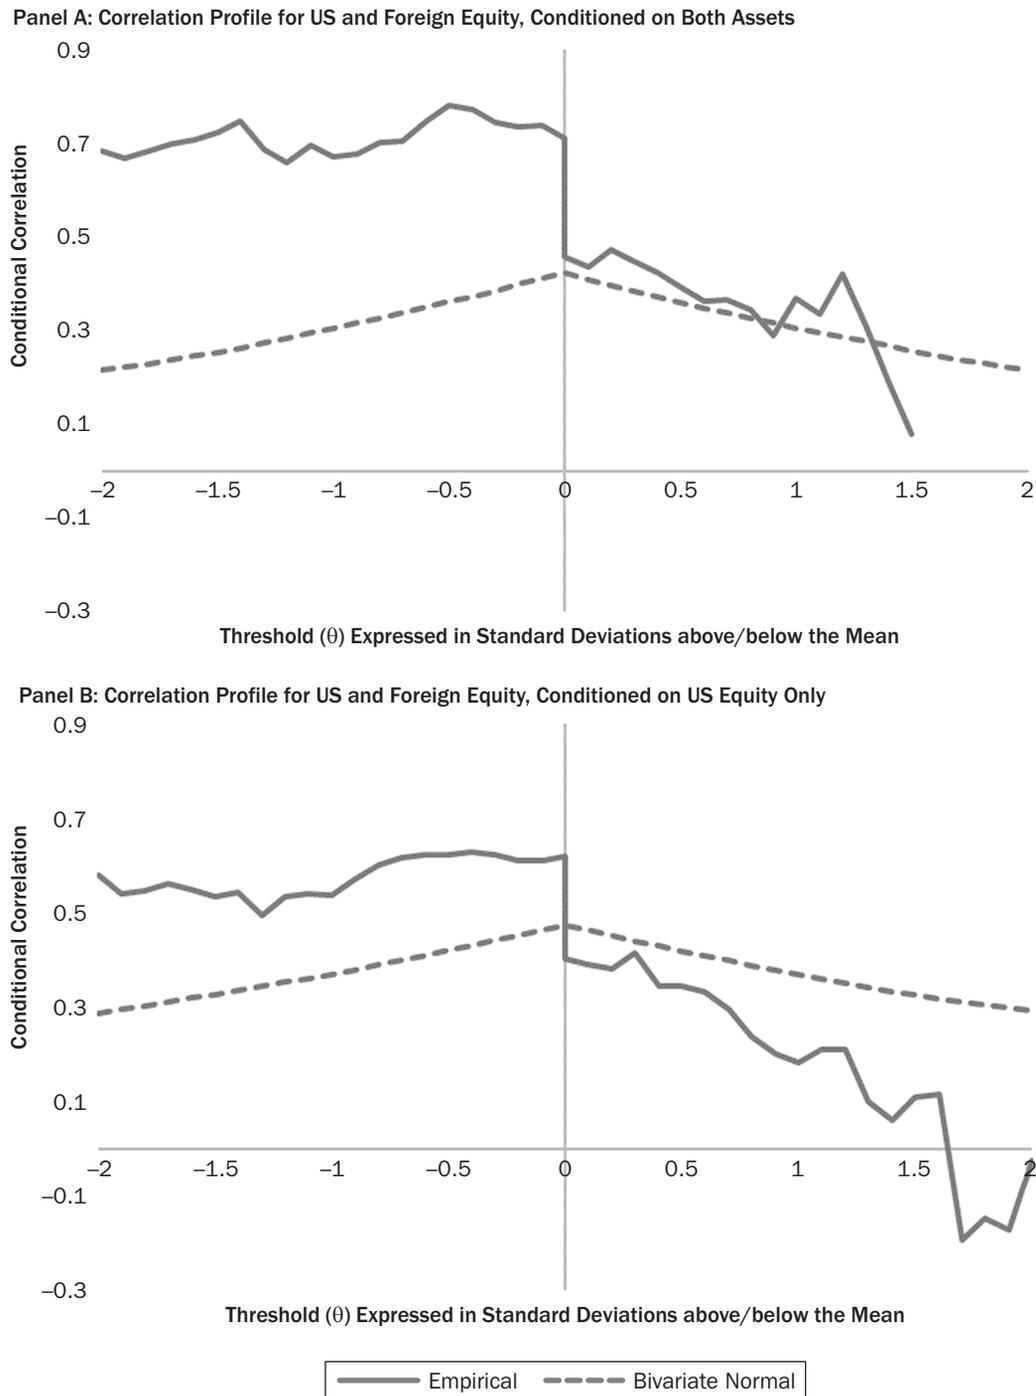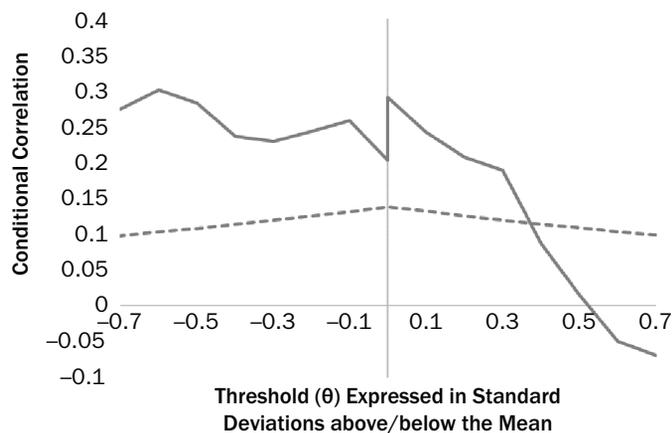
Empirical ----- Bivariate Normal

