Instantaneous Credit Risk Correlation

Hans Byström

Hans Byström
is a professor at the Department of Economics, Lund University in Lund, Sweden.
hans.bystrom@nek.lu.se

While, traditionally, the pricing and risk assessment of credit has been more of an art than a science, the situation is rapidly changing. The credit world is now becoming as sophisticated as the equity world and many of the issues since long dealt with by equity investors, such as taking short positions and creating all kinds of derivatives, are now quickly being replicated by credit investors.

As a result, the behavior of *individual* credits is now quite well understood, and the quality of models describing such stand-alone credits is much better than it was just a few years ago. However, the level of understanding and modeling of *multiple* (dependent) credits is less satisfactory. The central parameter in the modeling of portfolios of credit-risky securities is the correlation between credit events affecting the portfolio, i.e., the default correlation. The problem, though, is that the default correlation is not only difficult to estimate, it is even quite hard to define. There are numerous ways of defining default correlation and a consensus on the "right" definition is currently lacking; should one look at the correlation between actual defaults, the correlation between default probabilities, the correlation between credit spreads or, perhaps, the correlation between asset (or equity) values?

In this article, we look at credit risk correlations or, more exactly, credit spread correlations. These correlations are more or less identical to default intensity or default probability correlations, since the size of the spread is directly linked to the likelihood of default (see, for e.g., van der Voort [2004] or Elizalde [2005]).[1] However, since the history of good-quality credit spread data (i.e., credit default swap (CDS) data) is very short, it is almost impossible to properly study the dynamics of credit spread correlations using the ordinary time-series sample (or Pearson) correlation measure. Moreover, the traditional correlation measure suffers from a number of well-known problems, such as extreme variability and "ghost" effects. Our suggestion is, therefore, to extend the cross-sectional correlation measure of Solnik and Roulet [2000] to the world of credits. By using cross-sectional variations, instead of time-series variations, we create instantaneous estimates of the correlation level among a set of credit-risky securities. These daily updated estimates of the credit correlation then lend themselves very well to time-series studies of the correlation dynamics.

The contribution of this article is twofold. First, we propose a novel approach to the estimation of credit risk correlation where, instead of looking at asset value based correlations, we compute daily correlations based on credit spreads. These credit risk correlations could be useful in pricing and relative-value analysis as well as in risk management and financial stability contexts. Second, by applying the approach to sector-wide and countrywide *iTraxx* CDS indexes, we try to shed some light on the co-movements of actual credit spreads. Among other things, this could give an estimate of potential diversification benefits in the credit market and how it has changed over time. As far as we know, this is the first study investigating this issue.

The next section gives a brief introduction to credit default swaps and to the *iTraxx* CDS indexes. The following section describes the proposed cross-sectional correlation measure, with the subsequent section presenting some empirical evidence on the dynamics of credit risk correlations among industry sectors and geographical regions (countries) worldwide. The next section discusses possible applications of the method, and the last section concludes our article.

CREDIT DEFAULT SWAPS AND CREDIT DEFAULT SWAP INDEXES

A CDS protects the buyer against credit losses caused by a particular underlying reference name. The credit protection buyer (the buyer of the CDS contract) pays the protection seller (the seller of the CDS contract) a periodic fee, the CDS spread. If a credit event strikes the reference name before the CDS contract matures, the CDS buyer typically delivers debt owed by the reference name to the CDS seller in return for a sum equal to the debt's nominal amount.

A CDS index, in turn, is an index created by pooling together several credit default swaps. CDS indexes are fairly new instruments and their main purpose is to provide investors with a quick and easy way of getting market-wide (positive or negative) credit risk exposure. The most well known CDS index, the *Dow Jones iTraxx* index, was created in 2004 from two older indexes and there are now several different *iTraxx* indexes covering different industry sectors and regions of the credit universe.[2] Each *iTraxx* index contains the most liquid CDS contracts in its particular segment.

