Systematic Risk in Corporate Bond Credit Spreads

MONICA PEDROSA AND RICHARD ROLL

MONICA PEDROSA is professor of applied economics and economic history at the Universidad Nacional de Educacion a Distancia in Madrid.

RICHARD ROLL is Allstate professor of finance and insurance at the Anderson School of the University of California at Los Angeles.

Large financial institutions around the world are investing heavily in systems for measuring market risks. They are responding to two forces. First, regulators and investors are insisting on better information about short-term fluctuations in firm value. Second, senior executives are focusing on *internal* risk monitoring and control systems, in order to know more precisely what they should be worried about every day.

Financial companies can be regarded as levered portfolios with long positions in assets of varying liquidity and short positions in diverse debt obligations. Daily movements in interest rates, equity prices, and exchange rates affect the overall value of the portfolio/firm. But the portfolio's composition is under the control of managers who may choose to restructure the portfolio frequently in response to market conditions or customer demand, and who employ a mutable menu of hedging instruments. As a consequence, portfolio risk estimated from historical public data such as equity returns can be seriously off the mark. Even insiders can find it difficult to cope with sizable swings in portfolio composition induced by daily trading activity.

An ideal risk measurement system would produce a probability distribution of returns conditional on the firm's current portfolio composition. If such a distribution were available in real time, widely employed regulatory constructs such as "value at risk" (VaR) could be calculated easily; VaR is a level of loss exceeded with a prespecified frequency.

The ideal is rarely attained. To explain the causes of this failure, we note that all assets and liabilities fall into one of three categories according to the difficulty of measuring and managing their risks:

1. Liquid positions for which liquid derivatives are available.
2. Liquid positions without related derivatives.
3. Illiquid assets and liabilities.

Few financial institutions expose themselves to market risks that can be inexpensively avoided. If in the normal course of business positions are maintained in the first category above, they are typically hedged with offsetting derivative contracts. For example, a dealer in long-term government securities will employ futures to lessen or even eliminate the risk of adverse movements in interest rates. A corporate bond dealer faces a more troublesome problem vis-à-vis risk, because interest rate futures can hedge only a portion of the price volatility. At present, there are no convenient derivatives for hedging changes in credit perceptions. Nonetheless, for corporate bonds that are relatively liquid, the firm can alter its credit exposure simply by liquidating the position.

Such an expedient is not available for firms with an illiquid portfolio. Prices are not usually available, so even for that part of an illiquid position theoretically hedgeable with a derivative. the appropriate hedge ratio is hard to estimate. Even worse, the firm cannot simply terminate the position. Risk management in this case devolves to the forlorn hope that diversification will accomplish the bulk of the job, and the remaining risk is typically ignored.

Since (1) hedgeable positions are hedged, and (2) illiquid positions are ignored, the risks of many financial firms agglomerate in the intermediate category of positions that can be marked reliably to market but have no available derivatives. Our article focuses on one of the most common risks of this type, the *systematic* components of credit spreads of fixed-income obligations subject to default. By "systematic," we refer to components of credit valuation that are shared by many individual bonds and hence tend to resist elimination by diversification.

To make the focus more specific, think about the corporate bond trading desk of a investment bank in New York or a commercial bank in London. At a given moment, flow traders will have many individual bond positions, long and short, accumulated as accommodations to customers. The net interest rate sensitivity of the aggregate position can be hedged in large part, although perhaps not perfectly, with interest rate derivatives. Changes in the credit quality of *individual* companies do not result in significant volatility either, provided that the position is well-diversified.

But the interest rate-hedged, well-diversified position is still risky because yield spreads move together.[1] As investors become more or less cautious or alter their beliefs about the general outlook for the economy, they reassess the probability of default for *all* bonds. Even a desk with the best intentions and skills will experience losses when spreads shift unexpectedly in the wrong direction.

In this situation, the probability distribution of systematic changes in credit spreads assumes paramount importance. Calculation of tail probabilities for VaR, internal trader control via position limits, regulator anxieties and actions, and investment performance all depend heavily on this distribution. Its functional form, its parameters, and its stability should be the primary focus of the risk manager's attention.

Theoretically, a credit spread is attributable entirely to the corporation's default option. Rather than make the payments promised on its debt, the corporation can choose to deliver its assets to the bondholders. It would exercise this option, again theoretically, whenever the present value of the remaining bond payments exceeds the value of the assets plus the unexercised option.

Valuing the default option is difficult for several reasons. The option is relatively long-term and has multiple possible exercise dates interspersed with cash emissions by the corporation. In the event of default, the legal system intervenes, subverting the contractual specifications of the bond indenture and creating additional randomness in the eventual cash settlement.

Nonetheless, although an exact option valuation formula remains elusive, some qualitative comparative statics are available. The option's value should depend directly on 1) term to expiration (bond maturity), 2) the likelihood of default (bond rating), and 3) perceived volatility of the firm's assets. The likelihood of default itself must depend on the difference between the value of the firm's assets and the value of its indebtedness. The last factor depends inversely on the level of interest rates.

Using a sample of dollar-denominated bonds, we study the empirical distributions of systematic credit spread changes. The distributions turn out to be mor complex than stationary-Gaussian. We provide evidence about the nature of the intertemporal instability and document the influences of credit quality, maturity, and industry, proxies for determinants of default option value.

I. DATA

Bloomberg Financial Services provides daily on-line indexes of bond credit spreads. Each index is an average of individual corporate bond yields within a maturity, industry, and credit rating category spread relative to a corresponding U.S. Treasury yield. The spreads are "adjusted" for embedded options (presumably for call or convertibility, not default), but we have no information about the adjustment methods. Also, we cannot verify how well the corporate bonds are matched in duration to their allegedly comparable Treasuries.

We secured indexes for five maturities: two, five, seven, ten, and thirty years; four industry groupings: Industrials, Utilities, Financials, and Yankees; and three credit ratings: AA, A, and BBB — a total of sixty different indexes.[2] The time period of the sample spans October 5, 1995, through March 26, 1997, 366 trading da observations. The latest date in the sample was just prioı to commencement of this project. The first date was selected mainly for convenience, although we had tertiary reasons: first, to collect a sufficient number of observations for statistical reliability, and second, to mimic practicing risk managers and regulators, who tend to focus on the recent past.[3]

Exhibit 1 plots the credit spread levels over the sample period.

II. TESTS OF TIME SERIES STATIONARITY

We begin the empirical analysis of credit spread indexes by recording evidence about their intertemporal stationarity, that is, about whether their time series are consistent with the presence of a unit root. If the existence of a unit root cannot be rejected, the series should be at least first-differenced before being used in any statistical model. Time series with unit roots will not usually produce reliable statistical results because asymptotic distributions are never achieved.

A standard unit root test involves the augmented Dickey-Fuller [1981] regression:

$$S_t = \alpha + \rho S_{t-1} + \sum_{j=1}^{p} \gamma_j (S_{t-j} - S_{t-j-1}) + \varepsilon_t$$

where S_t denotes the credit spread on date t; α, ρ, and γ_j are parameters to be estimated; and ε_t is a spherical disturbance. Although the number of lags, p, can be selected in advance, a common practice is to use the highest significant lag from the autocorrelation function of S.[4]

Allowing for a deterministic time trend in S, the null hypothesis of a unit root (non-stationarity) is rejected when the parameter ρ is significantly *less* than unity. Note that such a test is designed to err on the safe side; the point estimate of ρ might be considerably less than unity, but if the estimate is not *significantly* below unity, the time series is deemed dangerously close to having a unit root. It would then be safer to use first differences, after they too are checked for a unit root.

Exhibit 2 reports Dickey-Fuller t-test results of H_0: $\rho \geq 1$, for levels and first-differences of credit spreads.[5] The critical 90% value is −2.57. A smaller t-statistic rejects the null hypothesis; i.e., it is a one-sided test. For the credit spread *levels*, the BBB-rated, two-year Yankee is the sole series out of sixty that rejects the null hypothesis of a unit root at the 90% level.[6] A few other series are close, and one might very well surmise that they probably do not have unit roots. Indeed, it seems implausible that *any* yield spread would actually explode, as a unit root process could. The power of unit root tests is notoriously low against *near* unit root alternatives; nonetheless, first-differencing seems prudent for all these series.

The first-differences in Exhibit 2 display an entirely opposite pattern. Every series rejects the unit root hypothesis at the 90% level, most by a substantial margin. Even at the 99% level, whose critical value is −3.43, all but one series reject. First-differences appear to be safe.

As an aside, the number of lags (p) required in the Dickey-Fuller first-difference regressions vary widely across the series. Many series are strongly autocorrelated; the p's range from a minimum of one to a maximum of nineteen over the sixty series.

We also investigate whether the sixty credit spread series are cointegrated. The test procedure is to run a regression of one credit spread (level) on the fifty-nine others, and thereafter examine the regression residuals for the presence of a unit root. If the residuals still display a unit root, the sixty series are *not* cointegrated. Basically, this implies that non-stationarity within an individual series cannot be explained by the non-stationarity in other series.

In the cointegrating regression with the AA two-year Industrials selected arbitrarily as the regressand, the R^2 is 0.897, and the Durbin-Watson is 1.349. This already indicates probable cointegration. A formal Dickey-Fuller test on the regression residuals gives a t-test statistic of −6.31, which, compared against the 90% critical value of −4.70 and the 99% critical value of −5.52, is indeed strong evidence of cointegration. Since the residuals do not have a unit root, the sixty credit spreads share their non-stationarity in common.

The economic intuition for this result points to *common* option-theoretic characteristics of credit spreads. Even across industries, maturities, and rating categories, there must be underlying *systematic* state variables driving default option values. Indeed, we have already mentioned one such common variable, the (default-free) interest rate, whose movements change the present values of corporate liability promises. Other possible common factors include consensus perceptions about volatility and asset values.

Apparently, *all* of the non-stationarity in credit spread indexes can be attributed to these underlying common influences.

