Risk, Regimes, and Overconfidence

Mark Kritzman, WCMB
mkritzman@wcmbllc.com

Kenneth Lowry, State Street Global Markets
CELowry@statestreet.com

Dr. Anne-Sophie VanRoyen, WCMB and State Street Associates

Published in *The Journal of Derivatives*, Spring 2001

272-2 Revere Street Working Papers

Risk, Regimes, and Overconfidence

Mark Kritzman, Kenneth Lowry, and Anne-Sophie Van Royen

Mark Kritzman is a managing partner at Windham Capital Management Boston in Cambridge, MA.

Kenneth Lowry is vice president of State Street Bank in Boston, MA.

Anne-Sophie Van Royen is director of research at Windham Capital Management Boston and is senior quantitative strategist for State Street Associates in Cambridge, MA.

Investors typically think of risk as the uncertainty of wealth at the end of their investment horizon. By focusing on the dispersion of ending wealth, investors ignore the effect of interim losses, no matter how severe. Investors also measure risk as though returns come from a single regime, which may understate the likelihood and severity of interim losses.

The authors argue that the perception of risk as fully represented by the distribution of terminal wealth, together with the assumption of a single regime, leads to overconfidence. They apply first-passage probabilities to compare the risk of loss during an investment period with the risk of loss at the end of a horizon. Application of a methodology to measure risk based on quiet or turbulent regimes shows the extent to which the traditional measurement of risk understates exposure to loss.

The authors present a forecasting procedure to assess the relative likelihood of quiet and turbulent regimes, and show how to use this information to structure portfolios that are regime-sensitive.

When investors think of risk as the uncertainty of wealth at the end of their investment horizon, they implicitly ignore the effect of interim losses, no matter how severe. Investors also measure risk as though returns come from a single regime characterized by a lognormal distribution.[1] If, instead, returns are generated by two regimes, one quiet and the other turbulent, severe interim underperformance may be more likely to occur than is implied by the full-sample covariance matrix.

We argue that the perception of risk as depending only on the distribution of terminal wealth, together with the assumption of a single regime, leads to overconfidence. We propose a methodology for structuring portfolios that are regime-sensitive.

The methodology, which identifies multivariate outliers according to work by Chow, Jacquier, Kritzman, and Lowry [1999], is better than an approach that stratifies returns solely on the basis of volatility. We present evidence that belies the assumption of a single lognormal distribution, and we measure the extent to which this assumption, together with emphasis on ending wealth, understates interim risk. Regime-sensitive portfolios may be structured using a model for forecasting the relative probability of quiet and turbulent regimes.

I. RISK OF LOSS

In Exhibit 1 we compare the distinction between the risk of loss if we focus on ending wealth instead of on interim wealth. Each line represents the path of a hypothetical investment through four periods. The horizontal line represents a loss threshold, which in this example equals 10%. Only one of the five paths breaches the loss threshold as of the end of the horizon; hence we might conclude that the likelihood of a 10% loss equals 20%.

Four of the five paths breach the loss

EXHIBIT 1

Risk of Loss: Ending Wealth versus Interim Wealth

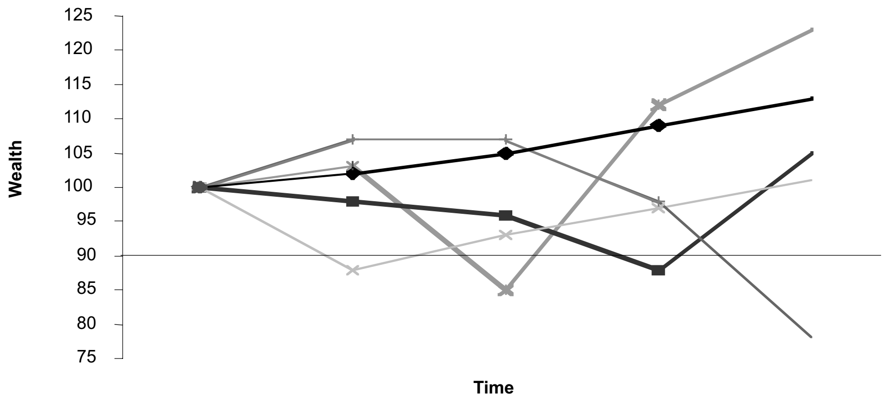

threshold at some point during or at the end of the investment period, however, although three of the four paths subsequently recover. If we care about the investment's performance along the way to the end of our horizon, we should conclude that the likelihood of a 10% loss equals 80%.