In this study, we use two groups of *iTraxx* indexes. On the one hand, we use six industry sector indexes within the *iTraxx* Europe index family and on the other hand, we use six regional *iTraxx* indexes chosen from Europe, Asia, and Australia. These indexes will serve as proxies for the general credit spread levels in the various industry sectors and geographical regions (countries).

CROSS-SECTIONAL CREDIT CORRELATION

One of the most widely accepted models of dependencies in the credit market is the one-factor Gaussian copula model. In a sense, the model is the market standard, similar to how the Black-Scholes model is the standard model in options pricing. The one-factor Gaussian copula model can be used to price collateralized debt obligation (CDO) tranches, and it assumes a multivariate normal joint distribution of the asset values of all the names in the CDO pool. Furthermore, it simplifies the problem by assuming that all the pair-wise credit risk correlations in the pool are identical.[3]

In this article, we suggest an alternative way of modeling credit risk correlations. Instead of relying on asset value correlations, we choose to look at credit spread correlations. These correlations are assumed to be closely related to the default probability correlations among the same names. To calculate these correlations, we have chosen to extend the cross-sectional correlation measure suggested by Solnik and Roulet [2000] to credits. One important advantage of this approach is that it produces instantaneous correlation estimates that are freshly updated on a daily basis. Moreover, just like the one-factor Gaussian copula method described above, our cross-sectional credit correlation approach simplifies the problem by assuming a single market-wide correlation number.[4]

Solnik and Roulet [2000] focus on the equity market and, as an alternative to the common time-series based correlation measure, they suggest a correlation measure based on the cross-sectional dispersion of equity returns. Basically, a large dispersion of equity returns around an average "world" return is supposed to be an indication of a low correlation among the equity markets. Solnik and Roulet [2000] then suggest a way of computing a correlation coefficient based on this dispersion measure. Such a correlation measure has the advantage that it only uses observations from one point in time, not extended (overlapping) time-periods.[5] In this way, one gets instantaneous correlation estimates that are perfectly suited for studies of time-varying correlations. Moreover, while traditional correlation coefficients normally focus on pair-wise dependencies, the idea of the cross-sectional correlation coefficient is to produce a "global" measure of correlation that captures the co-movements among a group of names in one single coefficient.

Solnik and Roulet [2000] start by defining a *world return*, calculated as the average return of all the markets, around which the market returns are distributed. If the markets are highly correlated, they will demonstrate similar returns at a certain time, and if they are weakly correlated, the returns will be widely spread out around the world return. This tendency for the returns to move or not to move together is summarized by the dispersion quantity.

Solnik and Roulet [2000] then continue by proposing how one can transform the dispersion measure into a correlation coefficient. First, they assume that the return, $R_{i,t}$, in a single market i at time t is driven by the world return, $R_{w,t}$, plus an independent error term, $e_{i,t}$

$$R_{i,t} = a_i + R_{w,t} + e_{i,t} \tag{1}$$

The time-varying standard deviation, $\sigma_e(t)$, of the error term will then serve as a measure of the cross-sectional dispersion.[6] Second, Solnik and Roulet [2000] identify the (possibly time-varying) volatility of the world return as $\sigma_w(t)$,[7] and third, they calculate the instantaneous cross-sectional correlation, $\rho_{i,w}(t)$, between market i and the world as

$$\rho_{i,w}(t) = \frac{\sigma_w(t)}{\sigma_i(t)} = \frac{1}{\sqrt{\left(1+\frac{\sigma^2_e(t)}{\sigma^2_w(t)}\right)}} \tag{2}$$

where $\sigma_e(t)$ is the dispersion of the market returns around the world return.[8] The measure, $\rho_{i,w}(t)$, can now be interpreted as the average level of global correlation at time t (since every market's correlation with the world index is the same).

In this article, we suggest an extension of the Solnik and Roulet [2000] framework to the calculation of correlations among credit spreads. We make no modifications at this point (in this article) and we simply replace all the equity returns above with credit spread changes. In the next section, we look at the resulting cross-sectional credit correlations for two sets of *iTraxx* CDS indexes.

