PORTFOLIO
MANAGEMENT
RESEARCH

The Journal of Portfolio Management

VOLUME 36 NUMBER 1 JPM.pm-research.com FALL 2009

The Myth of Diversification

David B. Chua, Mark Kritzman, and Sébastien Page

The Myth of Diversification

David B. Chua, Mark Kritzman, and Sébastien Page

David B. Chua
is an assistant vice president at State Street Associates in Cambridge, MA.
dbchua@statestreet.com

Mark Kritzman
is president & CEO at Windham Capital Management in Cambridge, MA.
mkritzman@windhamcapital.com

Sébastien Page
is a senior managing director at State Street Associates in Cambridge, MA.
spage@statestreet.com

Perhaps the most universally accepted precept of prudent investing is to diversify, yet this precept grossly oversimplifies the challenge of portfolio construction. Consider a typical investor who relies on domestic equities to drive portfolio growth. This investor will seek to diversify this exposure by including assets that have low correlations with domestic equities. Yet the correlations, as typically measured over the full sample of returns, often belie an asset's diversification properties in market environments when diversification is most needed, for example, when domestic equities perform poorly. Moreover, upside diversification is undesirable; investors should seek unification on the upside. Ideally, the assets chosen to complement a portfolio's main engine of growth should diversify this asset when it performs poorly and move in tandem with it when it performs well.

In this article, we will first describe the mathematics of conditional correlations assuming returns are normally distributed. Then, we will present empirical results across a wide variety of assets, which reveal that, unlike the theoretical profiles, empirical correlations are significantly asymmetric. We are not the first to uncover correlation asymmetries, but our empirical investigation updates prior research and extends the analysis to a much broader set of assets. Finally, we will show that a portfolio construction technique called full-scale optimization produces portfolios in which the component assets exhibit relatively lower correlations on the downside and higher correlations on the upside than mean-variance optimization.

CORRELATION MATHEMATICS

It has been widely observed that correlations estimated from subsamples comprising volatile or negative returns differ from correlations estimated from the full sample of returns.[1] These differences do not necessarily imply that returns are non-normal or generated by more than a single regime. Consider a joint normal distribution with equal means of 0%, equal volatilities of 15%, and an unconditional (full-sample) correlation of 50%. Suppose we condition correlations on both assets, as illustrated in Exhibit 1, mathematically,

$$\rho(\theta) = \begin{cases} \operatorname{corr}(x, y \mid x > \theta, y > \theta) & \text{if } \theta > 0 \\ \operatorname{corr}(x, y \mid x < \theta, y < \theta) & \text{if } \theta < 0 \end{cases} \quad (1)$$

where the variables x and y are observed values for each asset, $\rho(\theta)$ is the conditional correlation, and θ is the threshold. Longin and Solnik [2001] labeled these conditional correlations exceedance correlations.

Exhibit 2 shows the exceedance correlation profile for x and y. We generated this profile using the closed-form solution presented in the appendix. The profile's peak shows that when we truncate the sample to include only observations when both x and y return a positive value (or when they both return a negative value), the correlation decreases from

EXHIBIT 1
Stylized Illustrations of Exceedance Correlations

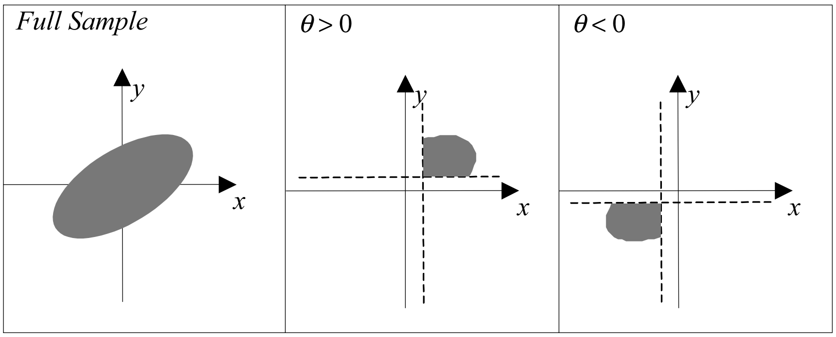

EXHIBIT 2
Correlation Profile for Bivariate Normal Distribution

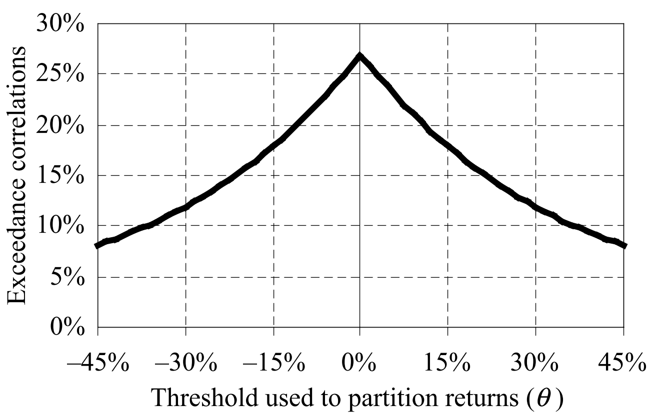

50% to 27%. Moving toward the tails of the distribution, the correlation decreases further. For example, if we focus on observations when x and y both return –15% or less, the correlation decreases to 18%. This effect, which is called the conditioning bias, may lead us to conclude falsely that diversification increases during extreme market conditions.

