RiskMetrics Group Research Monthly

Spread values

Christopher C. Finger
chris.finger@riskmetrics.com

November 2005

In a recent issue of *The Economist*, the Charlemagne column discusses the variety of social models that exist in Europe. By way of a clever analogy to national beverages, the column illustrates that European social welfare is not one type of system – as it is often viewed, if not caricatured – but at least four. Crucially, the different systems are a result, among other things, of what policy makers in each state value, and of how those values influence the decisions that make policy.

In this month's column, we take up the issue of credit spreads, and build on recent work explaining and straightening out the myriad of notions of spreads that exist. As with any topic with many definitions, it is tempting to try to conclude which definition is best. We argue that such an attempt is misguided here, and that it is best to invoke a bit of relativism in our discussion. Investors have settled on spread conventions because those conventions give investors information that they value. Risk managers, who may value different properties, may rightly settle on different, or at least additional, spreads. Values matter.

O'Kane and Sen (2005) provide an excellent survey of the notions of credit spreads that market participants use. Beyond the spread definitions, which we will recap shortly, the authors comment on the interpretation and usefulness of each. As interesting as the interpretations themselves are the questions that the authors use to evaluate the various spread definitions. Specifically, the authors expect that a good notion of spread should achieve at least one of the following:

- *Relative value*. Differentiate cheap from expensive securities within a broad pool.
- *Actual return*. Communicate an achievable return to an investor for bearing credit risk.

From a trading point of view, it is difficult to argue with either of O'Kane and Sen's requirements. However, from a risk point of view, there are other contributions we require from a spread notion. More specifically, we care not just about the information contained in a single day's spread observation, but also about the properties of these spreads as they move with time. What we value is different, and so our conclusions as to the most appropriate spread notions may be as well.

Spreads and more spreads

O'Kane and Sen begin with two spread notions that rely on the yield-to-maturity (YTM) of a corporate bond: the *yield spread* is defined as the difference between the YTM on the corporate bond and the YTM on an actual (benchmark) Treasury bond of comparable maturity; and the *interpolated spread* is the difference between the YTM of the corporate and the YTM for a hypothetical Treasury bond with the same maturity as the corporate. Though simple to define, and for that reason widely used, these two measures inherit many of the well known deficiencies of the YTM measure.

The YTM is the single rate that, when used to discount all of the bonds cashflows, successfully reprices the bond; it does not account for the actual term structure of interest rates. For this reason, the first two spread measures are at best crude indicators of relative value, and neither effectively differentiates bonds of different maturities. As an indicator of an actual return, the yield spread is slightly more attractive since it measures return against an actual bond; still, both spread notions are lacking here, since the YTM itself measure requires that we be able to lock in future reinvestment rates for the return to be achievable.

An improved bond spread measure is the *option adjusted spread (OAS)*.[1] As O'Kane and Sen point out, the name OAS is a bit misleading, since we apply it here to bonds with no optionality. So the name is somewhat of an odd convention, but the spread notion is well founded: the OAS for a bond is the constant spread we must add to the discount curve in order to recover the bond's price. Though a bit more complex than the simple yield measure, OAS does account for the term structure of interest rates, and as such is a useful relative value measure across bonds of differing maturities.

In all of the cases mentioned so far, the spread is a comparison of a corporate bond price to some base (or risk-free) curve. The choice of base curve is in a strict sense arbitrary, but in practice, is driven by what comparisons are appropriate for the user of the particular spread. The first set of spread measures are typically expressed versus yields on Treasury bonds; this owes to the traditional use of these measures by long-only bond investors, whose performance is benchmarked to Treasuries. In contrast, the OAS measure is more commonly expressed with respect to the Libor curve, as this is a better representation of the return (or cost of funding) that a bond trader can expect on cash positions.