CREDIT CORRELATIONS AROUND THE WORLD

In this section, we present some empirical evidence on the evolution of credit risk correlations among industry sectors and geographical regions worldwide. We look at six sector indexes within the *iTraxx* Europe index family as well as six regional *iTraxx* indexes chosen from Europe, Asia, and Australia. The results are important for anyone investing in credits and credit-linked instruments. We also briefly discuss the well-known sector-country diversification controversy, but in a credit setting.

The European *iTraxx* index is split up into *sector* sub-indexes (Autos, Industrials, TMT, Energy, Consumer, and (senior) Financials), and these six (five-year maturity) sub-indexes are the sectoral CDS indexes that we use in this study. For the regional (five-year maturity) indexes, we have chosen to look at *iTraxx* Europe, *iTraxx* Japan, *iTraxx* Australia, *iTraxx* Korea, *iTraxx* Greater China and *iTraxx* Rest of Asia.[9] Again, these are a total of six indexes. The various *iTraxx* spreads are plotted in Exhibits 1 and 2. The time period is June 21, 2004, to October 3, 2006, for the sectors and July 26, 2004, to January 17, 2007, for the regions.[10,11]

We then go on to present the cross-sectional credit correlations, i.e., the daily credit spread co-movements for the sectors and the regions in Exhibits 3 and 4, respectively.[12] As one can see from these two exhibits, the daily updated cross-sectional correlation measure jumps around significantly.[13] In order to simplify the visual analysis, we choose to smooth the correlation series using one month of daily estimates around the day of interest. The fat curves indicate these moving averages and, compared to the daily updated correlations, they are much smoother.[14] On average, the correlation level among the industry sectors (average = 0.84) is somewhat higher than that among the geographical regions (average = 0.75). Both time-series show a mean-reverting behavior, even though this is most obvious for the correlation among the geographical regions. None of the series show any significant trend in any direction, though. One reason for this might be, of course, that the time-periods are fairly short.

Now, the time-period for the sectors already ends in October 2006 and that makes it difficult to compare the current situation in the two segments of the credit market. We can, however, compare the situation in the past, for instance, around the crisis in the U.S. auto industry during the spring months of 2005, and around the market correction in May 2006. During the spring 2005 crisis, the cross-sectional correlation fell significantly among both the European corporate sectors and among the geographical regions. For the group of sector CDS indexes, this de-coupling is clearly an indication of sector-specific

EXHIBIT 1
iTraxx CDS Index Spreads: European Industry Sectors (basis points)

