

Value at Risk for Interest Rate-Dependent Securities

NUSRET CAKICI AND KEVIN R. FOSTER

NUSRET CAKICI
is professor of economics
at the City College of
New York.
ncakici@yahoo.com

KEVIN R. FOSTER
is assistant professor of
economics at City College
of New York
kfooster@ccny.cuny.edu

Value at risk modelers face difficult trade-offs in choosing between the two major VaR methodologies: historical/ Monte Carlo simulations or parametric models. The historical approach uses the recent history of the asset price (or, for derivatives, the price of the underlying), while the parametric approach imposes functional form assumptions upon the expected returns and uses history only to estimate the few governing parameters.

Each model has its limitations. The parametric model gains computational ease at the cost of unrealistic assumptions for the forward distribution, often even assuming normal errors despite a long history of demonstrated abnormal tail behavior (Mandelbrot [1963]). The historical model suffers the "rewind" problem; there is no guarantee that the future will replay the past in the same way. Randomization (Monte Carlo) experiments mitigate this problem but at computational cost. Duffie and Pan [1997] and Linsmeier and Pearson [2000] give excellent overviews.¹

Interest rate modeling has allowed for a wide degree of non-parametric flexibility in how the forward distribution of rate changes depends on the level. Whether one uses the classic Cox, Ingersoll, and Ross (CIR [1985]) specification that volatility increases with the square root of the level, or the non-parametric Ait-Sahalia [1996a] and Stanton [1997] models, it is widely believed that this relationship is important in accurately representing interest

rate changes. A VaR calculated by ignoring this dependence of volatility on the level of rates would miss this essential feature.

Parametric methods that explicitly model this dependence, such as CIR, assume normal errors (actually, non-central chi-squared in the discrete case, although this is built from Wiener processes), and so understate the actual kurtosis of the distribution, leading to inaccurate tail probability and VaR estimates.

Contemporary non-parametric methods for valuing derivatives, such as Stutzer [1996] and Stanton [1997], are increasingly adopted in option pricing, but are only just beginning to find prominence in VaR analyses. They can accurately reflect the thick tails, pronounced skewness, and excess kurtosis of financial asset price returns. Kernel density estimation allows modeling of this tail behavior in flexible yet workable ways.

Stutzer [1996] provides a way to construct a risk-neutral probability distribution using the maximum amount of information from the historical record of the underlying security. Replacement of the conventional log-normal assumption with this risk-neutral historical distribution allows a rich variety of effects on option values. Rather than squeezing the empirical variation into a single volatility parameter, Stutzer allows an entire dimension of variation to enter non-parametrically, but only a single dimension, while interest rate-dependent securities need more.

While we keep the desirable non-parametric features of Stutzer and Stanton, the method we demonstrate allows an extension of the historical record to two dimensions, to include the characteristic dependence of rate changes upon the present level of interest rates. This approach is not just a non-parametric stochastic volatility model, but allows all the higher moments to vary, so it may be termed a *stochastic higher moments kernel model*.

We call the multidimensional model suitable for VaR analysis the *historical higher moments VaR* (HHM-VaR). It allows a tremendously flexible, completely non-parametric dependence of the distribution of asset price changes upon the level of interest rates. This method uses a normal kernel with adjustable bandwidth (see Chapman and Pearson [2000] for the importance of bandwidth choice).

Between the twin perils of overparameterization or historical "rewind," this method is both feasible and desirable. Many competing measures of VaR suffer from additional model risk—the possibility that the model is misspecified and so the VaR bounds are incorrect. The HHM-VaR necessitates very few assumptions about the driving processes, requiring little prior information, which makes it an ideal component of a process. In many applications, a completely non-parametric baseline model that makes very few assumptions about implied future rate paths can form an enlightening comparison.

An attractive feature of the HHM-VaR is that it does not assume symmetry in error distributions; the asymmetry is determined entirely by the data on the underlying security. Any asymmetry in empirical returns is important to a risk manager, as it means that long and short positions do not cancel each other out by a VaR measure. Our method can be extended to evaluate portfolios of bonds, although adding additional factors may exacerbate worries of overfitting.

While there is debate about whether the Aït-Sahalia and Stanton models go too far in giving the maximum voice to the data, we believe that this HHM-VaR gives an interesting description of the risk that faces a particular portfolio. No portfolio manager should believe that there is a single true VaR somewhere, but rather there are a variety of different measures (based on particular statistical representations) that can be used to triangulate the true risk facing the position.

We test the performance of the HHM-VaR against two alternative specifications: a simple rolling normal VaR model and a restricted GARCH(1, 1) (generalized autocorrelated conditional heteroscedastic) model (see Bollerslev, Chou, and Kroner [1992] and Vlaar [2000]). The

historical higher-moments method of constructing VaR boundaries is indistinguishable from its theoretical value in 19 of the 30 comparisons, while the GARCH is indistinguishable in just 12.

The HHM-VaR produces hit rates that are generally more accurate than a GARCH specification, and even when the HHM-VaR is somewhat less accurate, it is never as far off as the GARCH can be. The worst miss for the HHM measure is 0.65 percentage points. The GARCH is at least that far away in 8 of 30 subgroups, and as far away as 1.67 percentage points.