To avoid the conditioning bias in our empirical analysis, we compared empirical correlation profiles with those expected from the normal distribution. Our goal was to determine whether correlations shift because returns are generated from more than a single distribution or if they differ simply as an artifact of correlation mathematics.

EMPIRICAL CORRELATION PROFILES

Prior studies have shown that exposure to different country equity markets offers less diversification in down markets than in up markets.[2] The same is true for global industry returns (Ferreira and Gama [2004]); individual stock returns (Ang, Chen, and Xing [2002], Ang and Chen [2002], and Hong, Tu, and Zhou [2007]); hedge fund returns (Van Royen [2002b]); and international bond market returns (Cappiello, Engle, and Sheppard [2006]). The only conditional correlations that seem to improve diversification when it is most needed are the correlations across asset classes. Kritzman, Lowry, and Van Royen [2001] found that asset class correlations within countries *decrease* during periods of market turbulence.[3] Similarly, Gulko [2002] found that stocks and bonds decouple during market crashes.

We extended the investigation of correlation asymmetries to a comprehensive dataset comprising equity, style, size, hedge fund, and fixed-income index returns. Exhibit 3 shows the complete list of the data series we investigated. For most asset classes, we generated correlation profiles with the U.S. equity market because it is the main engine of growth for most institutional portfolios. We also generated correlation profiles for large versus small stocks, value versus growth stocks, and various combinations of fixed-income assets.

To account for differences in volatilities, we standardized each time series as follows: (x – mean)/standard deviation. We then used the closed-form solution outlined in the appendix to calculate the corresponding normal correlation profiles.

Exhibit 4 shows the correlation profile between the U.S. (Russell 3000) and MSCI World Ex-U.S. equity markets, as well as the corresponding correlations we obtained by partitioning a bivariate normal distribution with the same means, volatilities, and unconditional correlation. Observed correlations are higher than normal correlations on the downside and lower on the upside; in other words, international diversification works during good times—when it is not needed—and disappears during down markets. When both markets are up by more than one standard deviation, the correlation between them is –17%. When both markets are *down* more than one standard deviation, the correlation between them is +76%. And in times of extreme crisis, when both markets are down by more than two standard deviations, the correlation rises to +93% compared to +14% for the corresponding bivariate normal distribution.

We investigated the pervasiveness of correlation asymmetry across several asset classes. For each correlation

EXHIBIT 3
Data Sources

	Source	Start Date	End Date
Equity Markets			
U.S.	Russell 3000 or MSCI	1/31/1970	2/29/2008
U.K.	MSCI	1/30/1970	2/29/2008
France	MSCI	1/30/1970	2/29/2008
Germany	MSCI	1/30/1970	2/29/2008
Japan	MSCI	1/30/1970	2/29/2008
World Ex-U.S.	MSCI	1/30/1970	2/29/2008
Style and Size Indices			
U.S. Value	Russell 3000 Value	1/31/1979	2/29/2008
U.S. Growth	Russell 3000 Growth	1/31/1979	2/29/2008
U.S. Large Cap	S&P 500	1/29/1988	2/29/2008
U.S. Small + Mid Cap	Russell 2500	1/31/1979	2/29/2008
Hedge Funds			
Event Driven	HFR	4/30/2003	2/29/2008
Relative Value Arbitrage	HFR	4/30/2003	2/29/2008
Convertible Arbitrage	HFR	4/30/2003	2/29/2008
Equity Market Neutral	HFR	4/30/2003	2/29/2008
Merger Arbitrage	HFR	4/30/2003	2/29/2008
Global Hedge Fund	HFR	4/30/2003	2/29/2008
Fixed-Income Indices			
U.S. Bonds	Lehman Aggregate	1/31/1979	2/29/2008
Government	Lehman	1/31/1979	2/29/2008
High Yield	Lehman	1/31/1979	2/29/2008
Mortgage-Backed Securities	Lehman	1/31/1979	2/29/2008
Credit BAA	Lehman	1/31/1979	2/29/2008

EXHIBIT 4
Correlation Profile between U.S. and World Ex-U.S., January 1979–February 2008

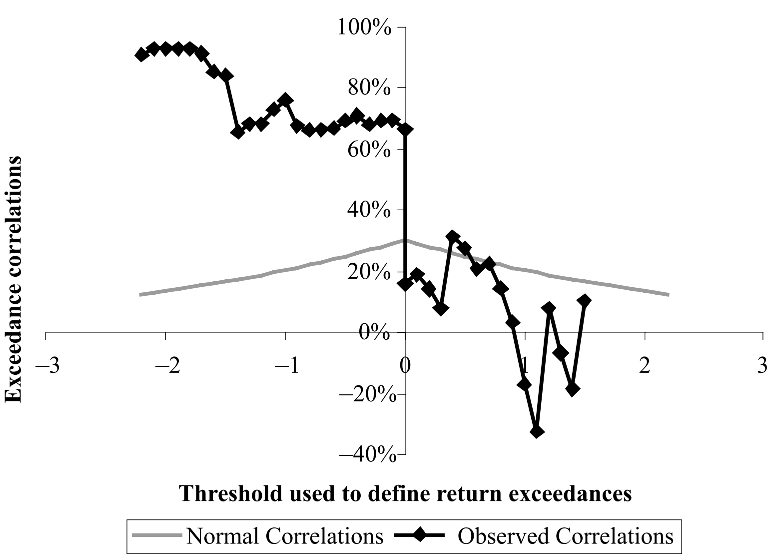

profile, we calculate the average difference between *observed* and *normal* exceedance correlations.