This example illustrates the distinction between the riskiness of ending wealth and interim wealth. To quantify this distinction, we use a statistic called first-passage time probability.[2] For a lognormal diffusion with an initial value of S at time 0, the probability that it will first reach a given critical value C at or before date T is given by:

$$N[(\ln(C/S) - \mu T)/(\sigma\sqrt{T})] + (C/S)^{2\mu/\sigma^2} \times N[(\ln(C/S) + \mu T)/(\sigma\sqrt{T})] \quad (1)$$

where:

N[ ] = cumulative normal distribution function;
ln = natural logarithm;
μ = continuous return; and
σ = continuous standard deviation.

We use the first-passage time statistic to estimate the probability that an investment will depreciate to a particular value over some horizon if it is monitored continually. Consider, for example, an investment that has an expected return of 10% and a standard deviation of 20%. Exhibit 2 compares the likelihood of a 10% loss over a one-, five-, and ten-year investment horizon if we monitor the investment only at the end instead of continually throughout each period.

There are two interesting points to note from Exhibit 2. First, the likelihood of a 10% loss during an investment horizon is much greater than the likelihood of a 10% loss at the end of the horizon. And second, while the likelihood of loss at the end of a horizon diminishes as the horizon is extended from one to ten years, the likelihood of loss during the horizon rises with its duration.

The critical question, of course, is whether we have the forbearance to weather severe interim losses and to sustain our exposure to risky assets. If not, our expectation for the strategy's long-term performance may be misguided. Moreover, it is often the case that this decision is not ours to make. If we are managing someone else's assets, we must be confident that our client will not force us to abandon the agreed-upon strategy should it experience severe losses along the way.

II. IDENTIFYING OUTLIERS

The common assumption is that returns are generated by a single regime characterized by a lognormal distribution. One of the implications of this assumption is that the correlations we estimate do not change as we sample different returns, as long as the samples are sufficiently large. If, however, returns are generated by different regimes, we must estimate regime-specific correlations if we are to accurately assess the risk of loss during an investment period.

We invoke the methodology of Chow, Jacquier, Kritzman, and Lowry (CJKL) [1999] to partition return samples into ordinary returns and unusual returns. They argue that ordinary returns are more likely to be associated with noise, and thus characteristic of quiet regimes, while unusual returns are more likely to be event-driven and characteristic of turbulent regimes.

EXHIBIT 2
Likelihood of a 10% Loss

Expected Return	10%		
Standard Deviation	20%		

	1 Year	5 Years	10 Years
End of Investment Horizon	15.33%	10.72%	5.81%
During Investment Horizon	41.81%	56.49%	58.85%

CJKL define an unusual return vector as an outlier and measure its unusualness in terms of a multivariate distance from the mean as shown in Equation (2).

$$d_t = (y_t - \mu) \sum^{-1} (y_t - \mu)' \tag{2}$$

where:

d_t = vector distance from multivariate average (dimension: 1 x k);
y_t = return series (1 x k);
μ = mean vector of return series y_t (1 x k); and
$\sum$ = covariance matrix of return series y_t (k x k).

CJKL assume the return series y_t is normally distributed with a mean vector μ and a covariance matrix $\sum$. Thus the distance d_t follows a chi square distribution with k degrees of freedom. The authors choose a confidence level, calculate the corresponding tolerance distance for the chi square distribution, and compare it with each d_t. If the observed distance d_t is greater than the tolerance distance, CJKL define that vector as an outlier, and if the distance is smaller, as an inlier.

For two return series that are uncorrelated and have equal variances, the d_t are the radii of circles that are centered on the average of the return pairs and whose perimeters pass through each return pair. Exhibit 3 illustrates this process for the return pair denoted by the star. The radius of the tolerance circle corresponds to the tolerance distance associated with the chosen confidence level.

CJKL's methodology assumes implicitly that returns are not generated by a single distribution, but rather from separate distributions representing a quiet regime and a turbulent regime. If so, the risk parameters estimated from the full sample of returns may understate a portfolio's exposure to first-passage risk. The risk parameters associated with the quiet and turbulent regimes identified by Equation (2) provide a more accurate description of first-passage risk.