I. JOINT DISTRIBUTION OF r_t AND Δr_{t+1}

The modeling of the dependence of the distribution of rate changes as a function of the level extends at least back to Cox, Ingersoll, and Ross [1985]. This model posits a Brownian motion process whose diffusion parameter (variance) is itself a function of the level, in addition to the standard dependence of the drift upon the level. The model, with a linear error correction term for the drift and a diffusion proportional to the square root of the level, is:

$$dr = \kappa(r - \theta)dt + (\sigma\sqrt{r})dW \quad (1)$$

where κ is the speed at which deviations from the long-term rate, θ , are corrected and $\sigma\sqrt{r}$ is the diffusion. CIR provide an exact solution so that the discrete-time difference of this continuous-time interest rate process is distributed as a non-central chi-square.

Aït-Sahalia [1996a] makes an important generalization by specifying the diffusion as depending non-parametrically upon the level of interest rates. That specification, however, maintains the assumption of a linear error correction factor, so that the rate process is given as:

$$dr = \kappa(r - \theta)dt + \sigma(r)dW \quad (2)$$

Finally, Stanton [1997] gives a further generalization by allowing the drift process to also depend non-parametrically upon the level, so:

$$dr = \mu(r)dt + \sigma(r)dW \quad (3)$$

so both $\mu(r)$, the drift, and $\sigma(r)$, the diffusion, can be flexibly estimated.

Stanton's model is criticized by Chapman and Pearson [2000] for making a low power estimation at very high levels of interest rates. This low power estimation will make the VaR bounds too tight at extreme values, but backtesting will reveal if this problem is important.

The empirical evidence has developed parallel to the theory, with various proposals to estimate or calibrate model parameters. Traditional methods of recovering the structure of interest rate changes from the historical record run regressions of the general form, $\Delta r_{t+1} = g(r_t, \epsilon_t)$, where $\Delta r_{t+1} = r_{t+1} - r_t$ and the function $g(r_t, \epsilon_t)$ may be as simple as an error correction model or may be expanded to include various GARCH specifications of the error term, lags of either r_t or ϵ_t , or further complications (see Chan et al. [1992]). A few approaches, such as that of Conley et al. [1997], address the identification issues forced by the discrete nature of the data collected.

The regression coefficients represent the central tendency only and do not allow for a heterogeneity of effects. Methods such as kernel regressions can allow a richer variety of effects, but these are still methods of summarizing the rich variation of the joint distribution, $f(r_t, \Delta r_{t+1})$, and compressing it into a few parameters.

Is there a need for such compression? The joint distribution can be estimated directly, and all the variation can be employed, without the need for compaction, to simulate bond and option prices. Stutzer [1996] uses this basic approach in the univariate case. We generalize to higher dimensions to estimate the joint distribution of the change in the interest rate, Δr_{t+1} and the rate itself, r_t . This joint distribution is estimated with a two-dimensional kernel.

A kernel estimate is basically like a histogram, but while a histogram stacks boxes to represent each data point, the kernel stacks "bumps," and a normal kernel uses Gaussian "bumps." The width of the bump (for the normal kernel, the variance) is controlled by the bandwidth choice, h . The two-dimensional kernel constructs the empirical probability density function $\hat{f}(r_t, \Delta r_{t+1})$, which will under appropriate regularity conditions converge to the true probability density function, $f(r_t, \Delta r_{t+1})$.

The kernel estimator will be appropriate as long as the generating distribution is stationary and ergodic, and the data-generating process satisfies some relatively weak assumptions of stochastic equicontinuity in order for a uniform law of large numbers result to obtain (Andrews [1994]). Stochastic equicontinuity of an empirical process,

$H_T(\theta)$, is defined as $\forall \epsilon > 0$ and $\forall \eta > 0$, $\exists \delta > 0$ such that:

$$\lim_{T \rightarrow \infty} P \left[\sup_{\theta, \theta' \in \Theta, \rho(\theta, \theta') < \delta} \|H_T(\theta') - H_T(\theta)\| > \eta \right] < \epsilon \quad (4)$$

There is a trade-off between the smoothness required of the function and the dimension of the random variable, here two. The kernel estimator at each point is constructed as:

$$H_T(x) = \frac{1}{h^2 T} \sum_{i=1}^T K \left(\frac{x - X_i}{h} \right) \quad (5)$$

where K is the normal (Gaussian) kernel. A normal kernel does not limit the estimated probability density function to be normal.

Since Chapman and Pearson have demonstrated the importance of bandwidth choice in estimating the drift and diffusion functions, we use an adjustable bandwidth kernel. This adjustable bandwidth estimator of the kernel density can be seen as a version of their recommended least squares cross-validation bandwidth choice for kernel regression functions.

The adjustable bandwidth starts with a pilot estimate of the density. This pilot estimate uses the identical and independently distributed (iid) bandwidth:

$$h = \hat{\sigma} T^{-1/6} \quad (6)$$

where $\hat{\sigma}$ is the standard deviation of the sample, and the two variates are prewhitened to have the same variance. This bandwidth choice will be acceptable for a wide range of functional forms.