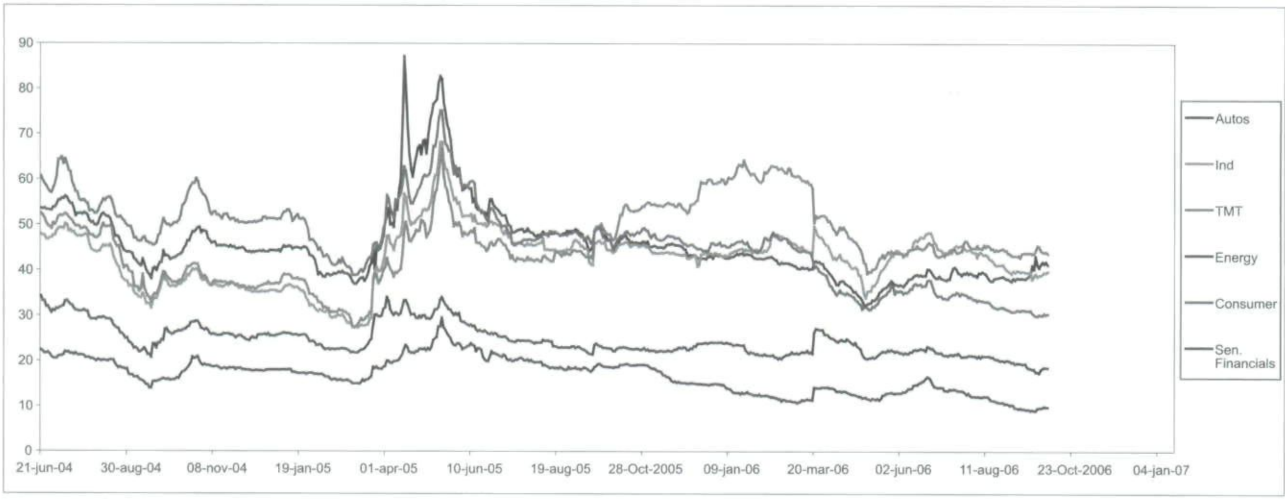

problems (the auto industry). For the country CDS index group, on the other hand, the cause of the drop in correlation is less obvious. Possibly, it is caused by the auto industry-heavy European and Japanese indexes being relatively more affected by the auto industry problems than the other regional indexes. If so, this is confirmed by looking at Exhibit 2 where the European, Japanese and, to some extent, the Korean (another car-producing country) index spreads are widening at the time of the crisis, while the other index spreads barely move. The effect of the May 2006 market correction on the co-movements in the CDS market can also be seen in Exhibits 3 and 4. Here, however, the fall in correlation is much smaller than for the 2005 auto industry crisis. In fact, it is hard to distinguish the fall from general day-to-day fluctuations. This holds for the sectors as well as for the regions and it is in line with the fact that the May 2006 correction pretty much affected all markets and asset classes.

Finally, there is not much we can add to the debate on sector vs. country diversification. As mentioned above, in neither of the two markets do we detect (from visual inspection) any trend in the correlation. Furthermore, in

EXHIBIT 2
iTraxx CDS Index Spreads: World-wide Regions (Countries) (basis points)

