

CORPORATE CREDIT LIMITS FOR FIXED INCOME PORTFOLIOS

Miikka Taurén^a and Thomas Philips^b

Fixed income portfolio managers and risk managers constantly grapple with the question of how to size their corporate credit trades. Their task is made more difficult by the fact that corporate credit events are rare, particularly among Investment Grade bonds, and that tail risk is not well captured by most multifactor risk models. In this article, we propose a simple, but effective, method for sizing credit trades based on their spread. In particular, we model the cross-sectional behavior of corporate spreads, estimate the expected shortfall of monthly spread returns, and use our results, along with some observations on the duration of corporate bonds, to derive a simple upper bound on the permissible exposure to any single issuer in a credit portfolio.

The method has been applied successfully to Investment Grade and High Yield fixed income portfolios in both developed and emerging markets, and has proven its worth in daily use by protecting portfolios against disproportionate idiosyncratic losses, while allowing portfolio managers sufficient flexibility to express their investment views with clarity. Its use is not confined to limits on issuers—our method is easily extended to create limits on a portfolio's exposures to individual industries, sectors, countries, and regions.



1 Introduction and overview

Fixed income portfolio managers and their risk managers constantly grapple with the problem

^aSenior Risk Manager, Fischer Francis Trees and Watts, 5 Aldermanbury Square, London, United Kingdom. Tel.: +44 207 063 7770; E-mail: miikka.taurén@bnpparibas.com.

^bGlobal Head of Investment Risk, BNP Paribas Investment Partners, Institutional Division, 200 Park Avenue, New York, NY 10166, USA. Tel.: +1 212 681 3016; E-mail: thomas.philips@fftw.com.

of setting appropriate issuer limits for corporate credit trades. However, in spite of its fundamental importance in the management of credit portfolios, most approaches to setting issuer limits are surprisingly crude. More often than not, issuer limits are based simply on credit ratings, with a maximum exposure being defined for each rating class, e.g., AAA: 5%, AA: 4%, A: 2%, and BBB: 1%. But credit ratings can be inaccurate, and can also be slow to reflect changes in credit quality, as exemplified by the AAA ratings accorded Enron prior to 2001, WorldCom prior

to 2007, and most European sovereigns prior to 2010.

In this article, we describe an alternative approach to setting corporate credit limits that is based on issuer spreads. It is rooted in the widely observed fact that many measures of spread risk, such as spread volatility and the distribution of spread changes, are proportional to spread levels. This observation can be combined with realistic assumptions about the duration of corporate bonds, achievable information ratios, and tolerable levels of loss to derive a useful upper bound on the exposure to a single issuer.

The remainder of this article is organized as follows. We first survey the literature on the empirical behavior of corporate spreads, and present an independent body of evidence that supports the proportional spread change models of Taurén (1999), Ben Dor *et al.* (2007), and Benzschawel and Lee (2011). We then use these results to guide the development of a position sizing methodology and a simple formula for determining the maximum exposure to a single issuer. Next, armed with some assumptions about the duration of corporate bonds and the information ratios of actively managed portfolios, we select appropriate values for certain parameters in the formula. Finally, we summarize our results and present some numerical examples that illustrate the use of our formula in setting credit limits in both Investment Grade and High Yield portfolios.

2 The behavior of corporate spreads

In early work, Taurén (1999) shows that the volatility of the option-adjusted spread (OAS) of U.S. investment grade corporate bonds is proportional to OAS, i.e.,

$$\sigma(s) \approx c \times s \quad (1)$$

for some constant c . If we define r , the *spread return*, to be the ratio of the one-period change in

spread to the initial spread, Equation (1) can be rewritten as

$$\sigma(r) = \sigma\left(\frac{s_{t+1} - s_t}{s_t}\right) \approx c \quad (2)$$

In their seminal paper, Ben Dor *et al.* (2007) show that this result holds both in U.S.A. and in Europe, across industries, rating classes and maturities, and for both bonds and Credit Default Swaps (CDS). In addition, they show that it is robust to the curve against which spreads are measured—in particular, it holds when spreads are measured relative to both the Treasury curve and the Libor curve. More recently, Benzschawel and Lee (2011) and Benzschawel *et al.* (2012) verify this relationship for corporate bonds across the entire gamut of agency ratings (see Figure 20 in both articles), and show that it applies to hard-currency sovereign bonds as well.

A useful rule of thumb emerges from these studies: over long horizons, the annualized volatility of the spread of a corporate bond is approximately 30% of its current level; i.e., if its current spread is 200 basis points (bps), its spread volatility will be about 60 bps per annum. If its spread tightens to 100 bps, its spread volatility will decline to about 30 bps per annum, and if it widens to 400 bps, its spread volatility will correspondingly double to about 120 bps per annum. In quiet periods, the constant of proportionality is somewhat smaller (say 20%), and in turbulent periods it is somewhat higher (say 50%), but over long horizons, it averages out to about 30%. This insight has been incorporated into the risk model embedded in Barclays POINT[®], a commercial risk management application that includes the product of duration and spread (or DTS) as a prominent risk factor. It has also been used by Benzschawel and Lee (2011) and Benzschawel *et al.* (2012) to identify bonds that are trading rich and cheap relative to their true probability of default.