We calculate these average differences for up and down markets, as follows:

$$\mu_{dn} = \frac{1}{n}\sum_{i=1}^{n}[\breve{\rho}(\theta_i) - \rho(\theta_i)] \quad \forall\, \theta_i < 0$$
$$\mu_{up} = \frac{1}{n}\sum_{i=1}^{n}[\breve{\rho}(\theta_i) - \rho(\theta_i)] \quad \forall\, \theta_i > 0 \tag{2}$$

The variables μ_{dn} and μ_{up} are the average differences for down and up markets, respectively; $\breve{\rho}(\theta_i)$ is the observed exceedance correlation at threshold θ_i; and $\rho(\theta_i)$ is the corresponding normal correlation. For up and down markets, we use n thresholds ($\theta_i > 0$ and $\theta_i < 0$, respectively) equally spaced by intervals of 0.1 standard deviations.[4] Exhibit 5 summarizes our results in detail and Exhibit 6 shows the corresponding ranking of correlation asymmetries defined as $\mu_{dn} - \mu_{up}$. We find that correlation asymmetries prevail across a variety of indices. The correlation profile for equity-market-neutral hedge funds raises questions about such claims of market neutrality. Only a few asset classes offer desirable downside diversification, or *decoupling*. And unfortunately, most of these asset classes—MBS, high yield, and credit—failed to diversify each other during the recent subprime and credit-crunch crisis of 2007–2008.

The last column of Exhibit 5 shows the percentage of paths—from a block bootstrap of all possible 10-year paths—for which asymmetry ($\mu_{dn} - \mu_{up}$) has the same sign as the full-sample result shown in the third column. We find a reasonable degree of persistence in the directionality of correlation asymmetries, which suggests historical correlation profiles should be useful in constructing portfolios.

PORTFOLIO CONSTRUCTION WITH ASYMMETRIC CORRELATIONS

How should we construct portfolios if downside correlations are higher than upside correlations? Research on conditional correlations has led

EXHIBIT 5
Average Excess Correlations versus Normal Distribution (%)

		Downside (μ_{dn})	Upside (μ_{up})	Difference ($\mu_{dn} - \mu_{up}$)	Prob. of Same Sign
Average excess correlations of U.S. Equities (Russell 3000) with	**Equity Markets**				
	U.K.	38.88	–23.18	62.05	78.35
	France	47.06	–21.33	68.39	74.46
	Germany	15.63	–15.15	30.78	100.00
	Japan	6.80	23.44	–16.65	59.74
	World Ex-U.S.	53.49	–4.26	57.75	95.67
	Hedge Funds				
	Event Driven	24.26	–0.14	24.12	NA
	Relative Value Arbitrage	65.76	–30.78	96.54	NA
	Convertible Arbitrage	–24.96	6.69	–31.65	NA
	Equity Market Neutral	66.44	–18.41	84.85	NA
	Merger Arbitrage	–40.52	–45.66	5.13	NA
	Global Hedge Fund	14.60	–31.87	46.47	NA
	Fixed-Income Indices				
	U.S. Bonds	10.33	6.68	3.65	12.12
	Government	26.08	–4.21	30.30	9.09
	High Yield	–14.06	–17.67	3.62	77.27
	Mortgage-Backed Securities	24.44	–0.71	25.16	59.31
	Credit BAA	–4.36	15.19	–19.56	82.68
Other average excess correlations	**Style and Size**				
	Value vs. Growth	34.99	–16.22	51.22	100.00
	Large vs. Smid	31.03	–21.84	52.88	100.00
	Fixed-Income Indices				
	Government vs. High Yield	10.21	48.15	–37.93	45.45
	Government vs. MBS	48.04	85.41	–37.37	48.92
	Government vs. Credit BAA	72.55	83.08	–10.52	51.52
	Credit BAA vs. MBS	–5.86	65.13	–70.99	40.26
	High Yield vs. MBS	24.29	23.37	0.92	65.91

to several innovations in portfolio construction. In general, conditional correlations lead to more conservative portfolios than unconditional correlations. For example, Campbell, Koedijk, and Kofman [2002] presented two efficient frontiers for the allocation between domestic stocks (S&P 500) and international stocks (FTSE 100). The first frontier uses the full-sample historical returns, volatilities, and correlation; the second substitutes the downside correlation for the full-sample correlation. Campbell, Koedijk, and Kofman found that for the same level of risk the downside-sensitive allocation to international stocks is 6% lower and cash increases from 0% to 10.5%. We present similar evidence of these shifts by deriving optimal allocations based on downside, upside, and full-sample correlations.

Exhibit 7 shows expectations and optimal portfolios for U.S. equities, World Ex-U.S. equities, and cash. Downside correlations lead to an increase in cash allocation from 3% to 9% and a higher allocation to World Ex-U.S. equities, while upside correlations lead to an all-equity portfolio with a higher allocation to the U.S. market.