Exhibit 1

Credit Spread Levels

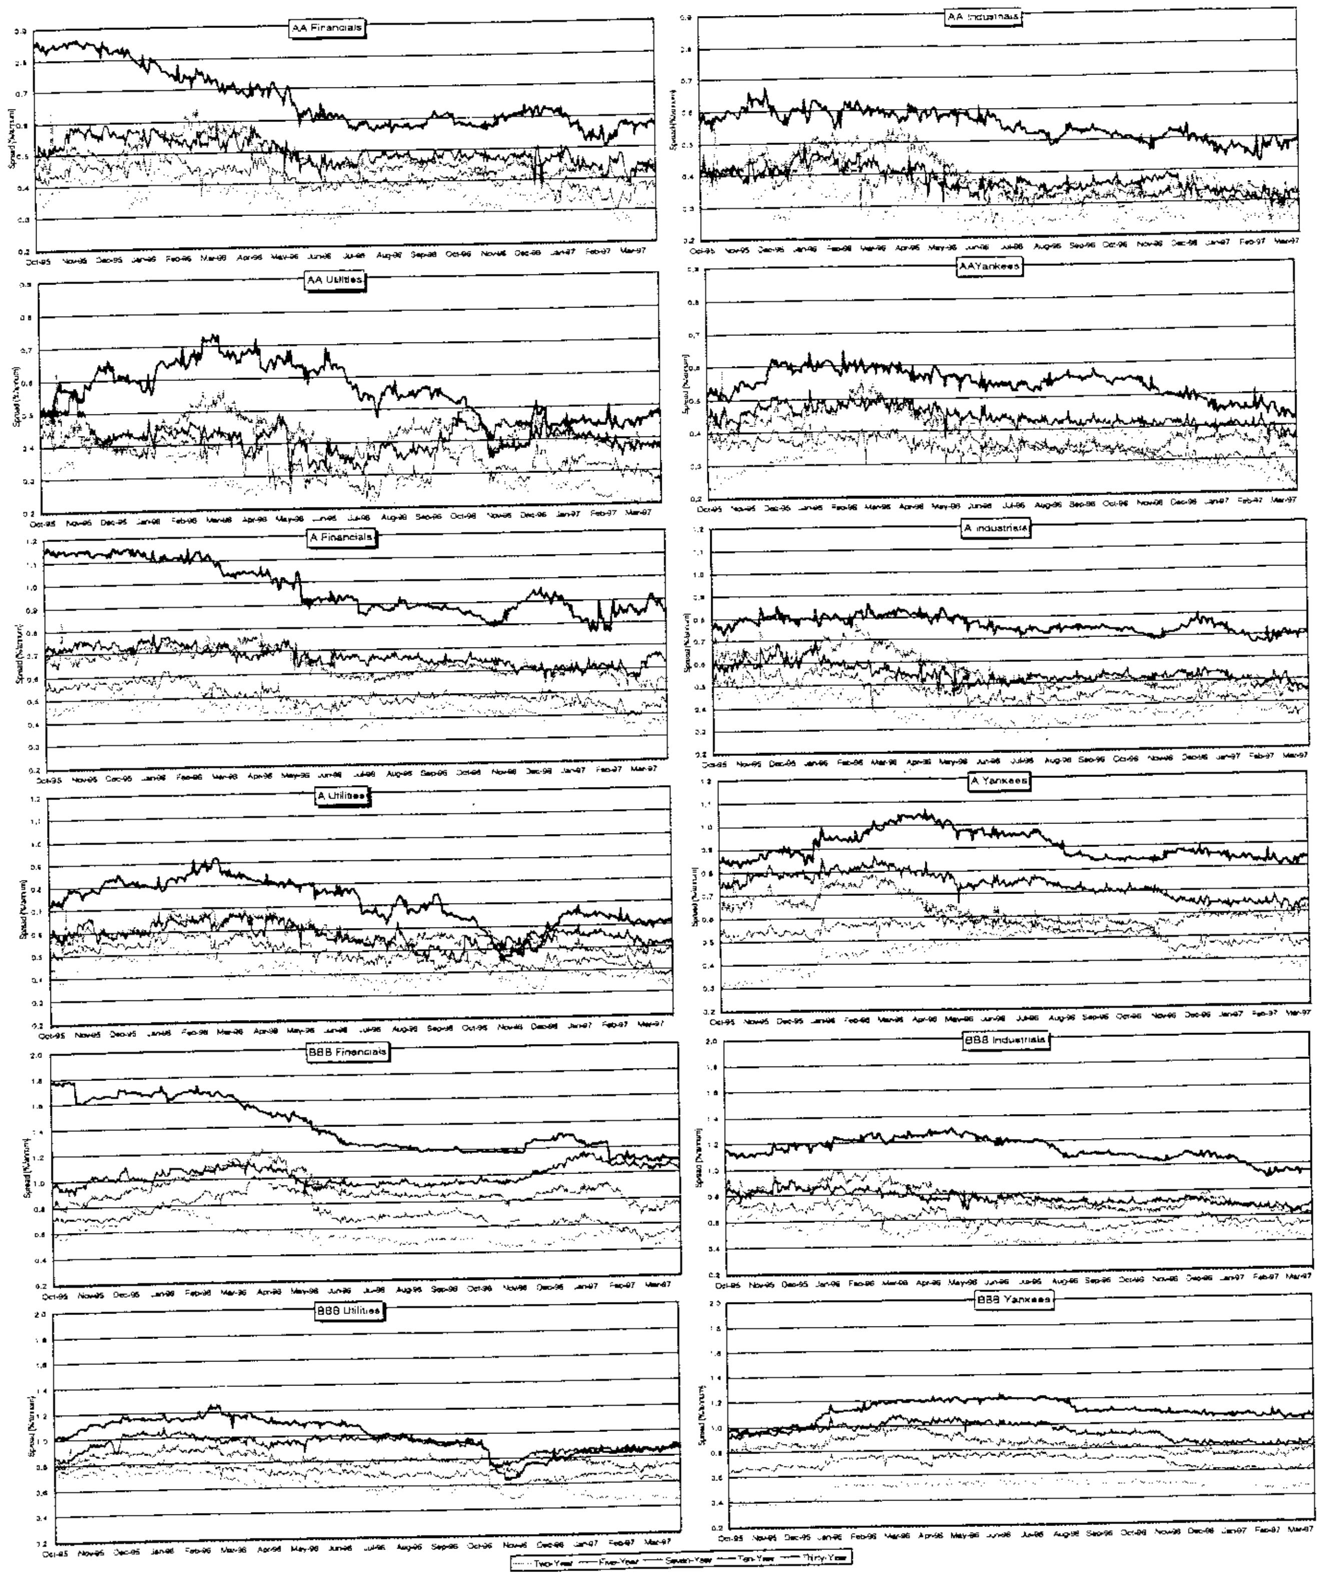

EXHIBIT 2

Unit Root Tests for Credit Spread Indexes ■ Levels and First-Differences

	Critical Value (90%) = −2.570									
Maturity (years)	2	5	7	10	30	2	5	7	10	30
	"AA — Level"					"AA — First-Difference"				
Industrials	−2.554	−1.778	−0.984	−0.928	−0.836	−4.977	−5.660	−4.853	−5.965	−6.206
Utilities	−1.999	−1.687	−2.198	−2.438	−1.104	−4.513	−5.590	−4.578	−5.607	−4.778
Financials	−2.150	−1.994	−1.159	−0.958	−1.486	−5.204	−6.086	−5.109	−6.460	−4.421
Yankees	−0.290	1.081	−0.770	0.085	−0.134	−6.500	−7.034	−7.558	−8.484	−4.609
	"A — Level""A — First-Difference"									
Industrials	−2.056	−0.961	−0.630	−0.584	−1.621	−4.566	−5.703	−7.957	−6.431	−5.382
Utilities	−1.442	−0.366	−1.084	−1.311	−1.431	−4.776	−5.943	−5.430	−6.491	−4.310
Financials	−1.011	−1.057	0.050	−1.604	−1.434	−4.878	−4.416	−6.788	−5.579	−3.785
Yankees	−2.199	−1.174	−1.279	−0.224	−1.404	−6.106	−5.372	−4.437	−5.205	−3.721
	"BBB — Level""BBB — First-Difference"									
Industrials	−2.451	−1.845	−0.403	−0.495	0.129	−5.321	−5.439	−4.954	−7.097	−5.014
Utilities	−0.539	−1.095	−1.381	−1.984	−1.250	−5.480	−5.108	−5.323	−4.433	−5.299
Financials	−1.122	−1.296	−1.261	−1.598	−0.413	−4.507	−3.112	−3.584	−5.058	−4.101
Yankees	−2.711	−1.080	−0.844	−0.420	−2.042	−4.193	−6.881	−4.664	−6.430	−3.773

Stationarity is an important topic for further investigation because many leading models of credit derivatives assume implicitly that spreads follow stationary processes; e.g., see Das and Tufano [1996] or Jarrow, Lando, and Turnbull [1997].

III. CHARACTERISTICS OF CREDIT SPREAD CHANGES

Continuously compounded percentage changes in spreads between successive trading days were computed as first-differences in the logs of spread levels. Exhibit 3 presents summary statistics for each spread change during this sample period. Exhibits 4-6 plot estimated volatility, skewness, and kurtosis.

Volatility displays a pronounced and consistent pattern across the four industry groups. Higher ratings and shorter maturities have greater volatilities. Although spreads are always higher for lower ratings, their percentage changes are typically smaller in magnitude. Similarly, spreads are usually (but not invariably) higher for longer maturities, but their percentage changes have smaller amplitudes. Yield spread percentage change volatility is quite large compared to, say, return volatility. Most of the volatilities in Exhibit 3 are well above 2% *per day*, and the shortest AA series exceed 5% per day; (approximately 80% per year.)

Spread changes are correlated. *All* the estimated 1,770 pairwise correlation coefficients among the sixty time series are positive. The mean is 0.581, and the minimum and maximum are 0.281 and 0.932, respectively. Correlations are higher for similar maturities, but they are roughly the same within and across industries and credit ratings.[7]

Evidence Against the Gaussian Distribution

If yield spread changes were normally distributed, the kurtosis statistics in Exhibit 3 would not differ materially from zero. There would be some sampling scatter, of course, but about as many negative as positive values. As Exhibit 6 shows, however, *every one* of the yield spread changes has excess kurtosis. All the estimates are positive, and many are very large. This implies a particular type of departure from the normal distribution that can have troubling implications for risk management.[8]

Excess kurtosis is caused by "thick" tail areas in the unconditional probability distribution. There are too many extreme observations relative to what would have occurred if the data were normally distributed. But such

EXHIBIT 3

Yield Spread Changes, (%/Day), Univariate Statistics ■ October 5, 1995-March 26, 1997

Rating	Maturity (years)	Mean	Standard Deviation	Excess Skewness	Kurtosis
			Industrials		
AA	2	-0.1401	8.575	-0.357	1.621
AA	5	-0.0981	6.835	-0.288	3.329
AA	7	-0.0668	6.094	0.101	1.993
AA	10	-0.0566	5.740	0.292	3.474
AA	30	-0.0367	3.519	-0.141	1.898
A	2	-0.1203	6.261	-0.333	2.100
A	5	-0.0814	5.046	-0.115	4.251
A	7	-0.0937	4.428	0.249	3.418
A	10	-0.0603	4.274	-0.064	3.770
A	30	-0.0163	2.597	-0.047	1.753
BBB	2	-0.0689	4.580	-0.464	1.877
BBB	5	-0.0671	3.785	-0.289	3.601
BBB	7	-0.0933	3.235	0.097	2.432
BBB	10	-0.0510	3.063	-0.002	2.938
BBB	30	-0.0537	1.926	0.061	1.504
			Utilities		
AA	2	-0.1062	7.403	0.178	2.256
AA	5	-0.1126	6.928	-0.336	2.984
AA	7	-0.0630	5.726	0.348	4.792
AA	10	-0.0741	5.354	0.111	4.252
AA	30	-0.0196	3.845	-0.196	1.265
A	2	-0.0695	5.096	-0.027	1.160
A	5	-0.0687	4.339	-0.219	0.827
A	7	-0.0452	4.246	0.044	6.215
A	10	-0.0328	3.884	0.255	2.020
A	30	-0.0442	2.888	0.055	2.154
BBB	2	-0.1212	3.162	-0.127	0.448
BBB	5	-0.0562	2.820	-0.420	2.446
BBB	7	-0.0299	2.840	0.072	8.499
BBB	10	0.0046	2.400	-1.332	12.120
BBB	30	-0.0440	2.124	-0.515	4.601
			Financials		
AA	2	-0.0283	5.689	-0.183	0.687
AA	5	-0.0330	4.965	-0.044	2.353
AA	7	-0.0308	4.841	0.526	6.427
AA	10	-0.0342	4.546	0.358	4.027
AA	30	-0.1089	2.751	-0.205	1.174
A	2	-0.0397	4.999	-0.012	1.154
A	5	-0.0666	4.112	0.171	2.064
A	7	-0.0703	3.730	0.352	4.477
A	10	-0.0233	3.113	-0.116	2.315
A	30	-0.0707	2.319	0.159	4.681
BBB	2	-0.0329	3.548	-0.304	1.277
BBB	5	-0.0738	3.122	-0.558	3.179
BBB	7	-0.0248	2.751	0.410	5.483
BBB	10	0.0290	2.098	0.186	2.587
BBB	30	-0.1284	1.618	-1.843	12.280

extremes are precisely the object of most interest in many risk systems such as in VaR calculations. Risk monitoring often attempts to predict large losses, extreme negative outcomes. Ignoring the excess kurtosis of yield spread changes can understate the true probability of a large loss, perhaps by an order of magnitude.

