

The Journal of Portfolio Management

VOLUME 33, NUMBER 2

www.ijpm.com

WINTER 2007

DTSSM (Duration Times Spread)

ARIK BEN DOR, LEV DYNKIN, JAY HYMAN,
PATRICK HOUWELING, ERIK VAN LEEUWEN
AND OLAF PENNINGA

Robeco is an international asset manager offering an extensive range of active investments, from equities to bonds. Research lies at the heart of everything we do, with a 'pioneering but cautious' approach that has been in our DNA since our foundation in Rotterdam in 1929. We believe strongly in sustainability investing, quantitative techniques and constant innovation.

ROBECO
The Investment Engineers

DTSSM (Duration Times Spread)

A new measure of spread exposure in credit portfolios.

Arik Ben Dor, Lev Dynkin, Jay Hyman, Patrick Houweling,
Erik van Leeuwen, and Olaf Penninga

ARIK BEN DOR

is an analyst, Quantitative
Portfolio Strategy, at
Lehman Brothers in
New York City.
abendor@lehman.com

LEV DYNKIN

is a managing director,
head of Quantitative
Portfolio Strategy at
Lehman Brothers in
New York City.
ldynkin@lehman.com

JAY HYMAN

is a senior analyst,
Quantitative Portfolio
Strategy at Lehman
Brothers.
jhyman@lehman.com

PATRICK HOUWELING

is a senior quantitative
researcher at Robeco Asset
Management in Rotterdam,
The Netherlands.
P.Houweling@robeco.nl

ERIK VAN LEEUWEN

is a senior portfolio
manager at Robeco Asset
Management in Rotterdam,
The Netherlands.
E.M.H.van.Leeuwen@robeco.nl

OLAF PENNINGA

is a senior portfolio
manager at Robeco Asset
Management in Rotterdam,
The Netherlands.
O.Penninga@robeco.nl

DTS is a registered service mark of Lehman Brothers.

Asset allocation in a portfolio or a benchmark is typically expressed in terms of percentage of market value. It is widely recognized that this is not sufficient for fixed-income portfolios, where differences in duration can cause two portfolios with the same allocation of market weights to have extremely different exposures to macro-level risks.

As a result many fixed-income portfolio managers have become accustomed to expressing their allocations in terms of contributions to duration—the product of the percentage of portfolio market value represented by a given market cell and the average duration of securities in that cell. This represents the sensitivity of the portfolio to a parallel shift in yields across all securities within this market cell. For credit portfolios in particular, the corresponding measure would be contributions to spread duration, measuring the sensitivity to a parallel shift in spreads.¹

Determining the set of active spread duration contributions to market cells or issuers is one of the primary decisions required of credit portfolio managers. Yet all spread durations are not created equal. Just as one could create a portfolio that matches the benchmark exactly by market weights, but clearly takes more credit risk (e.g., by investing in the longest-duration credits within each cell), one could match the benchmark exactly by spread duration contributions and still take more credit risk—by choosing the securities with the widest spreads within each cell.

These bonds presumably trade wider than their peer groups for a reason—that is, the market consensus has determined that they are more risky—and are often referred to as high-beta because their spreads tend to react more strongly than the rest of the market to a systematic

shock. Portfolio managers are well aware of this effect, but many tend to treat it as a secondary effect rather than as an intrinsic part of the allocation process.

To reflect the view that higher-spread credits represent greater exposures to sector-specific risks, we propose a simple risk sensitivity measure that uses spreads as a fundamental element in the credit portfolio management process. We represent sector exposures by contributions to *Duration Times Spread* (DTS), computed as the product of market weight, spread duration, and spread. For example, an overweight of 5% to a market cell implemented by purchasing bonds with a spread of 80 basis points and spread duration of three years is equivalent to an overweight of 3% using bonds with an average spread of 50 b.p. and spread duration of eight years.

To understand the intuition behind the new measure we propose, consider the return due strictly to change in spread— R_{spread} . Let D denote the spread duration of a bond and s its spread; the spread change return is then given by:²

$$R_{spread} = -D\Delta s \quad (1)$$

It is quite easy to see that this is equivalent to

$$R_{spread} = -Ds \frac{\Delta s}{s} \quad (2)$$

That is, just as spread duration is the sensitivity to an absolute change in spread (e.g., spreads widen by 5 b.p.), DTS (Ds) is the sensitivity to a relative change in spread. Note that this notion of relative spread change provides for a formal expression of the idea mentioned earlier—that credits with wider spreads are riskier since they tend to experience greater spread changes.

In the absolute spread change approach in Equation (1), the volatility of excess returns is approximated by:

$$\sigma_{return} \cong D\sigma_{spread}^{absolute} \quad (3)$$

while in the relative spread change approach of Equation (2), excess return volatility follows:

$$\sigma_{return} \cong Ds\sigma_{spread}^{relative} \quad (4)$$

Given that the two representations are equivalent, why should one of them be preferable over the other?

We provide ample evidence that the second approach, based on relative spread changes, has an advantage due to the stability of the associated volatility estimates. Using a large sample with over 560,000 observations over the

period September 1989–January 2005, we demonstrate that the volatility of spread changes (both systematic and idiosyncratic) is indeed linearly proportional to spread level. This relation holds for both investment-grade and high-yield credit, irrespective of the sector, duration, or time period. Additional analysis (to be published later) indicates that these results are not confined to the realm of U.S. corporate bonds but also apply to European corporate bonds and credit default swaps (CDX and iTraxx), regardless of the spread reference curve (e.g., Treasury or LIBOR). This explains why relative spread volatilities of spread asset classes are much more stable than absolute spread volatilities, both across different sectors and credit quality tiers and also over time.

The paradigm shift we advocate has many implications for portfolio managers, both in terms of the way they manage exposures to industry and quality factors (systematic risk) and in terms of their approach to issuer exposures (non-systematic risk). Throughout this article, we present evidence that the relative spread change approach offers increased insight into both of these sources of risk.

ANALYSIS OF SPREAD BEHAVIOR OF CORPORATE BONDS

How should the risk associated with a particular market sector be measured? Typically, for lack of any better estimate, the historical volatility of a particular sector over some previous time period is used to forecast its volatility for the coming period.³


For this approach to be reliable, we would like to find that these volatilities are fairly stable. Unfortunately, this is not always the case.

As an example, Exhibit 1 shows the 36-month trailing volatility of spread changes for various credit ratings in the Lehman Brothers Credit Index between September 1989 and January 2005. It is clear that spread volatility changed dramatically over the period, declining steadily until 1998 and then increasing significantly through 2005. The dramatic rise in spread volatility since 1998 was only partially a response to the Russian crisis and the Long-Term Capital Management debacle, as volatility has not reverted to its pre-1998 level.

If the investment-grade corporate universe is instead partitioned by spread levels, we find considerably more stable spread volatilities, as seen in Exhibit 2. After an initial shock in 1998, the volatilities within each spread bucket revert almost exactly to their pre-1998 level (beginning

EXHIBIT 1


Spread Change Volatility by Credit Rating (trailing 36 months: 9/89-1/05)



Source: Lehman Brothers.

EXHIBIT 2

Spread Change Volatility by Spread Range (trailing 36 months: 9/89-1/05)



Source: Lehman Brothers.

in August 2001, exactly 36 months after the Russian crisis occurred). In this respect, one could relate the results of Exhibit 1 to an increase in spreads—both across the market and within each quality group.

As Equation (4) suggests, a potential remedy to the volatility instability problem is to approximate the absolute spread volatility (basis points per month) by multiplying the historically observed relative spread volatility (% per month) by the current spread level (b.p.). This can help stabilize the process if relative spread volatility is more stable than absolute spread volatility. The results in Exhibit 2 point in this direction, and indicate a relationship between spread level and volatility.