Exhibit 5 presents the analogous results for US equity and US corporate bonds, which have a relatively low full-sample correlation. In this case, we observe that the two methods lead to entirely different conclusions. The correct approach, shown in Exhibit 5, Panel B, shows that US corporate bonds offer some desirable downside decoupling. The incorrect approach, shown in Panel A, suggests that the two asset classes have higher-than-normal correlation on the downside.
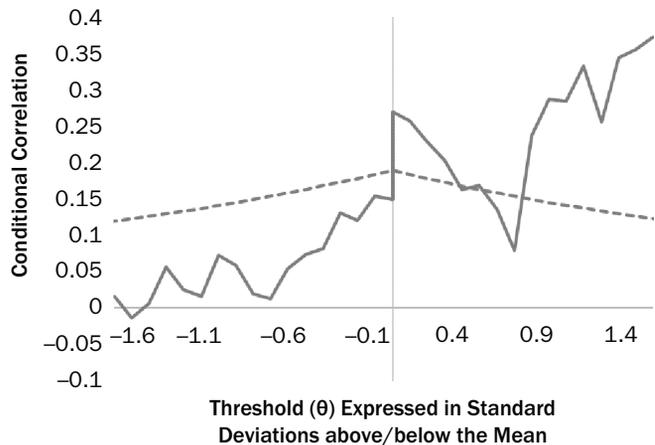
EXHIBIT 5
Correlation Profile for US Equity and Corporate Bonds

Panel A: Correlation Profile for US Equity and Corporate Bonds, Conditioned on Both Assets

Panel B: Correlation Profile for US Equity and Corporate Bonds, Conditioned on US Equity

Threshold (θ) Expressed in Standard Deviations above/below the Mean

Threshold (θ) Expressed in Standard Deviations above/below the Mean

——— Empirical   ---- Bivariate Normal

We present this example because it highlights how the single and double conditioning approaches can lead to different conclusions, not to suggest that US corporate bonds impart the most desirable correlation asymmetry to US equity. That designation belongs to Treasury bonds. Corporate bond returns have two components: a duration component and a credit component. The former tends to diversify equity exposure, whereas the latter does not. During some periods, the duration component overpowers the credit component, enabling corporate bonds to decouple from equities on the downside. Yet the double conditioning approach is blind to these periods, which are akin to outcomes in the upper-left quadrant of Exhibit 1. Double conditioning therefore produces higher downside correlation estimates by design, as is evident from Exhibit 5, Panel A. The single conditioning approach captures the full left side of the distribution and therefore produces a lower downside correlation estimate, as shown in Exhibit 5, Panel B.