The pilot estimate of the density is used to scale the bandwidth for calculation of the adjustable-bandwidth kernel, so that areas of high density use a small bandwidth while areas of low density use a higher bandwidth choice. Following Silverman [1986], we define local bandwidth factors, λ_i , by

$$\lambda_i = [\tilde{f}(X_i) / g]^{-\alpha} \quad (7)$$

where the pilot estimate is $\tilde{f}(X)$; g is the geometric mean:

EXHIBIT 1

First Four Central Moments of One-Day Interest Rate Changes—One-Month Eurodollar Rates



Kernels estimated with a two-dimensional adjustable bandwidth normal kernel. Each plots the central tendency and 90% error bands.

$$g = \exp\left\{\frac{1}{T} \sum \tilde{f}(X_i)\right\} \quad (8)$$

and α is the sensitivity parameter, which here is 0.5.

Then the adaptive bandwidth kernel is


$$\hat{f}(x) = \sum_{i=1}^T \frac{1}{h^2 \lambda_i^2 T} K\left(\frac{x - X_i}{h \lambda_i}\right) \quad (9)$$

where K is again the normal (Gaussian) kernel. The kernel is estimated on a 500×500 grid and converted to a conditional cumulative distribution at each discretization of r_t . Each conditional density is allowed to be different, for different levels of the interest rate, r_t .

Calculating the value at risk interest rate is straightforward. The percentiles of the conditional kernel density

EXHIBIT 2

First Four Central Moments of Ten-Day Interest Rate Changes—One-Month Eurodollar Rates



Kernels estimated with a two-dimensional adjustable bandwidth normal kernel.
Each plots the central tendency and 90% error bands.

are calculated at the gridpoints, and the actual VaR is interpolated. Specifically, to find the value at risk interest rate, we calculate a value, $V(r_t)$, so that there is probability of α that the interest rate τ days from now is higher than it:

$$\Pr\{r_{t+\tau} > V(r_t)\} = \alpha \quad (10)$$


or, for the other tail, $V''(r_t)$, so that:

$$\Pr\{r_{t+\tau} < V''(r_t)\} = 1 - \alpha \quad (11)$$

We calculate results for the ten-day ($\tau = 10$) and one-day ($\tau = 1$) cases.

EXHIBIT 3

First Four Central Moments of Ten-Day Interest Rate Changes—Ten-Year Treasury Rates



Kernels estimated with a two-dimensional adjustable bandwidth normal kernel.
Each plots the central tendency and 90% error bands.

II. STOCHASTIC KERNEL VAR MODEL

The data on interest rates downloaded from the Federal Reserve website incorporate a variety of maturities. We examine one-year, five-year and ten-year rates. The one-month and three-month eurodollar interest rates are for each business day from January 4, 1971, through February 15, 2002.² The short eurodollar rates are con-

verted to a continuously compounded yield to maturity basis. The longer-maturity rate series are on constant-maturity Treasury bonds, beginning January 2, 1962.

With a kernel specification, we use the maximum amount of data available. Other models such as the GARCH seem to generate results that depend crucially on the exact time period used, and therefore on the discretion of a particular modeler. If a risk manager wants

an analysis that is not so dependent upon these idiosyncrasies, the kernel method may be better.

The first four moments of a kernel estimate of the one-day change in the short-term (one-month) rate are shown in Exhibit 1 along with 90% error bands. Exhibit 2 shows the moments of the ten-day changes. In both it is clear that while the variance increases with the level of the interest rate, the higher moments also change. Clearly the variance increases with the level of interest rates, with a suggestion of either a break or perhaps a parabolic shape, as found in the Ait-Sahalia [1996b] study of interest rate volatility.

These changing moments do not appear to have a straightforward parametric form, suggesting that any simple model would miss important characterizations of the evolution of rate changes.³ A simple VaR calculation based on normally distributed errors would have to assume that the higher moments are constant, even if it accounts for the changing variance, and so lead to important inaccuracies in the estimation of the tail probabilities.

Exhibit 3 shows the moments of the ten-day changes for the ten-year rate.⁴ The level of interest rates has an important influence on all the conditional moments. The strong negative skew will influence the VaR bounds calculated below. Again, it is clear that a proper VaR model should take account of these effects.


The estimated kernel, at each discretization point of the interest rate level, is converted to a cumulative density, and the appropriate percentiles are linearly interpolated. These percentiles represent the VaR bounds, either as levels of interest rates or as corresponding bond prices. Since the distribution of rate changes is allowed to vary, depending on the level, the VaR bounds will also change dynamically, depending on conditions.

These dynamically changing VaR bounds, for a ten-day-ahead prediction period, are shown in Exhibit 4. These graphs show the 1% and 5% VaR bounds (in each direction, as well as the median) for the one-month eurodollar and the ten-year Treasury (constant-maturity) interest rates. The shorter-rate percentile bounds are more sensitive to the level, showing a little narrowed volatility for rates from 5% to 7%. For a one-day prediction period, the plots look substantially similar, although with tighter bounds (these graphs, as well as the results for other maturities, are omitted but are available from the authors).

The asymmetry in the bounds is particularly important, but it is difficult to see in Exhibit 4 (and better views are shown in exhibits to follow). Exhibit 4 is intended to show the intuition of how the bounds depend on the

EXHIBIT 4

Ten-Day Percentile Bounds— One-Month Eurodollar and Ten-Year Treasury




Central line is the median; the dashed lines are 5% and 95%, and the dotted lines are 1% and 99%. Bounds are calculated from an adjustable bandwidth normal kernel.

level of interest rates. The asymmetry may be crucial in evaluating the risk of a portfolio, as equal short and long positions will not carry equal levels of risk—the skewness and kurtosis (and higher moments) of the forward distribution can change the risk.