Exhibit 3 plots side by side the volatility of absolute and relative spread changes of all bonds in the Lehman Brothers Credit index rated Baa (relative spread changes are calculated simply as the ratio of spread change to the beginning-of-month spread level). The comparison


illustrates that while a modest stability advantage is gained by measuring volatility of relative spread changes, the improvement is not as great as we might have hoped, and the plot seems to show that even relative spread changes are quite unstable. This apparent instability, however, is due only to the dramatic events that took place in the second half of 1998.

When we recompute the two time series excluding the four observations representing the period 8/98–11/98, the difference between the modified time series is striking. From a low of 3 b.p. per month in mid-1997, absolute spread volatility increases steadily through a high of 16 b.p. per month in 2002–2003, growing by a factor of five. Relative spread volatility, however, increases more modestly over the same period, from 3% per month to 7% per month.

Another demonstration of the enhanced stability of relative spreads is seen when comparing of the volatilities of various market segments over distinct time periods. We

EXHIBIT 3

Absolute and Relative Spread Change Volatility of Baa Credit (trailing 36 months)



Actual data replaced with missing values for months that are excluded.
Source: Lehman Brothers.

have already identified 1998 as a critical turning point for the credit markets, due to the combined effect of the Russian default and the Long-Term Capital Management crisis. To what extent is volatility information prior to 1998 relevant in the post-1998 period?

In Exhibit 4, we plot pre-1998 volatility on the x axis, and post-1998 volatility on the y axis. We do this for two different measures of volatility: absolute spread volatility, and relative spread volatility.⁴ Each point on the graph represents a particular sector-quality cell of the Lehman Brothers Credit Index, which is divided into eight industry groups by three quality cells. Points along the diagonal line reflect identical volatilities in both time periods.⁵

Two clear phenomena can be observed here. First, most of the observations representing absolute spread volatilities are located quite far above the line, pointing to an increase in volatility in the second period of the sample, even though the events of 1998 are not reflected in the data. Relative spread volatilities, however, are quite stable, with almost all observations located on the 45-degree line or very close to it. This is because the pickup in volatility in the second period was accompanied by a similar increase in spreads.

Second, relative spread volatilities of various sectors are quite tightly clustered, ranging from 5% to a bit over 10%, while the range of absolute spread volatilities is much wider, ranging from 5 b.p. per month to more than 20 b.p. per month.


Our results so far clearly indicate that absolute spread volatility is highly unstable and tends to rise with increasing spread. Computing volatilities based on relative spread change however, generates a more stable time series. These findings have important implications for the appropriate way to measure excess return volatility and demonstrate the need to better understand the behavior of spread changes.

To analyze the behavior of spread changes, we first examine the dynamics of month-to-month changes in spreads of individual bonds. When spreads widen or tighten across a sector, do they tend to shift in parallel or shift proportionally? This key issue should determine how we measure exposures to systematic spread changes.

If spreads change in a relative fashion, then the volatility of systematic spread changes across a given sector of the market should be proportional to the average spread of that sector. This is true when comparing the risk of

EXHIBIT 4

Absolute and Relative Spread Change Volatility Before and After 1998



Source: Lehman Brothers.

different sectors at a given time, or when examining the volatility of a given sector at different times.

We also examine issuer-specific (or idiosyncratic) spread volatility. The dispersion of spread changes among the various issuers within a given market cell, or the degree to which the spread changes of individual issuers can diverge from those of the rest of the sector, also tends to be proportional to spread.

We investigate each of these issues using monthly spread data from the Lehman Brothers Credit Index historical database. The dataset includes more than 15 years, from September 1989 through January 2005, and provides monthly spreads, spread changes, durations, and excess returns for all bonds in the Credit Index. For the sections of our study that include high-yield bonds as well as investment-grade, we augment the dataset with historical data from the Lehman Brothers High Yield Index. A more detailed description of the dataset can be found in the appendix.

Dynamics of Spread Change

To understand why absolute spread volatility is so unstable, we first need to examine at a more fundamental level how spreads of individual securities change in a given month. One basic formulation of the change in spread of some bond i at time t is that the overall change is simply the sum of two parts, systematic and idiosyncratic:

$$\Delta s_{i,t} = \Delta s_{J,t} + \Delta s_{i,t}^{idiosyncratic} \quad \text{for } i \in J \quad (5)$$

where J denotes a peer group of bonds with similar risk characteristics (such as Financials rated Baa with duration of up to five years). This formulation is equivalent to assuming that spreads change in a parallel fashion across all securities in a given market cell J (captured by $\Delta s_{J,t}$).

Alternatively, if changes in spreads are proportional to spread level, we have (omitting the subscript t for simplicity):

$$\frac{\Delta s_i}{s_i} = \frac{\Delta s_J}{s_J} + \frac{\Delta s_i^{idio}}{s_i} \quad (6)$$

$$\text{or: } \Delta s_i = s_i \left[\frac{\Delta s_J}{s_J} \right] + \Delta s_i^{idio}$$

Equation (6) reflects the idea that systematic spread changes are proportional to the current (systematic) spread level and that the sensitivity of each security to a systematic spread change depends on its level of spread.

Higher-spread securities are riskier in that they are affected more by a widening or tightening of spreads relative to lower-spread securities with similar characteristics.

To analyze the behavior of spread changes across different periods and market segments, we estimate the parameters of two models based on Equations (5) and (6). The first model corresponds to the parallel shift approach shown in Equation (5):

$$\Delta s_{i,t} = \alpha_{J,t} + \varepsilon_{i,t} \quad (7)$$

The second model reflects the notion of a proportional shift in spreads as in Equation (6):

$$\Delta s_{i,t} = \beta_{J,t} \cdot s_{i,t} + \varepsilon_{i,t} \quad (8)$$

Comparing Equation (8) to Equation (6) reveals that the slope coefficient we estimate, $\beta_{J,t}$ corresponds to the proportional systematic spread change, $\Delta s_{J,t}/s_{J,t}$. These two models are nested in a more general model that allows for both proportional and parallel spread changes to take place simultaneously:

$$\Delta s_{i,t} = \alpha_{J,t} + \beta_{J,t} \cdot s_{i,t} + \varepsilon_{i,t} \quad (9)$$

Exhibit 5 shows changes in spreads experienced by the large issuers that constitute the Communications sector of the Lehman Brothers Corporate Index against their beginning-of-month spreads in January 2001.⁶ It is clear that this sectorwide rally was not characterized by a purely parallel shift; rather, issuers with wider spreads tightened by more.

Exhibit 6 reports the regression results when we fit the three general models of spread change to the data in this example. The results verify that spreads in the Communications sector in January 2001 changed in a proportional fashion. The slope estimate is highly significant, and the high R^2 (97.1%) indicates that the model fits the data well.⁷


The combined model that allows for a simultaneous parallel shift achieves only a slightly better fit (97.7%) and yields a somewhat unintuitive result; it shows that the sector widens by a parallel shift of 16 b.p. and simultaneously tightens by a relative spread change of -28%. We therefore estimate a fourth model, which is essentially a variant of the combined model:

$$\Delta s_{i,t} = \bar{\alpha}_{J,t} + \beta_{J,t}(s_{i,t} - \bar{s}_{J,t}) + \varepsilon_{i,t} \quad (10)$$

Normalizing spreads by subtracting the average spread level in Equation (10) yields slope coefficients and

EXHIBIT 5

Average Spreads and Spread Changes for Large Issuers in the Communications Sector (January 2001)



Source: Lehman Brothers.

EXHIBIT 6

Regression Estimates of Various Models of Spread Change—Communications

Model	Equation	Coefficients		T-stats		R ²	
		Shift (bp)	Slope (%)	Shift	Slope	Jan-05	Aggregate
Parallel	7	-45		-10.9		88.2%	16.9%
Relative	8		-21%		-23.2	97.1%	33.0%
Combined	9	16	-28%	2	-7.9	97.7%	
Combined with normalized spread	10	-45	-28%	-24.1	-7.9	97.7%	35.2%

R² values reported in the last column are based on 1,480 individual regressions (185 months × 8 sectors).

