ket options and options on the max or on the min, there are no dependable sources of implied correlations to use in pricing them. Correlations must be estimated from historical data, which leads to substantial estimation risk. In this article, Fengler and Schwendner describe a block bootstrap procedure that allows an investor to evaluate a multiasset option's exposure to parameter risk from imperfectly estimated correlations. The results are translated into minimum bid-ask spreads that are required to account for this additional source of risk.

Numerical option pricing methods discretize a continuous-time continuous-state state space and approximate the solution to the continuous problem with the results of local computations on all of the nodes of the discrete lattice. This procedure converges to the correct solution as the time and price intervals go to zero, but at the cost of larger and larger numbers of calculations. Many of the calculations on a fine lattice are irrelevant to valuing the derivative. They may occur at stock prices that have virtually no probability of being reached. Or the stock price may be reachable from the initial price, but be so far away from the strike price that the option is either (almost) certain to be out of the money at maturity, meaning it is worthless now, or (almost) certain to be exercised at maturity, meaning it can be priced as a forward now, with no need for any further calculations. Eliminating the superfluous calculations at these nodes can greatly improve the performance of a lattice valuation technique, and the possible increase in computation speed becomes much greater as the number of stochastic variables in the lattice increases. Here, Andricopoulos et al., show how to curtail the price range for a numerical valuation technique and demonstrate the substantial improvement in speed that it can produce.

INNOVATIONS

Securitization has been one of the most important capital market innovations in the last 30 years. Beginning with mortgage pass-throughs and CMOs, the creative use of collateralization and tranching to redistribute risk exposure has recently expanded significantly, as bonds, loans, and other securities exposed to credit risk have been securitized into CDOs, CLOs, and the like. In this article, Stone and Zissu describe the latest new product in this area, which combines both securitization and active asset management. The "Collateralized Fund Obligation" (CFO) is a securitization of a fund of hedge funds. The investor acquires the benefit of hedge fund returns that are (hopefully!) enhanced by active management and relatively uncorrelated with other asset classes, along with a much greater ability to manage the risk characteristics of the investment.

models. See, for instance, Andersen and Bollerslev [1998].

An alternative use of this regression is to provide a bias correction for volatility forecasts using the estimated regression parameters. In-sample, such a bias-corrected forecast will by construction be both unconditionally and conditionally unbiased, and can thus be expected to provide a better forecast of volatility than the original uncorrected volatility forecast (see Figlewski [1997]).

Day and Lewis [1993] apply this approach to the pricing of crude oil futures options. They estimate the regression parameters using the first half of a sample, and then use these to correct the volatility forecasts for the second half of the sample. Overall, they find only limited improvements in the corrected volatility forecasts for some of the models.

The bias of these models, however, and hence also the regression parameters used for bias correction are likely to change over time, a feature that will not be captured by a static regression approach. Moreover, the degree of potential improvement from such a bias correction will depend on how badly specified the conditional volatility model is in the first place.

We consider the bias in the estimation of value at risk (VaR) by the RiskMetrics approach of J.P. Morgan, which assumes daily returns are conditionally normal and uses an EWMA estimator to generate forecasts of daily conditional volatility (see "RiskMetrics" [1996]). The EWMA estimator has proven very effective in forecasting the volatility of financial asset returns over short horizons, often outperforming the forecasts of more sophisticated GARCH models. Several authors find VaR forecasts based on the standard EWMA estimator superior to those based on the GARCH model. See, for instance, Alexander and Leigh [1997] and Boudoukh, Richardson, and Whitelaw [1997].

Whatever its attractions, the RiskMetrics EWMA model is likely to be misspecified in three ways. First, the EWMA model assumes the true conditional variance is integrated, while it is likely that for most financial asset series the true conditional variance is highly persistent but stationary (a non-stationary conditional variance would imply that the unconditional variance of returns is infinite).

Second, as it is based on squared returns, the EWMA model implicitly assumes the conditional distribution of returns is normal. If, as is suggested by empirical evidence, the conditional distribution of returns is leptokurtic, the EWMA estimator will be inefficient, and will attach too much weight to extreme returns (see Baillie and DeGennaro [1990]; Guermat and Harris [2002]; and Harris and Shen [2003]).

Finally, it is unlikely that, within the class of EWMA estimators, the decay factor RiskMetrics uses is optimal for a particular series over a particular time period.[2]

If the RiskMetrics EWMA model is misspecified, it will generate conditional volatility forecasts that are conditionally biased. We estimate the conditional bias in the RiskMetrics approach using a realized volatility regression with daily returns, and use the estimated parameters to generate bias-corrected conditional volatility forecasts. The error term in this regression is highly non-normal, so the usual ordinary least squares (OLS) estimator will be inefficient. We therefore also use the robust least absolute deviations (LAD) estimator in order to estimate the regression parameters. To allow for parameter inconstancy (i.e., time variation in the degree of conditional bias), we adopt a rolling regression approach, using the most recently available data to estimate the regression parameters.

We apply this approach to the estimation of daily VaR for three equity portfolios (the S&P 500, FTSE 100, and Nikkei 225), three long bond portfolios (U.S., U.K., and Japan), and three currency portfolios (USD/GBP, USD/JPY, and GBP/JPY). For each portfolio, we generate out-of-sample estimates of one-day portfolio VaR at the 95.0%, 97.5%, and 99.0% confidence levels. We evaluate the VaR forecasts using tests of both unconditional and conditional coverage.

In most cases, VaR estimates based on the bias-corrected conditional volatility perform substantially better than those based on the unadjusted estimates of conditional volatility, particularly in terms of unconditional coverage, and particularly at high VaR confidence levels. Moreover, the LAD estimator provides significantly better VaR forecasts than the OLS estimator.