When attempting to apply this result to the setting of corporate credit limits, though, we need to be cognizant of the entire distribution of spread changes, and not just its volatility, as its tail can be very long. To further our understanding of the behavior of spreads, we therefore examine the cross-sectional and time-series properties of monthly spread changes and spread returns for the constituents of two disparate indices—the Barclays U.S. Corporate Index,¹ which covers the

domestic market for corporate bonds, and the J.P. Morgan CEMBI Index,² which is an index of dollar denominated emerging market corporate bonds. Exhibits 1 and 2 display the evolution of the number of issues and issuers contained in these indices with time. The two indices differ sharply not only in their regional exposure and their composition but also in their number of constituents and in the number of issues associated with each issuer.




Exhibit 1 The number of issues and issuers vs. time: J.P. Morgan CEMBI Index.

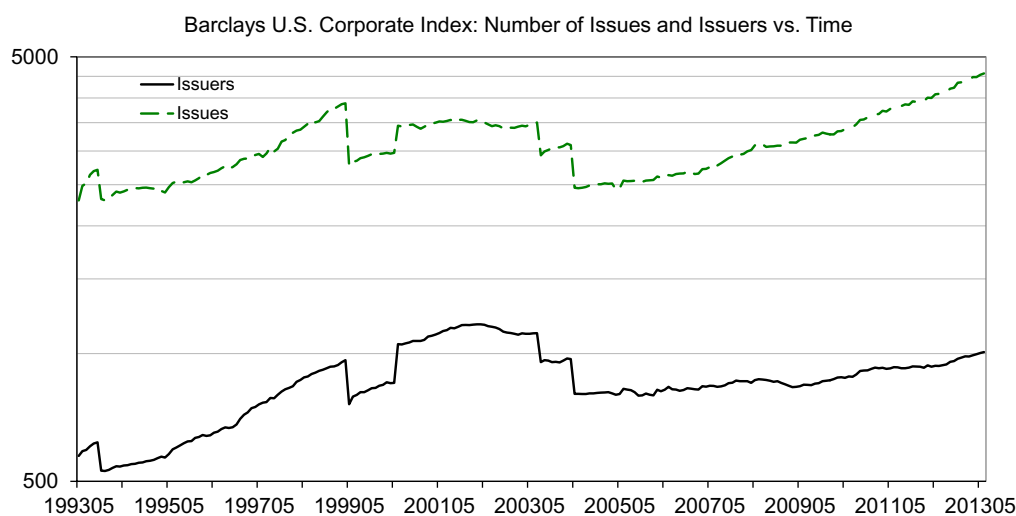


Exhibit 2 The number of issues and issuers vs. time: Barclays U.S. Corporate Index.

Our study is conducted as follows: each month, we obtain the constituents of both indices and their OASs. These OASs are computed by the index provider, and include the impact of their (proprietary) interest rate model and any optionality (call schedule, sinking fund, etc.) associated with each bond. For convenience, we denote the OAS of bond i at the start of month t by s_{it} . For each bond and in each month, we compute the spread return, or the ratio of the change in its OAS during the month to its beginning of month OAS, i.e.,

$$r_{it} = \frac{s_{i(t+1)} - s_{it}}{s_{it}} \quad (3)$$

If a bond has just entered the index, or if no OAS is available for it, it is excluded from this computation. In addition, we exclude bonds with less than 1 year to maturity, as they are excluded from most indices, and therefore from most investors' opportunity set. In spite of the fact these OASs are provided by large and sophisticated index providers with significant analytic and computational resources, and are widely distributed to virtually the entire universe of institutional investors, we identified errors in the data that required cleansing before any analysis could be conducted.

For example, all bonds in the Barclays U.S. Corporate Index have a spread of 0 from 1988 to 1992. These, naturally, are removed, as are bonds with very low (i.e., less than 20 basis points) or negative OAS, bonds whose OAS more than quadruples in a month, and bonds whose OAS more than doubles in one month and then reverts to within 10% or 5 basis points of their initial OAS in the subsequent month. A given bond may be excluded in some months, but not in others.

Given the vast number of bonds and issuers and their heterogeneity, as well as our lack of access to the analytic engines that were used to produce these OASs, these rules must necessarily be

arbitrary, and we cannot guarantee that they will catch all data errors. They were drawn up after extensive exploratory data analysis followed by detailed investigation of a selected subset of outliers in coordination with credit analysts with long investment histories. While we believe that the filters described here catch the most egregious errors, we openly acknowledge our inability to identify and cleanse all errors in our data, and therefore use both classical and outlier-resistant robust estimators when estimating parameters.

Our data for the Barclays U.S. Corporate Bond Index starts in May 1993, while that for the J.P. Morgan CEMBI Index starts in December 2001. Data for both indices ends in July 2013. For each index, and in each month, we first filter the raw data as described above and then compute the cross-sectional volatility and the range of monthly spread returns and monthly spread changes when trimming the data at two percentiles—the 1st and 99th percentile and the 5th and 95th percentile, both by issue and by issuer, with each issue being market value weighted when computing issuer level changes. In addition, we compute the standardized range of monthly spread returns (i.e., the ratio of the range of spread returns in a month to its cross-sectional standard deviation), as well as the expected shortfall of monthly spread returns and monthly spread changes at the same percentiles.³ We consciously do not DV01 weight individual issue OASs when computing issuer OASs as long maturity corporate bonds (say those with 10 years or more to maturity) tend to be less liquid than intermediate maturity corporate bonds (say those with 3–10 years to maturity), and their OASs are correspondingly less reliable.