In the next section, we extend our analysis of correlation asymmetry to present a comprehensive empirical study among six major asset classes. To do so concisely, we introduce a summary metric to capture the degree of correlation asymmetry for each asset pair. The metric we choose is the average difference between the empirical and expected downside correlations across threshold values for down markets, less the analogous quantity for up markets. We calculate these average differences for up and down markets as given by Equation 2 as proposed by Chua, Kritzman, and Page (2009).

$$\mu_{dn} = \frac{1}{n}\sum_{i=1}^{n}[\rho_{emp}(\theta_i) - \rho_{exp}(\theta_i)], \ for \blacksquare < 0$$

$$\mu_{up} = \frac{1}{n}\sum_{i=1}^{n}[\rho_{emp}(\theta_i) - \rho_{exp}(\theta_i)], \ for \blacksquare > 0 \qquad (2)$$

The terms μ_{dn} and μ_{up} are the average differences for up and down markets, respectively; $\rho_{emp}(\theta_i)$ is the empirical upside or downside correlation at threshold θ_i; and $\rho_{exp}(\theta_i)$ is the corresponding expected upside or downside correlation for a

Asset Class Returns and Standard Deviations

Asset Class	Return (% p.a.)	Standard Deviation (% p.a.)
US Equities	12.66	14.54
Foreign Developed Equities	11.10	16.38
Emerging Market Equities	13.51	22.36
Treasury Bonds	7.16	5.20
US Corporate Bonds	8.15	6.61
Commodities	6.26	19.06

NOTES: Annualized return is the arithmetic return of each asset class from January 1976 through December 2019, with the exception of emerging markets, for which data begin in January 1988. Standard deviation is the annualized standard deviation of monthly returns over the same period.

bivariate normal distribution. For both up and down markets, we iterate through n thresholds in equally spaced intervals of 0.1 standard deviations. We stop iterating when there are fewer than 30 observations beyond the threshold. We then take the difference between the upside and downside averages given by Equation 2 ($\mu_{dn} - \mu_{up}$) to arrive at a summary metric for correlation asymmetry. Positive values for this metric indicate the presence of an undesirable correlation profile (net downside unification or upside diversification). Negative values indicate a desirable profile (downside decoupling or upside unification). For US and foreign developed equities, this average difference is equal to 0.32, reflecting an undesirable correlation profile.

CORRELATION ASYMMETRY BETWEEN ASSET CLASSES

Research has shown that most asset universes offer less diversification during down markets than during up markets, including country equity markets, global industries, individual stocks, hedge funds, and international bonds. The correlation between stocks and bonds is often an exception to this pattern. Kritzman, Lowry, and Van Royen (2001) found that stock–bond correlations within countries decrease during periods of market turbulence. In this section, we build on this finding to analyze the pervasiveness of correlation asymmetry across the six major asset classes shown in Exhibit 6. Exhibit 6 shows the full-sample return and standard deviation of each asset class, and Exhibit 7 shows the full-sample correlation matrix.[3]

To measure correlation asymmetry for each pair of asset classes, we report in Exhibit 8 the summary metric ($\mu_{dn} - \mu_{up}$) as given by Equation 2. Panel A reports this summary metric following the double conditioning specification. Because the threshold is applied to both assets simultaneously, this matrix is symmetric. Panel B reports the summary metric for each asset class pair following our favored single conditioning approach. In this case, the matrix is not symmetric. The values presented in this panel capture the correlation asymmetry between the row asset and column asset when we condition on the returns of the row asset. Panel C reports the differences between Panels A and B.