A last spread measure from the O'Kane and Sen survey is the *credit default swap (CDS) spread*. As the authors point out, CDS spreads are not really a spread over anything, but rather the fair market cost of default protection on a class of debt of a given issuer. As a measure of the credit riskiness of an issuer, the CDS spread is cleaner than any of the various bond-related spread measures. Since the CDS spread is an actual price, it has no explicit base curve; however, implicitly, the base curve is best thought of as Libor, since this represents the cost to a seller of default protection of funding the potential protection payout. Last, it is important to note that the CDS and bond markets are distinct, and as such are subject to a different set of technical factors; while the bond and CDS spreads may represent credit risk for the same issuer, they could well diverge due to differences in legal documentation,[2] or to different supply and demand effects.

As the bond and CDS markets represent mostly the same risks, but can provide slightly different returns, it is desirable to assess relative value across the markets. This represents a significant challenge, as most of the traditional spread measures were conceived to assess relative value only among bonds. The most

[1] The OAS in this sense is also referred to as the Z-spread.

[2] For instance, CDS contracts can differ with respect to their definitions of the default event, as well as their definition of the bonds that they reference.

common method to assess the relative value of a bond and CDS is to compare the bond OAS to the fair CDS spread. If the CDS spread is greater than the bond OAS, there is said to be a positive basis, and the conclusion is that credit risk in the form of the CDS is cheaper than in the form of the bond: if we own a bond, we should consider selling it and taking the same risk through the CDS.

Unfortunately, this standard definition of the bond-CDS basis does not account for differences in convention[3] across the two products. More importantly, the family of spreads related to bonds are essentially different ways of discounting future (risky) cashflows, while the CDS spread is an actual price. When we write down mathematical formulas, the bond spreads appear in the denominators of our expressions (where the discount rates go), while the CDS spreads appear in the numerator (where the actual sizes of cashflows go). So in at least two senses, OAS and CDS spreads are different animals, and we should take care to properly interpret their differences. We will return to this point later.

Risk values

How should we evaluate these spread notions for our purposes? Naturally, we should begin by attempting to enumerate what it is that risk managers actually do with spreads. We could start with the many modeling challenges, but it is crucial to state first that risk managers must communicate with the trading operation. Thus, for all of the desirable properties that an alternate notion of spreads may display, there is a disincentive to adopting such a notion if it differs from those most familiar to traders. We want to speak the language of the market, so that we represent, communicate, and manage risk in a way that is consistent with how the institution actually *takes* risk.

We proceed by moving through the different characterizations of risk that a risk manager might produce. First among these are sensitivity measures: the effect of a small movement in spreads on the present value of our positions. Of the information offered by a risk system, sensitivity measures are often the first examined by traders, meaning that our first point is applicable here. Rightly or no, we may find that our risk system's sophisticated risk forecasting models hold no credibility if it cannot also produce sensitivity measures consistent with the market. Here, then, convention dominates over any advantages with respect to modeling. Further, the interpretation of sensitivities is typically at the security level; there is less focus on rolling these figures up to a portfolio level, and so consistency across spread definitions is subordinate to market convention.

Our second broad risk characterization takes the form of stress tests, or scenario analysis. While at first glance, we might make the same conclusions as with sensitivities, there may be in fact room for modeling concerns. In this context, the most common definition of a stress test is a question of the form "what happens if all spreads widen by fifty basis points?" Simple enough, but there is more to this question than meets the eye. We could produce a straightforward answer by choosing the conventional spread notion for each market, and interpreting the question as a request to widen this by the same amount for all positions (both bond and CDS) and all issuers. If this is in fact the intent of the question, then we have done

[3] Credit default swaps effectively always are entered at par, while bonds can trade at a premium or discount. Further, bond and CDS vary in coupon frequency and daycount definitions.