EXHIBIT 5

5% Value at Risk of Bonds over Ten-Day Horizon—One- and Three-Month



5% value at risk (in dollars) on a bond with current value of \$100 and the specified maturity over a ten-day period. Solid line: VaR from a long position (if interest rates rise). Dashed line: VaR from a short position.

While a GARCH model can be specified that allows this asymmetry, the choice of how to specify the asymmetry adds further discretion to the GARCH modeler's risk analysis (see Campbell and Hentschel [1992] and Wu

[2001]). The HHM-VaR handles this asymmetry naturally and flexibly, without undue dependence upon the modeler's choices. Instead, the data determine the nature and extent of the dependence of skewness and kurtosis on the level of rates.⁵

The actual value at risk, for a bond with a given maturity and a current value of 100 (i.e., a face value of $100\exp\{rT\}$), is calculated as:

$$VaR_{\alpha} = \left(100 \exp\left\{\left(r - \hat{r}_{\alpha}\right)T\right\}\right) - 100 \quad (12)$$

where r is the present level of the interest rate, $\hat{r}_{\alpha}$ is the critical α -percentile bound, and T is the time to maturity.

Exhibits 5, 6, and 7 show the 5% value at risk over a ten-day horizon for a bond with a \$100 current value and maturities ranging from one month to ten years. (The figures for other percentiles of value at risk are similar.)

In each graph, two values at risk are shown. The solid line shows the VaR for a long position (if interest rates rise), while the dashed line shows the VaR for a short position (if rates fall). With this simple model of the interest rate, more complex derivatives can be straightforwardly modeled.

The first graph in Exhibit 5 shows the short and long VaR on one-month Eurodollar bonds over a standard ten-day interval, depending on today's interest rate (on the x-axis). The dashed line is the VaR for a short bond position, representing the chance that rates will fall; the solid line is the VaR for a long bond position.



At low levels of rates, the one-month leaves only a few pennies at risk, but as rates rise the value at risk rises steeply, so that at interest rates above 10% the VaR is at least four times higher. This should convince any risk manager that the dependence of forward rate changes upon the present level is important even for very short-maturity instruments and should be carefully modeled. At these short rates, there is little pronounced skew to push the bounds apart.

The three-month rates, in the second graph in Exhibit 5, have a similar shape. In this case, the plot more dramatically shows that not only does the VaR increase dramatically at rates above 9% but also for rates from 4% to 6% the VaR may actually fall.

Exhibits 6 and 7 show the Treasury bond rates at longer maturities. The one-year is similar to the short instruments, while the longer maturities show a marked pivot of the long and short VaR. There is clearly more VaR from a short bond position when interest rates are high, and less when rates are low. The pivot point, for both

EXHIBIT 6

5% Value at Risk of Bonds over Ten-Day Horizon—One- and Five-Year




5% value at risk (in dollars) on a bond with current value of \$100 and the specified maturity over a ten-day period. Solid line: VaR from a long position (if interest rates rise). Dashed line: VaR from a short position.

the five-year and the ten-year, is about 7%. For ten-year bonds, at high interest rates the value at risk of a short position is almost \$2 more than the VaR of a long position.

This difference is caused by the skewness in the distribution, clearly shown in Exhibit 3. This skew is an important characterization of the temporal evolution of

EXHIBIT 7

5% Value at Risk of Bonds over Ten-Day Horizon—Ten-Year



5% value at risk (in dollars) on a bond with current value of \$100 and the specified maturity over a ten-day period. Solid line: VaR from a long position (if interest rates rise). Dashed line: VaR from a short position.

long-maturity interest rates, yet a model based on normally distributed errors would overlook it entirely. The skew means that a portfolio with balanced short and long positions would not be riskless by the VaR criterion.

III. EVALUATING PERFORMANCE OF STOCHASTIC KERNEL VaR MODELS

Exhibits 5, 6, and 7, showing the evolution of the value at risk as a function of the level of interest rates, give an overview of how these bounds are set, but they must be accurate predictors of the VaR over time as well. To test the performance of these HHM-VaR bounds, we backtest the ability of the HHM-VaR to characterize the realistic risks facing the bondholder. The performance of the HHM-VaR is tested against two primary alternative specifications: a simple rolling normal VaR and a restricted GARCH(1, 1).

For the rolling normal VaR, the last T data points are used to imply the standard deviation of rate changes going forward if they were normally distributed. We use $T = 100$. This backward-looking window is obviously not a predictor that would be based on rational expecta-

071

EXHIBIT 8

Ten-Day VaR Hit Rates—One-Month Eurodollar Interest Rates

	Theoretical Value	HHM	Rolling Normal	GARCH	Tail
Percent	0.01	0.004503	0.045713	0.014463	Lower
Number (of 7329)	73	33	335	106	
Percent	0.01	0.007777	0.039479	0.024424	Upper
Number	73	57*	289	179	
Percent	0.025	0.019512	0.061781	0.025788	Lower
Number	183	143	453	189*	
Percent	0.025	0.02838	0.053747	0.040251	Upper
Number	183	208*	394	295	
Percent	0.05	0.044481	0.083253	0.044344	Lower
Number	366	326	610	325	
Percent	0.05	0.054032	0.073556	0.066994	Upper
Number	366	396*	539	491	

Exhibits 8-12: The first column (in bold) is the theoretical probability and the number of times that the interest rate will be beyond the given VaR threshold. The remaining columns show the fraction and number of days the rate was actually beyond, as calculated by the historical higher moments (column 2), the rolling normal (column 3), and restricted GARCH(1, 1) (column 4). The last column indicates whether the comparison is for rates below the lower bound or those above the upper bound in the tail of errors. An asterisk indicates if the particular specification is not significantly different (at 95%) from the theoretical value. Because the rolling normal does not use the first 100 observations, the hit numbers are rescaled to be comparable.