Our focus on the cross-sectional volatility of spread returns stands in contrast to earlier work, which focuses on the time series variation of individual issuer OASs, typically over 3- and 5-year horizons. We link our results to these

earlier results by assuming that the spread return-generating process is ergodic,⁴ so that its time series properties are equal to its cross-sectional, or ensemble, properties. The validity of this assumption is easily tested by comparing our cross-sectional estimates of volatility to the time series estimates reported by Ben Dor *et al.* (2007), Benzschawel and Lee (2011), and Benzschawel *et al.* (2012). Averaged over issuers and time periods, both time series and cross-sectional estimates of the volatility of spread returns are approximately 10%/month or 30%/annum, which supports our assumption of ergodicity. Our cross-sectional estimation procedure is similar in spirit to that of Adrian (2007), who models the behavior of hedge fund returns using cross-sectional estimates of volatility and instantaneous covariance.

We acknowledge that there is no fundamental reason for the spread return-generating process to be ergodic, just as there is no fundamental reason for the time series volatility of spread returns to be independent of spread levels, and moving between cross-sectional and time series estimates brings with it the possibility of error. In the long run, the spread return-generating process must have a mean of zero—spreads do not grow to the sky or become reliably negative, and their mean

and variance must, therefore, exhibit mean reversion. In times of crisis, though, the distribution of spread returns is both volatile and unpredictable, and the assumption of ergodicity is likely not valid. But over the long run, corporate spreads do appear to be reasonably well modeled by a mean reverting ergodic process. Time variation in the mean is best thought of as the systematic component of spread return, whereas cross-sectional variation around the mean is best thought of as its idiosyncratic component. For reasons we next describe, we deliberately make the simplifying assumption that the spread return-generating process is ergodic with zero mean, and that all its cross-sectional variation is idiosyncratic. Empirically, this is a surprisingly good assumption, and greatly simplifies our analysis.

Our estimates of the mean, the volatility, the range, and the expected shortfall of spread returns are plotted in Exhibits 3–10, and are summarized in the tables that follow. We display information only by issuer, as the graphs are virtually identical when plotted by issue. More often than not, the magnitude of the mean monthly spread return is smaller than its volatility and much smaller still than its cross-sectional range. September 2008 provides a clear exception to this observation, but




Exhibit 3 Cross-sectional mean and volatility of spread returns vs. time: J.P. Morgan CEMBI Index.




Exhibit 4 Cross-sectional mean and volatility of spread returns vs. time: Barclays U.S. Corporate Index.




Exhibit 5 Inter-percentile range of spread returns vs. time: J.P. Morgan CEMBI Index.

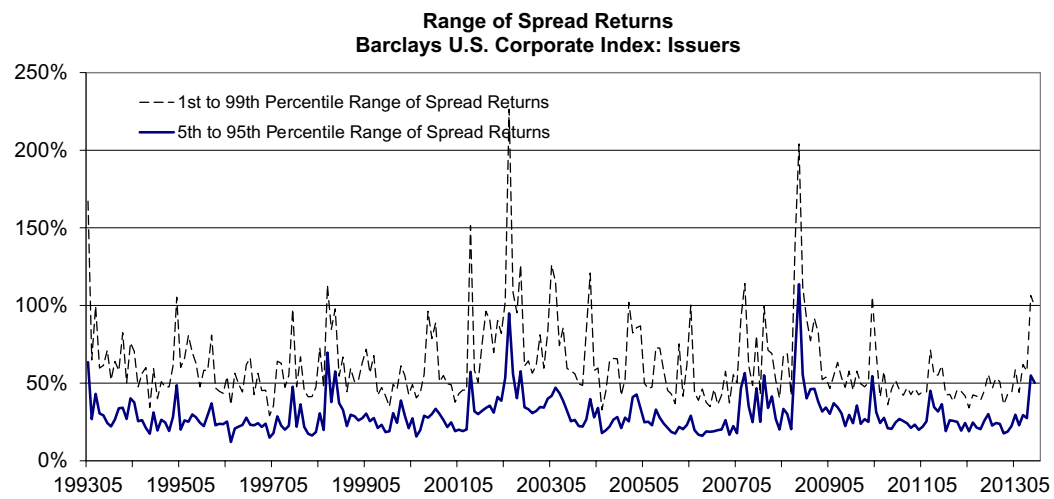


Exhibit 6 Inter-percentile range of spread returns vs. time: Barclays U.S. Corporate Index.




Exhibit 7 Standardized inter-percentile range of spread returns vs. time: J.P. Morgan CEMBI Index.




Exhibit 8 Standardized inter-percentile range of spread returns vs. time: Barclays U.S. Corporate Index.




Exhibit 9 Expected shortfall of spread returns vs. time: J.P. Morgan CEMBI Index.