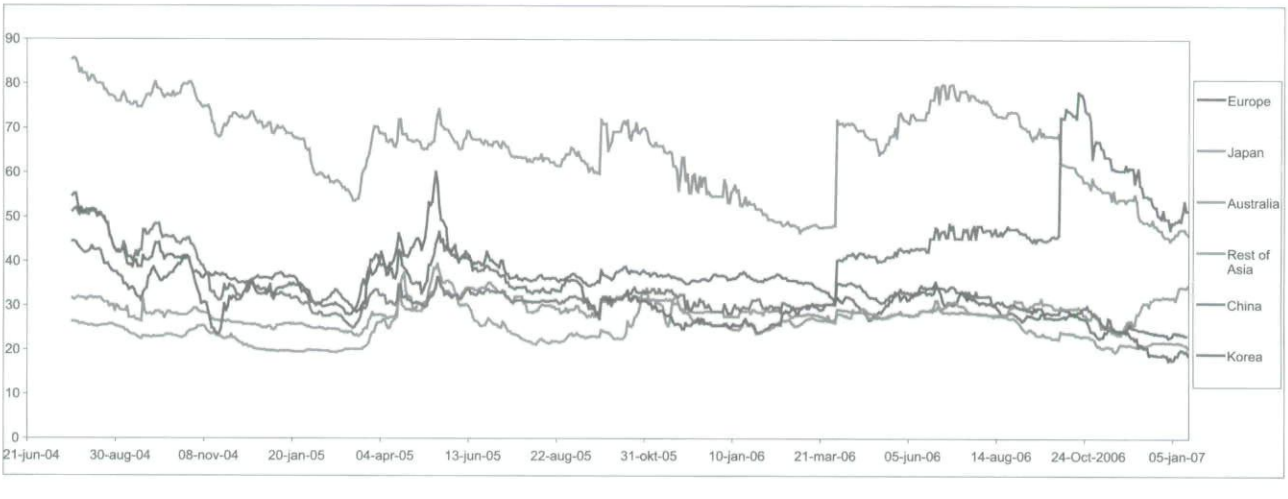

EXHIBIT 3

Cross-sectional Credit Correlation: European Industry Sectors

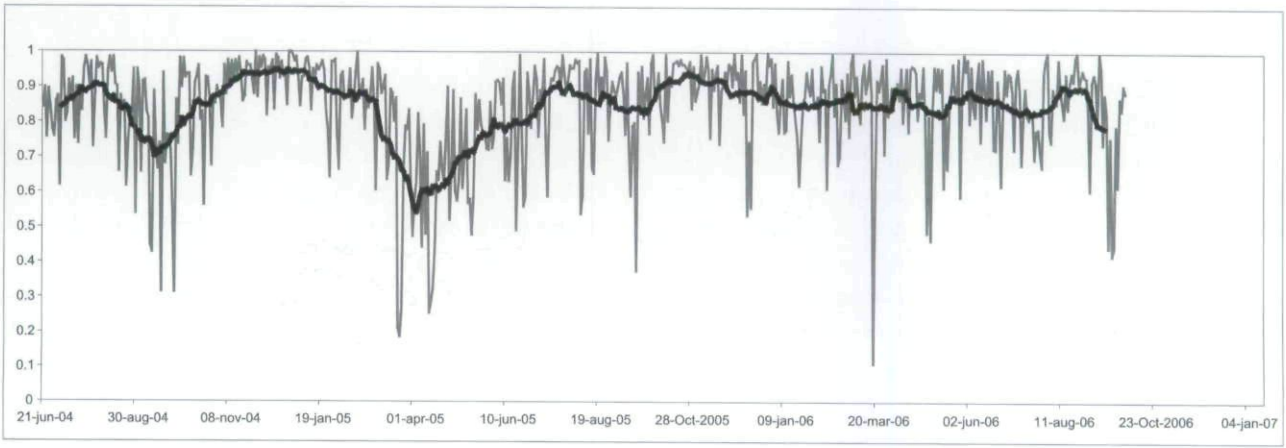

both markets, the correlation level is very high; the regional spreads co-move slightly less than the industry sector spreads but the (common) high level of correlation is more striking than any small and, most likely, economically non-significant difference. The high correlation simply tells us that neither country nor sector diversification seems to be as crucial for the credit investor as is often argued. The high correlation will also possibly make it more difficult for market-neutral hedge funds and the like to create alpha by taking sector- or country-wide positions in the credit market. Basically, skilled managers will find fewer opportunities to outperform their peers, and the role of chance will increase.

POSSIBLE AREAS WHERE OUR CREDIT CORRELATION MEASURE COULD BE USEFUL

Here follow some areas where our cross-sectional credit correlation measure could be particularly useful.

It could be useful in order to infer the anatomy of a credit-related crisis. By looking at the link between an overall CDS index spread and the credit spread correlation among the names/sectors in the index, one can possibly tell whether a credit crisis, identified as a sudden widening of the index spread, is sector-specific or systemic in its nature. We saw some evidence of this in the previous section. Basically, a market-wide crisis would cause the correlation among the names/sectors to increase and a

EXHIBIT 4

Cross-sectional Credit Correlation: World-wide Regions (Countries)