A critical feature of the CJKL methodology is that it takes into account not only unusually volatile returns but also returns that interact in strange ways. If we associate a turbulent regime only with returns that are unusually

EXHIBIT 3
Bivariate Outlier for Uncorrelated Return Pairs with Equal Variance

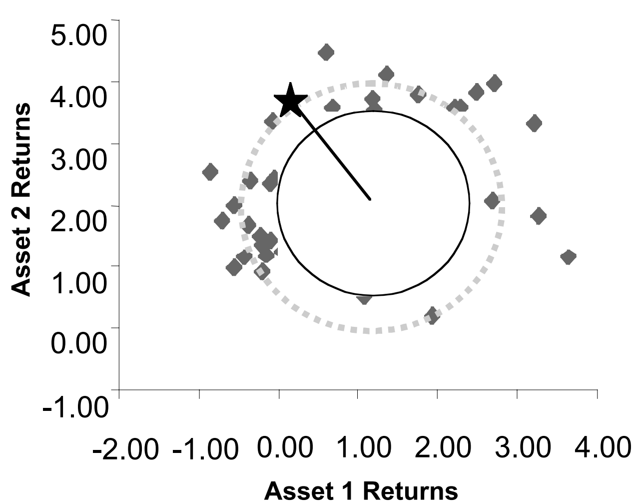

EXHIBIT 4

Cross-Country Correlations Within Asset Classes—Quiet Regime

	U.S. Bonds	U.K. Stocks	U.K. Bonds	German Stocks	German Bonds	Japanese Stocks	Japanese Bonds
U.S. Stocks		51%		42%		22%	
U.S. Bonds			27%		36%		21%
U.K. Stocks				48%		39%	
U.K. Bonds					68%		51%
German Stocks						31%	
German Bonds							63%
Japanese Stocks							

volatile, we might conclude falsely that there are multiple regimes.[3] An experiment demonstrates this point:[4]

- We draw 30,000 return pairs randomly from a correlated joint-normal distribution. Both assets have a 10% mean and 20% standard deviation. The correlation is equal to 60%.
- We rank the return pairs from highest to lowest on the basis of their average returns.
- We select the 10% highest and 10% lowest return pairs, ranked by their average.
- We select outliers using a 20% cutoff.
- We calculate the correlation of the full sample, the sample of high and low returns, and the sample of outliers.

This experiment produces the results as follows:

Full sample correlation:	60.38%
Correlation of high and low returns:	85.84%
Correlation of outliers:	60.48%

The experiment reveals that the correlation estimated from only the volatile returns is significantly biased upward, simply by virtue of the nature of this calculation. The correlation of multivariate outliers, by comparison, is unbiased.

Therefore, if we observe an increase in the correlation of returns from volatile periods, it does not necessarily follow that the returns come from multiple regimes or have an unusual distribution. The rise in the correlation may be simply an artifact of correlation mathematics.[5] A rise in the correlation of an outlier sample, however, is evidence that the returns come either from more than one regime or from an unusual distribution.

III. OVERCONFIDENCE

We apply the CJKL methodology to determine whether global stock and bond returns come from a single regime or from quiet and turbulent regimes. We find that cross-country correlations within asset classes are substantially higher during turbulent regimes than during quiet regimes, and that correlations between asset classes within countries

EXHIBIT 5

Cross-Country Correlations Within Asset Classes—Turbulent Regime (10% Outliers)

	U.S. Bonds	U.K. Stocks	U.K. Bonds	German Stocks	German Bonds	Japanese Stocks	Japanese Bonds
U.S. Stocks		78%		48%		40%	
U.S. Bonds			54%		41%		27%
U.K. Stocks				64%		60%	
U.K. Bonds					56%		55%
German Stocks						37%	
German Bonds							62%
Japanese Stocks							

EXHIBIT 6

Cross-Asset Class Correlations Within Countries—Quiet Regime

	U.S. Bonds	U.K. Stocks	U.K. Bonds	German Stocks	German Bonds	Japanese Stocks	Japanese Bonds
U.S. Stocks	45%						
U.S. Bonds							
U.K. Stocks			68%				
U.K. Bonds							
German Stocks					46%		
German Bonds							
Japanese Stocks							53%

are substantially lower during turbulent regimes than during quiet regimes, which is consistent with the presence of multiple regimes. We also measure the extent to which the assumption of a single regime understates a portfolio's exposure to risk if, in fact, there are multiple regimes.

Our sample includes monthly stock and bond returns in the U.S., U.K., Germany, and Japan from January 1986 through December 1999. This sample is but one pass through history and may not be representative of the future, but it is nonetheless illustrative of the point we wish to make. With this caveat in mind, we choose a cutoff that defines the 10% most unusual periods as outliers.