072

EXHIBIT 9

Ten-Day VaR Hit Rates—Three-Month Eurodollar Interest Rates

	Theoretical Value	HHM	Rolling Normal	GARCH	Tail
Percent	0.01	0.003913	0.038398	0.014641	Lower
Number (of 7923)	79	31	304	116	
Percent	0.01	0.008709	0.037246	0.01666	Upper
Number	79	69*	295	132	
Percent	0.025	0.020194	0.053373	0.025243	Lower
Number	198	160	423	200*	
Percent	0.025	0.02802	0.05862	0.033068	Upper
Number	198	222*	464	262	
Percent	0.05	0.046952	0.077051	0.05036	Lower
Number	396	372*	610	399*	
Percent	0.05	0.053768	0.083835	0.062603	Upper
Number	396	426*	664	496	

tions, and of course the window of how many previous data points are used may change, depending on a particular analyst's view of market conditions, but we intend the rolling backward-looking VaR to be a useful benchmark rather than the best possible measure.

The GARCH(p, q) model is based on the conventional specification of the evolution of the conditional variance, σ_t^2 , of rate innovations, that:

$$\sigma_t^2 = \kappa + \sum_{i=1}^p \gamma_i \sigma_{t-i}^2 + \sum_{j=1}^q \alpha_j \varepsilon_{t-j}^2 \quad (13)$$

As often is true in VaR modeling, our basic specifications improve little whether the parameters are estimated or specified *a priori* to the RiskMetrics values: $\kappa = 0$, $\gamma = 0.94$, $\alpha = 0.06$, and $p = q = 1$.⁶ The restriction is intended to reduce the risk of overfitting within-sample. This GARCH specification, however, does not take into account the predictable effects of the level of interest rates upon the variance and higher moments. These effects are predictable and (in other strands of literature) commonly accounted for. The innovation of our research is in showing a model that explicitly but flexibly takes this dependence seriously.

The typical analysis of the performance of a VaR measure counts the number of times a stated bound is hit. Gauging the accuracy of VaR numbers is difficult, however, without a well-specified loss function. While the hit rate, the number of times the actual loss is greater than the predicted VaR, is important, the costs of setting too high or too low a VaR are unclear. A model that tightens up the VaR without sacrificing the hit rate would still be an improvement (see Christoffersen [1998] for a discussion of possible loss functions). Our analysis will assume that misses of the VaR in either direction are penalized.

The comparisons of hit rates for a ten-day window are in Exhibits 8-12; each exhibit shows the results for a different maturity. The hit rates for a one-day window are substantially similar and so are omitted. The exhibits show both the lower tails of rates that were lower than predicted and the upper tail of rates that were higher than predicted with bounds set for levels of at 1.0%, 2.5%, and 5.0%.

While an α of 1% is most common, we present the results for a broader range of categories, as we believe a risk manager would not consider a VaR estimate that hits the 1% level in backtesting at the expense of inadequate power at other levels. To check the robustness of the model, we also show the higher levels.

An important feature of the HHM-VaR is that it does not assume symmetry in the error distributions. Therefore, we present the tails separately rather than hiding behind composite or two-tailed tests. Portfolios of long and short positions in bonds and derivatives will generally necessitate an accurate idea of the risk of both up and down movements in rates.

Each pair of rows in Exhibits 8-12 gives the percentage and the actual number of hits during backtesting, when the specified bound was breached. The first column shows the theoretical percent and number, which is what each method is trying to approximate. The second column gives the HHM-VaR hits, the third the hits for the rolling

normal VaR, and the fourth column the hits for the restricted GARCH(1, 1) specification.⁷ The rolling normal statistics are included to provide a baseline for how much information any model is adding, in order to better assess the marginal contribution.

An asterisk indicates if the model's VaR hit rates are statistically indistinguishable at a 95% level from the true value (see Kupiec [1995] and Christoffersen [1998]). The HHM-VaR measure is statistically indistinguishable from the true value in 19 of the 30 maturity-risk cases, while the GARCH meets this standard only in 12. In 6 cases of the 30, neither measure falls within a 95% confidence interval around the true value.

Exhibit 8 shows the estimates for the ten-day VaR of one-month eurodollar rates. The 1.00% risk on the lower tail of interest rates should be breached only 73 times during the backtesting period (= 1% of 7,329). The HHM-VaR sets a too-conservative bound that is breached only 33 times, which is just 0.45% of the periods; the GARCH sets a closer bound that is breached 106 times, at 1.45%. The rolling normal is clearly worse; it breaches the bound 335 times, which is 4.57%.