Table 1: Volatility (bp) of daily spread changes. October, 2004 through October, 2005

	Interp. spread		OAS		CDS
Bond	Treas.	Swap	Treas.	Swap	
DCX 8.000 2010-06-15	6.45	6.50	6.48	6.53	4.52
DCX 8.500 2031-01-18	6.93	6.99	7.14	7.21	5.63
DOW 5.000 2007-11-15	2.31	1.91	2.31	1.92	0.80
DOW 5.970 2009-01-15	1.81	1.66	1.82	1.67	0.89
DOW 6.125 2011-02-01	2.79	2.85	2.81	2.87	1.07
DOW 7.375 2029-11-01	2.41	2.47	2.48	2.55	2.48
SRJ 7.000 2007-06-15	6.36	6.36	6.37	6.37	1.47
SRJ 6.250 2009-05-01	5.40	5.42	5.42	5.44	1.79

our job. But what if we ask which notion of spread to use, and receive a response similar to “I don’t care, just widen them."

Knowing that this parallel spread stress test is the most commonly employed, we should ask if any of the notions of spreads more typically moves in this fashion. A simple exercise is to examine the historical volatilities of spread moves for a sample set of bonds; spread volatilities that are insensitive to maturity would indicate that parallel moves are more likely.

In Table 1, we see that the spread volatilities actually increase (though weakly) with maturity; moreover, this property holds true equally for all of the spread notions. For this stress test then, we might conclude that our choice of spread definition is not critical, but that stress tests which apply greater shifts to longer bonds might be more relevant. This is a question we should ask about more issuers, however.

Risk measures

We come finally to what is often the first consideration of the risk manager: the portfolio risk forecasts. In this case, there are really two separate modeling challenges: one is to forecast the overall portfolio risk; the second is to decompose the risk, and to estimate that portion due purely to spread movements. These two challenges bring to light a new set of considerations that were not relevant for purposes of sensitivity and stress analysis.

One such consideration is the statistical properties of the spreads that we choose. For both the sensitivity and stress analysis, we use only the current price information and a pricing model; there is no dependence on historical data. As we move to forecasting, we become reliant on historical data, and are thus concerned that the data we choose have statistical properties that are amenable to forecasting. An obvious casualty of this consideration is the yield spread. From time to time, the benchmark Treasury bond for a given corporate bond will change, introducing an

Table 2: Correlation (%) of daily spread to daily base rate changes. October, 2004 through October, 2005

	OAS		CDS
Bond	Treas.	Swap	Swap
DCX 8.000 2010-06-15	-21.5	-24.9	-11.1
DCX 8.500 2031-01-18	-16.9	-21.3	-7.6
DOW 5.000 2007-11-15	-20.4	-8.9	-9.1
DOW 5.970 2009-01-15	-19.8	-23.3	-8.2
DOW 6.125 2011-02-01	-31.2	-40.9	-5.9
DOW 7.375 2029-11-01	-20.5	-34.2	-3.3
SRJ 7.000 2007-06-15	-4.2	-1.1	4.4
SRJ 6.250 2009-05-01	-13.5	-16.1	3.2

artificial jump in the yield spread time series; it would be incorrect to interpret such a jump as an indication of greater risk.

A less obvious issue with the yield spread is that its time series are not homogeneous. In other words, it would be erroneous to assume that all historical changes in a specific yield spread are drawn from the same distribution. As suggested by Table 1, spread volatility is often greater for longer bonds; the time series for any specific bond spread contains information from when it had a longer maturity than it does today, and therefore was characterized by greater spread volatility. Thus, rather than analyzing spreads on individual bonds, it is likely a better choice to create spread curves on each day, and to apply our statistics on data for a constant maturity point.