Exhibit 7 also shows the impact of correlations on expected utility. For example, the portfolio constructed on downside correlations—holding everything else constant—has higher utility (0.0633) than the portfolios constructed on unconditional correlations (0.0630) and upside correlations (0.0627), if the downside correlations are realized.

Another approach for addressing correlation asymmetry is to dynamically change our correlation assumptions.

EXHIBIT 6

Correlation Asymmetries ($\mu_{dn} - \mu_{up}$) Ranked

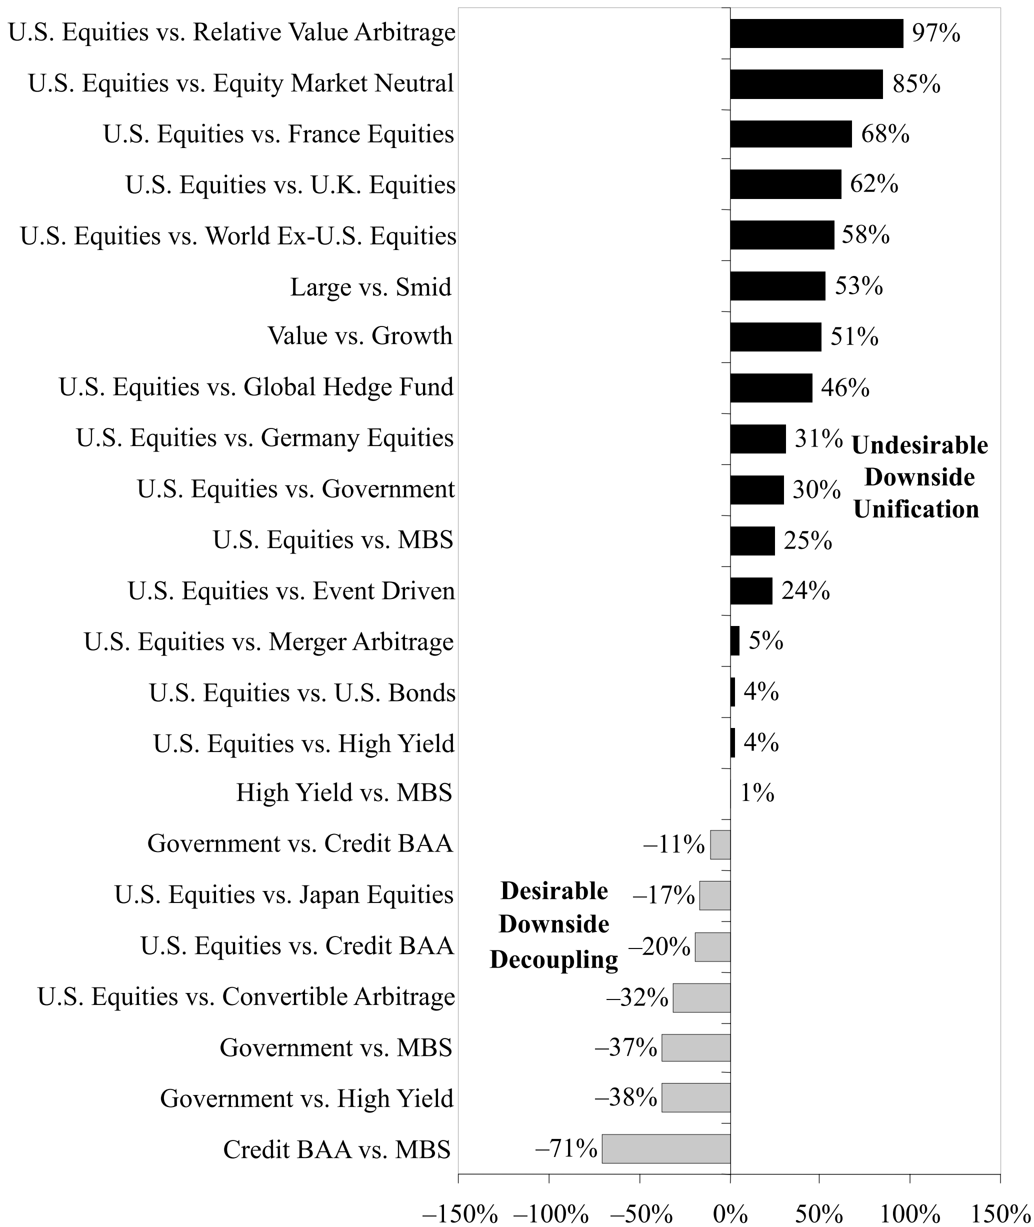

Previous research suggests that regime-switching models are ideally suited to handle correlation asymmetries in sample.[5] From a practitioner's perspective, the question remains whether regime shifts are predictable. For example, Gulko [2002] suggested a mean-variance regime-switching model that calls for the investor to switch to an all-bond portfolio for one month following a one-day crash. Unfortunately, the superiority of this strategy is unclear—the author's model relies on six events and assumes that bonds outperform stocks for one month after a market crash.

Our goal was not to develop active trading strategies, thus we did not try to predict regime shifts. Instead, we focused on strategic asset allocation with the argument that portfolios should be modeled after airplanes, which means they should be able to withstand turbulence whenever it arises, because it is usually unpredictable.