The 1.00% bound for the upper tail set by the HHM-VaR is broken a conservative 0.78% (57 times). This is not significantly different from the theoretical 1.00% value, unlike the 2.44% at which the GARCH bound is broken (179 hits).

For the 2.50% level, the GARCH is a clear winner for the lower tail, setting boundaries that are hit 189 times compared with 183 "true," while the HHM is again conservative and sets bounds that are hit only 143 times (1.95%). On the other tail, however, the HHM-VaR is better: 208 hits (2.84%) rather than the GARCH's 295 hits (4.03%).

At the 5.00% level, the HHM and GARCH bounds are essentially tied for estimating the lower bound, but the upper bound is set significantly tighter by the HHM-VaR. If an analyst were to set a level of rates above which only 5.00% of the ten-day windows would rise, it should be hit 366 times. The HHM-VaR hits this level 396 times (5.40%, which is indistinguishable from 5%), while the GARCH sets a level that is 491 times (6.70%)—clearly much worse.

This relative performance of the two chief measures holds for other VaR bounds and other maturities. As the GARCH specifies symmetric conditional errors, it cannot account for the skewness of rate changes. One boundary is typically set well by both models, but the GARCH may miss the other boundary by as much as a full percentage point. In no test at any maturity-risk partition

023

EXHIBIT 10

Ten-Day VaR Hit Rates—One-Year Treasury Bond Interest Rates

	Theoretical Value	HHM	Rolling Normal	GARCH	Tail
Percent	0.01	0.007605	0.034403	0.010908	Lower
Number (of 9993)	100	76	344	109*	
Percent	0.01	0.007405	0.0427	0.013509	Upper
Number	100	74	427	135	
Percent	0.025	0.022015	0.054235	0.021515	Lower
Number	250	220*	542	215	
Percent	0.025	0.024417	0.063746	0.022816	Upper
Number	250	244*	637	228*	
Percent	0.05	0.045832	0.079531	0.040528	Lower
Number	500	458*	795	405	
Percent	0.05	0.050335	0.092381	0.045832	Upper
Number	500	503*	923	458*	

does the HHM-VaR miss by over 0.65 percentage points.

As the graphs have shown, since skewness and excess kurtosis are persistent characteristics of the data, and vary in intensity depending on the level of interest rates, any quantification of the risk of a portfolio must take account of these effects. Since these effects are predictable consequences of the level of interest rates, the GARCH model, where heteroscedasticity is a function of lagged deviations, must struggle to keep up.

Exhibit 9 shows the same comparison for the three-month eurodollar rates. These show that even when the HHM-VaR misses, it is rarely by far. The worst performance of the HHM-VaR of all the targets for the three-month bonds, where the 1.00% lower tail should be hit 79 times but the conservative HHM bound is hit only 31 (0.39%), is 48 misses. The three-month GARCH model misses by at least 48 in three of six maturity-risk comparisons.

In every comparison at each maturity, the HHM-VaR is either better or only slightly worse, while the GARCH is at times far from the true value. Sometimes the HHM-VaR and GARCH targets are on opposite sides of the theoretical level; one is a bit too conservative and the other a bit too lax. This suggests that, while risk managers might not throw away GARCH-set VaR targets, the HHM-VaR may be another useful tool to gauge the riskiness of portfolios.

The comparisons for the one-year constant-maturity Treasury bond interest rates are in Exhibit 10. The HHM measure is statistically indistinguishable from the

true value at the higher risk levels, and at the 1% level the GARCH wins for the lower tail. In four of the six comparisons, the HHM is so close to its theoretical value as to be statistically indistinguishable; in three of six, the GARCH is very close. Since these longer-maturity interest rates are less volatile, there is less of a difference between the GARCH and HHM-VaR performance.

Exhibit 11 provides the evaluations for the five-year constant-maturity Treasury bond rates. The two measures trade places for best at the 1.0% level; the GARCH is better for the lower tail and the HHM-VaR for the upper. In all the remaining categories, the HHM-VaR is indistinguishable from the theoretical value. The GARCH misses the lower tails of both the 2.5% and 5.0% as well.

The ten-year maturity comparisons are in Exhibit 12. For each risk level, the HHM is adequate for the upper tail but misses significantly in the lower tail. The GARCH is quite good for the 2.5% level and the upper tail of the 5.0% level. In both the 2.5% and 5.0%, the GARCH and HHM miss conservatively, while the 1.0% bound is breached too often with the GARCH but still set conservatively by the HHM-VaR.

To summarize the results of the backtesting, the historical higher moments method of constructing VaR boundaries is statistically indistinguishable from the theoretical value in 19 of the 30 comparisons, while the GARCH is statistically indistinguishable in just 12. When the GARCH misses, it can miss very wide of the mark. The worst miss for the HHM measure is 65 basis points.