A second desirable statistical property derives specifically from our goal of decomposing the total portfolio risk. We define pure interest rate risk as the risk we see by holding credit spreads constant while allowing the base interest rate curve (or curves) to evolve according to our risk forecast. Likewise, we define spread risk as the risk we see by leaving the base curves constant, but allowing spreads to move according to our forecast. Technically, these definitions do not rely whatsoever on the statistical properties of our spreads, but the usefulness of the risk decomposition does. If spreads are strongly correlated to the base curve, the distinction between interest rate and spread risk is meaningless; the two types of risk are really the same, and one is not likely to occur without the other. What we would prefer, therefore, is a decomposition of risk across factors that are close to independent, where knowing that interest rates move against us does not give us any information about what to expect from spreads. To the extent possible, then, for risk decomposition, we should seek out notions of spreads that are least correlated with base interest rate curves.[4]

We present spread-to-base rate correlations for a

[4]That spreads be roughly independent of the base rate is also desirable for stress testing, as it makes the scenario of the base rate moving and the spread staying constant (or vice versa) more realistic.

number of bonds in Table 2. Interestingly, the correlations of OAS to the base rate are moderately negative. This is actually a common finding, indicating that bond yields can be "sticky", that is, that they react slowly when interest rates move. Unfortunately, it is difficult to ascertain whether such an effect is real, or is an artifact of the illiquidity of the corporate bond market. Spreads on CDS exhibit little correlation with base rates, making them potentially attractive as a mechanism to decompose risk. These spreads do not relate explicitly to bonds, though, so there is a tradeoff to consider; we will return to this point in the next section.

Building bridges

Beyond the pure statistical properties of the time series we choose to analyze, we must also consider the more subtle question of how many series we should analyze in the first place, and particularly whether the different markets for credit (chiefly, bonds and default swaps) constitute two or one sources of information.

This brings us back to an issue we raised earlier – what is the appropriate bond spread to which to compare CDS spreads – albeit this time from a risk perspective. As we mentioned before, while OAS and CDS spreads quantify the same fundamental concept – the amount of compensation an investor receives for bearing credit risk – there are technical differences between the two spread notions that keep them from being strictly comparable. This means that we cannot view the traditional spread basis as a strict arbitrage indicator, nor can we view volatilities of these series in the same light. What both the trader and risk manager need is a bridge.

It is tempting to extend the OAS framework, and use the information contained in this spread to price CDS products. The idea, loosely, is that the OAS added to the base curve gives us a discount curve for cashflows promised by the issuer. Having the issuer (risky) discount curve is not enough to price the CDS, however, since we quickly run into the problem of the relationship between credit and interest rates. More specifically, a default swap added to a bond position provides a guarantee that at default, we receive par for a fixed-coupon bond, regardless of the level of interest rates. To replicate such a profile without a default swap, we need an interest rate swap which allows us to swap the fixed bond coupons for a floating rate, but which also cancels in the event of a default. Without the price of such a product,[5] we have no way of pricing a CDS precisely given bond pricing information. This leaves us with some approximate bridges, but nothing that uses the OAS framework to rigorously infer a CDS spread from a bond price.

A better bridge, then, is one that models the credit risk in the bond and CDS explicitly. We introduce the notion of a term structure of default probabilities, and price the bond and CDS by weighing the two distinct possibilities (default or no default) at each point in time. Under the specific Hull-White version of this approach, each instrument is priced by taking the difference of the instrument's risk-free value (which we obtain by discounting all cashflows by our base curve) and the expected loss due to default (which we obtain by multiplying the default probability by the loss in the event of default at each time point).[6] Thus, using the base interest rate curve and the curve

[5] Or, more or less equivalently, a credit-risky floating rate note or an asset swap

[6] Note that this approach does require us to assume that interest rates evolve independently of credit quality.

Figure 1: Spreads for DaimlerChrysler 8.0% 2010-06-15

of default probabilities, we can price both the bond and the CDS in a consistent framework, capturing the specific cashflow profiles (in default and non-default) of each. The default probabilities are then the bridge from bond to CDS pricing and vice versa. Further, it is clear what is meant by the assumption that default probabilities are the same for bonds and CDS.[7]

We can now utilize the bridge to compute a true bond-implied CDS: we fit a default probability curve that prices our bond accurately, then apply this curve to the CDS to find an implied fair CDS spread level.[8] The result is something that is directly comparable to the market CDS level. We show an example of the OAS, bond-implied CDS equivalent, and actual CDS spread in Figure 1.