Exhibit 10 Expected shortfall of spread returns vs. time: Barclays U.S. Corporate Index.

it is clear that the mean does not dominate the cross-sectional volatility of spread returns in most months (see Exhibits 13 and 14 as well). In addition, rapid reversion in both the cross-sectional mean and volatility is particularly evident in the period immediately following the crisis. Finally, in spite of the fact that the two indices are completely disparate in their construction and in the types of bonds they contain, the monthly spread returns of their constituents exhibit very similar patterns of behavior. By way of contrast, we

plot the same inter-percentile range of monthly changes in OAS (the numerator of Equation (3)) for both indices in Exhibits 11 and 12; these are noticeably more variable, differ significantly in magnitude across indices, and, in fact, must be plotted on a logarithmic scale to accommodate the entire range of their time variation.

The information in Exhibits 3 through 10 is summarized in Exhibits 13 and 14. Over the entire period, as can be seen from these tables, the




Exhibit 11 Inter-percentile range of spread changes vs. time: J.P. Morgan CEMBI Index.




Exhibit 12 Inter-percentile range of spread changes vs. time: Barclays U.S. Corporate Index.

Exhibit 13 Summary statistics for monthly issuer spread returns: J.P. Morgan CEMBI Index: Entire period.

Parameter estimated	Conditioning	Average	Hodges–Lehmann average
Cross-sectional monthly mean	Including all observations	1.2%	−0.3%
	Excluding 1% largest and smallest observations	1.2%	−0.3%
	Excluding 5% largest and smallest observations	1.2%	−0.3%
Cross-sectional monthly Std. deviation	Including all observations	12.1%	11.2%
	Excluding 1% largest and smallest observations	10.9%	10.0%
	Excluding 5% largest and smallest observations	7.8%	7.2%
Cross-sectional monthly standardized range	Including all observations	7.9 σ	7.7 σ
	Excluding 1% largest and smallest observations	6.3 σ	6.3 σ
	Excluding 5% largest and smallest observations	4.6 σ	4.5 σ
Monthly expected shortfall	99% ES	43.6%	41.0%
	95% ES	27.7%	25.7%

cross-sectional standard deviations of monthly spread returns are approximately 12%, 10%, and 7% for both indices when untrimmed, trimmed by 1%, and trimmed by 5% on each side, and the corresponding expected shortfalls at the 99% and 95% confidence limits are approximately 45% and 26%, respectively. The 1st to 99th and the 5th to 95th percentile range ranges of spread returns

are approximately 7 and 5 standard deviations wide, respectively. All averages are computed using both the standard estimator and a robust Hodges–Lehmann (1963) estimator.⁵

The long-tailed nature of spread returns is very evident in the fact that trimming only a small percentage of observations (1% or 5%) on each side

Exhibit 14 Summary statistics for monthly issuer spread returns: Barclays U.S. Corporate Index: Entire period.

Parameter estimated	Conditioning	Average	Hodges–Lehmann average
Cross-sectional monthly mean	Including all observations	1.3%	0.2%
	Excluding 1% largest and smallest observations	1.3%	0.2%
	Excluding 5% largest and smallest observations	1.3%	0.2%
Cross-sectional monthly Std. deviation	Including all observations	11.3%	10.6%
	Excluding 1% largest and smallest observations	8.5%	7.9%
	Excluding 5% largest and smallest observations	6.1%	5.7%
Cross-sectional monthly standardized range	Including all observations	14.4 σ	14.1 σ
	Excluding 1% largest and smallest observations	7.3 σ	7.3 σ
	Excluding 5% largest and smallest observations	4.9 σ	4.9 σ
Monthly expected shortfall	99% ES	46.8%	44.5%
	95% ES	25.6%	23.9%

leads to a significant drop in the cross-sectional volatility and the range of spread changes and spread returns. If spread changes and returns were normally distributed, trimming 1% and 5% of the observations at each end of the distribution would lower the estimated volatility only by 1.5% and 11.5%, respectively, and not the 10–40% that we actually observe. Somewhat counter intuitively, it is the constituent of the Barclays U.S. Corporate Index, and not the J.P. Morgan CEMBI Index, that appears to have the longer tails, as exemplified by the particularly sharp decline in their cross-sectional volatility when only 1% of each month's observations are trimmed at both extremes. Some exploratory data analysis of the distribution of spread returns (e.g., histograms, Q–Q plots) suggests that it is unimodal, though with some skew and with exceptionally long tails. It does not seem to be drawn from any well-known family of distribution. In addition, both indices have sufficiently many outliers that parameter estimates from a standard maximum likelihood estimator for, say, a normal distribution, are either unrealistic or unstable, or both.

During the credit crisis (which we proxy by the period August 2008 to March 2009), the cross-sectional standard deviation and expected shortfall of spread *returns* increase by about 80%, but the standardized range changes hardly at all, as can be seen from Exhibits 15 and 16. It is striking that these quantities, as well as the cross-sectional mean, revert to their pre-crisis pattern of behavior in very short order. By way of comparison, it can be seen from Exhibits 17 through 20 that the expected shortfall and the cross-sectional standard deviation of spread *changes* increase dramatically: they increase by a factor of seven for the Barclays U.S. Corporate Index and by a factor of three for the J.P. Morgan CEMBI Index. This is driven in part by the fact that the U.S. Corporate Index covers a longer period of time, during which credit spreads were lower and more stable in U.S.A. than in emerging markets, but also by the fact that the 2008 credit crisis affected developed markets far more than it did emerging markets.