I. THEORETICAL FRAMEWORK

Suppose the return on an asset between time $t-1$ and time t, denoted r_t, is drawn conditionally from a distribution with mean zero and conditional variance, $\sigma_t^2 = E[r_t^2 \mid \Omega_{t-1}]$, where Ω_{t-1} is the information set at time $t-1$.[3] We assume the data are generated by an arbitrary stochastic volatility model, and that a conditional variance model is used to generate a forecast of the variance at time t, $\hat{\sigma}_t^2$, employing data available up to and including time $t-1$.

From the definition of the conditional variance, we have:

$$r_t^2 = \sigma_t^2 + \nu_t \tag{1}$$

where ν_t is a random forecast error with $E(\nu_t \mid \Omega_{t-1}) = 0$.

The regression between realized volatility, measured by the squared return, and the forecast conditional variance, is:

$$r_t^2 = a + b\hat{\sigma}_t^2 + \nu_t \tag{2}$$

A necessary and sufficient condition for $\hat{\sigma}_t^2$ to be unconditionally unbiased is that $a = (1-b)E(\sigma_t^2) = (1-b)\sigma^2$, where σ^2 is the unconditional variance of r_t (see, for instance, Taylor [1999]). A necessary condition for $\hat{\sigma}_t^2$ be *conditionally* unbiased is that, in addition, $a = 0$ and $b = 1$. This is not a sufficient condition; for $\hat{\sigma}_t^2$ to be conditionally unbiased, we must have $E(r_t^2 - \hat{\sigma}_t^2 \mid \Omega_{t-1}) = 0$. The regression given by Equation (2) tests whether $E(r_t^2 - \hat{\sigma}_t^2 \mid \hat{\sigma}_t^2) = 0$ and $\hat{\sigma}_t^2$ is only a subset of Ω_{t-1}.

The property of conditional unbiasedness is known as *efficiency* in the forecasting literature. When the only conditioning variable is the forecast itself, we call this *weak efficiency* (see Nordhaus [1987]). Application of Equation (2) to evaluation of economic forecasts in general is due to Theil [1966] and Mincer and Zarnovitz [1969].

Equation (2) is widely used in the volatility modeling literature to evaluate the explanatory power of a particular conditional variance model, and to compare competing models, using the R-squared statistic. To evaluate daily volatility forecasts, realized volatility may be measured by squared daily returns or by the sum of squared intraday returns.

Equation (2) can also be used to correct conditional variance forecasts. If for a particular model, $b < 1$, we know, a priori, that the conditional variance forecasts will be too dispersed in the sense that high forecasts are on average too high and low forecasts are too low. This can be seen by subtracting the realized conditional variance from both sides of (2), which shows that when $b < 1$, the forecast error is negatively related to the forecast. It should therefore be possible to refine the forecast by adjusting for this systematic bias.

If we have parameter estimates, $\hat{a}$ and $\hat{b}$, we can define a new bias-corrected forecast given by:

$$\hat{\hat{\sigma}}_t^2 = \hat{a} + \hat{b}\hat{\sigma}_t^2 \tag{3}$$

We could then expect the bias-corrected forecast, $\hat{\hat{\sigma}}_t^2$, to perform better empirically than $\hat{\sigma}_t^2$, as it will incorporate the same information as $\hat{\sigma}_t^2$, but will by construction be both conditionally and unconditionally unbiased. When the bias correction is applied to out-of-sample forecasts, the extent to which the bias is reduced will depend on how constant the parameters a and b are over time and how precisely they are estimated.

Estimation of the parameters of Equation (2) may not be straightforward. When the conditional volatility forecast $\hat{\sigma}_t^2$ is derived from an estimated model (such as implied volatility or GARCH), it will be subject to measurement error, so the OLS estimators of the parameters in Equation (3) will be biased. This problem can be circumvented using instrumental variables (see Andersen and Bollerslev [1998]; Christensen and Prabhala [1998]; and Barndorff-Nielsen and Shephard [2001]). The error term in Equation (2) is also likely to be highly non-normal, so OLS may not be an efficient estimator of the regression parameters.

II. EMPIRICAL ILLUSTRATION

The RiskMetrics approach assumes returns are conditionally normal, and uses an EWMA model to forecast the conditional variance. The EWMA estimator is given by:

$$\hat{\sigma}_t^2 = \lambda\hat{\sigma}_{t-1}^2 + (1-\lambda)r_{t-1}^2 \tag{4}$$

which by recursive substitution can be written as:

$$\hat{\sigma}_t^2 = (1-\lambda)\sum_{i=1}^{\infty}\lambda^{i-1}r_{t-i}^2 \tag{5}$$

As in practice only a finite number of observations are used, a correction must be made to ensure that the weights of the EWMA estimator sum to unity for a small sample size. This correction involves setting the variance forecast at date 1 to zero and then multiplying the variance forecast at date t by $1/(1-\lambda^{t-1})$. The EWMA variance forecast is therefore given by:

$$\hat{\sigma}_t^2 = 0 \qquad \text{for } t = 1$$

$$\hat{\sigma}_t^2 = \frac{(1-\lambda)}{(1-\lambda^{t-1})}\sum_{i=1}^{t}\lambda^{i-1}r_{t-i}^2 \qquad \text{for } t = 2, \ldots, T \tag{6}$$

For daily data, RiskMetrics uses a value of λ equal to 0.94 and discards the first 75 EWMA forecasts. Once the conditional volatility of portfolio returns is estimated, the VaR of a portfolio is then calculated by multiplying the negative of the appropriate quantile of the standard normal distribution by the square root of the estimated conditional variance from the EWMA model:

EXHIBIT 1
Summary Statistics

	Mean	Standard Deviation	Skewness	Excess Kurtosis
US Bonds	0.036%	0.547%	-0.281	1.823
UK Bonds	0.037%	0.483%	-0.026	3.361
Japanese Bonds	0.026%	0.395%	-0.923	9.476
S&P 500	0.043%	1.026%	-0.238	4.427
FTSE 100	0.037%	1.031%	-0.114	2.974
Nikkei 225	-0.019%	1.421%	0.205	3.945
USD/GBP	-0.002%	0.572%	-0.201	2.782
USD/JPY	0.003%	0.691%	0.504	4.845
GBP/JPY	0.005%	0.710%	0.309	3.440

$$V\hat{a}R_t = -\hat{\sigma}_t\, c(\alpha) \tag{7}$$

where $c(\alpha)$ is the α% quantile of the standard normal distribution, and $1 - \alpha$ is the VaR confidence level.

Because the EWMA model is misspecified, its forecasts will be conditionally biased. While correcting for the conditional bias in the EWMA variance forecasts could be expected to improve the RiskMetrics VaR forecasts, it will not yield optimal forecasts (see, for instance, Baillie and DeGennaro [1990]).[4]

Data

We use daily returns for several equity, bond, and exchange rate portfolios. The equity indexes are the S&P 500, the FTSE 100, and the Nikkei 225. The bond indexes are the long bond return indexes for the U.S., U.K., and Japan, and the exchange rates are the USD/GBP, USD/JPY, and GBP/JPY. Data are obtained from Datastream for December 30, 1987-November 18, 2003, which is the longest common sample available (a total of 4,145 observations), and used to calculate continuously compounded daily returns.

Exhibit 1 reports summary statistics for the nine series. The mean return in each case is very close to zero. All nine series are highly non-normal, with very significant leptokurtosis.

Estimation Method

The EWMA model in Equation (6) is used to estimate the conditional variance of each return series with $\lambda = 0.94$. The variance forecast for $t = 1$ is set to zero, and the first 75 observations are discarded. This yields a sample of 4,069 conditional variance forecasts. For each series, on each day, t, we estimate Equation (2) using daily returns and volatility forecasts for $t - 500$ to $t - 1$, a regression sample of 500 observations.[5]

Note that because the EWMA model is not estimated, the measurement error problem does not arise. Also, as noted above, the error term in the regression given by (2) is likely to be highly non-normal, in which case OLS estimates of the regression parameters will be inefficient. Therefore, we also estimate the regression parameters using the least absolute deviation (LAD) estimator, which chooses estimates of a and b that minimize the sum of the absolute value of the errors, rather than the sum of squared errors (see Huber [1981]).

Although OLS is inefficient, it will yield corrected volatility forecasts that are by construction unconditionally unbiased. This is because OLS imposes the restriction that $mean(r^2) = \hat{a}_{OLS} + \hat{b}_{OLS} mean(\hat{\sigma}_t^2)$, where $mean(r^2)$ and $mean(\hat{\sigma}_t^2)$ are the sample mean of r^2 and $\hat{\sigma}_t^2$ estimated over the regression sample. Taking expectations conditional on $\hat{a}_{OLS}$ and $\hat{b}_{OLS}$ yields $\hat{a}_{OLS} = (1 - \hat{b}_{OLS})\sigma^2$, which is analogous to the condition required for unconditional unbiasedness.

The LAD estimator does not impose this restriction, so volatility forecasts based on the LAD estimators of a and b will in general be unconditionally biased. To ensure that the LAD estimator generates estimates of a and b that lead to volatility forecasts that are unconditionally unbiased, we therefore impose the restriction that $\hat{a}_{LAD} = (1 - \hat{b}_{LAD})\hat{\sigma}^2$, where $\hat{\sigma}^2$ is the sample variance estimated over the regression sample.

For both the OLS and LAD estimators, we use the estimated regression parameters to compute the bias-corrected EWMA forecast for day t, given by Equation (3). We then compute three VaR forecasts for day t, an original VaR forecast based on the uncorrected EWMA forecast of conditional volatility, and VaR forecasts based on the OLS- and LAD-corrected forecasts of conditional volatility. This is repeated for every day in the sample, yielding a total of 3,568 VaR forecasts for each model for each of the nine portfolios.

VaR Evaluation Procedure

The evaluation of VaR forecasts is not straightforward. As with the evaluation of volatility forecasting models, a direct comparison between the forecast VaR and the actual VaR cannot be made, because the latter is unobservable. A variety of evaluation methods have been proposed. Examples are Kupiec [1995]; Christofferson [1998]; and Lopez [1999].

To evaluate the VaR forecasts of the different models we estimate, we use the unconditional coverage statistic of Kupiec [1995] and the independence statistic of Christofferson [1998].

Unconditional Coverage

The most basic requirement of a VaR model is that the proportion of times that the VaR forecast it generates is exceeded (the number of exceptions) should on average be equal to one minus the nominal VaR confidence level; in other words, the model should provide the correct unconditional coverage. To test the null hypothesis that the unconditional coverage is equal to the nominal significance level, Kupiec [1995] derives a likelihood ratio (LR) statistic based on the observation that the probability of observing N exceptions in a sample of size T is governed by a binomial process and is given by $(1-p)^{T-N}p^N$.

The LR statistic is computed as:

$$LR_U = -2\ln[(1-p)^{T-N}p^N] + 2\ln[(1-N/T)^{T-N}(N/T)^N] \tag{8}$$

where p is the desired significance level, which is equal to one minus the VaR confidence level. Under the null hypothesis of correct unconditional coverage, the LR_U statistic has a chi-squared distribution with one degree of freedom.