So how do we make use of this information? The relative value implications are clear, but what about for risk? We return again to our theme of values. For the long-only investor, positions in bonds and CDS are for the most part all intended to express a view on the credit of an issuer, rather than on the relative value between different instruments. It is thus more appropriate to identify a single set of risk factors for each issuer than to treat bonds and CDS separately. The single set can be either from bond or CDS information, or indeed from equity information by way of a structural credit model, depending on which market is most liquid and affords the better data source; we then use the bridge to convert bond spread forecasts to CDS, or vice-versa. Beyond giving a consistent risk assessment using the best available data, this approach reduces the number of dimensions in our risk model, with benefits to both statistical forecasting and computational efficiency.

Arbitrage or relative value investors have a different set of needs. A typical trade here might involve buying a bond and simultaneously buying credit protection through a CDS. The motivation may be that

[7]We admit that life is not that simple. Since the default probability inferred here contains information about risk, liquidity, and other pricing premia, what we are really saying is that the total risk premium, including credit, is consistent across the instrument types.

[8]Similarly, we can run this process in the opposite direction to obtain a CDS-implied OAS or other bond spread measure.

the basis between the two instruments appears historically wide, and the investor expects to profit as the basis returns to typical levels; alternately, the motivation may be that the basis is negative, and the net position pays the investor a cashflow stream (roughly, the difference between the bond spread and the cost of CDS protection) but does not have any overall credit risk. In either case, a risk model that used only a single risk factor to describe all positions on a single issuer would not be prudent: such a model would, by construction, not predict any risk in the basis. It would be particularly dangerous in the case of the negative basis trade, where it would predict that there is no risk at all. Despite the trade looking like free money, there is always the chance that the mark-to-market value of the trade moves against the investor. Only by admitting two distinct sources of risk (the bond and CDS market spreads) does our risk model recognize this possibility.

For the relative value portfolio, then, the simplification to a single source of risk is not justified. We may still make use of the bridge, though, in order to establish model consistency. Since, as we have argued previously, bond and CDS spreads are not exactly the same quantity, it would not be appropriate to make the same modeling assumptions about how the two evolve. The effect of such an assumption would be that a bond modeled from CDS data would have a different distribution (not simply a different volatility) from a bond modeled from OAS data. To avoid introducing such an inconsistency, we prefer to state our modeling assumption in terms of a single quantity. Whether this quantity is the bond OAS, the CDS spread, or even default probabilities is a topic for another time; for now, we bear in mind that a consistent approach is crucial to any credit investor.

Conclusions

As we have considered spread risk, we have noted that at times the values of the risk manager overlap with those of the trader, while at times they take us to topics that the trader does not consider, or at least does not focus on. The challenge of building a spread risk framework is to balance these two possibilities. Thus, when we express risk of individual positions through market sensitivities, it is incumbent on us to speak the language of the market, and to express those sensitivities in the conventions that, based on their values, traders have established. On the other hand, to forecast risk, we need a consistent framework across multiple instrument types, while limiting the number of spread factors we introduce. Here, the risk managers values – model parsimony, statistical forecasting power, consistency of assumptions – drive many of our decisions. Values matter.

Further reading

- — (2005). Choose your poison. *The Economist*, October 1, pg. 32.
- Hull, J. and White, A. (2000). Valuing credit default swaps I: No counterparty default risk, *Journal of Derivatives*, **8**(1): 29–40.
- O'Kane, D. and Sen, S. (2005). Credit spreads explained, *Journal of Credit Risk*, **1**(2): 61–78.
- Stern, E. (2005). Measuring spread risk consistently across instruments, Presentation, RiskMetrics Group European Client Conference.