The behavior of spread returns for the two indices is surprisingly similar and stable over time, and

Exhibit 15 Summary statistics for monthly issuer spread returns: J.P. Morgan CEMBI Index: 8/2008–3/2009.

Parameter estimated	Conditioning	Average	Hodges–Lehmann average
Cross-sectional monthly mean	Including all observations	8.7%	0.1%
	Excluding 1% largest and smallest observations	8.7%	0.1%
	Excluding 5% largest and smallest observations	8.7%	0.1%
Cross-sectional monthly Std. deviation	Including all observations	21.2%	19.4%
	Excluding 1% largest and smallest observations	19.4%	16.8%
	Excluding 5% largest and smallest observations	13.5%	10.8%
Cross-sectional monthly standardized range	Including all observations	7.2 σ	7.3 σ
	Excluding 1% largest and smallest observations	6.5 σ	6.5 σ
	Excluding 5% largest and smallest observations	4.6 σ	4.5 σ
Monthly expected shortfall	99% ES	72.7%	73.2%
	95% ES	49.0%	44.5%

Exhibit 16 Summary statistics for monthly issuer spread returns: Barclays U.S. Corporate Index: 8/2008–3/2009.

Parameter estimated	Conditioning	Average	Hodges–Lehmann average
Cross-sectional monthly mean	Including all observations	9.8%	8.0%
	Excluding 1% largest and smallest observations	9.8%	8.0%
	Excluding 5% largest and smallest observations	9.8%	8.0%
Cross-sectional monthly Std. deviation	Including all observations	20.4%	18.6%
	Excluding 1% largest and smallest observations	16.0%	14.4%
	Excluding 5% largest and smallest observations	11.9%	10.6%
Cross-sectional monthly standardized range	Including all observations	12.6 σ	12.9 σ
	Excluding 1% largest and smallest observations	6.7 σ	6.7 σ
	Excluding 5% largest and smallest observations	4.6 σ	4.6 σ
Monthly expected shortfall	99% ES	79.3%	75.7%
	95% ES	44.9%	41.08%

stands in sharp contrast to the very obvious differences in their composition, size, and spread levels, suggesting that the relationship between spread levels and spread risk (as measured by the volatility of spread returns) is universal,

and proves supporting evidence, but no causal explanation, for our assumption that the spread return-generating process is ergodic. That said, the lack of a causal explanation for the behavior of spread returns does not preclude the use of

Exhibit 17 Summary statistics for monthly issuer spread changes: J.P. Morgan CEMBI Index: Entire period.

Parameter estimated	Conditioning	Average	Hodges–Lehmann average
Cross-sectional monthly mean	Including all observations	5.1 bp	–2.3 bp
	Excluding 1% largest and smallest observations	5.1 bp	–2.3 bp
	Excluding 5% largest and smallest observations	5.1 bp	–2.3 bp
Cross-sectional monthly Std. deviation	Including all observations	131.3 bp	111.4 bp
	Excluding 1% largest and smallest observations	105.0 bp	80.4 bp
	Excluding 5% largest and smallest observations	52.0 bp	33.6 bp
Cross-sectional monthly standardized range	Including all observations	9.7 σ	9.7 σ
	Excluding 1% largest and smallest observations	7.9 σ	7.7 σ
	Excluding 5% largest and smallest observations	5.2 σ	5.1 σ
Monthly expected shortfall	99% ES	587.3 bp	458.6 bp
	95% ES	300.1 bp	237.7 bp

Exhibit 18 Summary statistics for monthly issuer spread changes: Barclays U.S. Corporate Index: Entire period.

Parameter estimated	Conditioning	Average	Hodges–Lehmann average
Cross-sectional monthly mean	Including all observations	1.2 bp	–0.3 bp
	Excluding 1% largest and smallest observations	1.2 bp	–0.3 bp
	Excluding 5% largest and smallest observations	1.2 bp	–0.3 bp
Cross-sectional monthly Std. deviation	Including all observations	27.9 bp	17.9 bp
	Excluding 1% largest and smallest observations	18.0 bp	12.5 bp
	Excluding 5% largest and smallest observations	11.1 bp	8.3 bp
Cross-sectional monthly standardized range	Including all observations	16.9 σ	16.5 σ
	Excluding 1% largest and smallest observations	8.3 σ	8.2 σ
	Excluding 5% largest and smallest observations	5.2 σ	5.2 σ
Monthly expected shortfall	99% ES	127.2 bp	78.8 bp
	95% ES	60.5 bp	40.1 bp

the proportional spread model in a variety of risk management applications.

Ben Dor *et al.* (2007) make much the same point: they present strong empirical evidence,

but no theoretical model, for the approximately constant time series volatility of spread returns across a wide range of fixed income asset classes, but do not let this deter them from putting their observation to good use in modeling the returns

Exhibit 19 Summary statistics for monthly issuer spread changes: J.P. Morgan CEMBI Index: 8/2008–3/2009.