Independence

If a VaR model accurately captures the conditional distribution of returns and its dynamic properties such as time-varying volatility, exceptions should be unpredictable, and hence independently distributed over time. To test the independence of the exceptions of a VaR model, Christofferson [1998] derives an LR statistic, computed as:

$$LR_I = 2(\ln L_A - \ln L_0) \tag{9}$$

where

$$L_A = (1-\pi_{01})^{T_{00}}\pi_{01}^{T_{01}}(1-\pi_{11})^{T_{10}}\pi_{11}^{T_{11}}$$
$$L_0 = (1-\pi)^{T_{00}+T_{10}}\pi^{T_{01}+T_{11}}$$
$$\pi_{ij} = T_{ij} / (T_{i0} + T_{i1})$$
$$\pi = (T_{01} + T_{11}) / (T_{00} + T_{01} + T_{10} + T_{11})$$

and T_{ij} is the number of times state i is followed by state j, where state 0 describes an actual portfolio loss of less than the estimated VaR and state 1 describes an actual return greater than the estimated VaR. Under the null hypothesis that exceptions are independently distributed, the LR_I statistic has a chi-squared distribution with one degree of freedom.

III. RESULTS

For each of the nine series, Exhibit 2 reports summary statistics for the uncorrected EWMA forecast of volatility and the estimated regression parameters a and b and the corrected forecast of volatility using both the OLS and LAD estimators. For all nine series, the average estimated value of b, using both OLS and LAD, is considerably lower than unity, and the average estimated value of a is greater than zero, implying that the RiskMetrics EWMA model systematically generates forecasts of conditional volatility that are conditionally biased.

For all nine series, $\hat{b}_{LAD}$ is higher on average than $\hat{b}_{OLS}$, and $\hat{a}_{LAD}$ is lower than $\hat{a}_{OLS}$. The average value of $\hat{b}_{OLS}$ ranges from 0.32 for the U.S. bond portfolio to 0.66 for the Nikkei 225 portfolio. The average value of $\hat{b}_{LAD}$ ranges from 0.44 to 0.71 for the same two portfolios. Across all nine portfolios, the average estimated value of b is 0.53 using OLS and 0.60 using LAD.

Exhibit 2 also shows considerable time variation in the estimated regression parameters. Somewhat counter-intuitively, LAD leads to more volatile estimates of a and b than OLS (a result of imposing the restriction that ensures

EXHIBIT 2

Summary Statistics of Estimated Regression Parameters and Uncorrected and Corrected Volatility

Panel A: US Bonds

	σ_t	$a_{t,OLS}$	$b_{t,OLS}$	$\sigma_{t,OLS}$	$a_{t,LAD}$	$b_{t,LAD}$	$\sigma_{t,LAD}$
Mean	0.533%	0.019	0.315	0.544%	0.015	0.444	0.556%
Standard deviation	0.130%	0.006	0.287	0.095%	0.010	0.405	0.106%
Minimum	0.279%	-0.007	-0.561	0.339%	-0.013	-0.991	0.378%
Maximum	1.115%	0.033	1.418	1.179%	0.051	1.572	1.229%
5th percentile	0.362%	0.000	-0.194	0.419%	0.000	-0.213	0.436%
95th percentile	0.791%	0.000	0.721	0.734%	0.000	0.987	0.436%
Average abs correction	-	-	-	9.582%	-	-	8.025%

Panel B: UK Bonds

	σ_t	$a_{t,OLS}$	$b_{t,OLS}$	$\sigma_{t,OLS}$	$a_{t,LAD}$	$b_{t,LAD}$	$\sigma_{t,LAD}$
Mean	0.477%	0.012	0.473	0.491%	0.009	0.582	0.494%
Standard deviation	0.138%	0.005	0.250	0.100%	0.007	0.306	0.116%
Minimum	0.249%	-0.006	-0.073	0.345%	-0.011	-0.214	0.313%
Maximum	1.174%	0.022	1.470	1.402%	0.030	1.428	1.364%
5th percentile	0.303%	0.000	0.021	0.393%	0.000	0.121	0.381%
95th percentile	0.740%	0.000	0.824	0.677%	0.000	1.042	0.381%
Average abs correction	-	-	-	10.229%	-	-	8.210%

Panel C: Japanese Bonds

	σ_t	$a_{t,OLS}$	$b_{t,OLS}$	$\sigma_{t,OLS}$	$a_{t,LAD}$	$b_{t,LAD}$	$\sigma_{t,LAD}$
Mean	0.361%	0.005	0.635	0.377%	0.003	0.697	0.379%
Standard deviation	0.182%	0.002	0.198	0.169%	0.003	0.265	0.172%
Minimum	0.106%	-0.008	0.180	0.171%	-0.005	-0.232	0.148%
Maximum	1.294%	0.010	2.065	1.458%	0.012	1.285	1.379%
5th percentile	0.173%	0.000	0.262	0.225%	0.000	0.225	0.225%
95th percentile	0.730%	0.000	0.854	0.707%	0.000	1.034	0.225%
Average abs correction	-	-	-	10.333%	-	-	8.263%

the unconditional unbiasedness of the LAD-corrected volatility estimator).[6] For both OLS and LAD, the average corrected volatility estimate is on average similar to the average uncorrected volatility forecast, reflecting the fact that the uncorrected, OLS-corrected, and LAD-corrected volatility forecasts are all unconditionally unbiased.