EXHIBIT 11

Ten-Day VaR Hit Rates—Five-Year Treasury Bond Interest Rates

074

	Theoretical Value	HHM	Rolling Normal	GARCH	Tail
Percent	0.01	0.007405	0.035718	0.011108	Lower
Number (of 9993)	100	74	357	111*	
Percent	0.01	0.009507	0.039158	0.01381	Upper
Number	100	95*	391	138	
Percent	0.025	0.024117	0.054639	0.021415	Lower
Number	250	241*	546	214	
Percent	0.025	0.024017	0.06243	0.025918	Upper
Number	250	240*	624	259*	
Percent	0.05	0.048034	0.079429	0.041029	Lower
Number	500	480*	794	410	
Percent	0.05	0.049335	0.09056	0.048234	Upper
Number	500	493*	905	482*	

EXHIBIT 12

Ten-Day VaR Hit Rates—Ten-Year Treasury Bond Interest Rates

075

	Theoretical Value	HHM	Rolling Normal	GARCH	Tail
Percent	0.01	0.006401	0.031857	0.013303	Lower
Number (of 9998)	100	64	319	133	
Percent	0.01	0.008602	0.033475	0.013503	Upper
Number	100	86*	335	135	
Percent	0.025	0.020904	0.054814	0.023705	Lower
Number	250	209	548	237*	
Percent	0.025	0.022104	0.051173	0.024905	Upper
Number	250	221*	512	249*	
Percent	0.05	0.043509	0.080502	0.044909	Lower
Number	500	435	805	449	
Percent	0.05	0.046309	0.0804	0.046809	Upper
Number	500	463*	804	468*	

The GARCH is more than 65 basis points away in eight subgroups and as far away as 167 basis points.

The HHM-VaR is particularly good in comparisons of the shorter maturities, which are the more volatile. The two eurodollar short maturities are the most flattering comparison. The HHM-VaR is not significantly different from the theoretical value in 7 of 12, while the GARCH is not significantly different 3 times. The three longer maturities show HHM wins in 12 of 18 while GARCH wins in 9.

IV. CONCLUSION

Our work demonstrates the importance of constructing VaR bounds that vary dynamically depending on the level of the interest rate. This dependence, which is crucial for interest rate-dependent securities, has been modeled non-parametrically by authors such as Ait-Sahalia [1996a] and Stanton [1997]. The non-parametric method can accurately reflect the thick tails, pronounced skewness, and excess kurtosis of financial asset price returns.

An adjustable-bandwidth two-dimensional kernel is a flexible yet workable estimator of the VaR. It allows all the higher moments to vary, not just the first two, so it is not stochastic volatility but stochastic higher moments with a kernel—the HHM-VaR.

We have shown how to construct one-day and ten-day HHM-VaR bounds that depend upon the present level of interest rates. Since the variance, skewness, kurtosis, and higher moments of the distribution are allowed to vary non-parametrically with the level of rates, the HHM-VaR produces dynamically changing boundaries of risk. These VaR bounds are shown to be significantly better measures by a variety of criteria.

The HHM-VaR produces hit rates that are generally more accurate than a GARCH specification; even when the HHM-VaR is somewhat less accurate, it is never as far off as the GARCH can be. The worst miss for the HHM measure is 0.65 percentage points. The GARCH is at least that far away in 8 of 30 subgroups and as far away as 1.67 percentage points.

The methods of updating are very different for the GARCH and the HHM-VaR approaches. The GARCH imposes a smooth relationship, where a sudden shock enters with a particular weight and begins a gradual process that changes the forward risk measures. The HHM-VaR doesn't impose such smoothness. A shock today that changes the level of the interest rate will immediately move the forward distribution to that implied by the new level, without any memory of previous history.

While a careful risk manager may not wish to entirely disregard GARCH-based VaR projections in favor of the HHM-VaR, the results of this study make a compelling case for the use of kernel-based estimators of dynamically varying moments. The method of Barone-Adesi and Giannopoulos [2001] is another non-parametric attempt to take seriously this dependence upon past return volatility, but our approach for interest rates takes advantage of our knowledge that the distribution of interest rate changes is dependent upon the present level of the rate.

Complications remain that are specific to modeling interest rates. An analyst evaluating portfolios of bonds must attend to the comovements along the yield curve, so further research must examine these complexities. The basic HHM-VaR methodology provides an interesting direction for yield curve modeling. Further studies can be made to determine the utility of kernel methods in accurately describing the rate variation. Generalizing interest rate modeling to include an extreme value approach (Bali [2001]) also seems interesting and promising.

The historical higher moments value at risk (HHM-VaR) methodology can improve the estimation and prediction of interest rate percentiles for risk analysis. A kernel estimator can accurately model the way the level of interest rates shifts the entire distribution of forward changes—not only variance, but all the higher moments as well. This kernel estimator is ideally suited for many VaR applications that require an estimator that entails a minimal amount of prior information, yet is still flexible and realistic, in order to eliminate the model risk of other measures. The kernel estimator can incorporate a wide variety of other possible representations including jump and diffusion processes.

ENDNOTES

The authors gratefully acknowledge support from the Schweger Fund, and are thankful for the comments of Turan Bali.

¹A third methodology choice, stress testing, is as much art as science, and so is not addressed here (but see Longin [2000]).

²Due to vagaries in the data collection in the early 1970s, the data sets include slightly different numbers of observations.

³Of course, there are direct implications for the form of the diffusion process, but in this exercise we remain agnostic about the form of the jump component.

⁴The conditional moments for the other maturities are similar, and are therefore not reported here.

⁵Also note that, in applying this method to other price series, there is no necessary reason for this particular state variable, the level of interest rates. The higher moments of different price series may depend on other state variables. This analysis is meant only to be illustrative of the broad technique.