Parameter estimated	Conditioning	Average	Hodges–Lehmann average
Cross-sectional monthly mean	Including all observations	69.5 bp	8.3 bp
	Excluding 1% largest and smallest observations	69.5 bp	8.3 bp
	Excluding 5% largest and smallest observations	69.5 bp	8.3 bp
Cross-sectional monthly Std. deviation	Including all observations	414.7 bp	449.8 bp
	Excluding 1% largest and smallest observations	347.3 bp	379.6 bp
	Excluding 5% largest and smallest observations	201.7 bp	194.7 bp
Cross-sectional monthly standardized range	Including all observations	9.7 σ	9.7 σ
	Excluding 1% largest and smallest observations	8.2 σ	8.2 σ
	Excluding 5% largest and smallest observations	5.1 σ	5.1 σ
Monthly expected shortfall	99% ES	1908.3 bp	1922.4 bp
	95% ES	986.8 bp	1056.6 bp

Exhibit 20 Summary statistics for monthly issuer spread changes: Barclays U.S. Corporate Index: 8/2008–3/2009.

Parameter estimated	Conditioning	Average	Hodges–Lehmann average
Cross-sectional monthly mean	Including all observations	37.8 bp	38.8 bp
	Excluding 1% largest and smallest observations	37.8 bp	38.8 bp
	Excluding 5% largest and smallest observations	37.8 bp	38.8 bp
Cross-sectional monthly Std. deviation	Including all observations	210.6 bp	212.3 bp
	Excluding 1% largest and smallest observations	126.8 bp	126.5 bp
	Excluding 5% largest and smallest observations	67.7 bp	65.6 bp
Cross-sectional monthly standardized range	Including all observations	17.4 σ	17.4 σ
	Excluding 1% largest and smallest observations	9.7 σ	9.6 σ
	Excluding 5% largest and smallest observations	5.6 σ	5.6 σ
Monthly expected shortfall	99% ES	1022.6 bp	1026.5 bp
	95% ES	453.6 bp	452.2 bp

of corporate bonds. In a similar spirit, we derive a simple but effective rule for setting corporate credit limits based on the robustness of the expected shortfall of spread returns across long periods of time and a wide range of corporate bonds.

3 Application to corporate credit limits

We propose to set limits on the sizes of individual issuers in a portfolio in accordance with the following principle: *Limit the allocation to any single issuer so that, in the course of a month,*

when its OAS return tightens or widens by the 1st or 99th percentile Expected Shortfall, the impact on portfolio performance is not disproportionately large relative to the annual return target of the portfolio.

This principle is broadly applicable—it is as valid for a long only portfolio that is measured against a benchmark (in which case it is the relative allocation to an issuer that must be bounded) as it is for a trading book (in which case it is the absolute allocation to an issuer that must be bounded).

From Exhibits 13 and 14, we see that the 1st or 99th percentile Expected Shortfalls correspond to proportional spread tightenings and widenings of a little over 40% for both the Barclays U.S. Corporate Index and the J.P. Morgan CEEMBI Index. For convenience and computational tractability, we have settled on 40% as a uniform measure of expected shortfall across all credit indices at a 99% confidence level. While we express our rule in terms of the cross-sectional expected shortfall of spread returns, it is just as valid to express it in terms of the time series volatility of spread returns. Expressed in the language of Ben Dor *et al.* (2007), it is a 4 sigma rule—*size trades so that a 4 sigma move in spreads does not induce a disproportionately large loss in a portfolio.*

While there is no universally right answer to the question of what constitutes a large loss for a portfolio, we suggest that the threshold be set at 10% of its annual excess return target. This is not so large that a portfolio becomes unsalvageable in the event of a loss of this magnitude, and not so small that it limits a portfolio manager's ability to express his or her views. If no excess return target is available, but a tracking error target is, we assume that the portfolio has an Information Ratio of 0.5, so that its annual excess return target is half its annualized tracking error, or the standard deviation of its benchmark relative excess returns.

The return of a bond to a 40% change in its spread is easily computed. If we ignore the impact of convexity, the return of a bond is well approximated by the product of its effective duration and its change in yield, which we proxy by its change in OAS. If we denote the active weight (i.e., the difference between the portfolio weight and the benchmark weight) of the i th issuer at time t by w_{it} , and the weighted average duration of this issuer's bonds by D_{it} , we can express our credit limit as:

$$\begin{aligned} |w_{it}| \times D_{it} \times s_{it} \times 0.4 \\ \leq 0.10 \times \text{Target Excess Return} \\ = 0.05 \times \text{Target Tracking Error.} \end{aligned} \quad (4)$$

This can be further simplified by observing that much corporate issuance is done around the 5-year tenor, and so we set $D_{it} = 5$ for all issuers at all times in Equation (4). By doing so, we reap an additional benefit: it prevents the buildup of large short duration positions that can severely impact performance in the event of a default. If the portfolio's duration is significantly greater than 5 years (as is, for example, the case for portfolios that are designed to immunize pension liabilities), we suggest setting D_{it} to the average duration of all corporate bonds that are eligible for inclusion in the portfolio. Equation (4) is easily rearranged to read

$$\begin{aligned} |w_{it}| \leq \frac{0.050}{s_{it}} \times \text{Target Excess Return} \\ = \frac{0.025}{s_{it}} \times \text{Target Tracking Error.} \end{aligned} \quad (5)$$