Source: Lehman Brothers.

R² that are identical to those generated by the combined model, but now the intercept $\bar{\alpha}_{J,t}$ represents the average spread change in the sample. This model expresses the month's events as a parallel tightening of -45 b.p. coupled with an additional relative shift, with a slope of -28%, that represents how much more spreads move for issuers with above-average spreads, and how much less they move for issuers with below-average spreads.

We conduct an analysis similar to that in Exhibit 6 using individual bond data in all eight sectors and 185 months included in the sample. Our hypothesis that the relative model provides in general an accurate description of the dynamic of spread changes has several testable implications. First, the aggregate R² for the relative model should be significantly better than that of the parallel model, and almost as good as that of the combined model.

Second, we would like to find that the slope factor is statistically significant (as indicated by the t -statistic) in most months and sectors. Third, the realizations of the slope and the parallel shift factor in the combined model with normalized spread should be in the same direction, especially whenever the market experiences a large move. That is, in all significant spread changes, issues with wider spreads experience larger moves in the same direction.

We find support for all three implications. The last column of Exhibit 6 reports the aggregate R^2 for these regressions across all sectors and months. The relative model explains twice as much variation in spreads (33%) as the parallel shift model (16.9%) and almost as much as the less restrictive combined model. Only about a third of spread movements are explained because in many months there is little systematic change in spreads, and spread changes are largely idiosyncratic. Still, the slope factor is statistically significant 73% of the time.

Exhibit 7 shows that large spread changes are accompanied by slope changes in the same direction (the correlation between the two is 80%). That is, bonds that trade at wider spreads will widen by more in a widening

and tighten by more in a rally. There are essentially no examples of large parallel spread movements when the slope factor moves in the opposite direction. This clear linear relationship between the shift and slope factors serves as an additional validation of the relative model.


Systematic Spread Volatility

The security-level analysis has established that systematic changes in spreads are proportional to the systematic level of spread consistent with Equation (6). We now go one step further, and examine the relation between systematic spread volatility and the level of spreads. To do this, we would like to partition our dataset by spread level, separately measure the volatility of each spread bucket, and examine the relation between spread level and spread volatility.

The nature of the dataset presents several challenges, however. First, it is far from homogeneous—it includes bonds from different industries, credit qualities, and maturities. Second, the spreads of corporate bonds exhibited large variation over the course of the period studied, so

EXHIBIT 7

Regression Coefficients for Shift and Slope Factors



Source: Lehman Brothers.

that the populations of any fixed spread buckets vary substantially from one time period to another. Our goal was to design a partition fine enough so that the bonds in each cell share similar risk characteristics, yet coarse enough so that our cells are sufficiently well populated over the course of the period to give statistically meaningful results.

The credit index is first partitioned rather coarsely by sector (Financials, Industrials, and Utilities) and then further subdivided by duration (short, medium, and long). To ensure that every sector-duration cell is well-populated each month, we do not use prespecified duration levels but rather divide each sector into three equally populated duration groups.⁸

In the last step, bonds in each sector-duration cell are assigned to one of several spread-level buckets. To allow a detailed partitioning of the entire spread range while minimizing the number of months that a bucket is sparsely populated, the spread break points differ from sector to sector. In addition, the Financial and Industrial sectors are divided into six spread buckets, while the Utilities sector has only five spread buckets (a more detailed description


of the partition and sample population can be found in the appendix).

The systematic spread change in cell J in month t can be represented simply as the average spread change across all bonds in that bucket in month t . Therefore, for each cell in the partition, we compute every month the median spread, the average spread change, and the cross-sectional standard deviation of spread change. This procedure produces 51 distinct time series datasets; each consists of a fairly homogeneous set of bonds for which we have monthly spreads and spread changes. We then calculate the time series volatility of these systematic spread changes.⁹ Similarly, the spread level for bucket J is calculated as the time series average of the monthly median spread (rather than the average spread).

The relation between the volatility of systematic spread changes and spread level is plotted in Exhibit 8, where each observation represents one of the 51 buckets in the partition. Exhibit 8 illustrates a clear relation between spread volatility and spread level. Higher spreads are accompanied by higher volatilities for all sector-duration cells. Relatively minor differences can be seen between Industrials and the

EXHIBIT 8

Time Series Volatility of Systematic Spread Changes versus Spread Level (9/89-1/05)



Source: Lehman Brothers.

other two broad sectors. Similarly, duration does not seem to have any significant systematic effect on the results.¹⁰

Nonetheless, the results shown in Exhibit 8 do not perfectly corroborate our hypothesis of proportional spread volatility, which would predict that all our observations (or at least all observations within a given sector) should lie along a diagonal line that passes through the origin, of the form:

$$\sigma_{spread}^{absolute}(s) \equiv \theta s \quad (11)$$

While the points at the left side of Exhibit 8 seem to fit this description, the points to the right, representing higher spread levels, do not seem to continue along this line. Rather, volatility seems to flatten out beyond the 200 to 250 b.p. range. Is it possible that spread volatility does not continue to grow linearly when spreads increase beyond a certain point?

Before we reject our hypothesis, we should question the significance of these few highest-spread observations.


The 250–300 b.p. spread region represents the boundary between investment-grade and high-yield bonds. For a good part of the time period of our study, these spread cells are very lightly populated by our investment-grade bond sample. Because we exclude any cell with fewer than 20 bonds, the summary results for these cells may be less robust than desired.

To examine the relation between systematic spread change volatility and spread level beyond the 250 b.p. level, we repeat the analysis including all bonds rated Ba and B. This increases the sample size by roughly 35%, to 565,602 observations. We employ the same sector x duration x spread partition, with the addition of several spread buckets to accommodate the widening of the spread range (the number of cells increases to 66).

Exhibit 9 plots the relation between systematic spread volatility and spread level using both investment-grade and high-yield data. The linear relation between the two now extends out through spreads of 400 b.p. As before, the three observations that represent the highest-spread bucket

EXHIBIT 9

Systematic Spread Volatility versus Spread Level (investment-grade + high-yield bonds)



Monthly observations for all bonds rated Aaa-B (9/1989-1/2005).
Source: Lehman Brothers.

in the Industrials sector (circled) have somewhat lower-than-expected spread volatility. Once again, we suspect the statistical relevance of these most extreme data points.

The simple linear model of Equation (11) provides an excellent fit to the data, with θ equal to 9.1% if we use all observations or 9.4% if we exclude the three circled outliers. Hence, the results suggest that the historical volatility of systematic spread movements can be expressed quite compactly, with only minor dependence on sector or maturity, in terms of a relative spread change volatility of about 9% per month. That is, spread volatility for a market segment trading at 50 b.p. should be about 4.5 b.p. per month, while that of a market segment at 200 b.p. should be about 18 b.p. per month.

Idiosyncratic Spread Volatility

To study the spread dependence of idiosyncratic spread volatility, we use the same partition as for our study of systematic spread volatility. Instead of the average spread change experienced within a given cell in a given month, we examine the dispersion of spread changes within each cell.

The idiosyncratic spread change of bond i in market cell J at time t is defined as the difference between its spread change and the average spread change for the cell in that month:

$$\Delta s_{i,t}^{idio} = \Delta s_{i,t} - \Delta s_{J,t} \quad (12)$$


The volatility of idiosyncratic spread changes is then exactly equal to the cross-sectional standard deviation of total spread changes.¹¹

Exhibit 10 shows a scatterplot of the cross-sectional volatility for all months and spread buckets including high-yield bonds. This plot clearly shows the general pattern of volatilities increasing with spread, as well as the relative paucity of data at the higher-spread levels.