The cross-country correlations within asset classes are shown in Exhibit 4 for the quiet regime and in Exhibit 5 for the turbulent regime (10% outlier sample). These correlations are from a U.S. dollar perspective. We omit the other correlations in order to focus on the correlations of interest.

Exhibits 6 and 7 show the asset class correlations within countries for the quiet and turbulent regimes. Again, we omit the other correlations in order to focus attention on the correlations of interest.

Exhibit 8 shows the average correlation across countries within asset classes and across asset classes within countries associated with a shift from a quiet regime to a turbulent one. We include results from the perspective of four currencies to show that the results are not caused by exchange rate differences.

Exhibit 8 shows that the average correlation across countries (top right bar) is about 12% higher in turbulent regimes than it is during quiet regimes. The average correlation across asset classes (top left bar) is by contrast nearly 25% lower during turbulent regimes than it is during quiet regimes. There is nearly a 40 percentage point difference between country correlations and asset class correlations, depending on which regime prevails.

Although these results are consistent with the presence of multiple regimes, they are not a sufficient condition. An alternative explanation is that these results are generated by a single regime with a very unusual distribution. In either case, the consequence is the same: We will underestimate risk if we assume a single and lognormal distribution.

EXHIBIT 7

Cross-Asset Class Correlations Within Countries—Turbulent Regime (10% Outliers)

	U.S. Bonds	U.K. Stocks	U.K. Bonds	German Stocks	German Bonds	Japanese Stocks	Japanese Bonds
U.S. Stocks	0%						
U.S. Bonds							
U.K. Stocks			29%				
U.K. Bonds							
German Stocks					19%		
German Bonds							
Japanese Stocks							47%

E X H I B I T 8

Changes in Correlations—Quiet Regime →Turbulent Regime

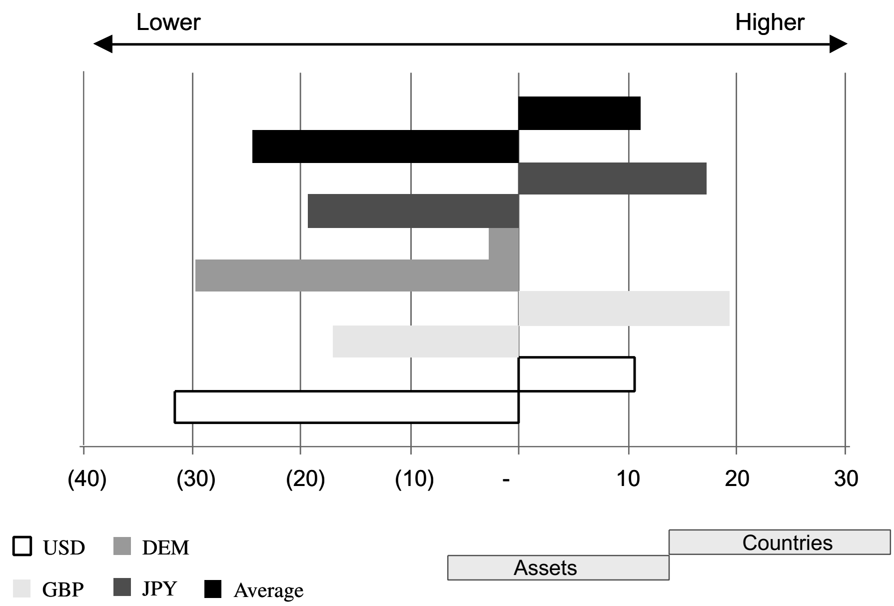

To what extent does the assumption of a single lognormal distribution understate risk when instead returns are generated by a quiet regime or a turbulent regime? To answer this question, we compare the results of a Monte Carlo simulation with the results of a non-parametric bootstrapping simulation.

We begin by estimating the optimal weights of the global stocks and bonds in our sample, based on equilibrium returns, the full-sample covariance matrix, and risk aversion equal to 2.0. We estimate equilibrium returns as those that are proportional to the assets' systematic risk, assuming the full-sample covariance matrix, a 60/40 stock/bond mix equally weighted across countries as the reference portfolio, a riskless return of 4%, and an expected return for the reference portfolio of 10%. The optimal weights are subject to constraints of greater than or equal to 0% and less than or equal to 100%. They are: U.S. stocks 12.5%, U.S. bonds 0.0%, U.K. stocks 20.8%, U.K. bonds 8.9%, German stocks 18.4%, German bonds 9.9%, Japanese stocks 19.3%, and Japanese bonds 10.2%.