⁶Certainly a determined econometrician could search to find some combination of parameters that improve VaR bounds (see Chen, Hong, and Stein [2001] for a fascinating endeavor), but part of the appeal of the historical higher moments kernel lies exactly in the fact that it does not need to be tweaked. Estimating GARCH parameters that balance the twin risks of overfitting and poor backtesting is an art. The kernel method does not require a strong prior belief and so can provide a baseline figure that is not idiosyncratic depending on the particular modeler's judgment. See Christoffersen, Hahn, and Inoue [2001] for an overview.

⁷Several other specifications were tested. Rolling normal VaRs based on predictions of either the absolute level of the interest rates (rather than the change) as well as rolling normal VaRs based on the percent change are not much different from the rolling normal shown. The percent change can give a rough approximation to the notion that rate change volatilities depend on the level.

REFERENCES

- Aït-Sahalia, Yacine. "Nonparametric Pricing of Interest Rate Derivative Securities." *Econometrica*, 64 (3) (1996a), pp. 527-560.
- _____. "Testing Continuous-Time Models of the Spot Interest Rate." *Review of Financial Studies*, 9 (2) (1996b), pp. 385-426.
- Andrews, Donald W.K. "Empirical Process Methods in Econometrics." Cowles Foundation Discussion Paper No. 1059, 1994.
- Bali, Turan G. "An Extreme Value Approach to Estimating Volatility and Value at Risk." *Journal of Business*, 76, 2001.
- Barone-Adesi, Giovanni, and Kostas Giannopoulos. "Nonparametric VaR Techniques: Myths and Realities." *Economic Notes*, 30 (2) (2001), pp. 167-181.
- Bollerslev, Tim, Ray Y. Chou, and Kenneth F. Kroner. "ARCH Modeling in Finance: A Review of the Theory and Empirical Evidence." *Journal of Econometrics*, 52 (1992), pp. 5-59.
- Campbell, John Y., and Ludger Hentschel. "No News is Good News: An Asymmetric Model of Changing Volatility in Stock Returns." *Journal of Financial Economics*, 31 (1992), pp. 281-318.
- Chapman, David A., and Neil D. Pearson. "Is the Short Rate Drift Actually Nonlinear?" *Journal of Finance*, 55 (1) (2000), pp. 355-388.
- Chen, Joseph, Harrison Hong, and Jeremy C. Stein. "Forecasting Crashes: Trading Volume, Past Returns, and Conditional Skewness in Stock Prices." *Journal of Financial Economics*, 61 (2001), pp. 345-381.
- Christoffersen, Peter F. "Evaluating Interval Forecasts." *International Economic Review*, 39 (4) (1998), pp. 841-862.
- Christoffersen, Peter, Jinyong Hahn, and Atsushi Inoue. "Testing and Comparing Value-at-Risk Measures." *Journal of Empirical Finance*, 8 (2001), pp. 325-342.
- Conley, Timothy G., Lars Peter Hansen, Erzo G.J. Luttmer, and José A. Scheinkman. "Short-Term Interest Rates as Subordinated Diffusions." *The Review of Financial Studies*, 10 (3) (1997), pp. 525-577.
- Cox, John C., Jonathan E. Ingersoll, Jr., and Stephen A. Ross. "A Theory of the Term Structure of Interest Rates." *Econometrica*, 53 (2) (1985), pp. 385-408.
- Duffie, Darrell, and Jun Pan. "An Overview of Value at Risk." *The Journal of Derivatives*, 4(3) (1997), pp. 7-49.
- Kupiec, Paul H. "Techniques for Verifying the Accuracy of Risk Measurement Models." *The Journal of Derivatives*, Winter 1995.
- Linsmeier, Thomas J., and Neil D. Pearson. "Value at Risk." *Financial Analysts Journal*, 56 (2) (2000), pp. 47-67.
- Longin, François M. "From Value at Risk to Stress Testing: The Extreme Value Approach." *Journal of Banking and Finance*, 24 (2000), pp. 1097-1130.
- Mandelbrot, Benoit. "The Variation of Certain Speculative Prices." *Journal of Business*, 36 (4) (1963), pp. 394-419.
- Silverman, B.W. *Density Estimation for Statistics and Data Analysis*. New York: Chapman and Hall, 1986.
- Stanton, Richard. "A Nonparametric Model of Term Structure Dynamics and the Market Price of Interest Rate Risk." *Journal of Finance*, 52 (5) (1997), pp. 1973-2002.
- Stutzer, Michael. "A Simple Nonparametric Approach to Derivative Security Valuation." *Journal of Finance*, 51 (1996), pp. 1633-1652.
- Vlaar, Peter J.G. "Value at Risk for Dutch Bond Portfolios." *Journal of Banking and Finance*, 24 (2000), pp. 1131-1154.
- Wu, Guojun. "Determinants of Asymmetric Volatility." *The Review of Financial Studies*, 14 (3) (2001), pp. 838-859.

To order reprints of this article, please contact Ajani Malik at amalik@iijournals.com or 212-224-3205.

©Euromoney Institutional Investor PLC. This material must be used for the customer's internal business use only and a maximum of ten (10) hard copy print-outs may be made. No further copying or transmission of this material is allowed without the express permission of Euromoney Institutional Investor PLC. Copyright of Journal of Portfolio Management is the property of Euromoney Publications PLC and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.