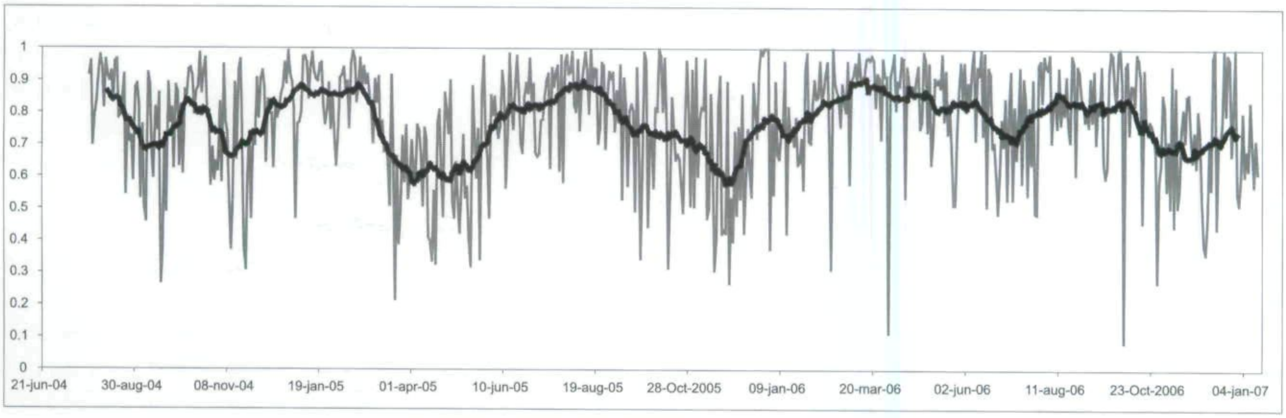

sector-specific crisis would cause the correlation to fall. In other words, by looking at the cross-sectional credit spread correlation-level (or perhaps its monthly moving average), together with the spread-level movements, one can tell more about the anatomy of a crisis. This is not a new idea (see Belsham, Vause, and Wells [2005]) but we believe that our instantaneous correlation measure could capture correlation changes much quicker than traditional historical correlation measures, without suffering from some of the drawbacks of market-implied correlations.

Our approach can be compared to that of Chan-Lau and Lu [2006]. Just like us, Chan-Lau and Lu [2006] look at the European *iTraxx* CDS index market but instead of focusing on the overall index spreads, they suggest a way of extracting idiosyncratic and systemic risk components from quoted CDS index *tranche* spreads. CDS index tranches are nothing but single-tranche synthetic CDOs and, currently, there are traded *iTraxx* tranches for the European *iTraxx* CDS index.[15] One drawback of the Chan-Lau and Lu [2006] approach is that it only is applicable to markets with traded index tranches (i.e., the European corporate market).[16] Our approach, however, is applicable to a range of different credit markets around the world.

Now, if we compare our result with those of Chan-Lau and Lu [2006], we see some differences. During the spring 2005 auto industry crisis, for example, our method sends a strong signal of the crisis being idiosyncratic rather than systemic in its character. The Chan-Lau and Lu [2006] method, on the other hand, is more ambiguous since both the idiosyncratic and the systemic risk component increase significantly at the time of the crisis. The reason for this discrepancy could be supply/demand effects in a still fairly illiquid *iTraxx tranche* market, making the Chan-Lau and Lu [2006] method less useful. On the other hand, it could also be due to the small number of sectors (six) in our cross-sectional analysis, which could cause one or two individual sectors to have a disproportionately large effect on the overall correlation-level. Finally, it could also be caused by our different definitions of idiosyncrasy; in our setup, idiosyncratic means sector-specific, while in the set-up of Chan-Lau and Lu [2006], it means firm-specific.

Belsham, Vause, and Wells [2005], in turn, studied the U.S. *CDX* CDS index market from 2004 to 2005. In a sense, the Belsham et al. study [2005] foregoes us just as well as Chan-Lau and Lu by looking at both tranche spreads and (market-implied) correlations. Belsham et al. [2005] find that different tranche spreads react differently to the developments in the auto industry in 2005. This is further confirmed by a sharp fall in market-implied (equity tranche) correlation.[17] Even though the results in Belsham et al. [2005] are similar to ours, only a more comprehensive comparison of the various approaches could shed light on their relative performance. In doing so, it should be remembered that the implied correlations in Belsham et al. [2005] are inherently forward-looking (as are those of Chan-Lau and Lu [2006]), while our correlations only indirectly represent the future through the spreads' link to the default probability over the coming five-year period.

It could be useful for rating agencies in their ratings of CDOs. Default correlation estimates are very important for the rating of CDOs and other multi-name structured products, and we therefore propose a possible inclusion of our simple dispersion-based approach in the rating of such products. Just like us, in the example in the previous section, Standard & Poor's, Moody's and Fitch all focus on industry sectors in their models (S&P [2001]; Fender and Kiff [2004], Fitch [2006]) and they all typically divide the default correlations in the underlying pool into intra-industry and inter-industry correlations. Furthermore, the inter-industry default correlations are often assumed to be zero or close to zero. This clearly seems to be at odds with our findings; even if credit spread correlations are much higher than corresponding default event (or default-time) correlations, it seems drastic to assume the latter to be zero, considering the very high spread correlations (default probability correlations) observed in the previous section.