The last row in each panel of Exhibit 2 shows the average absolute difference between the corrected and uncorrected volatility forecasts, as a percentage of the average uncorrected volatility forecast. For both OLS and LAD, the average absolute correction is of the order of 10%; LAD results in a marginally lower correction than OLS. For both OLS and LAD, the corrected volatility forecast has significantly lower variance than the uncorrected volatility forecast; the LAD-corrected forecast is slightly more volatile than the OLS-corrected forecast.[7]

Exhibit 3 reports the unconditional coverage for the RiskMetrics EWMA model for the nine portfolios at the 99% VaR confidence level for VaR estimates based on the uncorrected EWMA forecasts, and the bias-corrected EWMA forecasts using both OLS and LAD. It also reports the LR_U and LR_I statistics to test the null hypotheses of correct unconditional coverage and independence, respectively.

For all nine portfolios, the unconditional coverage of the uncorrected VaR forecasts is higher than the nominal significance level of 1%. In all cases, the null hypothesis of correct unconditional coverage can be rejected at the 5% level using the LR_U statistic, and in all but one case at the

EXHIBIT 2 (continued)

Summary Statistics of Estimated Regression Parameters and Uncorrected and Corrected Volatility

Panel D: S&P 500

	$\hat{\sigma}_t$	$\hat{a}_{t,OLS}$	$\hat{b}_{t,OLS}$	$\hat{\hat{\sigma}}_{t,OLS}$	$\hat{a}_{t,LAD}$	$\hat{b}_{t,LAD}$	$\hat{\hat{\sigma}}_{t,LAD}$
Mean	0.955%	0.043	0.508	0.972%	0.036	0.565	0.989%
Standard deviation	0.415%	0.021	0.189	0.374%	0.025	0.283	0.386%
Minimum	0.352%	-0.056	0.020	0.464%	-0.040	-0.292	0.473%
Maximum	2.597%	0.101	2.048	3.194%	0.182	1.312	2.774%
5th percentile	0.436%	0.000	0.151	0.526%	0.000	0.022	0.536%
95th percentile	1.782%	0.000	0.806	1.644%	0.000	0.907	0.536%
Average abs correction	-	-	-	8.713%	-	--	7.491%

Panel E: FTSE 100

	$\hat{\sigma}_t$	$\hat{a}_{t,OLS}$	$\hat{b}_{t,OLS}$	$\hat{\hat{\sigma}}_{t,OLS}$	$\hat{a}_{t,LAD}$	$\hat{b}_{t,LAD}$	$\hat{\hat{\sigma}}_{t,LAD}$
Mean	0.979%	0.032	0.625	1.000%	0.025	0.644	1.016%
Standard deviation	0.417%	0.014	0.210	0.401%	0.026	0.374	0.412%
Minimum	0.425%	-0.022	-0.086	0.529%	-0.056	-0.603	0.505%
Maximum	2.991%	0.092	1.373	3.374%	0.146	1.373	3.225%
5th percentile	0.530%	0.000	0.177	0.597%	0.000	-0.014	0.600%
95th percentile	1.843%	0.000	0.902	1.820%	0.000	1.102	0.600%
Average abs correction	-	-	-	6.367%	-	-	6.787%

Panel F: Nikkei 225

	$\hat{\sigma}_t$	$\hat{a}_{t,OLS}$	$\hat{b}_{t,OLS}$	$\hat{\hat{\sigma}}_{t,OLS}$	$\hat{a}_{t,LAD}$	$\hat{b}_{t,LAD}$	$\hat{\hat{\sigma}}_{t,LAD}$
Mean	1.423%	0.075	0.663	1.469%	0.065	0.710	1.472%
Standard deviation	0.480%	0.033	0.154	0.417%	0.061	0.253	0.433%
Minimum	0.500%	-0.046	0.119	0.734%	-0.044	-0.217	0.606%
Maximum	3.954%	0.157	2.368	4.732%	0.270	1.291	4.058%
5th percentile	0.762%	0.000	0.397	0.960%	0.000	0.210	0.888%
95th percentile	2.388%	0.000	0.857	2.230%	0.000	1.024	0.888%
Average abs correction	-	-	-	7.576%	-	-	6.305%

1% level. The null hypothesis of independence can be rejected LR_I at the 1% level for the U.S. equity portfolio and at the 5% level for the Japanese bond portfolio.

Using the OLS-corrected VaR forecasts improves the unconditional coverage in all cases. The null hypothesis is now not rejected for two of the portfolios, and is rejected at only the 10% level for one portfolio. The null hypothesis of independence is now not rejected for the U.S. equity portfolio, but is rejected for the USD/GBP portfolio.

Using the LAD-corrected VaR forecasts offers a considerable improvement over OLS. There are large reductions in both the LR_U and LR_I statistics for most series. Only for the S&P 500 series is the null hypothesis of independence rejected, and then only at the 10% level. At the 99% VaR significance level, therefore, use of bias correction substantially improves both the unconditional and conditional coverage of the RiskMetrics VaR approach. Moreover, it is clear that the LAD estimator performs significantly better than the OLS estimator.