To obtain a single measure of idiosyncratic spread volatility for each bucket, we pool all observations of idiosyncratic risk in a given market cell J over all bonds and all months, and compute the standard deviation. This pooled measure of idiosyncratic spread volatility per market cell is plotted in Exhibit 11 against the median spread of the cell.

EXHIBIT 10


Volatility of Idiosyncratic Spread Change versus Spread Level



Monthly calculations (9/1989-1/2005) for all bonds rated Aaa-B, computed separately by sector, duration, and spread bucket ($N = 7,250$).
Source: Lehman Brothers.

EXHIBIT 11

Pooled Idiosyncratic Spread Volatility versus Spread Level



Each observation represents the standard deviation of idiosyncratic spread changes aggregated across all sample months separately by sector, duration, and spread bucket for all bonds rated Aaa-B (9/1989-1/2005).

Source: Lehman Brothers.

The linear relation between spread and spread volatility is strikingly clear. Observations that represent buckets populated almost exclusively by high-yield bonds exhibit more variation than those representing investment-grade bonds, but follow the exact same pattern. The regression results indicate a zero intercept, but the estimated slope coefficient (the relative volatility of idiosyncratic yield change) is somewhat higher than estimated previously, 11.5% versus 9.4%.

Stability of Spread Behavior

We have established that spread volatility is linearly proportional to the level of spread. What is the extent of time variation in the spread slope or the change in spread volatility as spreads vary?

For each bucket we compute the yearly systematic spread volatility and corresponding average spread level (using 12 months of average spread change).¹² We then regress these estimates of systematic spread volatility against

an intercept and a spread slope factor. We follow the same approach for idiosyncratic spread volatility, except that we use the monthly cross-sectional volatility estimates.


Panels A and B of Exhibit 12 present the yearly spread slope estimates and the corresponding adjusted R^2 . The results are plotted separately for systematic and idiosyncratic volatility. The estimated coefficients are all highly significant, with t -statistics ranging between 15 and 30 for both systematic and idiosyncratic spread volatility.

Not surprisingly, Exhibit 12 reveals that including high-yield data generally increases the spread estimate for both systematic and idiosyncratic volatility. The spike in volatility caused by the 1998 Russian crisis is evident in the high estimate of spread slope in 1998 (except for idiosyncratic volatility with high-yield). Excluding 1998, the spread slope estimates are remarkably stable despite the limited number of observations in the estimation.

Panel B of Exhibit 12 reveals that the regressions have better and more stable explanatory power when high-yield securities are included. When we analyze investment-

EXHIBIT 12

Yearly Regression of Spread Volatility Against Spread Level (1/1990-12/2004)



Source: Lehman Brothers.

grade data only, the R^2 of our regressions goes as low as 40% for systematic volatility and 30% for idiosyncratic volatility. When we include high-yield data, the regression results are much better, achieving R^2 values of consistently over 70% for systematic volatility and 60% for idiosyncratic volatility.

Overall, this pattern confirms that relative spread changes characterize both investment-grade and high-yield credit.

A NEW MEASURE OF EXCESS RETURN VOLATILITY

What are the implications of spread proportionality? Which measure—duration times spread, or spread duration—is more appropriate for representing the risk of credit securities? We show here that excess return volatility increases linearly with DTS, consistent with the formulation in Equation (4). Furthermore, portfolios with very different spreads and spread durations but with similar DTS exhibit the same excess return volatility.

For example, a portfolio with a weighted spread of 200 b.p. and spread duration of two years is as risky as a portfolio with a spread of 100 b.p. and spread duration of four years. We also show that DTS generates better estimates of future excess return volatility than those calculated by spread duration.

DTS, Spread Duration, and Excess Returns

If the volatility of both systematic and idiosyncratic spread changes is proportional to the level of spread, the volatility of excess returns should be linearly related to DTS, with the proportionality factor equal to the volatility of relative spread changes over the corresponding period [see Equation (4)].


To examine this prediction, each month bonds are assigned to quintiles according to their DTS value. Each quintile is further subdivided into six buckets based on spread. Every month the average excess returns and median DTS are calculated, and then the time series volatility of excess returns and average DTS are calculated separately for each bucket. This formulation yields two empirical predictions:

1. Excess return volatility should increase linearly with DTS, where the ratio of the two (or slope) represents the volatility of relative spread changes previously estimated.
2. The level of excess return volatility should be approximately equal across spread buckets with a similar DTS value.

The results of the analysis, presented in Exhibit 13, strongly support both empirical predictions, even though we do not control for industry, quality, maturity, or any other effect.

EXHIBIT 13

Excess Return Volatility versus DTS (9/1989-1/2005: investment-grade credit)



Bonds are first divided into DTS quintiles and then further subdivided into six buckets by spread level.
 Source: Lehman Brothers.

First, it is clear that excess return volatility increases with the level of DTS and that a straight line through the origin provides an excellent fit. This is confirmed by a regression of the excess return volatility on average DTS, which finds a fit of 98% and an insignificant intercept. The slope estimate is 8.8%, which is in line with the estimated slope from the analysis of systematic spread volatility.

Second, consistent with prediction (2), observations representing the same DTS quintile but with differing spread levels exhibit very similar excess return volatilities. The one exception to this is in the highest DTS quintile, where the subdivision by spread causes wide variations in DTS as well. As a result, the points no longer form a tight cluster, although they do continue to exhibit the same general relation between DTS and volatility.

To demonstrate the significance of the second result, Exhibit 14 reports the average spread and spread duration for all 30 buckets. It illustrates the extent of the variation in spreads and corresponding spread duration across buckets with almost identical DTS. For example, the top and

bottom spread buckets in the second DTS quintile (shown in boldface) exhibit very close DTS values of 299 and 320. Yet they have very different spread and spread duration characteristics; bonds in the top bucket have average spread duration of 5.48 and trade at a spread of 54 b.p., while bonds in the bottom cell have spread duration of 2.53 and a spread of 127 b.p. Hence, a portfolio of high-spread bonds with short duration can be as risky as a portfolio of low-spread bonds with high duration, as long as they both have roughly the same duration times spread.¹³

A Comparison of Excess Return Volatility Forecasts

A natural step to extend the analysis is to examine which approach provides a better forecast of the excess return volatility of a portfolio:

1. Spread duration \times historical volatility of absolute spread change.
2. DTS \times historical volatility of relative spread change.

EXHIBIT 14

Summary Statistics by DTS and Spread Buckets

Panel A:		DTS buckets				
Spread Sub-buckets	Spread	Low	2	3	4	High
	Low	41	54	64	77	97
	2	52	68	79	94	116
	3	60	78	88	106	135
	4	69	87	98	118	156
	5	79	99	112	135	184
	High	100	127	143	172	246

Panel B:		DTS buckets				
Spread Sub-buckets	Spread Duration	Low	2	3	4	High
	Low	2.38	5.48	7.20	9.53	11.15
	2	2.19	4.24	6.12	7.17	10.62
	3	2.17	3.80	5.50	6.51	9.78
	4	2.17	3.54	4.96	6.09	9.09
	5	2.09	3.25	4.43	5.72	8.23
	High	1.65	2.53	3.52	4.53	6.91

Source: Lehman Brothers.

To directly compare the forecasting accuracy of the two measures, for every month we compute the realized excess return of each of the 24 buckets in the Lehman Brothers Credit Index (8 sectors \times 3 credit ratings). The carry component (spread/12) is stripped from the realized excess return, and the random part is then divided by one of the two forecasts of excess return volatility.¹⁴

If the projected excess return volatility is an unbiased estimate of the “true” volatility, then the time series volatility of these standardized excess return realizations should be very close to 1.0.

Our premise is that relative spread change volatility is a more timely measure than absolute spread change volatility, as it can react almost instantaneously to a change in market conditions. Hence we expect the sample time series standard deviation of excess returns to be closer to 1.0 under approach 2 than under approach 1. A volatility measure that adjusts more quickly for changing market conditions should also generate less extreme realizations (realizations that fall above or below two or three standard deviations) than a measure that is slower to react.