The typical rating agency approach is to extract the degree of default dependency among names from their asset (or equity) values. Our approach, instead, is based entirely on the names' credit spreads. Considering the fast growth of the credit derivative market, it seems reasonable to start incorporating spread information in the rating models. Furthermore, our cross-sectional correlation measure is applicable to any multi-name structured product whose underlying names have quoted credit spreads.[18] Examples of such products are the standardized *iTraxx* Europe tranches (that reference 125 individual names within six industry sectors) as well as typical single-tranche CDOs and basket default swaps.

It could be useful for pricing tranches of non-standard (bespoke) CDOs. The market practice of quoting implied correlations of standard CDO tranches is useful as a convention for market participants. However, if tranches of non-standard pools of names are to be priced, these implied standard correlations are less useful. A possible solution to price tranches of a CDO whose pool of names is similar to, but not identical to, let's say, the standard *iTraxx* Europe index portfolio, is to modify, or perturb, the implied correlation using our cross-sectional credit correlation. Since it is straightforward to calculate the spread correlation for any pool of traded names, one could simply use the difference in credit spread correlation (the sign as well as the size) between the standard pool and the bespoke pool as a (crude) indicator of how to perturb the market-implied default-time correlation from the standardized CDO market in order to use it to price the bespoke CDO.

It could be useful for relative value analysis of new generations of structured products. Above, we mentioned that the typical CDO rating methodology relies on asset (or equity) value correlations between the names in the pool. Now, in recent CDO generations, many types of structured products are used as collateral; asset-backed securities (structured finance CDOs or SFCDOs), mortgage-backed securities (CMOs) and CDO tranches (CDO^2s) are some examples. This complicates the situation since asset value correlations, or for that matter the asset values themselves, of the underlying structured products are nowhere to be found. The same holds for CDOs that reference pools of loans to privately held firms, sovereign debt CDOs, and novel products such as collateralized foreign exchange obligations (CFXOs) and microfinance CDOs (MiCDOs). This is a potential area where our cross-sectional correlations could serve useful for relative value analysis.

It could be useful, more generally, for traders in the credit derivatives market. Different views on default correlation are a trigger for trading in so called "correlation products"; i.e., products whose value depends on the default correlation among a set of credits. Examples of such products are CDOs, such as the *iTraxx* tranches, as well as basket products, such as n^{th}-to-default basket default swaps. Consequently, in liquid markets such as the *iTraxx* tranche market, where the tranches are quoted using market-implied (default-time) correlations, it is straightforward to envision trading strategies exploiting the differential between changes in the market-implied default-time correlations and changes in the corresponding cross-sectional credit spread correlations.

CONCLUSION

We have demonstrated a method of estimating instantaneous credit risk correlations using credit spreads. By relying on cross-sectional information instead of time-series information, we are able to create freshly updated estimates of the credit risk correlation on a daily basis. In this way, it is possible to follow the daily (or monthly, if preferred) movements in the correlation among a set of credits and this can be used both for pricing and trading purposes in the quickly growing credit derivatives market and for risk management purposes. For instance, the correlation measure could be compared to daily quoted CDS index tranche correlations in order to detect trading opportunities. Or, alternatively, the market-wide credit correlation level, combined with the market-wide CDS index spread level, could serve as an indicator of whether a crisis in the credit market is idiosyncratic or systemic.

In the empirical part of the article, we present some evidence on the dynamics of the credit spread correlation throughout the history of the CDS index market. We compare the correlation dynamics amongst a set of industry sectors with that amongst a set of geographical regions. Our findings indicate that the correlation levels are fairly similar among the sectors and the regions. Furthermore, in both cases, the general correlation level is quite high, which limits possible gains from diversification among countries or industries.