We draw several conclusions from Exhibit 8:

- When paired with US equities as the conditioning asset, Panel B reveals that emerging market equities, foreign developed equities, and commodities are less desirable complements. They exhibit correlation asymmetry that is unfavorable to investors. On the other hand, Treasury and corporate bonds exhibit a favorable correlation profile.
- When paired with Treasury bonds as the conditioning asset, Panel B shows that corporate bonds offer the most favorable correlation profile and commodities the least.

[3] We use the following benchmark indexes as proxies for each asset class. For US equities, we use the S&P 500 Total Return Index. For foreign developed equities, we use the MSCI World ex-US Total Return Index. For emerging market equities, we use the MSCI Emerging Markets Total Return Index. For Treasury bonds, we use the Barclays US Treasury Total Return Index. For US corporate bonds, we use the Barclays US Credit Total Return Index. For commodities, we use the S&P GSCI Commodity Total Return Index. We procured all data from Datastream.

EXHIBIT 7
Asset Class Correlations

Asset Class		a	b	c	d	e
a	US Equities					
b	Foreign Developed Equities	0.67				
c	Emerging Market Equities	0.66	0.72			
d	Treasury Bonds	0.08	0.02	−0.15		
e	US Corporate Bonds	0.30	0.24	0.22	0.86	
f	Commodities	0.18	0.30	0.29	−0.09	0.02

NOTES: Correlation coefficients are derived from the full sample of monthly returns from January 1976 through December 2019. Emerging markets correlations are pairwise and derived from the period starting in January 1988, when data became available.

EXHIBIT 8
Correlation Asymmetry among Major Asset Classes

		a	b	c	d	e	f
Panel A: Excess Downside Minus Excess Upside Correlation, Conditioned on Both Assets' Returns							
a	US Equities		0.34	0.35	0.00	0.14	0.39
b	Foreign Developed Equities	0.34		0.31	−0.17	0.19	0.35
c	Emerging Market Equities	0.35	0.31		−0.17	0.24	0.10
d	Treasury Bonds	0.00	−0.17	−0.17		−0.14	−0.04
e	US Corporate Bonds	0.14	0.19	0.24	−0.14		0.17
f	Commodities	0.39	0.35	0.10	−0.04	0.17	
Panel B: Excess Downside Minus Excess Upside Correlation, Conditioned on Row Asset Only							
a	US Equities		0.32	0.32	−0.35	−0.19	0.27
b	Foreign Developed Equities	0.35		0.37	−0.25	0.07	0.31
c	Emerging Market Equities	0.29	0.31		−0.11	0.10	0.05
d	Treasury Bonds	0.02	−0.08	0.05		−0.19	0.20
e	US Corporate Bonds	0.12	0.07	0.27	−0.23		0.37
f	Commodities	0.47	0.56	0.46	−0.01	0.21	
Panel C: Difference between Panels A and B							
a	US Equities		−0.02	−0.03	−0.35	−0.33	−0.12
b	Foreign Developed Equities	0.01		0.06	−0.08	−0.12	−0.04
c	Emerging Market Equities	−0.06	0.00		0.06	−0.14	−0.05
d	Treasury Bonds	0.02	0.09	0.23		−0.05	0.24
e	US Corporate Bonds	−0.02	−0.13	0.03	−0.09		0.20
f	Commodities	0.08	0.20	0.36	0.02	0.05	

NOTES: We estimate the empirical and expected correlation profile for each asset pair in the same manner as in Exhibits 4 and 5. We estimate the empirical correlation profile using data starting in January 1976 and ending in December 2019, except for emerging market equities, which start in January 1988. Emerging market equity values are estimated pairwise over this time period with each other asset class. We only estimate values for thresholds for which at least 30 observations are available above or below the threshold. We then average across all thresholds to estimate the excess upside and downside correlation.