To illustrate this effect, consider a typical situation that induces excess kurtosis, a mixture of two normal distributions. Intuitively, we might imagine that credit spreads fluctuate with only a modest amplitude during "quiet" epochs while they move dramatically during "exciting" times. For simplicity of illustration, assume that spread changes are Gaussian, but that the standard deviation jumps randomly from one regime to another.

Exhibit 7 plots the excess kurtosis and standard deviations from such a mixture of two normal distributions, one distribution with a standard deviation of 1.0 and the second distribution with a standard deviation of 3.0. The excess kurtosis and standard deviation of the mixture is plotted against the probability that the lower standard deviation regime will occur on a given date.

As the graph makes evident, excess kurtosis is produced at all mixing probabilities, but it is maximized when the probability of the higher standard deviation is relatively low. In practical terms, if a "usual" regime with low volatility is interrupted *infrequently* by an "eventful" regime with high volatility, one will observe significant excess kurtosis.

Implications for Risk Management

Calculations of loss probabilities, such as VaR, will be seriously compromised by a risk manager who relies inappropriately on an assumption of normality and ignores the possibility of shifting regimes. Let's take a typical example from Exhibit 7, say, a probability of the low volatility regime of 0.85; this corresponds to an unconditional standard deviation of about 1.48 and an excess kurtosis of 5.06. The 1% VaR of a normal distribution with volatility 1.48 is -3.45 below the mean of the distribution. But the 1% VaR of an 85%/15% mixture of low and high volatility regimes, *with*

EXHIBIT 3 (CONTINUED)

Yield Spread Changes, (%/Day), Univariate Statistics
■ October 5, 1995-March 26, 1997

Rating	Maturity (years)	Mean	Standard Deviation	Excess Skewness	Kurtosis
			Yankees		
AA	2	-0.0184	5.223	-0.012	0.628
AA	5	-0.1278	5.528	-0.197	2.943
AA	7	-0.0662	5.460	0.370	6.402
AA	10	-0.0498	4.130	0.201	2.958
AA	30	-0.0484	3.133	-0.054	1.854
A	2	0.0481	4.154	0.108	1.321
A	5	-0.0361	3.468	0.031	2.817
A	7	-0.0103	3.549	0.158	7.183
A	10	-0.0339	2.510	0.087	2.522
A	30	-0.0012	1.968	0.274	2.607
BBB	2	0.0730	3.594	0.162	1.433
BBB	5	0.0016	2.749	0.370	3.273
BBB	7	-0.0001	2.721	0.553	11.940
BBB	10	-0.0224	2.004	0.209	2.949
BBB	30	0.0163	1.668	0.035	2.527

the same volatility, 1.48, has a VaR –4.50 below the mean.

To put this into practical terms, if a VaR of $100 million had been calculated with the present internal model under these conditions, it would have been about 30% too small. The true 1% VaR would have been about $130 million.

The VaR error induced by non-normal thick tails becomes worse, percentagewise, with a lower cutoff level; e.g., the error for a 1% VaR is generally worse than the error for a 5% VaR. Also, the error generally increases with kurtosis, but the effect is complex; it depends on the number of different regimes generating the data, the probabilities of each regime, and the persistence of a regime.

On a more reassuring note, the effect of excess kurtosis is likely to be somewhat attenuated by diversification. A relatively large portfolio will usually have less kurtosis than many of its individual constituents. This diversification benefit will, however, be less marked when individual returns are highly correlated. Since the correlations among credit spread changes are relatively large, particularly for similar maturities, one should expect that a portfolio of corporate bonds will retain a troubling degree of excess kurtosis.

Turning finally to skewness, Exhibits 3 and 5 reveal that slightly more than half of the sixty yield spread changes have positive skewness. Since price changes have the opposite skewness of yield changes, positive yield skewness implies that the left tail of the loss distribution (for a long position) contains more probability than a normal distribution. Again, a mindless calculation of VaR will be understated. On the other hand, for those cases of negative yield spread skewness, VaR can actually be overstated.

The degree of skewness varies by industry; twelve of the fifteen Yankee spreads are positively skewed, but only five of the fifteen industrial spreads. There is a slightly greater tendency toward positive skewness in the longer maturities.

IV. MODELING CREDIT SPREAD CHANGES WITH RANDOMIZED GAUSSIAN MIXTURES

Thick tails occur under many probability distributions, but there is a strong a priori reason to think that a Gaussian law *ought* to be a good model for credit spread changes, at least over a short horizon. The reason stems from the Gaussian's position as the limiting distribution for sums of independent random increments with finite variance.[10] Since the actions of many market participants are reflected in bond prices and hence in credit spreads, some more or less bell-shaped limiting distribution is likely to occur.

If, in addition, markets undergo periodic changes in turbulence, the simplest and most straightforward model would be a mixture of several Gaussian laws. We undertake an excursion into this terrain by estimating the likely number of distinct Gaussian laws, their respective probabilities of occurrence, and their associated parameters for the sixty credit spread change indexes in our sample. The technique assumes that each member of the mixture has a fixed probability of occurrence every day. It does not attempt to ascertain whether regimes persist once they occur. That topic is reserved for the next section.

Denote by g(s) the unconditional density function for s, the log first-difference in the yield spread.[11] The simple mixtures model can be expressed as

$$g(s) = \sum_{k=1}^{K} \pi_k f_k(s)$$

EXHIBIT 4

Volatility of Yield Spread Changes (% Day) ■ October 5, 1995-March 26, 1997

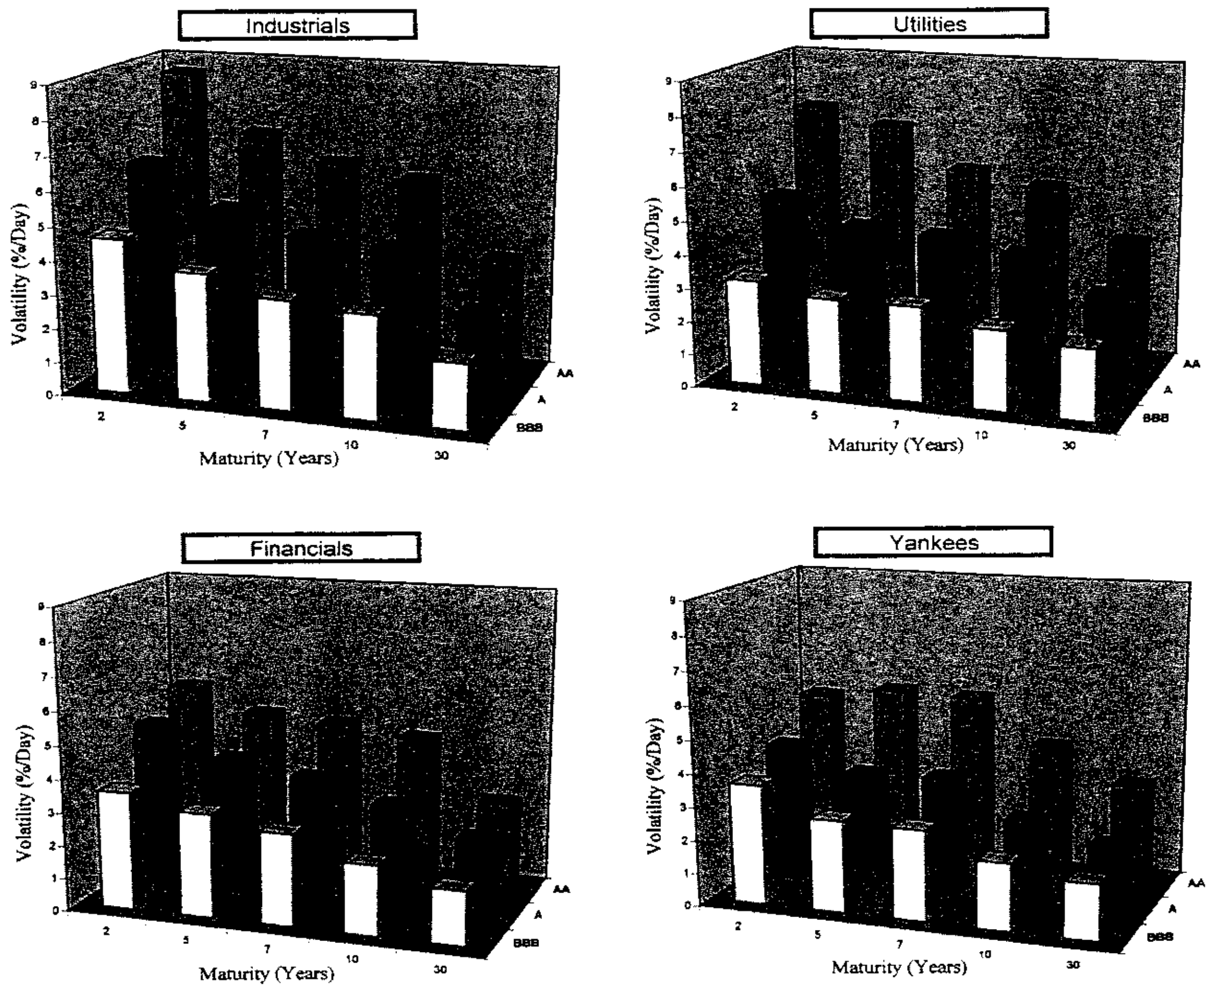

where K is the number of distinct Gaussian elements of the mixture, $f_k(s)$ is the k-th Gaussian density, with mean μ_k and standard deviation σ_k, and π_k is the probability of a regime whose density is $f_k(s)$, $\Sigma_k \pi_k = 1$.

Let s_t denote a sample observation of the yield spread change on date t, a single observation drawn from g(s). Conditional on K, the likelihood function for a sample of size T is

$$L_k(s_1,\ldots,s_T) = \prod_{t=1}^{T} \left\{ \sum_{k=1}^{K} \pi_k f_k(s_t) \right\}$$

which can be maximized with respect to the 3K − 1 parameters, μ_k, σ_k, and π_k, k = 1, …,K. A conditional likelihood value can be calculated for K = 1, 2, …, until a K is found for which there is no evidence that K + 1 improves the overall fit.