Exhibit 4 reports the results for the 97.5% VaR confidence level. Again, for all nine portfolios, the unconditional coverage of the uncorrected VaR forecasts is higher than the nominal significance level of 2.5%, significantly so in seven of nine cases. Using the OLS-corrected VaR forecasts improves the unconditional coverage in all cases. The null hypothesis of correct conditional coverage is rejected for only two portfolios, and only at the 10% level. The null hypothesis of independence using the OLS bias-

EXHIBIT 2 (continued)

Summary Statistics of Estimated Regression Parameters and Uncorrected and Corrected Volatility

Panel G: USD/GBP

	$\hat{\sigma}_t$	$\hat{a}_{t,OLS}$	$\hat{b}_{t,OLS}$	$\hat{\hat{\sigma}}_{t,OLS}$	$\hat{a}_{t,LAD}$	$\hat{b}_{t,LAD}$	$\hat{\hat{\sigma}}_{t,LAD}$
Mean	0.522%	0.014	0.520	0.537%	0.011	0.618	0.548%
Standard deviation	0.185%	0.005	0.171	0.155%	0.008	0.274	0.168%
Minimum	0.210%	0.005	-0.054	0.286%	-0.017	-0.193	0.285%
Maximum	1.401%	0.029	0.961	1.397%	0.042	1.302	1.545%
5th percentile	0.295%	0.000	0.179	0.376%	0.000	0.076	0.360%
95th percentile	0.904%	0.000	0.749	0.858%	0.000	1.048	0.360%
Average abs correction	-	-	-	9.358%	-	-	8.314%

Panel H: USD/JPY

	$\hat{\sigma}_t$	$\hat{a}_{t,OLS}$	$\hat{b}_{t,OLS}$	$\hat{\hat{\sigma}}_{t,OLS}$	$\hat{a}_{t,LAD}$	$\hat{b}_{t,LAD}$	$\hat{\hat{\sigma}}_{t,LAD}$
Mean	0.659%	0.022	0.506	0.683%	0.021	0.552	0.696%
Standard deviation	0.223%	0.007	0.199	0.184%	0.014	0.302	0.196%
Minimum	0.304%	-0.014	0.030	0.426%	-0.009	-0.399	0.383%
Maximum	2.223%	0.040	1.451	2.651%	0.066	1.256	2.300%
5th percentile	0.404%	0.000	0.096	0.509%	0.000	-0.023	0.481%
95th percentile	1.069%	0.000	0.745	0.993%	0.000	0.924	0.481%
Average abs correction	-	-	-	10.428%	-	-	9.647%

Panel I: GBP/JPY

	$\hat{\sigma}_t$	$\hat{a}_{t,OLS}$	$\hat{b}_{t,OLS}$	$\hat{\hat{\sigma}}_{t,OLS}$	$\hat{a}_{t,LAD}$	$\hat{b}_{t,LAD}$	$\hat{\hat{\sigma}}_{t,LAD}$
Mean	0.709%	0.027	0.505	0.736%	0.022	0.599	0.740%
Standard deviation	0.217%	0.017	0.290	0.183%	0.019	0.340	0.195%
Minimum	0.286%	-0.023	-0.345	0.429%	-0.016	-0.444	0.370%
Maximum	1.968%	0.084	1.764	2.086%	0.086	1.410	2.085%
5th percentile	0.432%	0.000	-0.181	0.517%	0.000	-0.014	0.504%
95th percentile	1.078%	0.000	0.808	1.037%	0.000	1.024	0.504%
Average abs correction	-	-	-	9.313%	-	-	7.950%

corrected VaR forecasts is rejected for three series.

Using the LAD-corrected VaR forecasts, the unconditional coverage is further improved. The null hypothesis of correct unconditional coverage is now rejected for only one series. The null hypothesis of independence using the LAD bias-corrected VaR forecasts is again rejected for three series.

At the 97.5% VaR significance level, therefore, use of the proposed bias correction again substantially improves the unconditional coverage of the RiskMetrics VaR approach, but does not have a significant effect on the independence of the VaR exceptions.

Exhibit 5 reports the results for the 95% VaR confidence level. The uncorrected VaR forecasts provide approximately correct unconditional coverage for all portfolios except the Japanese equity portfolio, where the null hypothesis of correct unconditional coverage is rejected at the 1% level. The null hypothesis of independence is rejected for five of the nine portfolios, however.

Using the OLS-corrected VaR forecasts leads to some deterioration in the unconditional coverage, and in several cases the null hypothesis of correct unconditional coverage is rejected. The OLS correction does improve the independence of the VaR exceptions in six of the nine portfolios, and the null hypothesis of independence is rejected in only four cases. Using the LAD-corrected VaR forecasts leads to a further deterioration in the unconditional coverage, but a further improvement in the conditional coverage with the null hypothesis of independence rejected in only three cases.

EXHIBIT 3

Unconditional Coverage—LR_U and LR_I Statistics for Uncorrected and Bias-Corrected 99% One-Day VaR Forecasts

	Uncorrected			OLS Bias-Corrected			LAD Bias-Corrected		
	Coverage	**LR_U**	**LR_I**	**Coverage**	**LR_U**	**LR_I**	**Coverage**	**LR_U**	**LR_I**
US Bonds	1.794%	18.377***	0.000	1.766%	17.208***	0.012	1.541%	9.068***	0.000
UK Bonds	1.654%	12.862***	0.001	1.457%	6.607**	0.073	1.401%	5.161**	0.116
Japanese Bonds	1.822%	19.578***	4.361**	1.682%	13.898***	5.348**	1.626%	11.860***	0.933
S&P 500	1.850%	20.810***	6.995***	1.682%	13.898***	2.693	1.626%	11.860***	2.982*
FTSE 100	1.570%	9.962***	0.00	1.317%	3.299*	0.000	1.289%	2.763*	0.000
Nikkei 225	1.990%	27.423***	1.417	1.541%	9.068***	1.178	1.541%	9.068***	0.027
USD/GBP	1.570%	9.962***	1.093	1.345%	3.878**	4.747**	1.233%	1.824	2.403
USD/JPY	1.345%	3.878**	0.000	1.177%	1.070	0.410	1.149%	0.765	0.000
GBP/JPY	1.457%	6.607***	0.073	1.261%	2.271	0.000	1.233%	1.824	0.000
Average	1.672%	14.384	1.549	1.470%	7.911	1.607	1.404%	6.021	0.718