- Treasury bonds are the only asset that is universally favorable as a complement. Panel B shows that Treasury bonds impart beneficial asymmetry when returns are conditioned on any other asset class, although the benefit imparted to commodities is negligible. That said, interest rates are near zero as we write this article in February 2021. We therefore cannot expect Treasury bonds to rally as much during equity sell-offs as they have in the past.
- Because it ignores observations in the upper-left quadrant, the joint conditioning approach employed in Panel A understates the favorable correlation profile that Treasury and corporate bonds impart to the three equity asset classes.

It also overstates the diversification benefits that commodities impart to fixed-income asset classes.

IMPLICATIONS FOR PORTFOLIO CONSTRUCTION

Having demonstrated that correlation asymmetry is prevalent among major asset classes, we now turn to the question of what investors should do about it. There are two ways to account for correlation asymmetry in portfolio management. Investors can

1. reallocate their portfolios dynamically in anticipation of regime shifts, increasing their exposure to safe-haven assets that offer downside diversification when they expect conditions to deteriorate (see, e.g., Kritzman, Page, and Turkington 2012), or
2. place greater weight on downside correlations when setting policy weights, thereby constructing a static portfolio that is more resilient to downturns.

Investors who pursue the first approach must monitor market conditions and predict which correlations are most likely to prevail in the future. On the other hand, investors who pursue the second approach seek to prepare rather than predict. They build portfolios as one might design a house on the seashore: for routine use during balmy conditions, but sufficiently resilient to weather a hurricane if one should strike. This approach requires that the designer strike an optimal balance between resilience to storms and utility during fair weather conditions—a windowless, concrete bunker would provide maximum protection but would not be particularly appealing on a sunny day. We propose that investors strike the analogous balance when constructing portfolios.

There are several ways investors can construct portfolios that account for correlation asymmetry. One is to perform mean–variance optimization using only the downside, rather than full sample, correlations. However, this approach would be optimal only during extreme conditions. Another would be to blend the full-sample correlations with the downside correlations, but this leaves the critical choice of the blending ratio. Furthermore, in a portfolio with more than two assets, both methods require that the investor select a subset of the assets on whose returns the correlations will be conditioned.

We propose instead that investors use full-scale optimization to account implicitly for asymmetric correlations a well as other peculiarities of the multivariate return distribution. Full-scale optimization, introduced by Cremers, Kritzman, and Page (2005), identifies the optimal portfolio for any return distribution and any specification of investor preferences. Whereas mean–variance optimization yields an approximation of the in-sample solution if the return distribution is elliptical and investors have preferences that can be described by mean and variance, full-scale optimization yields the true optimal portfolio for a given return sample. Rather than relying on parameters such as means and covariances to approximate the distribution, full-scale optimization relies on numerical search algorithms to solve for the weights that maximize the given utility function precisely. To demonstrate how this technique accounts for correlation asymmetries, we employ a kinked utility function with the kink located at a 25% loss. This utility function is given by Equation 3.

$$U_{kinked}(R) = \begin{cases} \ln(1+R), & \text{for } r \geq k \\ \ln(1+R) - \omega(k-R), & \text{for } r < k \end{cases} \qquad (3)$$

The term $U_{kinked}(R)$ is expected utility, R is the return of the portfolio, k is the location of the kink (in this case, negative 25%), and ω is the slope of the linear utility function below the kink. With this utility function, investor satisfaction drops precipitously when returns fall below the kink. When returns fall above the kink, investor satisfaction conforms to a log-wealth utility function. This utility function is designed to express a strong aversion to losses below the kink.