Next, we assume investment in this portfolio with monthly rebalancing and calculate its return each month. It has a mean of 1.0626% and a standard deviation of 3.1171%. We then perform the simulations as described below.

Monte Carlo Simulation

- Randomly select monthly returns from a lognormal distribution based on empirical mean and standard deviation. Repeat 30,000 times.
- Calculate monthly cumulative wealth for each return sequence.
- Calculate frequency with which threshold is breached.

Bootstrapping Simulation

- Estimate the density function of the empirical returns.
- Randomly select monthly returns from the empirical distribution, replacing the selected return before each draw. Repeat 30,000 times.
- Calculate monthly cumulative wealth for each return sequence.
- Calculate frequency with which threshold is breached.

The Monte Carlo simulation assumes that returns come from a single regime characterized by a lognormal distribution. The bootstrapping simulation assumes that the returns come from the actual distribution associated with the sample of returns, which could be a blend of multiple regimes.[6]

Exhibit 9

Misestimation of Risk During Investment Horizon

Loss During 1 Year	Theoretical Probability	Empirical Probability
5%	19.9%	23.7%
10%	4.8%	6.9%
15%	0.9%	1.8%
20%	0.1%	0.4%

Loss During 2 Years	Theoretical Probability	Empirical Probability
5%	22.3%	26.7%
10%	6.5%	9.5%
15%	1.7%	3.4%
20%	0.4%	1.0%

Loss During 5 Years	Theoretical Probability	Empirical Probability
5%	23.0%	28.3%
10%	7.2%	11.2%
15%	2.1%	4.6%
20%	0.5%	1.8%

Loss During 10 Years	Theoretical Probability	Empirical Probability
5%	23.1%	28.5%
10%	7.3%	11.5%
15%	2.1%	5.0%
20%	0.5%	2.1%

Exhibit 9 presents the results. It shows how many of the 30,000 simulation paths breach the threshold in at least one of the month-ends during the investment period, depending on whether we assume the theoretical lognormal distribution or the empirical distribution. There are more breaches during the investment horizon assuming the empirical distribution than there are associated with the theoretical assumption of a single lognormal distribution. These differences measure the extent to which we underestimate exposure to first-passage risk by assuming a single lognormal distribution.

IV. DETECTING REGIME SHIFTS

We have developed an econometric model to forecast the relative likelihood that the next period will be quiet or turbulent. We first investigate the serial dependence of multivariate distances, and find that log differences are serially correlated. Thus we transform the multivariate distances D_t as follows:

$$d_t = \log(D_t) - \log(D_{t-1}) \qquad (3)$$

where:

d_t = log difference of multivariate distance, and
D_t = multivariate distance.

We assume that the extent to which d_t is correlated from one month to the next depends upon the current state. We model the random process that determines the current state as a two-state Markov chain. Let S_t be a random variable that assumes only two values, 1 or 2, corresponding to state 1 or 2 at time t. The probability that S_t takes on a particular value depends only on last period's state, S_{t-1}:

$$P[S_t = i \mid S_{t-1} = j, S_{t-2} = k...] = P[S_t = i \mid S_{t-1} = j]$$

We also assume that the probabilities of switching from one state to another are time-invariant and are collected in a transition matrix P. It is written as follows, with p_{ij} equal to the probability that state i will be followed by state j:

$$P = \begin{bmatrix} p_{11} & p_{21} \\ p_{12} & p_{22} \end{bmatrix} = \begin{bmatrix} p_{11} & 1 - p_{22} \\ 1 - p_{11} & p_{22} \end{bmatrix}$$

Every period there is a probability π_t of entering state 1 ($S_t = 1$) and a $1 - \pi_t$ probability of entering state 2 ($S_t = 2$). The elements of P are conditional probabilities, while π_t and $1 - \pi_t$ are unconditional. The relation between the two is expressed as follows:

$$\begin{bmatrix} \pi_t \\ 1 - \pi_t \end{bmatrix} = \begin{bmatrix} p_{11} & p_{21} \\ p_{12} & p_{22} \end{bmatrix} \times \begin{bmatrix} \pi_{t-1} \\ 1 - \pi_{t-1} \end{bmatrix}$$

The regime-switching model is written:

State 1 (with probability π_t): $d_t = c_1 + \phi_1 d_{t-1} + \varepsilon_t$
State 2 (with probability $1 - \pi_t$): $d_t = c_2 + \phi_2 d_{t-1} + \varepsilon_t$
with ε_t follows $N(0, \sigma^2)$