*, **, *** *denote significance at the 10%, 5%, and 1% levels.*

EXHIBIT 4

Unconditional Coverage—LR_U and LR_I Statistics for Uncorrected and Bias-Corrected 97.5% One-Day VaR Forecasts

	Uncorrected			OLS Bias-Corrected			LAD Bias-Corrected		
	Coverage	**LR_U**	**LR_I**	**Coverage**	**LR_U**	**LR_I**	**Coverage**	**LR_U**	**LR_I**
US Bonds	3.111%	5.077**	2.541	2.971%	3.064*	0.508	2.719%	0.681	1.392
UK Bonds	3.027%	3.812*	0.846	2.522%	0.007	0.225	2.410%	0.119	0.003
Japanese Bonds	2.915%	2.393	8.381***	2.578%	0.089	11.764***	2.466%	0.017	7.296***
S&P 500	3.167%	6.014**	0.539	2.719%	0.681	1.793	2.635%	0.260	0.840
FTSE 100	2.971%	3.064*	2.196	2.494%	0.000	0.261	2.550%	0.037	2.482
Nikkei 225	3.672%	17.595***	1.926	2.971%	3.064*	3.778*	3.027%	3.812*	3.471*
USD/GBP	3.335%	9.258***	2.046	2.719%	0.681	5.357**	2.635%	0.260	3.851**
USD/JPY	2.663%	0.379	0.126	2.158%	1.793	0.068	2.186%	1.505	0.845
GBP/JPY	2.971%	3.064*	0.508	2.382%	0.206	0.432	2.382%	0.206	0.000
Average	3.092%	5.628	2.123	2.613%	1.065	2.687	2.557%	0.766	2.242

*, **, *** *denote significance at the 10%, 5%, and 1% levels.*

At the 95% VaR significance level, therefore, use of the proposed bias correction worsens the unconditional coverage, but offers some improvement in terms of independence.

IV. CONCLUSION

A wide range of models have been proposed to capture the well-established feature of time-varying conditional volatility in financial markets, including GARCH, stochastic volatility, rolling sample variance, and EWMA estimators. As we note, these models are at best only approximations of the true data-generating process for financial asset returns. Many will be seriously misspecified, resulting in forecasts of the true conditional volatility that are conditionally biased.

We use a regression framework to evaluate the conditional bias in the RiskMetrics approach to the estimation of VaR that assumes returns are conditionally normal and uses an EWMA model with a decay factor of 0.94 in order to forecast volatility. Using daily data on nine portfolios of equities, bonds, and foreign currency, we find that the RiskMetrics EWMA conditional volatility forecasts are seriously biased, with estimated slope parameters of the

EXHIBIT 5

Unconditional Coverage—LR_U and LR_I Statistics for Uncorrected and Bias-Corrected 95% One-Day VaR Forecasts

	Uncorrected			OLS Bias-Corrected			LAD Bias-Corrected		
	Coverage	LR_U	LR_I	Coverage	LR_U	LR_I	Coverage	LR_U	LR_I
US Bonds	4.877%	0.115	0.034	4.765%	0.423	0.107	4.456%	2.301	0.191
UK Bonds	4.652%	0.928	6.057**	4.204%	5.019**	4.521**	4.148%	5.774**	4.865**
Japanese Bonds	4.737%	0.530	10.480***	3.924%	9.365***	17.464***	4.092%	6.586**	15.325***
S&P 500	5.017%	0.002	0.473	4.288%	3.991**	1.717	4.260%	4.320**	1.818
FTSE 100	5.437%	1.398	4.967**	4.709%	0.650	4.286**	4.568%	1.439	6.666***
Nikkei 225	6.278%	11.389***	2.564	5.297%	0.651	2.441	5.521%	1.977	1.585
USD/GBP	5.325%	0.778	7.098***	4.905%	0.069	5.784**	4.793%	0.327	2.680
USD/JPY	4.793%	0.327	0.201	3.812%	11.515***	0.007	3.672%	14.546***	0.296
GBP/JPY	5.269%	0.535	3.588*	4.288%	3.991**	0.031	4.204%	5.019**	0.081
Average	5.154%	1.778	3.940	4.466%	3.964	4.040	4.413%	4.699	3.723

*, **, *** *denote significance at the 10%, 5%, and 1% levels.*

order of 0.5 to 0.7. This suggests that significant improvements may be obtained by correcting the EWMA conditional forecasts using the estimated regression parameters.

To allow for parameter inconstancy, we use a rolling regression, estimating the bias-correction parameter using the most recently available data. We use the bias-corrected conditional volatility estimates to forecast one-day VaR for the nine portfolios at various VaR confidence levels.

On balance, use of the bias-corrected forecasts yields a substantial improvement, particularly in the unconditional coverage of the VaR estimates, and particularly at higher VaR confidence levels. Moreover, owing to nonnormality in the realized volatility regression, the robust LAD estimator provides significantly better VaR forecasts than the OLS estimator.

It would be straightforward to apply the bias correction procedure to other methods of estimating VaR, such as historical simulation and extreme value theory allowing for time-varying conditional volatility (see Hull and White [1998] and McNeil and Frey [2000]). Moreover, it could be applied using other volatility forecasting models.

We have used a conditional volatility model that we expect a priori to be seriously misspecified (but choose it because it is one widely used by practitioners). It should thus not be surprising that the volatility forecasts it generates are severely biased and that there are potential gains from bias correction. Although more sophisticated models such as GARCH are likely to be better specified, there is still likely to be room for improvement in the VaR forecasts that they generate.

ENDNOTES

[1]For a review of GARCH models, see Bollerslev, Engle, and Nelson [1994]. For an empirical evaluation of these models, see Hansen and Lunde [2001]. For a review of the stochastic volatility literature, see, for example, Taylor [1994].