The decision rule examines the sequence of likelihood ratios, L_{K+1}/L_K, and terminates when the test statistic is no longer significant at a prespecified level (which we chose a priori as 95%). The critical value (7.81) does not vary with K because the reduction in degrees of freedom equals the number of additional restrictions imposed.

The likelihood function is non-linear in the parameters. A simultaneous analytic solution for all parameters becomes intractable (or at least very tedious,

EXHIBIT 5

Skewness of Yield Spread Changes ■ October 5, 1995-March 26, 1997

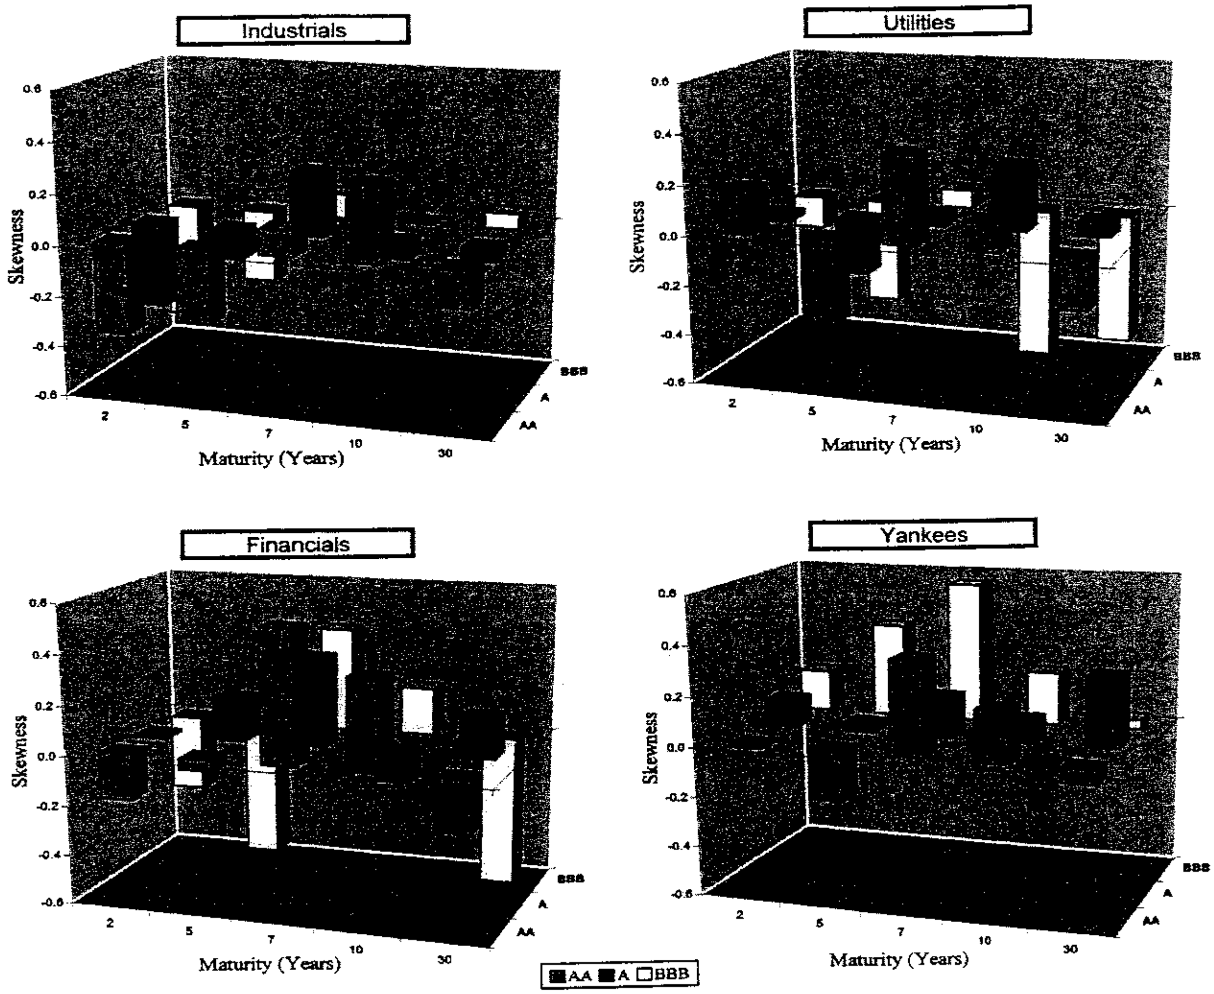

for K > 1. Consequently, we use a numerical algorithm for solving a system of non-linear equations.[12]

The results are presented in a series of figures organized to reveal interesting features of the data.[13] To begin with, Exhibit 8 reports the log likelihood ratio statistic for whether a single Gaussian distribution is adequate, as opposed to a mixture of at least two Gaussians. The critical 95% level is 7.81, and the adequacy of a single distribution is rejected for fifty-nine of the sixty time series. The only exception is for the BBB-rated, two-year Utilities, the series with the lowest excess kurtosis (see Exhibit 3.)

Exhibit 9 plots the sixty likelihood ratio test statistics for a two-distribution mixture against the computed kurtoses. The cross-sectional relation is apparent, but it is certainly not perfect and appears to be slightly non-linear.

Exhibit 10 plots the estimated probabilities of occurrence and standard deviations (percent/day) for the mixture of two normals (K = 2). There is a consistent pattern: a relatively high probability of a low-volatility regime (i.e., the quiet regime), and a smaller probability of a high-volatility regime (i.e., the exciting regime). In the low-volatility regime, estimated standard deviations decrease with maturity and increase with credit quality.

A similar pattern can be discerned in the Industrials sector during the high-volatility regime, but the other three sectors have an opaque pattern, if any at all. Also, one might be tempted to think that Utilities are less likely to experience high-volatility regimes, but

EXHIBIT 6

Excess Kurtosis of Yield Spread Changes ■ October 5, 1995-March 26, 1997

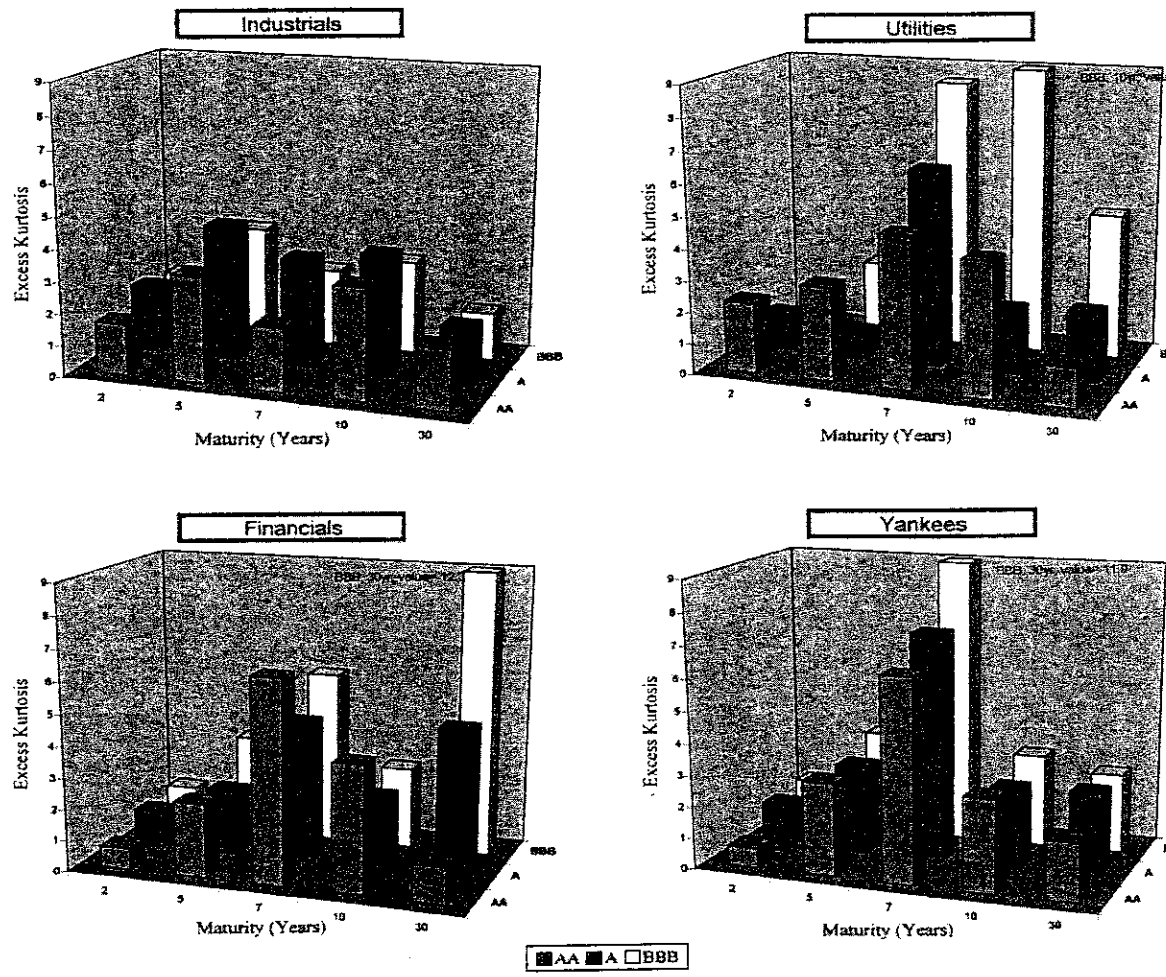

there are a few maturity/credit quality exceptions to this tendency.

A pictorial example of the technique is presented in Exhibit 11. Panel A plots the two separate elements of the Gaussian mixture estimated for the A-rated seven-year Financials along with the overall mixed density. Panel B compares the estimated mixture with an ordinary Gaussian having the same mean and variance and with a non-parametric density estimated from the sample.[14]

The non-parametric density estimate might be regarded as a curve fit through the sample observations, so differences with either the ordinary Gaussian or the mixed Gaussian reveal areas where these parametric estimates fail to fit the data perfectly. Clearly, the mixture provides a substantially better overall fit.

Panel C shows that the mixture has much more probability in the tails than the ordinary Gaussian, as would be expected. The mixture deviates from the non-parametric density tail areas primarily because the latter are multimodal. We think this non-monotonicity in the non-parametric density is evidence of sampling error that is prone to be exaggerated in the tail areas. Consequently, the mixture estimator might very well provide more reliable tail probability calculations and a better tool for risk managers.