EXHIBIT 7
Mean-Variance Optimization with Conditional Correlations

Inputs	U.S.	World Ex-U.S.	Cash
Expected Return	8.00%	10.00%	3.50%
Volatility	15.08%	16.45%	0.88%
Unconditional Correlation	57.17%		
Down-Down Correlation	67.30%		
Up-Up Correlation	44.60%		

Optimal Portfolios	U.S.	World Ex-U.S.	Cash
Unconditional Correlation	25%	72%	3%
Down-Down Correlation	14%	77%	9%
Up-Up Correlation	32%	68%	0%

Expected Utility	Full Sample	Down-Down Market	Up-Up Market
Unconditional Correlation	0.0643	0.0630	0.0659
Down-Down Correlation	0.0640	0.0633	0.0649
Up-Up Correlation	0.0642	0.0627	0.0661

Strategic investors, such as pension plans, are just like airline pilots—their goal is not to predict the unpredictable, but their portfolios should weather the storms.

FULL-SCALE OPTIMIZATION

Our approach, called full-scale optimization (Cremers, Kritzman, and Page [2005] and Adler and Kritzman [2007]), identifies portfolios that are more resilient to turbulence because they have better correlation profiles. It does not seek to exploit correlation asymmetries directly, but instead maximizes expected utility. By so doing, our approach constructs portfolios in which assets diversify each other more on the downside and move together more on the upside than portfolios derived from mean-variance analysis.

In contrast to mean-variance analysis, which assumes that returns are normally distributed or that investors have quadratic utility, full-scale optimization identifies the optimal portfolio given any set of return distributions and any description of investor preferences. It therefore yields the truly optimal portfolio in sample, whereas mean-variance analysis provides an approximation to the in-sample truth.

We apply mean-variance and full-scale optimization to identify optimal portfolios assuming a loss-averse investor with a kinked utility function. This utility function changes abruptly at a particular wealth or return level and is relevant for investors who are concerned with breaching a threshold. Consider, for example, a situation in which an investor requires a minimum level of wealth to maintain a certain standard of living. The investor's lifestyle might change drastically if the fund penetrates this threshold. Or the investor may be faced with insolvency following a large negative return, or a particular decline in wealth may breach a covenant on a loan. In these and similar situations, a kinked utility function is more likely to describe an investor's attitude toward risk than a utility function that changes smoothly. The kinked utility function is defined as

$$U(x)=\begin{cases}\ln(1+x), & \text{if } x\geq\theta\\ \upsilon\times(x-\theta)+\ln(1+\theta), & \text{if } x<\theta\end{cases} \tag{3}$$

where θ indicates the location of the kink and υ the steepness of the loss aversion slope.

In Exhibit 8, we illustrate the full-scale optimization process with a sample of stock and bond returns. We compute the portfolio return, x, each period as $R_S \times W_S + R_B \times W_B$, where R_S and R_B equal the stock and bond returns, respectively, and W_S and W_B equal the stock and bond weights, respectively. We use Equation 3 to compute utility in each period, with $\theta = -3\%$ and $\upsilon = 3$.

We then shift the stock and bond weights until we find the combination that maximizes expected utility, which for this example equals a 48.28% allocation to stocks and a 51.72% allocation to bonds. The expected utility of the portfolio equals 0.991456. This approach implicitly takes into account all features of the empirical sample, including possible skewness, kurtosis, and any other peculiarities of the distribution, such as correlation asymmetries.

To minimize the probability of breaching the threshold, full-scale optimization avoids assets that exhibit high downside correlation, all else being equal. Exhibit 9 shows an example using fictional distributions generated by Monte Carlo simulation. It shows an optimization between a fictional equity portfolio and a fictional market neutral hedge fund. Mean-variance is oblivious to the hedge fund's highly undesirable correlation profile because it focuses on the full-sample correlation of 0% as a proxy for the hedge fund's diversification potential. It invests half the portfolio in the hedge fund and, as a consequence, doubles the portfolio exposure to losses greater than –10%. In contrast, full-scale optimization chooses not to invest in the hedge fund at all.[6]

Using empirical data, we identified optimal country portfolios by allocating across the U.S., U.K., France, Germany, and Japan, and we also optimized across the entire sample of 20 asset classes presented in Exhibit 3.[7] In each case, we evaluated the mean-variance efficient portfolio with the same expected return as the true utility-maximizing portfolio. For purposes of illustration, we set the kink (θ) equal to a one-month return of –4%.

Before we optimized, we scaled each of the monthly returns by a constant in order to produce means that conform to the implied returns of equally weighted portfolios. This adjustment does not affect our comparisons because we apply it to both the mean-variance and full-scale optimizations; hence, the differences we find arise solely from the higher moments of the distributions and not their means.

We measure the degree of correlation asymmetry (ξ) in each of the optimal portfolios as

$$\xi = \sum_i w_i \bar{\rho}_i^{dn} - \sum_i w_i \bar{\rho}_i^{up} \tag{4}$$

where w_i is asset i's weight in the portfolio, $\bar{\rho}_i^{dn}$ is the weighted average of correlations between asset i and the other assets in the portfolio when the portfolio is down, and $\bar{\rho}_i^{up}$ is the weighted average of correlations between asset i and the other assets in the portfolio when the portfolio is up.