EXHIBIT 10
Hamilton-Markov Model Estimates

R^2 = 0.6747

σ^2 = 0.6527

Log-Likelihood = –189.3627

Variable	Coefficient	t-statistic	t-probability
Results for State 1			
Constant c_1	0.935474	5.961824	0.000000
Coefficient ϕ_1	–0.342242	-1.645128	0.101957
Results for State 2			
Constant c_2	–0.264555	-4.407919	0.000019
Coefficient ϕ_2	–0.512586	-6.113402	0.000000

Estimation of the unknown parameters p_{11}, p_{22}, p_{12}, c_1, ϕ_1, c_2, ϕ_2, and σ^2 is accomplished through maximization of the log-likelihood function, using the EM algorithm.[7] Exhibit 10 reports the results of the two-state model. The dependent variable is the log difference in multivariate distance.

All coefficients are significant at a 10% confidence level. Explanatory power is high ($R^2 = 67.4\%$). A simple autoregressive model that assumes only one state yields by comparison an R^2 equal to 38%.[8] This result suggests that it is more realistic to model returns as generated by two states rather than one state.

The estimated transition matrix equals:

$$P* = \begin{bmatrix} 17.40\% & 23.23\% \\ 82.60\% & 76.77\% \end{bmatrix}$$

This result shows that state 1 tends to be much more short-lived than state 2. The unconditional probability that we enter state 1 at any given time is equal to 21.9% $[(1 - p_{22})/(2 - p_{11} - p_{22})]$ compared to 78.1% for state 2.

The presence of two regimes as defined by CJKL does not require that a forecasting model of these regimes involve two states. Nonetheless, it would be interesting to see if there is a correspondence between the two regimes and the two states of the Markov switching model. In fact, we do find a correspondence between state 1 and the turbulent regime and between state 2 and the quiet regime.

Specifically, when a quiet regime prevails, the probability that the forecast is generated by state 2 exceeds the unconditional probability of 78.1% in 91% of the quiet periods. When a turbulent regime prevails, the probability that the forecast is generated by state 1 exceeds the unconditional probability of 21.9% in 57% of the turbulent periods. These frequencies suggest that state 1 corresponds to the turbulent regime (the outliers), while state 2 corresponds to the quiet regime (the inliers).

We also calculate the correlation between the multivariate distances and the probability that the forecasts are generated by state 1. This correlation equals 58.1%, providing additional evidence of the correspondence between the states of the forecasting model and the regimes.

We use the state-dependent coefficients, the estimated probabilities, and the transition matrix to obtain a forecast of next month's log difference and then add it to the current month's log distance. The exponential of this value provides a forecast of next month's multivariate distance. We assume that multivariate distances follow a chi-square distribution, and we estimate the probability that next month's distance will correspond either to a quiet observation or a turbulent one, as defined by CJKL.[9] For example, given a cutoff of 13.4, a forecasted distance equal to 14.2 based on eight assets implies a 92.5% probability that next month will be an outlier.[10]

Although we lack sufficient data to test our forecasting methodology out-of-sample, we can use a simple test proposed by Merton [1981] to determine whether our methodology adds information relative to an unconditional forecast within-sample. We sum the percentage of periods our methodology forecasts a quiet regime when in fact there was a quiet regime and the percentage of periods in which our methodology forecasts turbulence when there was indeed turbulence. A score of 2.0 would indicate perfect forecasting ability, while a score of 1.0 would indicate a total absence of forecasting ability.

Our forecasting methodology achieves a score of 1.28, which implies with 96.33% confidence that it improves the unconditional forecast.[11]

V. REGIME-SENSITIVE PORTFOLIOS

The results of the Merton test suggest that our regime-forecasting methodology helps to distinguish quiet regimes from turbulent regimes one period ahead. This is but one approach, however, for forecasting the relative likelihood of quiet and turbulent regimes. It might be the case that economic or financial variables can add explanatory power. In any event, to the extent that we can distinguish the relative likelihood of quiet and turbulent regimes, here is how we use such information to structure portfolios that are regime-sensitive. We begin by rewriting expected utility in matrix notation, as shown in Equation (4).

$$EU = w'\mu - \lambda\, (w' \textstyle\sum w) \tag{4}$$

where:

EU = expected utility;
w = vector of asset weights;
μ = mean of expected return vector;
λ = risk aversion; and
$\sum$ = covariance matrix.