Before we examine the results of the forecasting, it may be helpful to look at the conditions under which the

volatility forecasts generated by the two measures differ. If we explicitly write the expression for the ratio of the two measures at month t for some bucket J :

$$\begin{aligned}
 \text{Vol ratio}_{J,t} &= \frac{\sigma_t(\frac{\Delta s_J}{s_J}) \times \sum_{i \in J} D_{i,t} \times s_{i,t}}{\sigma_t(\Delta s_J) \times \sum_{i \in J} D_{i,t}} \\
 &\cong \frac{\theta \times \sum_{i \in J} D_{i,t} \times s_{i,t}}{\theta \times \bar{s}_{J,t} \times D_{J,t}} \\
 &\cong \frac{\sum_{i \in J} D_{i,t} \times (s_{J,t} + s_{i,t}^{idio})}{\bar{s}_{J,t} \times D_{J,t}} \\
 &\cong \frac{s_{J,t} \times D_{J,t}}{\bar{s}_{J,t} \times D_{J,t}} = \frac{s_{J,t}}{\bar{s}_{J,t}} \quad (13)
 \end{aligned}$$

The volatility measure based on relative spread changes reflects the current spread level of bucket J , while the volatility measure based on absolute spread changes reflects the time-weighted average spread the bucket has exhibited over the volatility estimation period (denoted $\bar{s}_{J,t}$).

If, for example, the systematic spread level of bucket J over the estimation period is unchanged, the ratio would be equal to 1.0. Otherwise, it would be above or below 1.0, depending on whether the current spread is above or below the historical average. Using a shorter period for estimating spread change volatility will not necessarily reduce the difference between the two measures, if the long-term historical spread is a better reflection of the current spread environment than the recent past.


Exhibit 15 displays the mean and standard deviation of the time series of normalized residuals separately for each volatility measure (each observation represents one of the 24 buckets). The normalized residuals are generated from the entire history of returns available at the beginning of each month. In addition, Figure 15 shows the mean and standard deviation of the normalized residuals when the absolute spread change volatility is calculated over the previous 36 months only. This corresponds to the approach many investors take in periods of exceptionally low or high volatility, namely, to rely only on recent data.

A comparison of the three sets of observations reveals that using absolute spread changes produces downward- (upward-) biased estimates of volatility for the entire available history (the previous 36 months). As a result, the average standard deviation of normalized excess returns using the entire history is above 1.0 (1.14) and using the partial history is below 1.0 (0.92). The observations generated using relative spread changes are evenly spread around 1.0, and the average standard deviation of standardized excess returns is 1.01. A close examination of the results does not suggest any relation between the deviation from 1.0 and the sector-quality bucket.

These findings support our empirical prediction and are also consistent with the analysis of the ratio of the two volatility measures. Excess return volatility estimates based on absolute spread changes are very sensitive to the length of the estimation period; they may overreact when using too few data points and can be slow to adjust when using a long history. What is the optimal estimation period is not clear ex-ante when using absolute spread changes.

EXHIBIT 15

Mean and Standard Deviation of Normalized Excess Return Realizations



Conditional volatility estimates are computed monthly by sector and credit quality based on the entire available history or previous 36 months, using monthly spread changes observations (9/89-1/05).

Source: Lehman Brothers.

A longer estimation period is always desirable for proportional spread changes, however, because it improves the accuracy of the proportionality factor, while at the same time the volatility estimate adjusts instantaneously because of multiplication by the current spread level.¹⁵

The second empirical prediction states that there should be a lower percentage of extreme realizations (positive or negative) in the case of relative rather than absolute spread change volatility. Exhibit 16 plots a histogram of the standardized excess return realizations for all sector-quality cells for the two volatility measures. The standard normal distribution is also displayed.

Not surprisingly, the histogram reveals that both volatility estimators generate distributions that are negatively skewed (−2.67 and −1.35 for the relative and absolute spread change-based volatility measures). With respect to the percentage of outliers, 7.06% of the observations in the distribution based on absolute spread changes are located beyond two standard deviations from the mean. In the case of the distribution based on relative spread changes, the same figure is almost half, at 4.03%.

THE SCOPE OF DTS

What do our findings imply for the level of spread volatility as spreads approach zero? Taking our results at

face value, one might say there is no lower bound for volatility and that spread volatility should decline to almost zero for very low-spread securities. Spread volatility, however, is not driven solely by changes in credit risk but also by non-credit risk-based factors. Non-credit risk-based spread changes can occur because of noise (i.e., pricing errors); demand/supply imbalance (for example, when securities enter or exit the Lehman Brothers Corporate Index); and other factors.

Spread volatility (systematic or idiosyncratic) can therefore be represented as the sum of two terms: a constant term that reflects non-credit risk-based spread volatility, and a second term that represents spread volatility due to changes in credit risk (which may be approximated by a linear function of spread) as follows:


$$\sigma(\Delta s) = \sqrt{\sigma_{\text{non credit risk}}^2 + \theta^2 s^2} \quad (14)$$

Equation (14) makes it clear that for sufficiently high spreads, the second term dominates the first, and spread volatility can be approximated well by a linear function of spread, as we find for corporate bonds. As spreads tighten and approach zero, the first term dominates, and spread volatility should converge to some minimum structural level.

Agency debentures provide a natural framework to examine the behavior of spread volatility for very low spreads.

EXHIBIT 16

Distribution of Standardized Excess Returns



Based on observations (9/1992—1/2005) grouped across all sectors and credit ratings.
Source: Lehman Brothers.

Because of market perception that securities issued by the three major agencies are backed by the U.S. government, these securities typically trade at very low spreads. Between September 1989 and April 2005, the median agency spread stayed between 20 and 50 b.p. except for a few distinct months. We follow the same approach as for corporates. Each month, bonds are partitioned according to beginning-of-month spread level. Average spread change and median spread level are computed separately for each bucket. We then examine the relation between the time series volatility and average (median) spread level of each bucket.

The sample spans roughly the same time period as the corporate dataset (September 1989–April 2005) and includes all Aaa-rated, non-callable debentures in the Lehman Brothers Agency Index.¹⁶ As before, we discard extreme observations (in either the top or bottom percentile of the spread distribution). Since the total number of observations (73,000) is about 17% of the corporate sample size, we use only eight spread buckets.

The results are presented in Exhibit 17. For comparison, we also show the spread volatility of long-duration financials which share many of the characteristics as agencies.

The plot in Exhibit 17 illustrates that spread volatility is roughly constant for spreads below 20 b.p., and the level of structural systematic volatility is about 2.5–3.0 b.p. per month. Above 20 b.p., the relation takes the usual linear

shape and fits nicely with that of long-financials. A regression of spread volatility against spread level reveals a flatter slope than we estimated for corporates (5.7% versus 9%), consistent with Equation (14).¹⁷

An analysis of idiosyncratic volatility indicates in a similar fashion that volatility increases moderately as spreads increase from 20 b.p. to 80 b.p. and indicates a structural volatility level of 4.0–4.5 b.p. per month. The fact that idiosyncratic structural volatility is higher than the corresponding systematic level is to be expected, as pricing noise should be more pronounced for individual securities.

To complete the analysis, we partition the sample into 12 DTS buckets and plot the excess return volatility of each bucket against its DTS (Exhibit 18). Similar to corporate bonds, excess return volatility increases linearly with DTS (the estimated slope from the regression is 9.8% versus 8.8% for corporates). As the DTS approaches zero, however, there is a clear flattening of the relation, and volatility does not decline further. Indeed, the regression yields a significant intercept of 3 b.p., which is consistent with our previous estimate for the structural level of systemic volatility.