Exhibit 10 shows the correlation asymmetry for country allocation portfolios using data from January 1970 to February 2008. It reveals the following:

1. Downside correlations are significantly higher than upside correlations.
2. Full-scale optimization provides more downside diversification and less upside diversification than mean-variance optimization.
3. Full-scale optimization reduces correlation asymmetry by more than half compared to mean-variance optimization.

EXHIBIT 8
Full-Scale Optimization

	Returns		Weights		Portfolio	
	Stocks	Bonds	Stocks	Bonds	Return	Utility
1993	10.06%	16.16%	48.28%	51.72%	13.21%	0.1241
1994	1.32%	–7.10%	48.28%	51.72%	–3.03%	–0.0315
1995	37.53%	29.95%	48.28%	51.72%	33.61%	0.2898
1996	22.93%	14.00%	48.28%	51.72%	18.31%	0.1682
1997	33.34%	14.52%	48.28%	51.72%	23.61%	0.2119
1998	28.60%	11.76%	48.28%	51.72%	19.89%	0.1814
1999	20.89%	–7.64%	48.28%	51.72%	6.13%	0.0595
2000	–9.09%	16.14%	48.28%	51.72%	3.96%	0.0388
2001	–11.94%	7.26%	48.28%	51.72%	–2.01%	–0.0203
2002	–22.10%	14.83%	48.28%	51.72%	–3.00%	–0.0305
					Utility	**0.9915**

EXHIBIT 9
Impact of Correlation Asymmetries on Full-Scale and Mean-Variance Portfolios

	Expected Return	Volatility	Unconditional Correlation	Downside Correlation	Upside Correlation	Asymmetry
Equity Portfolio	8%	12%	2%	85%	–30%	114%
Hedge Fund	8%	12%				

	Optimal Portfolios	
	Mean-Variance	Full-Scale
Equity Portfolio	50%	100%
Hedge Fund	50%	0%
Exposure to Loss (Prob. <–10%)	0.250%	0.125%

EXHIBIT 10
Weighted-Average Correlations for Country Allocation, January 1970–February 2008

Weighted-Average Correlation	Mean-Variance Optimal	Full-Scale Optimal
Upside	9.34%	11.07%
Downside	25.25%	17.96%
Correlation Asymmetry (ξ)	**15.91%**	**6.89%**

EXHIBIT 11

Full-Scale vs. Mean-Variance Optimization and Correlation Asymmetries (%)

		Correlation Asymmetry (ζ)			Utility	Turnover
Country Allocation		**MV**	**FS**	**FS Advantage**	**Gain (%)**	**Required (%)**
Full Sample	1/1970–2/2008	15.91	6.89	**9.02**	4.33	14.80
Subsamples	1/1970–12/1979	19.35	3.25	**16.10**	9.96	22.47
	1/1980–12/1989	7.02	1.75	**5.27**	23.40	22.40
	1/1990–12/1999	25.77	18.71	**7.06**	1.57	19.07
	1/1999–2/2008	0.52	0.36	**0.16**	3.59	9.67
Multi-Asset Class						
Full Sample	2/1988–2/2008	7.45	–0.41	**7.86**	29.11	91.91
Subsamples	2/1988–12/1998	58.42	49.67	**8.75**	16.35	19.52
	1/1999–2/2008	13.30	–3.46	**16.77**	37.62	91.21
	4/2003–2/2008*	3.58	–6.99	**10.58**	254.11	64.97

*Note: *Sample includes hedge funds.*
Proxy for U.S. Equities is MSCI USA.

Exhibit 11 shows the utility gain obtained by moving from the mean-variance to the full-scale optimal portfolio, as well as the turnover required to do so. Also, it shows results for subsamples and for the multi-asset optimization. In all cases, full-scale optimization improves correlation asymmetry when compared to mean-variance optimization. In some cases, full-scale optimization turns correlation asymmetry from positive to negative, which means the portfolio has lower downside correlations than upside correlations. Utility gains are highest for the multi-asset optimization. When hedge funds are included, the gain reaches 254.11%.

Note that utility gains are not perfectly correlated with improvements in correlation asymmetry—in some cases, large improvements in correlation asymmetry lead to a relatively small utility gain. For example, looking at country allocation, full-scale optimization improves correlation asymmetry by a greater amount in the 1970–1979 subsample than in the 1980–1989 subsample (16.10% versus 5.27%), while the improvement in utility is lower (9.96% versus 23.40%). This result occurs because the kinked utility function puts a greater premium on downside diversification than upside unification. Also, when we take into account other features of the distribution, all asymmetry improvements are not created equal in terms of expected utility. For example, Statman and Scheid [2008] found that return gaps provide a better definition of diversification because they include volatility. Full-scale optimization addresses this issue by using the entire return distribution. It is even possible to observe deterioration in correlation asymmetry associated with an increase in utility. But overall, our results show that to the extent that correlations are an important driver of utility—as opposed to volatilities, or returns, or other features of the joint distribution—an improvement in correlation asymmetry will lead to an improvement in utility.