In order to differentiate the probabilities we wish to assign to the quiet and turbulent regimes, we replace the full-sample covariance matrix $\sum$ with:

$$p_t \textstyle\sum_Q + (1 - p_t) \textstyle\sum_T \tag{5}$$

where:

p_t = predicted probability of an inlier at time t;
$\sum_Q$ = quiet regime covariance matrix; and
$\sum_T$ = turbulent regime covariance matrix.

Substituting these two covariance matrixes into the standard equation for the expected utility of a portfolio yields:

$$EU = w'\mu - \lambda\, (p_t w' \textstyle\sum_Q w + (1 - p_t)\, w' \textstyle\sum_T w) \tag{6}$$

In order to reduce risk of loss during an investment horizon we should reestimate the parameters of Equation (6) each period and adjust the portfolio weights accordingly.

IV. SUMMARY

We argue that investors with ostensibly long horizons may not have the forbearance to withstand severe interim losses, and we show that the likelihood of an interim loss is much greater than the likelihood of an end-of-period loss.

We use a methodology to partition returns into quiet and turbulent regimes on the basis of a sample's covariances as well as variances. There is evidence that this approach is better than one that relies only on volatility to stratify returns.

We present evidence from global stock and bond returns that is consistent with the notion that returns are generated by multiple regimes, rather than a single lognormal distribution. Moreover, we estimate the extent to which the assumption of a single lognormal distribution understates a portfolio's first-passage risk if instead our return sample is generated by multiple regimes.

Finally, we propose a methodology to assess the relative likelihood of a quiet or turbulent regime one month ahead using a switching regression model. This information can be used to structure regime-sensitive portfolios.

APPENDIX

Covariance of Truncated Returns

We consider n random variables X distributed following an N(O, V) multinormal distribution. The CJKL methodology separates the sample into two parts based on both the distance from the multivariate mean and the interaction between returns. Inliers are defined as all observations for which $X'V^{-1}X \leq b$ and outliers as $X'V^{-1}X > b$.[12] The methodology, known as elliptic truncation, differs from other forms of truncation in that it partitions all returns on the basis of their degree of unusualness, not on a given threshold level.[13]

We demonstrate that the covariance matrix of the truncated distribution takes a simple form. It equals the covariance matrix of the full distribution multiplied by a scale factor. This signifies that: 1) the CJKL methodology produces conditional moments, including correlations, that are unaffected by the conditioning bias reported in the literature; and 2) unless extreme and usual events are generated by two different distributions, the truncated correlation should not spuriously indicate different regimes.

We focus on the moment-generating function (MGF) of the distribution of truncated series. The MGF of a random variable with density f corresponds to:

$$\begin{aligned} M_X(t) &= E(e^{tx}) = \int e^{tx} f(x)dx \\ &= \int (1 + tx + \frac{1}{2!}t^2x \;+ \ldots +) f(x)dx \\ &= 1 + tm_1 + \frac{1}{2!}t^2 m_2 + \ldots + \end{aligned}$$

where m_r is the r-th moment of X around 0. Thus the n-th moment around 0 is given by the value of the n-th derivative of M_X with respect to t, taken at 0:

$$M_x'(t) = E(xe^{tx}) \text{ and } M_X''(t) = E(x^2e^{tx})$$

Therefore:

$$M_x'(0) = E(x),\ M_x''(0) = E(x^2)$$
and $$\mathrm{Var}(X) = M_x''(0) - M_x'(0)^2$$

Because $X' V^{-1} X$ is distributed following a chi-square distribution with n degrees of freedom, we obtain the expression for the MGF of the truncated outliers:[14]

$$M_{\mathrm{Trunc}}(t) = \frac{1}{\Pr(\chi_n^2 > b)} (2\pi)^{\frac{-1}{2n}} \mid V \mid^{-1}$$

$$\int_{X'V^{-1}X > b} \exp\left(-\frac{1}{2} X' V^{-1} X + t' X\right) dX$$

We make the transformation:

$$Z' = x'H' \text{ with } H'H = V^{-1}$$

Therefore:

$$\begin{aligned} M_{\mathrm{Trunc}}(t) &= \frac{1}{\Pr(\chi_n^2 > b)} (2\pi)^{\frac{-1}{2n}} \exp\left(\frac{t' Vt}{2}\right) \\ &\quad \int_{z'z > b} \exp\left[-\frac{1}{2}(Z - Ht)'(Z - Ht)\right] dZ \\ &= \frac{1}{\Pr(\chi_n^2 > b)} (2\pi)^{\frac{-1}{2n}} \exp\left(\frac{t' Vt}{2}\right) \Pr(\chi_n^2 (t' Vt) > b) \\ &= \frac{1}{\Pr(\chi_n^2 > b)} \sum_{j=0}^{\infty} \frac{\left(\frac{1}{2} t' Vt\right)^j}{j!} \Pr[(\chi_{n+2}^2 > b] \end{aligned}$$

It follows that the covariance matrix of the outliers is simply equal to cV in which

$$c = \frac{\Pr(\chi_{n+2}^2 > b)}{\Pr(\chi_n^2 > b)}$$

All elements of the covariance matrix of the outliers are multiplied by a scalar. This implies that, if returns come from a single distribution, truncated correlations are left unchanged.