For mixtures of three Gaussian laws, Exhibit 12 plots the likelihood ratio test statistics, and Exhibit 1

EXHIBIT 7
Excess Kurtosis and Volatility for Mixtures of Two Distributions

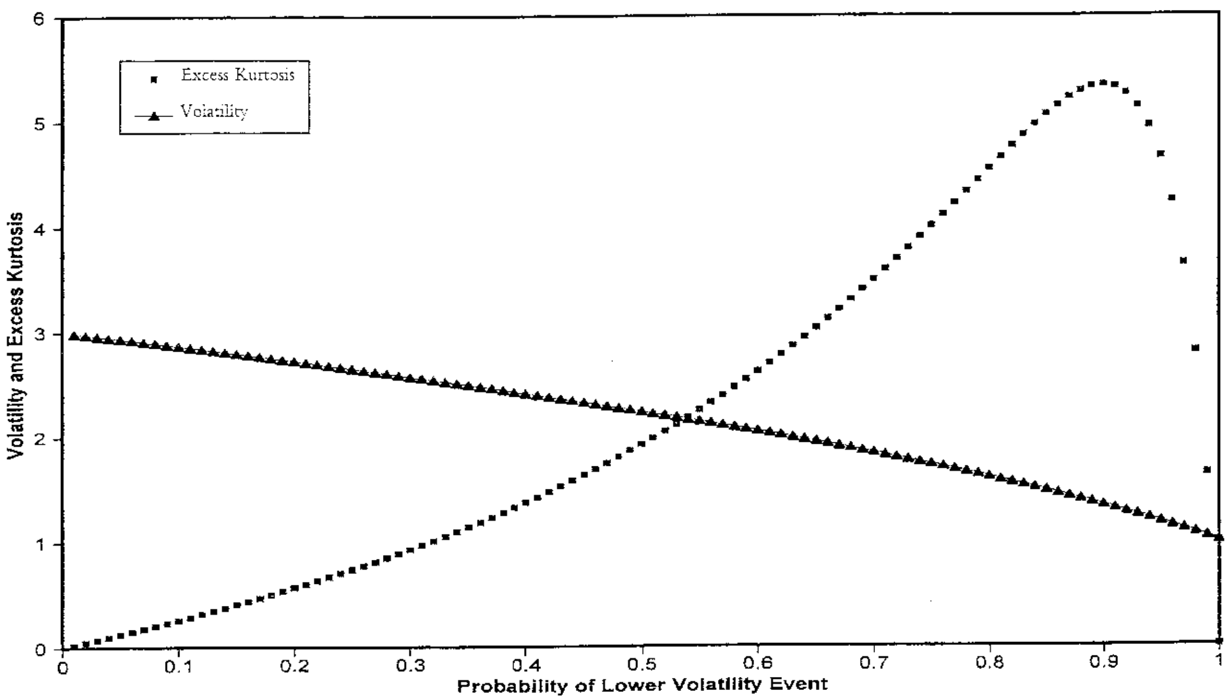

plots estimated mixture probabilities and standard deviations.[15] The pattern of the lowest volatility regime appears to be almost identical to that reported in Exhibit 10, which assumes just two components in the mixture. The higher-volatility regime of the two-component mixture has now been split further into two parts. The first part, labeled an "average" regime in Exhibit 13, has an estimated volatility somewhat greater than the "quiet" regime.

The third component comes in two distinct forms. For many series, it is a regime of minuscule volatility, undoubtedly associated with trading days on which credit spreads did not change at all. The associated probability of this mixture component is essentially the frequency of zero change observations. There are some series, however, with very large estimated volatilities, very "exciting" epochs indeed.

Exhibit 14 reports the likelihood ratio test statistics for the necessity of a four-distribution mixture, as compared to a mixture of at least three Gaussians. Missing bars indicate cases for which a three-element mixture is not necessary. At the 95% critical level, four components are required for only two of the sixty time series: the A-rated, five-year Industrials and the A-rated ten-year Financials.

V. PERSISTENCE OF REGIME SHIFTS

Although the previous results suggest that credit spreads display intertemporal changes in volatility, they provide no direct evidence about the validity of our working assumption that regime shifts occur randomly on a daily basis. Indeed, we regard this assumption as a priori implausible, and made it merely for convenience in the previous section.

The polar opposite of a daily regime switch is a slow evolution of volatility over time, as exemplified by a GARCH process.[16] GARCH provides a conditional prediction of volatility at date t based on the specification:

$$\sigma_t^2 = \alpha_0 + \sum_{j=1}^{p} \alpha_j s_{t-j}^{\ 2} + \sum_{j=1}^{q} \varphi_j \sigma_{t-j}^{\ 2}$$

where s_t is the observed credit spread (log) change on

EXHIBIT 8

Likelihood Ratio Test Values for a Mixture of Two Normals

Critical Value (5%) = 7.81

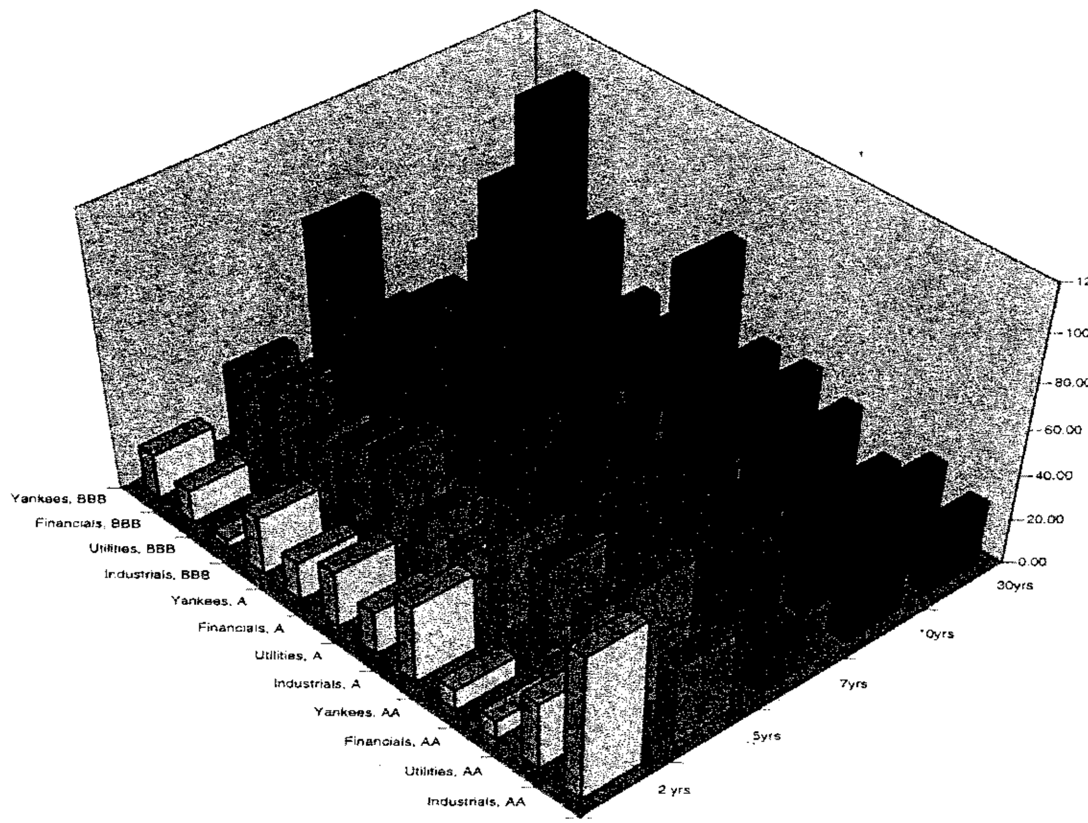

Type/Rating	2 yrs	5 yrs	Maturity 7 yrs	10 yrs	30 yrs
Industrials, AA	64.82	36.16	16.00	42.78	28.52
Utilities, AA	30.35	56.63	57.18	49.22	34.94
Financials, AA	8.50	33.57	50.09	50.28	23.69
Yankees, AA	9.45	42.66	49.37	25.17	37.45
Industrials, A	34.63	41.25	40.85	52.28	44.16
Utilities, A	19.17	24.95	49.76	19.80	43.23
Financials, A	25.20	26.57	47.00	24.89	78.47
Yankees, A	17.24	44.67	59.49	22.54	36.36
Industrials, BBB	25.66	35.80	34.28	42.04	35.62
Utilities, BBB	3.67	29.35	64.53	74.21	59.14
Financials, BBB	14.60	39.37	49.87	27.75	108.70
Yankees, BBB	19.75	37.16	82.38	29.59	64.82

date t, σ_t is the estimated (conditional) volatility of spread changes for date t and the α's and φ's are parameters obtained with the method of maximum likelihood. In general, the orders p and q of the process can be varied to obtain the best fit; a parsimonic exploratory initial choice is $p = q = 1$. Results for this

EXHIBIT 9
Likelihood Ratio Test Statistics versus Kurtosis

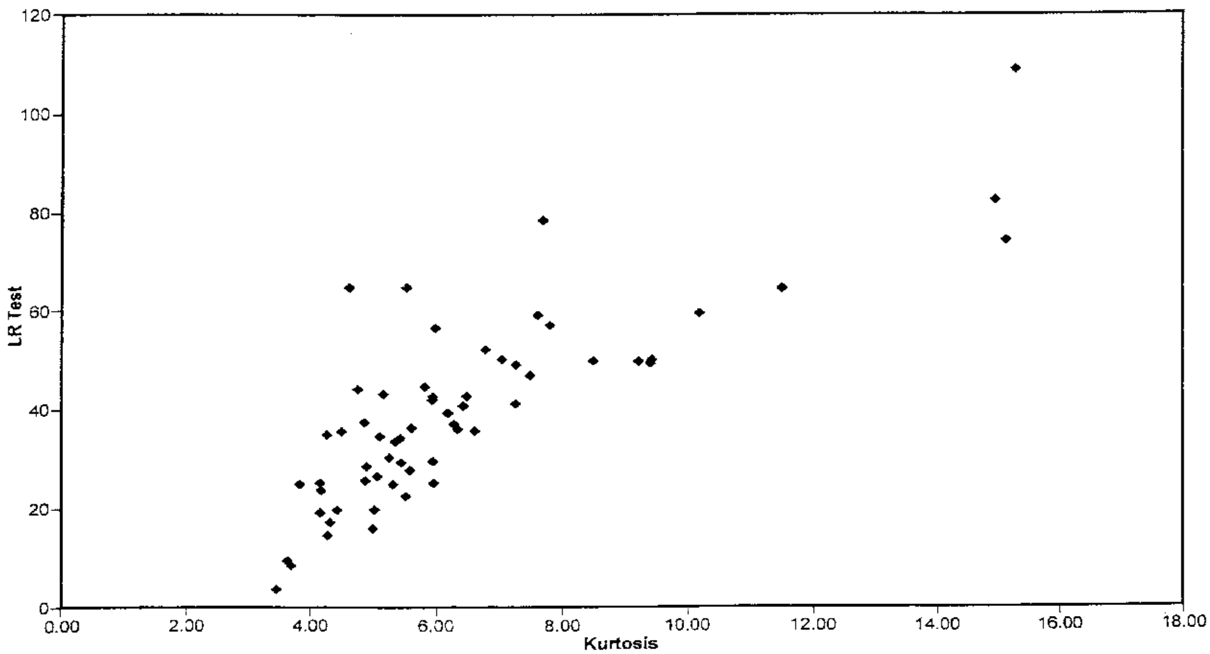

Mixture of Two Normals: Probabilities

specification are reported in Exhibit 15.

The intercept term, α_0, is always positive and significant as it should be. The "ARCH" coefficient α_1 on the lagged squared observation is significantly positive for fifty-nine of the sixty credit spread series; the only exception is BBB thirty-year Financials. This implies persistence in volatility at a single day's lag.