With a little final polishing to accommodate quoting conventions, the formula becomes exceptionally intuitive. Issuer weights, Target returns, and Target tracking errors are typically quoted in percentage points (e.g., an active weight of 1%). Spreads, on the other hand, are typically quoted

in basis points. If we multiply the numerator by 10,000, Equation (5) simplifies to

$$\begin{aligned} |w_{it}(\%)| &\leq \frac{500}{s_{it}(bp)} \\ &\quad \times \text{Target Excess Return } (\%) \\ &= \frac{250}{s_{it}(bp)} \\ &\quad \times \text{Target Tracking Error } (\%). \quad (6) \end{aligned}$$

In practice, this formula works well, provided that it is not applied to bonds are not deeply distressed.⁶ The limits it suggests make intuitive sense to both portfolio managers and risk managers, and have been used to good effect in a wide range of portfolios, both in developed and in emerging markets. The fact that the limit can be computed on the back of an envelope without any need for a specialized calculator makes it particularly useful for portfolio managers, who can apply it when they are executing trades with recourse only to simple mental arithmetic.

Nonetheless, there are a few loose ends that need tying up. First, as spreads expand and compress, the bound moves. This can induce unwarranted trading if a portfolio manager seeks the fullest possible expression of his insights and initiates positions at their maximum allowable size. Our pragmatic solution to this problem is to grandfather positions and limits. Once a position is put on, it need not be taken down if its spread expands, and can be taken up only as far as the initial limit if they compress, with pragmatic exceptions being made for significant spread and ratings changes, such as a transition from Investment Grade to High Yield, or vice versa.

Next, as spreads tighten, the upper bound on exposure increases without limit. This is misguided: bonds have default risk as well as market

risk, and, in addition, can drop sharply in value in the event of an LBO bid. It is therefore wise to place an absolute limit on the exposure to any issuer. For Investment Grade portfolios, whose Target Tracking Errors are typically on the order of 1%, we suggest that this limit be 1% for issuers rated BBB/Baa2 or lower, and 5% for all other issuers regardless of their rating or their spread. We set a separate limit for issuers near the boundary between Investment Grade and High Yield because when issuers get downgraded and cross this boundary, their bonds fall out of many indices, and forced selling by investors with ratings constraints can easily depress the price of a bond by 5% and render it far less liquid. This limit is of greatest value when spreads are tight and investors are reaching for yield.

Traditionally, many Investment Grade investors and portfolio managers have used bond ratings to guide their sizing decision, in spite of the fact that rating agencies have proven more reactive than predictive when companies undergo change. Even so, a simple ratings based limit can capture a significant portion of the wisdom embedded in our spread based limit, and our recommended ratings based limits follow. The maximum allowable exposure to an issuer is 5%, dropping by 0.5% for each one-notch drop in rating except at the very bottom, where it levels out.

$$|w_{it}(\%)| \leq \begin{bmatrix} \text{AAA} : 5\% \\ \text{AA+} : 4.5\% \\ \text{AA} : 4\% \\ \text{AA-} : 3.5\% \\ \text{A+} : 3\% \\ \text{A} : 2.5\% \\ \text{A-} : 2\% \\ \text{BBB+} : 1.5\% \\ \text{BBB} : 1\% \\ \text{BBB-} : 1\% \end{bmatrix}. \quad (7)$$

Exhibit 21 Issuer limits based on Equation (7).

Bond	Rating (S&P/ Moody's/Fitch)	Target annual TE (%)	OAS (bp)	5-yr. CDS spread (bp)	Issuer limit based on OAS (%)	Issuer limit based on 5-yr CDS (%)
MSFT 4% 2/8/2021	AAA/Aaa/AA+	2	68	36	5	5
KMI 5% 2/15/2021	BB/Ba2/BB+	2	306	172	1.6	2
MSFT 4% 2/8/2021	AAA/Aaa/AA+	4	68	36	5	5
KMI 5% 2/15/2021	BB/Ba2/BB+	4	306	172	3.3	4

The final limit for Investment Grade portfolios is the tighter of these two limits, i.e.,

$$\begin{aligned}
 |w_{it}(\%)| &\leq \min \left(\text{Ratings-Based Limit}, \frac{500}{s_{it}(bp)} \right. \\
 &\quad \left. \times \text{Target Excess Return}(\%) \right) \\
 &= \min \left(\text{Ratings-Based Limit}, \frac{250}{s_{it}(bp)} \right. \\
 &\quad \left. \times \text{Target Tracking Error}(\%) \right). \tag{8}
 \end{aligned}$$

High Yield investors and portfolio managers tend to be more attuned to risk limits than to ratings-based limits, and in this case, we suggest that the auxiliary limit be the target tracking error, or twice the target excess return. As in the case of Investment Grade portfolios, the limit used in practice is the tighter of this and the spread-based limit, so that the final limit for High Yield portfolios becomes

$$\begin{aligned}
 |w_{it}(\%)| &\leq \min \left(2, \frac{500}{s_{it}(bp)} \right) \\
 &\quad \times \text{Target Excess Return}(\%) \\
 &= \min \left(1, \frac{250}{s_{it}(bp)} \right) \\
 &\quad \times \text{Target Tracking Error}(\%). \tag{9}
 \end{aligned}$$

The auxiliary bound is equivalent to a lower bound of 250 bp on spreads, which is realistic for the High Yield market. While our limits are

described in terms of OAS, there is much to recommend the use of 5-year CDS spreads when they are available. CDS tend to be more liquid than bonds, and their spreads are quoted directly, obviating the need to compute the OAS of a bond. That said, it is not always possible to buy or sell protection on all levels of the capital structure of the firm, necessitating care with the substitution. This is particularly true of financials. As an aside, we note that sovereign CDS spreads prove useful when setting limits for local currency bonds in emerging markets, as the notion of a spread to a nearly risk-free asset is far more nebulous for such bonds. Exhibit 21 illustrates the credit limits suggested by Equations (8) and (9) for two issuers, one Investment Grade and the other High Yield, and with two very different levels of target tracking error.