In general, the stability of our results is not surprising given the reasonable level of stability shown in the underlying correlation profiles (Exhibit 5). Although our main goal is to solve the problem of utility maximization in sample, Adler and Kritzman [2007] provided a robust demonstration that full-scale optimization outperforms mean-variance optimization out of sample. Our findings help explain their results. We suggest that a significant portion of this out-of-sample outperformance comes from improvements in correlation profiles. In other words, conditional correlations matter—and mean-variance optimization fails to take them into consideration.

CONCLUSION

We measured conditional correlations to assess the extent to which assets provide diversification in down markets and allow for unification during up markets. We first derived conditional correlations analytically under the assumption that returns are jointly normally distributed in order to measure the theoretical bias we should expect from conditional correlations. We then measured conditional correlations from empirical returns, which revealed that correlation asymmetry is prevalent across a wide range of asset pairs. Finally, we turned to portfolio construction. We showed that conventional approaches to portfolio construction ignore correlation asymmetry, while full-scale optimization, which directly maximizes expected utility over a sample of returns, generates portfolios with more downside diversification and upside unification than alternative approaches to portfolio formation.

APPENDIX

Conditional Correlation for Bivariate Normal Distributions

Let $X = (x, y) \sim N(0, \Sigma)$ be a bivariate normal random variable, where x and y have unit variances and unconditional correlation ρ. Then the correlation of x and y conditional on $x < h$ and $y < k$ is given by

$$\operatorname{corr}(x, y \mid x<h,\ y<k)=\frac{\operatorname{cov}(x, y \mid x<h,\ y<k)}{\sqrt{\operatorname{var}(x \mid x<h,\ y<k)\ \operatorname{var}(y \mid x<h,\ y<k)}}$$

The variances and covariance of these conditional random variables can be written in terms of the moments $m_{ij} = E[x^i y^j \mid x < h, y < k]$ as

$$\begin{aligned} \operatorname{var}(x \mid x<h,\ y<k)&=m_{20}-m_{10}^2 \\ \operatorname{var}(y \mid x<h,\ y<k)&=m_{02}-m_{01}^2 \\ \text{and}\quad \operatorname{cov}(x, y \mid x<h,\ y<k)&=m_{11}-m_{10}\ m_{01} \end{aligned}$$

Let $L(h, k)$ denote the cumulative density of our bivariate normal distribution,

$$L(h,\ k)=\frac{1}{2\pi\sqrt{1-\rho^2}}\int_{-\infty}^{k}\int_{-\infty}^{h}\exp\left(-\frac{x^2-2\rho xy+y^2}{2(1-\rho^2)}\right)dx\,dy$$

As shown in Ang and Chen [2002], the first and second moments can be expressed as

$$\begin{aligned} L(h,k)\,m_{10}&=\psi(h,k;\rho)+\rho\,\psi(k,h;\rho) \\ L(h,k)\,m_{01}&=\psi(k,h;\rho)+\rho\,\psi(h,k;\rho) \\ L(h,k)\,m_{20}&=L(h,k)-\chi(k,h;\rho)-\rho^2\chi(h,k;\rho) \\ L(h,k)\,m_{02}&=L(h,k)-\chi(h,k;\rho)-\rho^2\chi(k,h;\rho) \\ \text{and}\quad L(h,k)\,m_{11}&=\rho\,L(h,k)+\rho\,h\,\psi(h,k;\rho)+\rho\,k\,\psi(k,h;\rho) \\ &\quad-\Lambda(h,k;\rho) \end{aligned}$$

where $\psi(\cdot)$, $\Lambda(\cdot)$, and $\chi(\cdot)$ are given by

$$\begin{aligned} \psi(h,k;\rho)&=-\phi(h)\,\Phi\left(\frac{k-\rho h}{\sqrt{1-\rho^2}}\right) \\ \Lambda(h,k;\rho)&=-\frac{\sqrt{1-\rho^2}}{\sqrt{2\pi}}\,\phi\left(\frac{\sqrt{h^2-2\rho hk+k^2}}{\sqrt{1-\rho^2}}\right) \\ \chi(h,k;\rho)&=-k\psi(k,h;\rho)+\frac{\rho}{(1+\rho^2)}\Lambda(h,k;\rho) \end{aligned}$$

and $\phi(\cdot)$ and $\Phi(\cdot)$ are the PDF and CDF, respectively, of the univariate standard normal distribution.

ENDNOTES

[1]See, for example, Longin and Solnik [2001], Kritzman, Lowry, and Van Royen [2001], and Van Royen [2002b].

[2]See, for example, Ang and Bekaert [2002], Kritzman, Lowry, and Van Royen [2001], Baele [2003], and Van Royen [2002a] on regime shifts; Van Royen [2002a] and Hyde, Bredin, and Nguyen [2007] on financial contagion; and Ang and Bekaert [2002], Longin and Solnik [2001], Butler and Joaquin [2002], Campbell, Koedijk, and Kofman [2002], Cappiello, Engle, and Sheppard [2006], and Hyde, Bredin, and Nugyen [2007] on correlation asymmetries.

[3]Kritzman, Lowry, and Van Royen [2001] condition on a statistical measure of market turbulence rather than return thresholds.

[4]Exceedance correlations are computed for any threshold with more than three events.

[5]See, for example, Ramchand and Susmel [1998], Ang and Bekaert [2002], and Ang and Chen [2002].

[6]Note that full-scale optimization will not always produce the most concentrated portfolio. It might invest a significant proportion of a portfolio in an asset with high full-sample correlation, but negative downside correlation, while mean-variance might ignore the asset altogether.