[2]In fact, the choice of 0.94 for the decay factor in the RiskMetrics EWMA model is based on the average estimate of the decay factor (using an out-of-sample mean square error criterion) for a large number of assets.

[3]It is assumed that the mean return is equal to zero, which is a common assumption in modeling the variance of short-horizon asset returns, justified by the fact that the mean return of an asset is typically several orders of magnitude lower than its standard deviation (see, for example, Figlewski [1997] or Jorion [2000]). It would be straightforward to accommodate a non-zero mean return.

[4]That the EWMA model generates forecasts of the conditional variance that are unconditionally unbiased can be seen by taking the unconditional expectation of both sides of the second equation in (6) and noting that $E[\sigma_t^2] = E[r_t^2] = \sigma^2$ by the law of iterated expectations. See also Figlewski [2003], who examines the impact of estimation error on VaR.

[5]We also tried regression window lengths of 250 and 1,000 observations. While there are some differences in the performance of different window lengths, these differences do not appear to be systematic. Results for 250, 500, and 1,000 window lengths, using a common evaluation sample, are available from the authors on request.

[6]The unrestricted LAD estimator, as expected, leads to less volatile estimates of *a* and *b*, but yields worse VaR forecasts because is it is unconditionally biased.

[7]The higher variance of the estimated regression parameters using LAD does not translate into significantly higher variance of the corrected volatility forecast because the negative correlation between $\hat{a}$ and $\hat{b}$ is higher for LAD than for OLS.

REFERENCES

Alexander, C., and C. Leigh. "On the Covariance Matrices Used in Value at Risk Models." *The Journal of Derivatives*, 4 (1997), pp. 50-62.

Andersen, T., and T. Bollerslev. "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts." *International Economic Review*, 39 (1998), pp. 885-905.

Baillie, R., and R. DeGennaro. "Stock Returns and Volatility." *Journal of Financial and Quantitative Analysis*, 25 (1990), pp. 203-214.

Barndorff-Nielsson, O., and N. Shephard. "Econometric Analysis of Realised Volatility and its Use in Estimating Stochastic Volatility Models." Working paper, University of Oxford, 2001..

Bollerslev, T. "Generalized Autoregressive Conditional Heteroskedasticity." *Journal of Econometrics*, 31 (1986), pp. 307-327.

Bollerslev, T., R. Engle, and D. Nelson. "ARCH Models." In R. Engle and D. McFadden, eds., *Handbook of Econometrics*. Amsterdam: North-Holland, 1994, pp. 2961-3038.

Boudoukh, J., M. Richardson, and R. Whitelaw. "Investigation of a Class of Volatility Estimators." *The Journal of Derivatives*, 4 (1997), pp. 63-71.

Christensen, B., and N. Prabhala. "The Relation Between Implied and Realized Volatility." *Journal of Financial Economics*, 37 (1998), pp. 125-150.

Christoffersen, P. "Evaluating Interval Forecasts." *International Economic Review*, 39 (1998), pp. 841-862.

Day, T., and C. Lewis. "Forecasting Futures Market Volatility." *The Journal of Derivatives*, 1 (1993), pp. 33-50.

Engle, R. "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation." *Econometrica*, 50 (1982), pp. 987-1007.

Figlewski, S. "Estimation Error in the Assessment of Financial Risk Exposure." Working paper, New York University, 2003.

——. "Forecasting Volatility." *Financial Markets, Institutions and Instruments*, 6, No. 1 (1997).

Guermat, C., and R. Harris. "Robust Conditional Variance Estimation and Value-at-Risk." *Journal of Risk*, 4 (2002), pp. 25-41.

Hansen, P., and A. Lunde. "A Forecast Comparison of Volatility Models: Does Anything Beat a GARCH(1, 1)?" Working paper, Brown University, 2001.

Harris, R., and J. Shen. "Robust Estimation of the Optimal Hedge Ratio." *Journal of Futures Markets*, 23 (2003), pp. 799-816.

Huber, P. *Robust Statistics*. New York: Wiley, 1981.

Hull, J., and A. White. "Incorporating Volatility Updating into the Historical Simulation Method for Value-at-Risk." *Journal of Risk*, 1 (1998), pp. 5-19.

Jorion, P. *Value at Risk*. New York: McGraw-Hill, 2000.

Kupiec, P.H. "Techniques for Verifying the Accuracy of Risk Measurement Models." *The Journal of Derivatives*, 3 (1995), pp. 73-84.

Lopez, J.A. "Regulatory Evaluation of Value-at-Risk Models." *Journal of Risk*, 1 (1999), pp. 37-63.

McNeil, A., and R. Frey. "Estimation of Tail-Related Risk Measures for Heteroscedastic Financial Time Series: An Extreme Value Approach." *Journal of Empirical Finance*, 7 (2000), pp. 271-300.

Mincer, J., and V. Zarnovitz. "The Valuation of Economic Forecasts." In J. Mincer, ed., *Economic Forecasts and Expectations*. Cambridge: National Bureau of Economic Research, 1969.

Nordhaus, W. "Forecasting Efficiency: Concepts and Applications." *Review of Economics and Statistics*, 69 (1987), pp. 667-674.

"RiskMetrics." Technical Document, 4th ed. J.P. Morgan, 1996.

Taylor, J. "Evaluating Volatility and Interval Forecasts." *Journal of Forecasting*, 18 (1999), pp. 111-128.

——. "Modeling Stochastic Volatility." *Mathematical Finance*, 4 (1994), pp. 183-204.

Theil, H. *Applied Economic Forecasting*. Amsterdam: North Holland, 1966.