DTS Across Seniority Classes

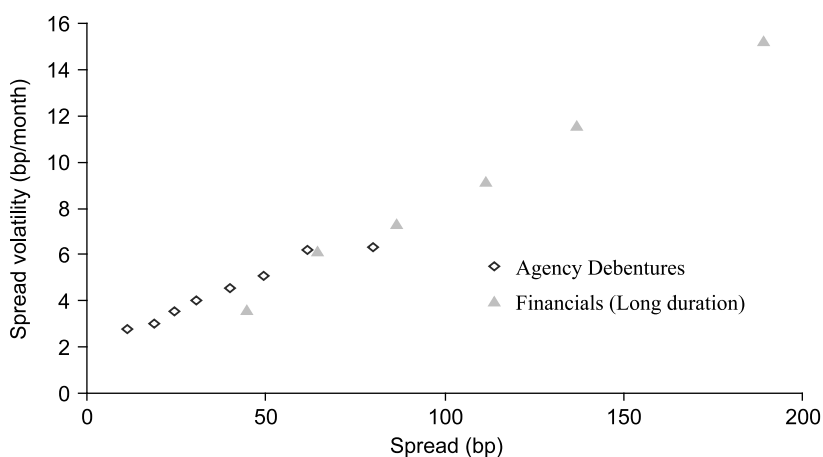
Probably some of the most convincing evidence in support of the DTS concept is that portfolios that are remarkably different in terms of their spread and spread duration, but with a similar product of the two (DTS), exhibit the same excess return volatility. Underpinning this result is the issue of whether credit risk is fully captured by spreads. If spreads incorporate on average all publicly available information related to credit risk, then all portfolios with similar DTS should have the same level of excess return volatility.

We re-examine this issue in the context of debt seniority by looking at portfolios of bonds from different seniority classes (e.g., senior notes, debentures) but with a very similar DTS. If spreads already incorporate the likelihood of default and the recovery value in such a case, all portfolios should exhibit the same excess return volatility. Such a result would provide further support for our earlier findings.

While credit ratings naturally lend themselves to cross-sectional comparisons,

EXHIBIT 17


Systematic Spread Volatility Versus Spread Level of Agencies (9/1989–4/2005)



Source: Lehman Brothers.

EXHIBIT 18

Excess Return Volatility versus DTS (9/1989-4/2005): agencies



Source: Lehman Brothers.

EXHIBIT 19

Summary Statistics for Senior and Subordinated Portfolios

Portfolio Composition		# of issuers	Ratio of Excess Return Volatility			Difference in Excess Returns (% / month)		
'Senior'	'Subord'		P ₂₅	P ₅₀	P ₇₅	P ₂₅	P ₅₀	P ₇₅
Senior debt	Subordinated debt	47	0.83	1.10	1.41	-0.15	-0.06	-0.01
Senior Notes	Notes	353	0.79	0.94	1.08	-0.08	-0.01	0.04
Sr. Debentures + Debentures	Sub. Debentures	46	0.80	0.94	1.04	-0.05	0.02	0.05
Sr. Debentures + Sr. Notes + Senior debt	Debentures + Sub. Debentures + Notes + Subordinated debt	535	0.80	0.93	1.08	-0.13	-0.04	0.01

Source: Lehman Brothers.

constructing portfolios based on debt seniority is more of a challenge. The classification of a bond as senior or subordinated depends on its payment priority in case of a default. The recovery value of any bond is affected by other claims on the same issuer that are more or less senior to that bond. Across issuers, however, the same seniority class does not necessarily imply similar recovery value in case of a default. Furthermore, even for a given issuer it is not always clear if a certain claim is senior to another claim (e.g., a debenture versus a senior note). Thus simply

grouping bonds into portfolios according to seniority class is inappropriate.¹⁸

To address these issues, we perform a more detailed analysis at the issuer level (identified by ticker). Each month, we construct two portfolios for each issuer, SENIOR and SUBORD, which include all the securities (often just a single one) defined as senior and subordinated. Months when only one of the portfolios is populated are eliminated.

We first compute the market-weighted duration times spread and excess return for each portfolio and the DTS ratio of the SENIOR portfolio to the SUBORD portfolio. We then match the DTS of the SENIOR portfolio to that of the SUBORD portfolio (i.e., the DTS is scaled up or down) and adjust the excess return accordingly. Hence for every issuer, we have a time series of excess returns for two portfolios with the same DTS each month.¹⁹

Using this approach for portfolio construction has clear advantages over the cross-sectional technique. First, it controls for any issuer-specific effect. Second, it accurately captures the relative seniority of different claims. Third, the fact that by construction the two portfolios have the same DTS has testable implications: The ratio of excess return volatility of the two portfolios should be 1.0 on average. In addition, any difference in excess return should be relatively small and reflect only idiosyncratic risk (for example, one portfolio may include bonds that on average are smaller and older than bonds in the second portfolio and are therefore less liquid).

Exhibit 19 presents the 25th percentile, 50th percentile, and 75th percentile of the ratio of excess return volatility for the SENIOR and SUBORD portfolios as well as the difference in average excess returns. As the ordering among seniority classes is not always clear, the table presents these statistics for different compositions of the SENIOR and SUBORD portfolios.

For example, the second row reports senior notes under SENIOR and notes under SUBORD. There were 353 different issuers with these portfolios populated over some time period. The median ratio of excess returns volatilities is 0.94, indicating no significant difference between

the two portfolios. One-quarter of the issuers exhibited ratios below 0.79, and one-quarter of the issuers ratios above 1.08, with the remaining half falling between these values. The typical performance of the two portfolios is also very similar and the median difference is 1 basis point per month (i.e., the SUBORD portfolio underperforms).

The results reported for other portfolio compositions are similar (in particular the last row, which represents the most inclusive case), and do not indicate the two portfolios exhibit different risk characteristics.

To examine the relation between duration times spread and excess return volatility across various seniority classes, each SENIOR and SUBORD portfolio (constructed for each issuer) is assigned to one of the DTS quintiles each month. We then calculate the weighted excess return and DTS for each quintile (separately by seniority class). The two aggregate portfolios in each quintile have the exact same DTS since at the issuer level the DTS of the SENIOR and SUBORD portfolios is equal by construction. As before, we compute the time series volatility of excess returns and the average DTS of the ten portfolios.

Exhibit 20 presents the results of the analysis comparing senior notes to notes and senior debentures to subordinated debentures. The scatterplot shows that the linear relation between excess return volatility and DTS is preserved, and that the slope does not depend on the seniority level. In both cases, there is an almost exact match between the volatilities of the SENIOR and SUBORD portfolios.

We obtain similar results for other compositions of the two portfolios that are reported in Exhibit 19.

SUMMARY AND IMPLICATIONS FOR PORTFOLIO MANAGERS

We have presented a detailed analysis of the behavior of spread changes. Using our extensive corporate bond database, which spans 15 years and includes well over 560,000 observations, we demonstrate that spread changes are proportional to the level of spread. Systematic changes in spread across a sector tend to follow a pattern of relative spread change; that is, bonds trading at wider spreads experience greater spread changes. The systematic spread volatility of a given sector (if viewed in terms of absolute spread changes) is proportional to the median spread in the sector; the non-systematic spread volatility of a particular bond or issuer is proportional to its spread as well. These findings hold irrespective of sector, duration, or time period.