To determine whether full-scale optimization addresses correlation asymmetry effectively, we need a way to measure the degree of correlation asymmetry in a portfolio and whether it is of the desirable or undesirable variety. For this purpose, we define the metric ξ in Equation 4:

$$\xi = \sum_{i=1}^{n}\sum_{j=1}^{n} w_i w_j (\mu_{i,j}^{dn} - \mu_{i,j}^{up}) \tag{4}$$

where n is the number of assets, w_i is the weight of asset i in the portfolio, w_j is the weight of asset j in the portfolio, and $(\mu_{i,j}^{dn} - \mu_{i,j}^{up})$ is the correlation asymmetry summary metric for assets i and j given by Equation 2, conditioned on asset i. Larger values of ξ indicate that the portfolio has excess downside correlation, which is undesirable. Of course, investors do not derive utility by reducing correlation asymmetry; they derive utility by growing wealth. Full-scale optimization does not maximize favorable correlation asymmetry directly, but as we will demonstrate, utility is well served by pairing assets that unify on the upside and diversify on the downside. Mean–variance optimization cannot account for these kinds of asymmetries because it implicitly assumes that correlations are symmetric.

Exhibit 9 shows a full-scale optimal portfolio and a corresponding mean–variance optimal portfolio with the same expected return of 7%. It also shows the degree of undesirable correlation asymmetry for each portfolio, ξ, as well as the utility of each portfolio as given by the kinked utility function.

Exhibit 9 reveals that, as we would expect, the mean–variance optimal portfolio has a lower standard deviation than the full-scale optimal portfolio. However, the mean–variance portfolio suffers a larger average loss when its returns fall below the threshold of negative 25%. The full-scale portfolio is able to achieve this reduction in downside exposure by reducing undesirable correlation asymmetry that is invisible to the mean–variance utility function. The changes in weights are intuitive. The full-scale optimal portfolio has larger allocations to US equities and Treasury bonds, the pair with the most desirable correlation asymmetry profile during the sample period. It holds almost no allocation to commodities or corporate bonds, the pair with the least desirable correlation asymmetry profile during the sample period.

SUMMARY

In this article, we debunk the fallacy that diversification is always beneficial to investors and that correlations are symmetric on the upside and downside. Although diversification is desirable on the downside, investors would prefer that all assets rise in concert and should seek unification on the upside. When measuring conditional correlations, it is important to adjust for the correlation changes that arise naturally as an artifact of correlation math. To detect correlation asymmetry properly, we must (1) condition returns on a single asset of interest rather than both assets and (2) compare the empirical upside and downside correlations to the values we would expect if returns emanated from a single, bivariate normal distribution. Unfortunately, when we make these adjustments, we conclude that most pairs of asset

EXHIBIT 9
Mean–Variance and Full-Scale Optimal Portfolios

	Expected Return	Standard Deviation	Full-Scale Optimal Portfolio	Mean– Variance Optimal Portfolio
US Equities	7.5%	14.1%	64.3%	42.3%
Foreign Developed Equities	7.8%	16.4%	9.7%	16.6%
Emerging Market Equities	8.5%	22.4%	4.1%	10.0%
Treasury Bonds	4.9%	4.4%	21.5%	12.4%
US Corporate Bonds	5.3%	4.9%	0.0%	10.8%
Commodities	6.5%	20.4%	0.5%	7.8%
Expected Return			7.0%	7.0%
Standard Deviation			10.9%	10.6%
Likelihood of Loss > 25%			0.2%	0.2%
Average Loss > 25%			27%	30%
Correlation Asymmetry (ξ)			4%	15%