4 Conclusion and summary

We have proposed a simple credit spread-based methodology for sizing corporate credit positions. These credit limits have been in use for close to 3 years across a wide range of portfolios in both developed and emerging markets, and have proved their worth by enabling a sensible level of diversification, neither underdiversifying and concentrating risk in a few positions, nor overdiversifying and creating unmanageably large numbers of line items in portfolios. By changing the constant of proportionality, they are easily extended to apply to sectors, countries, and

industries, and will prove as useful in controlling these broader exposures as they have proven in controlling issuer risk.

Acknowledgments

The authors would like to thank Antonio Silva of Barclays Capital, as well as the (anonymous) referee for their many constructive comments, which greatly improved both the content and the presentation of this paper. Any errors that remain are our own.

Notes

- ¹ All data for the Barclays U.S. Corporate Index were downloaded from the Barclays Live Web site (live.barcap.com) using their POINT® application.
- ² All data for the J.P. Morgan CEMBI Index were downloaded from the MorganMarkets® Web site (jpm.com).
- ³ We define the expected shortfall to be the average of the absolute value of the spread return (or spread) conditioned on it falling outside a specified inter-percentile range. Our definition is identical to that of Acerbi and Tasche (2002) for arbitrary random variables, except that the classical definition includes only losses, or exceedences on the left tail of the distribution, while we consciously ignore any asymmetry in the distribution of spread returns, and include exceedences on both tails. It is vital in our application that both tails be included, as the portfolio may be either long or short credit. The left tail of the distribution of spread returns has a negative impact on the portfolio if it is short credit, whereas the right tail has a negative impact if it is long credit.
- ⁴ A random process is said to be ergodic if every member of the process carries with it the complete statistics of the entire process, or for which the time average of one sequence of events is the same as the ensemble, or cross-sectional, average. Ergodicity is a subtle concept: All ergodic processes, for example, are stationary, but not all stationary processes are ergodic. The spread return-generating process is perhaps best thought of as a wide-sense (or covariance) stationary ergodic process that exhibits mean-reversion and not as a strict-sense stationary process.
- ⁵ The Hodges–Lehmann estimate of the mean of a sequence $\{x_i\}$ is given by $\text{median}\{(x_i + x_j)/2\}$. It is unbiased and translation invariant, is robust to outliers (it

has a breakdown point of 29%) and has high efficiency for symmetric distributions (0.955 for Normal random variables).

- ⁶ For bonds that are deeply distressed (bonds that are trading at a price lower than, say, 80), and which are consequently have very high yields and OASs, our simple duration approximation is inaccurate, and, in fact, is too conservative. The impact of convexity must be taken into account when modeling the relationship between OAS and price. While this is easily done, we have deliberately refrained from doing so: default risk becomes elevated as the price of a bond drops away from par, and our formula for computing credit limits loses its appealing back-of-the-envelope simplicity. Worse still, as default nears, bonds can suddenly become illiquid, and their bid-offer spread can widen by an order of magnitude. The statistics of spread returns utterly fail to capture these aspects of bond risk, and experienced portfolio managers tread lightly when trading distressed debt.

References

- Acerbi, C. and Tasche, D. (2002). “Expected Shortfall: A Natural Coherent Alternative to Value at Risk,” *Economic Notes* **31**, 379–388.
- Adrian, T. (2007). “Measuring Risk in the Hedge Fund Sector,” *Current Issues in Economics and Finance, Federal Reserve Bank of New York* **13**, 1–7.
- Ben Dor, A., Dynkin, L., Hyman, J., Houweling, P., van Leeuwen, E., and Penninga O. “DTS (Duration Times Spread) A New Measure of Spread Exposure in Credit Portfolios,” *Journal of Portfolio Management* **3**, 77–100.
- Benzschawel T. and Lee, C.-Y. (2011). “Market Implied Default Probabilities: Update,” *Fixed Income Research, Citigroup*, September 9.
- Benzschawel, T., Lee, C.-Y., Hawker, B., and Craft, D. (2012), “Market Implied Default Probabilities and Credit Indexes,” *Journal of Indexes* **62**, 24–31.
- Hodges, J. L. and Lehmann, E. L. (1963). “Estimation of Location Based on Ranks,” *Annals of Mathematical Statistics* **34**, 598–611.
- Taurén, M. (1999). “The Pricing of Credit Risk: Theory and Evidence.” Ph.D. thesis, Indiana University, Bloomington.

Keywords: Fixed income; portfolio management; credit limits; position limits; trade sizing; position sizing; acceptable loss; credit risk management