[7]We exclude the World Ex-U.S. asset class because it is redundant with the country choices.

REFERENCES

Adler, T., and M. Kritzman. "Mean-Variance versus Full-Scale Optimisation: In and Out of Sample." *Journal of Asset Management*, Vol. 7, No. 5 (2007), pp. 302–311.

Ang, A., and G. Bekaert. "International Asset Allocation with Regime Shifts." *Review of Financial Studies*, Vol. 15, No. 4 (2002), pp. 1137–1187.

Ang, A., and J. Chen. "Asymmetric Correlations of Equity Portfolios." *Journal of Financial Economics*, Vol. 63, No. 3 (2002), pp. 443–494.

Ang, A., J. Chen, and Y. Xing. "Downside Correlation and Expected Stock Returns." Working Paper, Columbia University, 2002.

——. "Downside Risk." *Review of Financial Studies*, Vol. 19, No. 4 (2006), pp. 1191–1239.

Baele, L. "Volatility Spillover Effects in European Equity Markets: Evidence from a Regime-Switching Model." Working Paper No. 33, United Nations University, Institute for New Technologies, 2003.

Billio, M., R. Casarin, and G. Toniolo. "Extreme Returns in a Shortfall Risk Framework." Working Paper No. 0204, GRETA, 2002.

Butler, K.C., and D.C. Joaquin. "Are the Gains from International Portfolio Diversification Exaggerated? The Influence of Downside Risk in Bear Markets." *Journal of International Money and Finance*, Vol. 21, No. 7 (2002), pp. 981–1011.

Campbell, R., K. Koedijk, and P. Kofman. "Increased Correlation in Bear Markets." *Financial Analysts Journal*, Vol. 58, No. 1 (2002), pp. 87–94.

Cappiello, L., R.F. Engle, and K. Sheppard. "Asymmetric Dynamics in the Correlations of Global Equity and Bond Returns." *Journal of Financial Econometrics*, Vol. 4, No. 4 (2006), pp. 385–412.

Chow, G., E. Jacquier, M. Kritzman, and K. Lowry. "Optimal Portfolios in Good Times and Bad." *Financial Analysts Journal*, Vol. 55, No. 3 (1999), pp. 65–73.

Cremers, J.-H., M. Kritzman, and S. Page. "Optimal Hedge Fund Allocations: Do Higher Moments Matter?" *Journal of Portfolio Management*, Vol. 31, No. 3 (2005), pp. 70–81.

Ferreira, M.A., and P.M. Gama. "Correlation Dynamics of Global Industry Portfolios." Working Paper, ISCTE Business School, Lisbon, 2004.

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

Hong, Y., J. Tu, and G. Zhou. "Asymmetries in Stock Returns: Statistical Tests and Economic Evaluation." *Review of Financial Studies*, Vol. 20, No. 5 (2007), pp. 1547–1581.

Hyde, S., D. Bredin, and N. Nguyen. "Correlation Dynamics between Asia-Pacific, EU and US Stock Returns." In *Asia-Pacific Financial Market: Integration, Innovation and Challenges*, *International Finance Review*, 8 (2007), pp. 39–61.

Kritzman, M., K. Lowry, and A.-S. Van Royen. "Risk, Regimes, and Overconfidence." *Journal of Derivatives*, Vol. 8, No. 3 (2001), pp. 32–43.

Lin, W.-L., R.F. Engle, and T. Ito. "Do Bulls and Bears Move Across Borders? International Transmission of Stock Returns and Volatility." *Review of Financial Studies*, Vol. 7, No. 3 (1994), pp. 507–538.

Longin, F., and B. Solnik. "Is the Correlation in International Equity Returns Constant: 1960–1990?" *Journal of International Money and Finance*, Vol. 14, No. 1 (1995), pp. 3–26.

——. "Extreme Correlation of International Equity Markets." *Journal of Finance*, Vol. 56, No. 2 (2001), pp. 649–676.

Markowitz, H. "Portfolio Selection." *Journal of Finance*, Vol. 7, No. 1 (1952), pp. 77–91.

Marshal, R., and A. Zeevi. "Beyond Correlation: Extreme Co-Movements between Financial Assets." Working Paper, Columbia University, 2002.

Ramchand, L., and R. Susmel. "Volatility and Cross Correlation across Major Stock Markets." *Journal of Empirical Finance*, Vol. 5, No. 4 (1998), pp. 397–416.

Solnik, B., C. Boucrelle, and Y. Le Fur. "International Market Correlation and Volatility." *Financial Analysts Journal*, Vol. 52, No. 5 (1996), pp. 17–34.

Statman, M., and J. Scheid. "Correlation, Return Gaps, and the Benefits of Diversification." *Journal of Portfolio Management*, Vol. 34, No. 3 (2008), pp. 132–139.

Van Royen, A.-S. "Financial Contagion and International Portfolio Flows." *Financial Analysts Journal*, Vol. 58, No. 1 (2002a), pp. 35–49.

——. "Hedge Fund Index Returns." *Hedge Fund Strategies*, 36 (2002b), pp. 111–117.

To order reprints of this article, please contact David Rowe at d.rowe@pageantmedia.com or 646-891-2157.