ENDNOTES

The authors thank George Chow, Alan Marcus, Paul O'Connell, André Perold, Krishna Ramaswamy, Bruno Solnik, and Edouard Stirling; and participants at New York University's Risk Management: State of the Art Conference; the Fall Seminar of the Institute for Quantitative Research in Finance; and the Spring Seminar of Inquire Europe for their helpful comments.

[1]Lognormality characterizes the distribution of discrete returns and arises from the process of compounding. The continuous counterparts of discrete returns are normally distributed.

[2]The first passage probability is described in Karlin and Taylor [1975].

[3]For example, Odier and Solnik [1993] find that correlations rise during volatile markets, but they fail to address the issue of whether this rise is due to multiple regimes or to correlation mathematics. The pitfalls of using conditional correlations to detect regimes are now widely documented. Boyer, Gibson, and Loretan [1997], Forbes and Rigobon [1999], Ang and Chen [2000], and Longin and Solnik [2001] note that deviations from the unconditional correlation do not necessarily indicate different regimes.

[4]See the appendix for a formal demonstration.

[5]See Ang and Chen [2000, p. 3.]

[6]We obtain a smooth approximation of the density function of the empirical returns by applying non-parametric kernel smoothing, which involves weighted local averaging of the density function. See Härdle [1990].

[7]See Hamilton [1994], Chapter 22.

[8]Obtained with a two-lag autoregressive model.

[9]Chi-square is a simplifying assumption. We would prefer to consider the empirical distribution of distances but, given the small number of observations in the sample, this would lead to unreliable results.

[10]The cutoff corresponds to the 10% most unusual observations.

[11]This confidence value is estimated using a normal approximation based upon the mean and variance of a hypergeometric distribution. For a more detailed discussion of this evaluation procedure, see Henriksson and Merton [1981].

[12]b corresponds to the inverse of the chi-square distribution for the chosen degree of tolerance and the number of assets.

[13]Most studies focus on single partitioning ($X > c$) or partitioning based on volatility ($|X| > c$).

[14]See Johnson and Kotz [1972], Chapters 28 and 35.

REFERENCES

Ang, A., and J. Chen. "Asymmetric Correlations of Equity Portfolios." Working paper, Columbia University, 2000.

Boyer, M., B. Gibson, and M. Loretan. "Pitfalls in Test for Changes in Correlations." International Finance Discussion Paper 597, Board of Governors of the Federal Reserve System, 1997.

Chow, G., E. Jacquier, M. Kritzman, and K. Lowry. "Optimal Portfolios in Good Times and Bad." *Financial Analysts Journal,* May/June 1999.

Forbes, K., and R. Rigobon. "No Contagion, Only Interdependence: Measuring Stock Market Co-Movements." Working paper 7267, NBER, 1999.

Hamilton, J. *Time Series Analysis*. Princeton: Princeton University Press, 1994.

Härdle, W. *Applied Nonparametric Regression*. Econometric Society Monographs No. 19. Cambridge: Cambridge University Press, 1990.

Henriksson, R., and R. C. Merton. "On Market Timing and Investment Performance: II. Statistical Procedures for Evaluating Forecasting Skills." *Journal of Business,* Vol. 54, No. 4 (1981).

Johnson, N.L., and S. Kotz. *Continuous Multivariate Distributions*. New York: John Wiley, 1972.

Karlin, S., and H. Taylor. *A First Course in Stochastic Processes*, 2nd ed. New York: Academic Press, 1975.

Longin, F., and B. Solnik. "Extreme Correlations of International Equity Markets." Forthcoming, *Journal of Finance,* 2001.

Merton, R. C. "On Market Timing and Investment Performance: I. An Equilibrium Theory of Value for Market Forecasts." *Journal of Business,* Vol. 54, No. 3 (1981).

Odier, P., and B. Solnik. "Lessons for International Asset Allocation." *Financial Analysts Journal*, March/April 1993.