The coefficient of the lagged conditional variance, φ_1, is not statistically significant for the majority of the Industrial, Utility, and Yankee series. In contrast, most of the Financial series are significant, including *all* of the BBBs. In the utilities group, the AA and A two-year series have significant φ_1, but the A and BBB seven-year series are aberrations, significantly negative. Only one of the Yankee series displays a highly significant φ_1.

Since the ARCH coefficient α_1 is so uniformly significant, we also compute a GARCH model with p = 2 and q = 1 to test for second-order lagged ARCH effects (not reported, but available on written request.) The first-order coefficients remain virtually unchanged while the new second-order coefficient, α_2, is insignificant for most series. Only four of the sixty series have associated t-statistics for this coefficient greater than 2.0, while seven t-statistics are less than –2.0. Apparently, the volatility persistence in these series is quite simple and short-term. There is little evidence that it persists for more than a single trading day.

Since the volatility of credit spread changes seems to be driven mainly by the *observed* squared credit spread change on the previous trading day, the basic assumption in our Gaussian mixtures models appears to be appropriate for most credit spreads. Every day, the volatility of the regime is drawn anew (although it is at least partly determined by the ex post observation on the previous day).

For many series, there is no evidence of long-term persistence in the volatility; the lagged conditional volatility as measured by the GARCH coefficient φ_1 is not a significant determinant of the subsequent day's conditional volatility. There are a few series, however, mainly the Financials and particularly the lower-rated Financials, where persistence is significant. For these series, a simple mixtures process could be improved upon.

VI. SUMMARY AND CONCLUSIONS

Most large financial institutions have a plethora of

EXHIBIT 10

Mixture of Two Normals: Standard Deviation

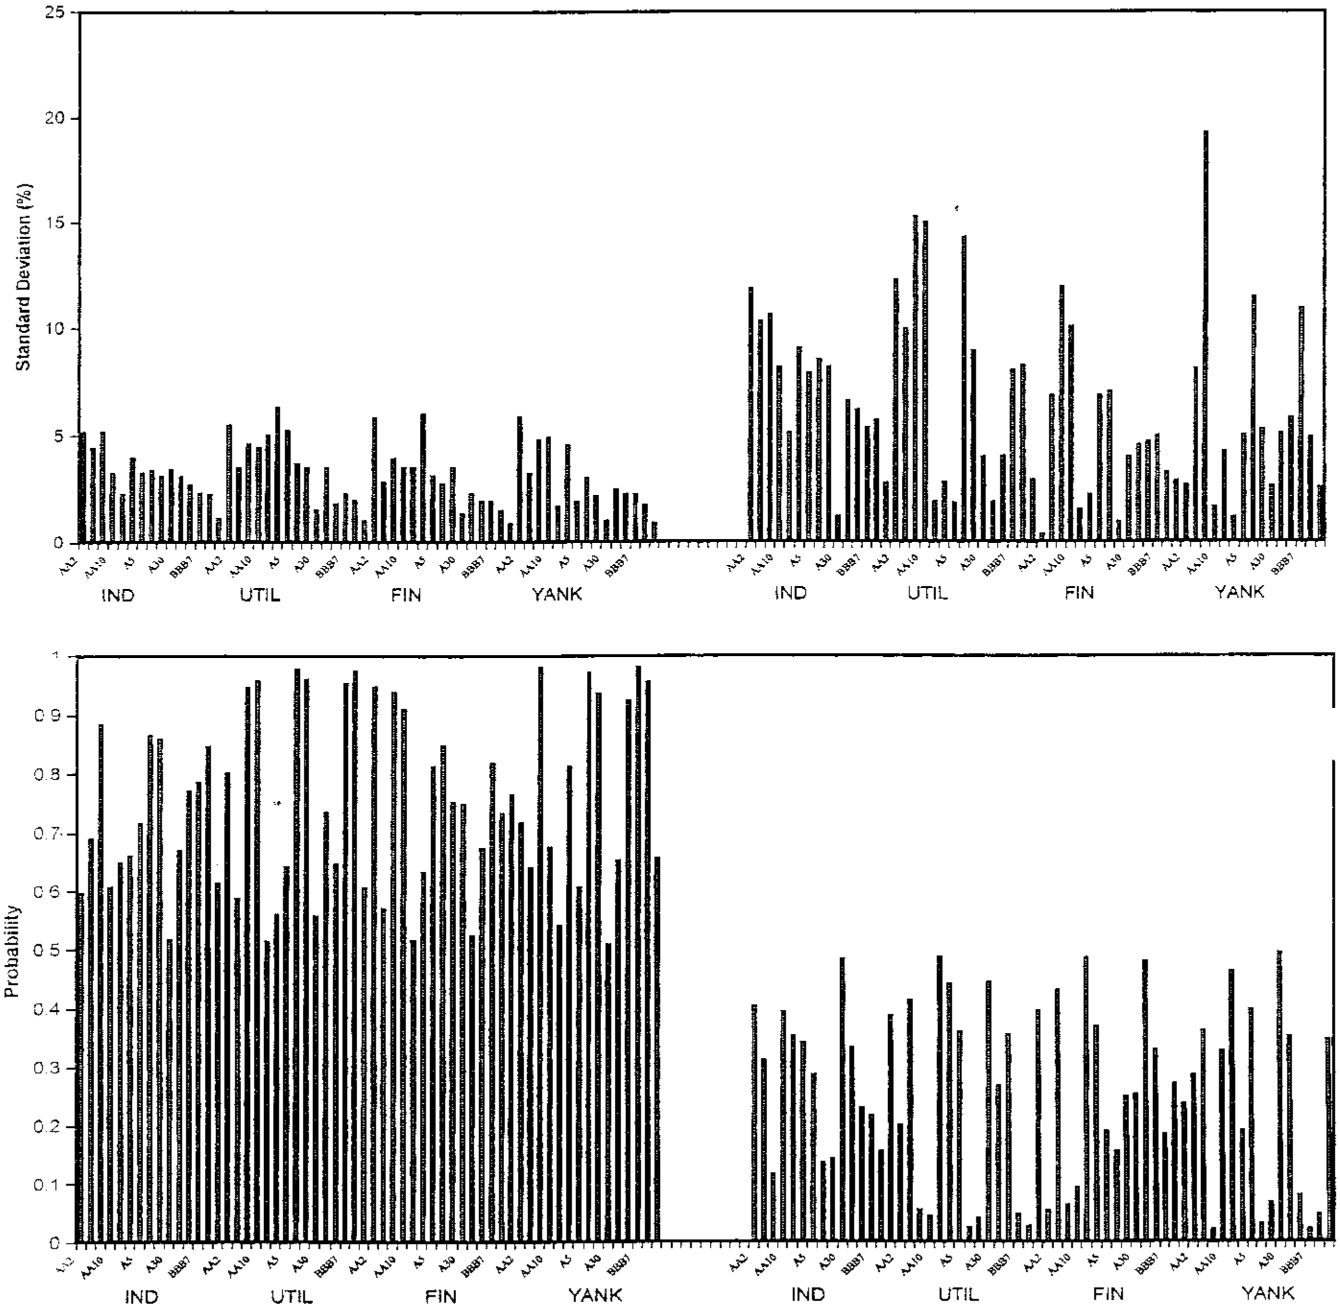

hedging instruments available for reducing the market risk of positions taken during the normal course of business. Typically, they employ hedges whenever possible while they diversify to reduce non-hedgeable risks. But there are some systematic (non-diversifiable) risks for which hedging instruments are illiquid or not even available. Among the most prominent of these are corporate bond credit yield spreads.

Corporate bond traders and investors can use interest rate derivatives to hedge part of the position risk in bond portfolios. Changes in individual company credit quality can be diversified. Unfortunately, average credit spreads still fluctuate dramatically, apparently driven common underlying influences on the default options

EXHIBIT 11

Estimated Mixture of Two Gaussian Densities for A-Rated 7-Year Financial Credit Spread Changes Compared with Non-Parametric Density Estimator and Gaussian Density with Same Mean and Standard Deviation

PANEL A: Density of Estimated Mixture and its Gaussian Components

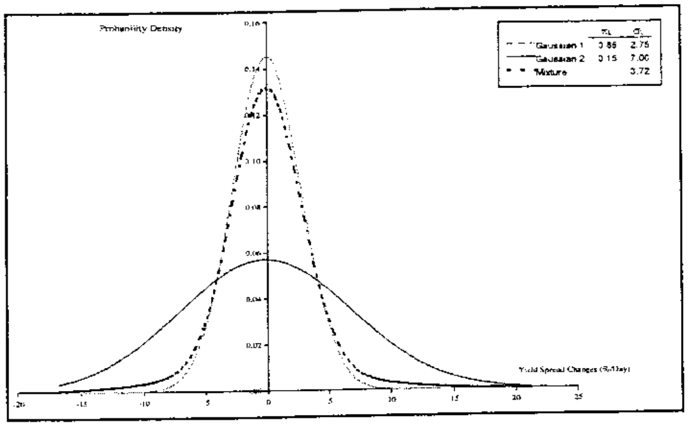

PANEL B: Comparison of Estimated Mixture with Gaussian Density having the same Mean and Standard Deviation and a Non-Parametric Density

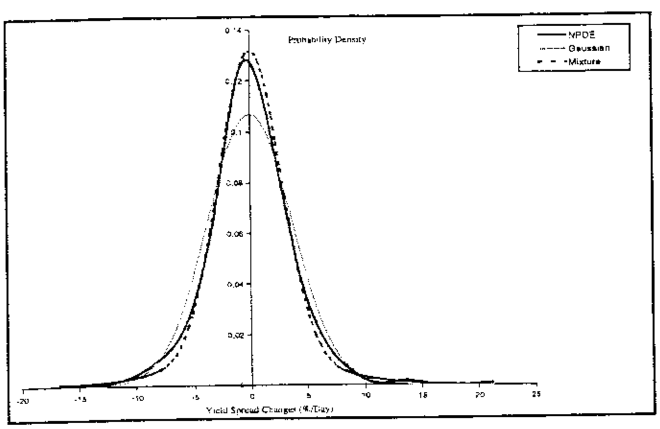

PANEL C: Right Tail of Estimated Mixture of Two Densities, Gaussian Density, and Estimated Non-Parametric Density

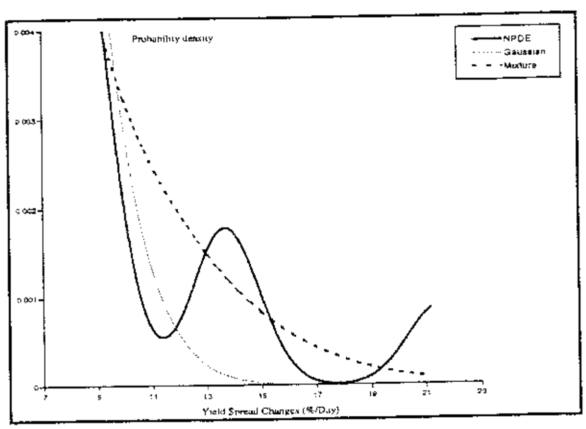

embedded in corporate bonds. This risk is not diversifiable, and there are few extant hedging instruments.