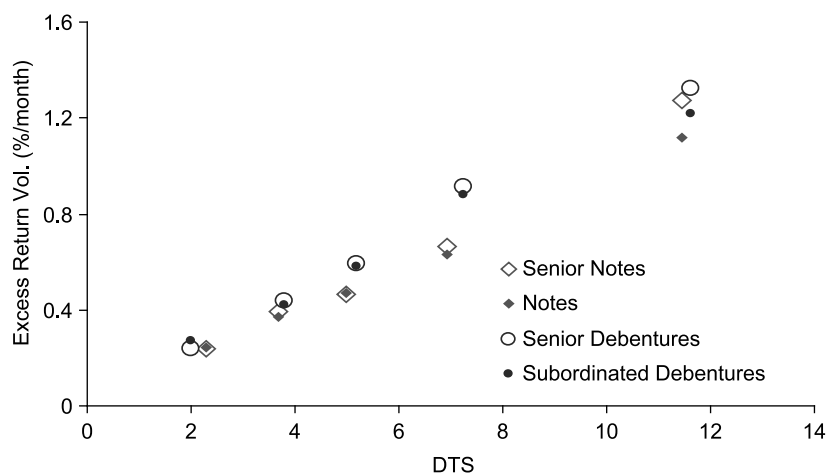
In a sense, these results are not altogether surprising. The lognormal models typically used to represent changes in interest rates assume that changes in yield are proportional to current yield levels. Models for pricing credit derivatives have used a similar lognormal model to describe changes in credit spreads (see Schönbucher [1999]).

An assumption of lognormal spread changes would imply two things: 1) That spread changes are proportional to spreads, and 2) that the relative spread changes are normally distributed. Our results can be seen as providing empirical evidence to support the first of these assumptions, but not necessarily the second.

There are several implications for a portfolio manager who wishes to act on these results. First, the best measure of exposure to a systematic change in spread within a given sector or industry is not the contribution to spread duration, but the contribution to duration times spread, or DTS. At many asset management firms, the targeted active exposures for a portfolio relative to its benchmark are expressed as contribution to duration overweights and underweights along a sector by quality grid—and reports on the actual portfolio follow the same format. In the relative spread change approach managers would express their targeted overweights and underweights in terms of contributions to DTS instead.

EXHIBIT 20

Excess Return Volatility versus DTS Across Seniority Classes



Source: Lehman Brothers.

Second, our finding that the volatility of non-systematic return is proportional to DTS offers a simple mechanism for defining an issuer limit policy that enforces smaller positions in more risky credits. Many investors specify some ad hoc weight cap by credit quality to control issuer-specific risk.²⁰ Alternatively, we can set a limit on the overall contribution to DTS for any single issuer.

For example, say the product of market weight \times spread \times duration must be 5.0 or less. Then, a position in issuer A, with a spread of 100 basis points and a duration of five years, could be up to 1.00% of portfolio market value, while a position in issuer B, with a spread of 150 and an average duration of ten years, would be limited to 0.33%.

Establishing issuer limits based on spreads has advantages and disadvantages compared to a ratings-based approach. One advantage, as described above, is that it is simple to specify a single uniform limit that requires increasing diversification with increasing risk. The key difference between the two approaches, though, concerns how often issuer limits are adjusted.

In a ratings-based framework, bond positions that are within policy on the date of purchase will tend to remain within policy unless they are downgraded. A spread-based constraint, by contrast, however, is by its very nature constantly adjusted as spreads change. One possible result is that as spreads widen, a position that was within policy when purchased can drift over the allowable DTS limit. Strict enforcement of this policy, requiring forced sales to keep all issuer exposures to stay within the limit, could become very distracting to managers, and incur excessive transaction costs as spreads trade up and down.

One possible solution would be to specify one threshold for new purchases and a higher one at which forced sales would be triggered. This could provide a mechanism that adapts to market events more quickly than the rating agencies without introducing undue instability.

Another possible disadvantage of the DTS-based issuer cap is that it allows for large positions in low-spread issuers and exposes the portfolio to credit torpedoes. This too would argue for using the DTS-based approach in conjunction with caps on market weights.

Third, there could be hedging implications. Say a hedge fund manager has a view on the relative performance of two issuers within the same industry, and would like to capitalize on this view by going long issuer A and short issuer B in a market-neutral manner. How do we define market neutrality?

A typical approach might be to match the dollar durations of the two bonds, or to go long and short credit default swaps of the same maturities with the same notional amounts. If issuer A trades at a wider spread than issuer B, however, our results would indicate that a better hedge against marketwide spread changes could be obtained by using more of issuer B, so as to match the contributions to DTS on the two sides of the trade.

Should portfolio management tools such as risk and performance attribution models be modified to view sector exposures in terms of DTS contributions and sector spread changes in relative terms? The answer in our opinion is yes, in both cases.

A risk model for any asset class is essentially a set of factors that characterize the main risks to which securities in that asset class are exposed. The risk of an individual security or portfolio depends on its risk loadings or sensitivity to the set of factors and their past realizations. For credit-risky securities, the traditional risk factors typically measure absolute spread changes in terms of a sector-by-quality partition that spans the universe of bonds. Specifying the risk factors in terms of relative spread changes instead has two important benefits.

First, such factors would exhibit more stability over time and allow better forward-looking risk forecasts. Second, the partition by quality would no longer be necessary to control risk, and each sector can be represented by a single risk factor. This would allow managers to express more focused views, essentially trading off the elimination of the quality-based factors with a more finely grained partition by industry.

Similarly, a key goal for attribution models is to match the allocation process as closely as possible. If and when a manager starts to state allocation decisions in terms of DTS exposures, performance attribution should follow suit.

One practical difficulty that may arise in the implementation of DTS-based models is an increased vulnerability to pricing noise. For the most part, models of portfolio risk and reporting of active portfolio weights rely largely on structural information. Small discrepancies in asset pricing give rise to small discrepancies in market values, but potentially larger variations in spreads. Managers who rely heavily on contribution-to-DTS exposures will need to implement strict quality controls on pricing.

Our investigation of the relation between duration times spread and excess return volatility has focused almost entirely on investment-grade credit in the U.S. We have

also included, however, some results from U.S. high-yield credit that show that the idea of proportional spread changes carries through to high-yield as well as agency debentures. Additional analysis we have conducted establishes that all our results are valid for European corporate bonds as well as for credit default swaps. Furthermore, the findings are robust to the return horizon (monthly or weekly) and to the reference curve used to calculate the spreads (Treasury or LIBOR).

Indeed, we believe that perhaps one of the most useful applications of DTS will be in the management of core-plus portfolios that combine both investment-grade and high-yield assets. Traditionally, investment-grade credit portfolios are managed according to contributions to duration, while high-yield portfolios are managed according to market value weights; using contributions to DTS across both markets could help unify this process.

Skeptics may point out that in high-yield markets, especially when moving toward the distressed segment, neither durations nor spreads are particularly meaningful, and the market tends to trade on price, based on an estimated recovery value. A useful property of DTS in that context is that in the case of distressed issuers, where shorter-duration securities tend to have artificially high spreads, DTS is fairly constant across the maturity spectrum, so that managing issuer contributions to DTS becomes roughly equivalent to managing issuer market weights.

We believe the duration times spread approach accurately represents the impact of spread changes on excess returns, and that acceptance of this result could have wide-ranging effects on portfolio management practice throughout the industry. We anticipate continued research in this area on several fronts, including extension to other asset classes, and implementation of DTS-based features into portfolio analytics offerings.

APPENDIX: DATA DESCRIPTION


The dataset used in the empirical analysis covers September 1989 through January 2005 (a total of 185 months). The sample includes all the bonds in the Lehman Brothers Credit Index excluding 1) zero-coupon bonds, 2) callable bonds, and 3) bonds with non-positive spreads. The final dataset provides a total of 416,783 observations. Exhibit A-1 shows the sample by sector and year.

We also extend the analysis to include high-yield bonds rated Ba and B (trading at a price above 80 to mitigate potential default effects), which increases the number of observations by roughly 35% (from 416,783 to 565,602).

Spread figures are model-driven and can exhibit extreme values (especially since the methodology for computing option-adjusted-spreads changed during the sample period). To mitigate the effect of outliers, we exclude observations where changes in spread fall above the 99th percentile or below the 1st percentile. As a result, monthly spread changes included in the analysis range from -60 b.p. to +78 b.p.