NOTES: Expected returns are illustrative views. Standard deviations are annualized and estimated from monthly data over the period from January 1988 through December 2019. We solve for the mean–variance optimal portfolio using these inputs as well as correlations estimated over the same period as standard deviations. We solve for the full-scale optimal portfolio using a nonlinear search function to maximize the kinked utility function given by Equation 3. For k and w, we select values of –25% and 100, respectively, although the results are robust to reasonable variation in these parameters. We constrain expected return to equal 7% in both optimizations. We apply the full-scale algorithm to a simulated sample consisting of 1,000 years of return observations for each asset class. We construct each year in this sample by drawing 12 monthly multivariate return observations at random, with replacement, from the empirical sample (1988–2019). We then re-mean the 1,000-year sample to reflect our expected returns. Likelihood of loss is the frequency of portfolio returns falling below the k value of –25% in the 1,000-year simulated sample. Average loss is the average portfolio return for the subsample of portfolio returns that falls below this threshold. Correlation asymmetry reflects the weighted average summary metric, as given by Equation 4, for each portfolio. Correlation asymmetries are calculated over the 1988–2019 period.

classes exhibit unfavorable correlation profiles. Diversification often disappears on the downside when it is most needed, and, like an unwieldy umbrella on a sunny day, it is often present on the upside when it imparts no benefit. Finally, we show how investors can use full-scale optimization to construct portfolios that account explicitly for asymmetric correlation profiles and other nonnormal features of the distribution to maximize expected utility.

REFERENCES

Ang, A., and G. Bekaert. 2015. "International Asset Allocation with Regime Shifts." *Review of Financial Studies* 15 (4): 1137–1187.

Ang, A., and J. Chen. 2002. "Asymmetric Correlations of Equity Portfolios." *Journal of Financial Economics* 63 (3): 443–494.

Baumeister, R. F., E. Bratslavsky, C. Finkenauer, and K. D. Vohs. 2001. "Bad Is Stronger Than Good." *Review of General Psychology* 5 (4): 323–730.

Chua, D., M. Kritzman, and S. Page. 2009. "The Myth of Diversification." *The Journal of Portfolio Management* 36 (1): 26–35.

Cremers, J.-H., M. Kritzman, and S. Page. 2005. "Optimal Hedge Fund Allocations." *The Journal of Portfolio Management* 31 (3): 70–81.

Garcia-Feijóo, L., G. R. Jensen, and R. R. Johnson. 2012. "The Effectiveness of Asset Classes in Hedging Risk." *The Journal of Portfolio Management* 38 (3): 40–55.

Gulko, L. 2002. "Decoupling." *The Journal of Portfolio Management* 28 (3): 59–66.

Hartmann, P., S. Straetmans, and C. G. de Vries. 2004. "Asset Market Linkages in Crisis Periods." *Review of Economics and Statistics* 86 (1): 313–326.

——. 2010. "Heavy Tails and Currency Crises." *Journal of Empirical Finance* 17 (2): 241–254.

Huang, J.-Z., M. Rossi, and Y. Wang. 2015. "Sentiment and Corporate Bond Valuations before and after the Onset of the Credit Crisis." *The Journal of Fixed Income* 25 (1): 34–57.

Kritzman, M., K. Lowry, and A. Van Royen. 2001. "Risk, Regimes, and Overconfidence." *The Journal of Derivatives* 8 (3): 32–42.

Kritzman, M., S. Page, and D. Turkington. 2012. "Regime Shifts: Implications for Dynamic Strategies." *Financial Analysts Journal* 68 (3): 22–39.

Leibowitz, M. L., and A. Bova. 2009. "Diversification Performance and Stress-Betas." *The Journal of Portfolio Management* 35 (3): 41–47.

Longin, F., and B. Solnik. 2001. "Extreme Correlation of International Equity Markets." *The Journal of Finance* 56 (2): 649–676.

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

Page, S. 2020. *Beyond Diversification: What Every Investor Needs to Know About Asset Allocation*. McGraw Hill.

Page, S., and R. A. Panariello. 2018. "When Diversification Fails." *Financial Analysts Journal* 74 (3): 19–32.

Van Oordt, M. R. C., and C. Zhou. 2012. "The Simple Econometrics of Tail Dependence." *Economics Letters* 116 (3): 371–373.