Using a sample of dollar-denominated average credit spreads categorized by industry, maturity, and rating, we have recorded some salient features of their probability distributions. The conclusions are:

1. The hypothesis of non-stationarity for credit spread levels cannot be rejected. Most of the series seem alarmingly close to having unit roots, and hence should not be used in statistical models without the utmost caution. First differences, however, can be used with confidence.
2. Our sixty credit spread time series appear to be cointegrated, suggesting that the observed non-stationarity is attributable to common underlying influences.
3. Changes in credit spreads exhibit substantial departures from the Gaussian distribution. Most seriously, changes have thick tails, which could compromise management calculations of probable loss such as value at risk, (VaR.) If VaR assumes normality, loss probabilities can be seriously understated.
4. Gaussian mixtures appear to provide reasonably good models for the thick-tailed distributions of credit

EXHIBIT 12

Likelihood Ratio Test Values for a Mixture of Three Normals

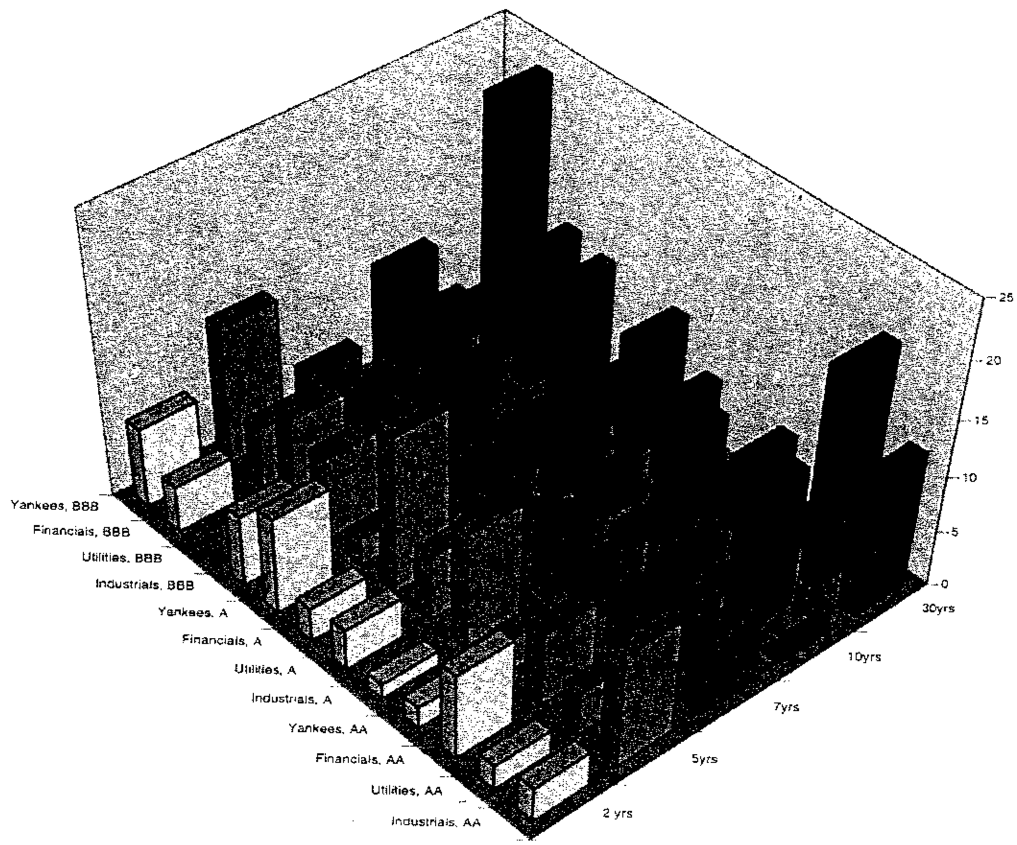

Critical Value (5%) = 7.81

Type/Rating	2 yrs	5 yrs	Maturity 7 yrs	10 yrs	30 yrs
Industrials, AA	3.12	10.42	7.05	0.48	12.19
Utilities, AA	2.25	4.03	14.40	12.62	19.29
Financials, AA	7.78	6.66	4.96	0.94	2.10
Yankees, AA	1.96	3.20	5.05	12.51	6.58
Industrials, A	1.32	10.50	6.44	1.78	1.33
Utilities, A	3.64	5.97	4.04	3.84	7.06
Financials, A	3.01	14.17	7.75	8.77	10.72
Yankees, A	8.66	2.55	7.07	4.86	4.32
Industrials, BBB	6.37	7.55	4.08	0.45	11.21
Utilities, BBB		8.62	7.46	4.85	11.93
Financials, BBB	4.18	5.54	4.19	7.91	23.25
Yankees, BBB	7.10	12.33	4.38	10.00	3.12

spread changes. We estimate mixing probabilities and parameters for each element of the mixture. Most credit spread series behave as if they were generated by two or three distinct regimes that occur randomly on a daily basis. In the two-regime model, quiet epochs are interspersed, but only rarely, with

EXHIBIT 13
Mixture of Three Normals: Standard Deviations

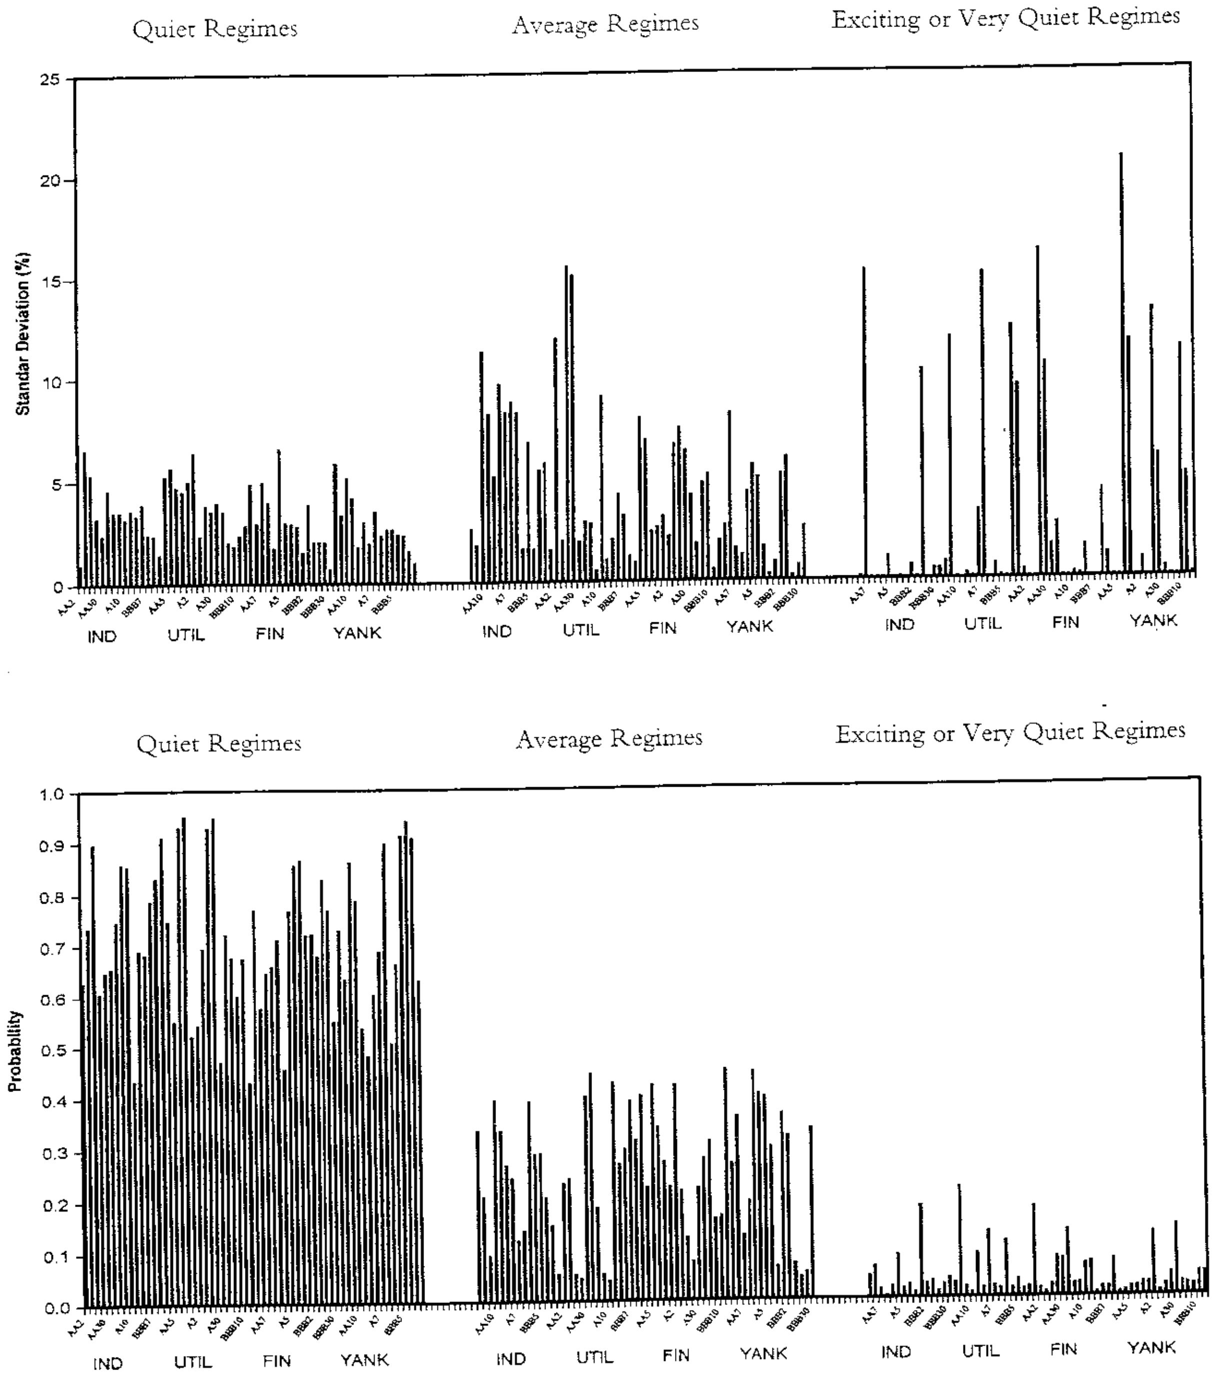

substantially more volatile periods.

5. GARCH models reveal some differences among the credit spreads across rating and industry classifications. In most cases, the persistence of volatility is rather simple; the conditional volatility on a trading day depends only on the observed squared credit spread on the preceding day. For a few series, however, there is evidence of longer-term volatility persistence.