EXHIBIT A - 1

Sample Population by Sector and Year



*Investment-grade bonds only; number of bonds is as of December of each year.
Source: Lehman Brothers.*

Figure A2 outlines the exact breakdown into spread buckets by industry and maturity we employ in analyzing the relation between spread volatility and spread level. Note that because spread has a tendency to widen with maturity, the population of the short-maturity bucket is concentrated in the lowest-spread bucket (denoted by 1) while the opposite holds for the long-maturity bucket.

We also report for each bucket the percentage of months during the sample period when the bond population exceeds 20. This statistic is of interest, because months with fewer than 20 observations are filtered out of any volatility calculation. The percentage of months with sufficient number of observations varies between 30% and 50% for Utilities and 50%-80% for Financial and Industrials.

EXHIBIT A - 2

Sample Partition by Sector, Duration, and Spread

Sector/ Maturity	Spread bucket / breakpoints					
	1	2	3	4	5	6
Financials	<0.50	0.50-0.75	0.75-1.00	1.00-1.25	1.25-1.50	>1.5
Short	16,881 (50.8%)	13,201 (82.7%)	9,351 (64.9%)	5,296 (46.5%)	2,677 (30.8%)	4,004 (37.3%)
Medium	5,839 (28.6%)	14,838 (65.4%)	11,156 (73.5%)	8,173 (61.6%)	5,133 (44.3%)	6,904 (48.1%)
Long	2,183 (18.9%)	12,875 (54.6%)	10,743 (81.1%)	8,174 (73.0%)	6,130 (58.9%)	11,993 (55.1%)
Industrials	<0.60	0.60-0.85	0.85-1.20	1.20-1.50	1.50-2.00	>2.00
Short	22,794 (84.9%)	13,705 (97.8%)	12,172 (78.9%)	7,670 (54.6%)	6,277 (48.6%)	6,167 (30.8%)
Medium	12,814 (70.3%)	14,621 (85.4%)	14,424 (96.2%)	9,109 (65.4%)	9,300 (54.6%)	9,131 (43.2%)
Long	9,212 (68.1%)	13,961 (81.6%)	16,248 (94.6%)	10,088 (69.7%)	11,010 (53.5%)	8,940 (40.5%)
Utilities	<0.55	0.55-0.75	0.75-1.15	1.15-1.50	>1.50	
Short	5,017 (46.5%)	3,233 (35.7%)	4,443 (48.6%)	2,388 (22.2%)	2,350 (16.8%)	
Medium	3,430 (41.1%)	3,552 (38.9%)	4,484 (41.1%)	2,699 (32.4%)	3,889 (23.2%)	
Long	3,030 (32.4%)	3,199 (40.5%)	4,457 (52.4%)	2,653 (25.4%)	2,350 (29.2%)	

Sample includes IG bonds only, between 9/89 and 1/05. Spread breakpoints, cell population, and percentage of months a bucket is populated by more than 20 bonds.

Source: Lehman Brothers.

ENDNOTES

¹Spread is the constant (absolute) shift to the zero-coupon discount curve in all scenarios that is required to ensure that the model value of the bond (average value over all scenarios) equals the observed market price.

²Spread change return is closely related to excess return, the return a corporate bond earns over a duration-matched Treasury bond. Excess return can be approximated by the sum of the spread change return and an additional component due to spread carry.

³This practice leads to perennial questions about how much history should be used in such estimation. A longer time period leads to more stable estimates of volatility; a shorter time period (or a weighting scheme that gives more weight to recent observations) makes the estimate less stable, but better able to adapt to fundamental changes in the marketplace. In either case, the large swings in volatility that the market can experience mean that we are always trying to catch up to market events, and there will always be some amount of lag between the time

of a volatility change and the time it is first reflected in our estimates.

⁴To enable the two to be shown on the same set of axes, both absolute and relative spread volatility are expressed in units of similar magnitudes. The interpretation is different, however. An absolute spread change of 10 represents a 10 b.p. parallel shift across a sector, while a relative spread change of 10 means that all spreads in the sector move by 10% of their current values (e.g., from 50 to 55, from 200 to 220).

⁵The eight sectors are: Banking, Finance, Basic Industry, Consumer Cyclical, Consumer Non-Cyclical, Communications, Energy, and Utility. Bonds are assigned to one of three quality cells: Aaa/Aa, A, and Baa.

⁶Large issuers are defined as issuers that have outstanding issues with market value in excess of 1% of the sector aggregate market value. In this case, 17 issuers represent 216 outstanding issues.

⁷Since we compare models with and without an intercept, Exhibit 6 reports uncentered R^2 values calculated using

the total sum of squares (without subtracting the average spread change) rather than a centered R^2 .

⁸We find that the distribution of spread duration varies significantly across time and therefore does not allow for a partition based on constant spread duration values.

⁹Despite our efforts to ensure uniform cell populations, some cells are very sparsely populated (or even empty) in some months. Months when a cell is populated by fewer than 20 bonds are not used in the analysis. As a robustness check, we repeat the analysis using the entire available time series of systematic spread changes and a weighted volatility estimate (where the weighting factor is the number of observations in each month). The results are essentially unchanged.

¹⁰Instead of a single spread coefficient, we also estimate an unconstrained model that allows the spread slope coefficient to vary by sector and duration as follows:

$$\sigma(\Delta s)_{i,d,s} = s_{i,d,s}(\beta + \beta_{Fin} \cdot I_{Fin} + \beta_{Ind} \cdot I_{Ind} + \beta_{Med} \cdot I_{Med} + \beta_{Long} \cdot I_{Long}) + \varepsilon_{i,d,s}$$

where i , d , and s denote the sector-duration-spread combination of each observation. I_{Fin} and I_{Ind} are dummy variables equal to 1 if i = Financials or Industrials, respectively, and zero otherwise. Similarly, I_{Med} and I_{Long} equal 1 if d = medium or long, respectively, and zero otherwise. The results confirm that relative spread volatility is not restricted to a single sector or maturity, and that there is roughly a 9 b.p. per month pickup in volatility for every 100 b.p. increase in spread.

¹¹To be consistent with Equation (6), idiosyncratic spread change should be defined as $\Delta s_{i,d}^{idio} = \Delta s_{i,d} - (s_{i,d}/s_{J,d})\Delta s_{J,d}$. As we are carrying out this test over relatively narrow spread buckets, however, there is very little difference in practice between the two definitions.

¹²Depending on the sample composition and population, this produces between 38 and 66 observations yearly for systematic volatility and 300 to 500 observations for idiosyncratic volatility. As before, only observations that represent buckets populated by at least 20 bonds during the entire year are included in the analysis.

¹³Our findings do not change when we repeat the analysis using other partitions.

¹⁴Although the carry component is time-varying, we analyze each month's excess return conditioned on the beginning-of-month spread. We can therefore treat the carry component as deterministic.

¹⁵A longer estimation period is always desirable as long as the proportionality factor is stable across periods, which we find to be the case.

¹⁶Including publicly issued debt of U.S. government agencies, quasi-federal corporations, and corporate or foreign debt guaranteed by the U.S. government (such as USAID securities).

¹⁷The results are unchanged when we exclude issues with a market value below \$300 million or non-U.S. agencies.

¹⁸When a bank is owned by a holding company, for example, owners of subordinated claims issued by the bank have priority in case of a default over owners of a senior claim issued by the holding company.

¹⁹Excess returns can be adjusted by the same scaling factor (the ratio of DTS of the two portfolios) because they are linearly related to DTS. To implement this in practice, we would need to take into account financing costs. As we do not form a trading strategy, however, but rather examine whether similar DTS portfolios exhibit similar excess return volatilities, we can ignore borrowing costs.

²⁰For example, an investment policy may specify that no more than 1% of the portfolio market value can be invested in securities of any single Baa-rated issuer, no more than 2% in any A-rated issuer, and no more than 4% in any Aa-rated issuer.

REFERENCE

Schönbucher, Philipp. "A Libor Market Model with Default Risk." Working paper, University of Bonn, 1999.

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