EXHIBIT 14
Likelihood Ratio Test Values for a Mixture of Four Normals

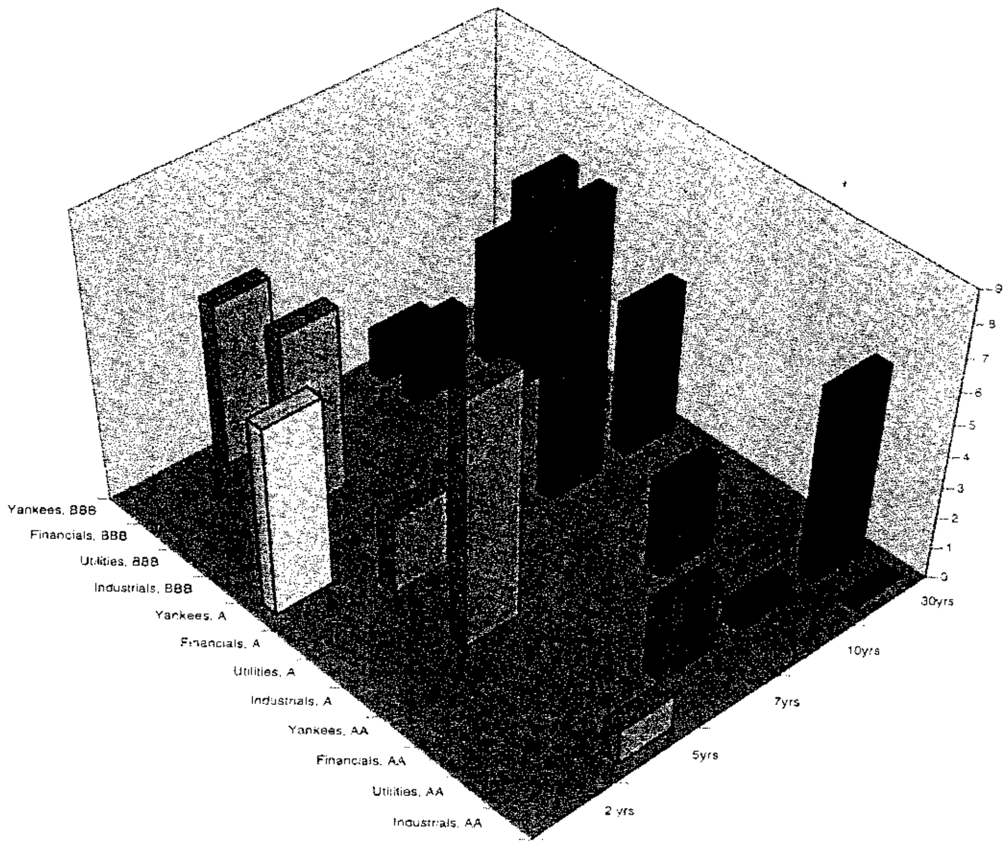

Critical Value (5%) = 7.81

Type/Rating	2 yrs	5 yrs	Maturity 7 yrs	10 yrs	30 yrs
Industrials, AA		1.14			0.09
Utilities, AA			2.46	0.71	6.07
Financials, AA					
Yankees, AA				3.04	
Industrials, A		7.87			
Utilities, A					
Financials, A		2.49		8.81	4.75
Yankees, A	6.01				
Industrials, BBB					2.23
Utilities, BBB		5.79			6.33
Financials, BBB				2.38	3.84
Yankees, BBB		5.19		1.49	

EXHIBIT 15

Persistence of Regime Shifts n Garch(1,1) Results

Rating	Maturity (years)	Industrials α0	Industrials α1	Industrials φ	Utilities α0	Utilities α1	Utilities φ	Financials α0	Financials α1	Financials φ	Yankees α0	Yankees α1	Yankees φ
Coefficients													
AA	2	14.968	0.097	0.702	16.169	0.206	0.505	11.128	0.239	0.419	11.273	0.297	0.301
AA	5	24.110	0.349	0.163	19.675	0.438	0.210	18.580	0.244	-0.017	20.579	0.299	0.038
AA	7	15.311	0.360	0.249	15.979	0.386	0.139	7.597	0.346	0.377	10.824	0.457	0.223
AA	10	19.733	0.465	-0.037	13.165	0.235	0.286	7.385	0.537	0.181	11.372	0.206	0.108
AA	30	5.343	0.382	0.214	11.244	0.273	-0.038	1.383	0.188	0.639	6.734	0.262	0.027
A	2	14.856	0.319	0.334	5.654	0.139	0.651	10.255	0.315	0.289	8.127	0.301	0.242
A	5	20.712	0.167	-0.006	12.211	0.329	0.048	10.471	0.327	0.058	10.357	0.239	-0.120
A	7	8.892	0.276	0.300	12.937	0.355	-0.063	4.599	0.346	0.374	6.260	0.500	0.039
A	10	9.005	0.275	0.216	7.857	0.138	0.339	4.224	0.324	0.222	4.349	0.207	0.077
A	30	3.768	0.323	0.134	2.873	0.201	0.465	0.396	0.263	0.684	1.719	0.240	0.332
BBB	2	13.490	0.327	0.037	4.556	0.268	0.288	3.971	0.209	0.487	5.662	0.298	0.273
BBB	5	10.025	0.213	0.080	6.374	0.240	-0.045	1.923	0.224	0.594	5.621	0.199	0.051
BBB	7	8.150	0.141	0.088	5.394	0.409	-0.053	1.721	0.188	0.605	2.389	0.495	0.269
BBB	10	4.342	0.293	0.246	1.097	0.224	0.621	1.445	0.196	0.472	2.639	0.165	0.161
BBB	30	1.615	0.192	0.370	0.845	0.110	0.703	0.526	0.007	0.793	1.894	0.300	0.037
T-Statistics													
AA	2	1.971	2.102	5.410	2.719	2.912	3.559	2.715	2.981	2.563	3.224	3.359	1.892
AA	5	4.350	3.586	1.181	4.644	4.156	1.927	4.465	2.974	-0.097	4.686	3.264	0.256
AA	7	3.729	3.808	1.854	4.572	3.726	1.128	3.691	3.702	3.311	4.153	4.272	1.987
AA	10	6.594	4.444	-0.548	3.149	2.781	1.572	4.668	4.615	1.904	3.533	2.641	0.532
AA	30	4.132	3.886	1.690	4.762	3.278	-0.250	2.587	3.102	6.209	4.606	3.185	0.175
A	2	3.458	3.524	2.462	2.182	2.516	5.076	3.510	3.531	1.955	3.547	3.370	1.525
A	5	3.929	2.305	-0.030	4.594	3.646	0.342	4.780	3.458	0.429	4.959	3.012	-0.752
A	7	3.432	3.099	1.916	8.185	3.588	-2.826	3.701	3.677	3.227	6.449	4.469	0.552
A	10	3.751	3.110	1.349	2.208	2.031	1.300	3.786	3.592	1.528	3.565	2.628	0.369
A	30	4.465	3.536	0.954	2.772	2.882	3.035	3.050	4.239	11.745	3.137	2.977	1.978
BBB	2	5.060	3.531	0.302	3.000	3.123	1.610	2.581	2.908	3.188	3.395	3.326	1.724
BBB	5	3.907	2.658	0.426	6.010	2.949	-0.412	2.814	3.273	5.744	3.663	2.493	0.240
BBB	7	2.867	2.039	0.321	8.628	3.832	-3.665	2.781	2.918	5.477	4.072	4.278	2.537
BBB	10	3.668	3.260	1.602	2.887	3.147	6.616	2.680	2.803	3.013	2.801	2.196	0.624
BBB	30	2.698	2.635	1.907	2.389	2.437	6.680	1.485	0.497	5.689	5.312	3.355	0.299

ENDNOTES

The authors thank Eduardo Schwartz and Walter Torous for their suggestions and comments.

[1]The yield spread is the difference between a bond's yield, adjusted for embedded options except default, relative to an otherwise equivalent default-free security such as a Treasury.

[2]We are grateful to J.P. Morgan for assembling, checking, and providing the data.

[3]For instance, the Basle Committee associated with the Bank for International Settlements recommends that VaR calculations use a sample spanning the most recent twelve months. See Jorion [1997, p. 50].

[4]See SHAZAM User's Reference Manual [1997, p. 168]. The SHAZAM software uses the highest lag significant at a 95% level.

[5]In an effort to assure symmetry, we adopt the commonplace procedure of using log levels and first differences.

[6]If the series were independent of each other, about six rejections out of sixty could be anticipated at the 10% level, even if every series has a unit root. Our single rejection is not enough. As reported later, however, the series are quite correlated, so the paucity of rejections is less suspicious than it might at first appear.

[7]The full matrix of individual correlation coefficients is available to interested readers upon request.

[8]A comprehensive treatment of excess kurtosis as a problem in VaR calculations is provided by Duffie and Pan [1997].

[9]Except, of course, at the extreme probabilities of zero and one.

[10]More generally, for increments without convergent integrals for the second moment, one might posit that a non-Gaussian member of the stable family would be appropriate.

[11]Lower case "s" denotes the first difference of upper case "S."

[12]Starting values were generated with the EM algorithm (see Dempster, Laird, and Rubin [1977]), while the rapid BHHH algorithm (Berndt, Hall, Hall, and Housman) [1974] was used to provide estimates and standard errors unless it did not converge, in which circumstance it was replaced by the slower but more reliably convergent EM algorithm. See Constrained Maximum Likelihood Application Manual [1997, pp. 6-37] and also Hamilton [1994, pp. 685-689].

[13]Tables with the exact values of all results are available from the authors upon request.

[14]The non-parametric density estimate employs a Gaussian kernel and the "normal" reference rule (Scott [1992, p. 131].)

[15]In Exhibit 12 and later analogous figures, a series is omitted if it has an insignificant likelihood ratio at an even lower value of K. In this instance, if there were no evide--e that a two-component mixture was required, a three-cor nent mixture was not estimated.

[16]GARCH is Generalized Auto-Regressive Conditional Heteroscedasticity; see Bollerslev [1986]. An intermediate case, not investigated here, could be something like a Markov switching model; see Hamilton [1994, ch. 22].

REFERENCES

Berndt, E.R., B.H. Hall, R.E. Hall, and J.A. Housman. "Estimation and Inference in Non-Linear Models." *Annals of Economic and Social Measurement*, 3, 4 (1974), pp. 653-666.

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

Das, Sanjiv T., and Peter Tufano. "Pricing Credit Sensitive Debt when Interest Rates, Credit Ratings and Credit Spreads are Stochastic." *Journal of Financial Engineering*, 5 (1996), pp. 161-198.

Dempster, A., N. Laird, and D. Rubin. "Maximum Likelihood Estimation from Incomplete Data via the EM Algorithm." *Journal of the Royal Statistical Society*, Series B, 39 (1977), pp. 1-38.

Dickey, R.A., and W.A. Fuller. "Likelihood Ratio Tests for Autoregressive Time Series with a Unit Root." *Econometrica*, 49 (1981), pp. 1057-1072.

Duffie, Darrell, and Jun Pan. "An Overview of Value at Risk." *Journal of Derivatives*, 4, 3 (Spring 1997), pp. 7-49.

GAUSS. *Constrained Maximum Likelihood Application Manual*, Version 3.2.32. Maple Valley, WA: Aptech Systems, 1997.

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

Jarrow, Robert A., David Lando, and Stuart M. Turnbull. "A Markov Model for the Term Structure of Credit Risk Spreads." *Review of Financial Studies*, 10, 2 (Summer 1997), pp. 481-523.

Jorion, Philippe. *Value at Risk*. Chicago: Irwin, 1997.

Scott, David W. *Multivariate Density Estimation*. New York: Wiley, 1992.

SHAZAM. *User's Reference Manual, Version 8.0*. New Yc McGraw-Hill.