ENDNOTES

Financial assistance from *Handelsbankens Forskningsstiftelser* and *Stiftelsen Bankforskningsinstitutet* is gratefully acknowledged.

[1]Of course, this assumes that the recovery rates of the various credits remain constant over time.

[2]The indexes are managed and administrated by *International Index Company* (IIC [2004]).

[3]In the marketplace, *iTraxx* index tranches (essentially synthetic CDOs) are sometimes quoted using "implied correlations" backed out from such a one-factor Gaussian copula model.

[4]In the one-factor Gaussian copula method, this number is the default-time correlation (or the asset value correlation), and here it is the credit spread correlation.

[5]It is well known that the ordinary sample correlation measure has the disadvantage of being very insensitive to temporary changes in the correlation and of suffering from "ghost" effects and spurious dependencies.

[6]Behind this equation lies an assumption of a β_i equal to

one for each market, i, and an assumption of independent multivariate normally distributed deviations of the individual market returns around the world return, $e_{i,t} \sim N(0, \sigma_e(t))$.

[7]The model therefore relies on an accurate and instantaneous estimate of the world return volatility, $\sigma_w(t)$, at each point in time. Solnik and Roulet [2000] choose to model the volatility as constant over time and argue that this gives sufficiently accurate correlation estimates that do not differ significantly from those produced by other volatility estimators. We follow Solnik and Roulet by assuming $\sigma_w(t)$ to be constant over time.

[8]Since all the β_i are equal to one and can be written as $\rho_{i,w}\sigma_i/\sigma_w$.

[9]The time difference between Europe and Asia does not seem to affect the results; when replacing today's *iTraxx* Europe quotes with yesterday's quotes, we get close to identical results.

[10]The data was downloaded through the web page of *International Index Company* and no sector data was available after October 2006.

[11]The large jumps in the spread series are caused by the re-balancing of the *iTraxx* index in March and September each year. Without any loss of stringency, these spikes, which are to some extent predictable, are ignored in this study.

[12]The large negative spikes in the correlation series are also caused by the re-balancing of the *iTraxx* index in March and September each year. Again, these spikes are of no interest to us and can, of course, easily be removed from the series.

[13]Most likely, the reason for this relatively volatile behavior is the rather small number of spreads (six) in the two cross-sections of indexes.

[14]One should recall that ordinary sample correlations jump around a lot more even when several months of data are used in the estimation.

[15]Junior (equity) and senior CDO tranches, respectively, react in opposite ways to a change in the default correlation, which is exploited in the method of Chan-Lau and Lu [2006].

[16]Eventually, this will, of course, change when tranches are introduced in other CDS index markets as well.

[17]Here, one problem of using market-implied correlations, instead of our credit spread correlations, is immediately noticeable; i.e., the need to determine which of the quoted correlations to use (due to the correlation skew, they are all different).

[18]Of course, to get statistically significant correlation estimates, we cannot have too few names in the product. Six or seven names is probably a minimum.

REFERENCES

Belsham, T., N. Vause, and S. Wells. "Credit Correlation: Interpretation and Risks." *Bank of England, Financial Stability Review* (2005), pp. 103–115.

Chan-Lau, J.A., and Y. Lu. "Idiosyncratic and Systemic Risk in the European Corporate Sector: A CDO Perspective." Working Paper, IMF, 2006.

Elizalde, A. "Do We Need To Worry About Credit Risk Correlation?" *The Journal of Fixed Income*, Vol. 15, No. 3 (2005), pp. 42–59.

Fender, I., and J. Kiff. "CDO Rating Methodology: Some Thoughts on Model Risk and its Implications." BIS Working Papers No. 63, Bank for International Settlements, 2004.

Fitch. "Global Rating Criteria for Collateralized Debt Obligations." Credit Products Criteria Report, Fitch, October 2006.

IIC. "Dow Jones iTraxx CDS Indices Europe." Product Description, International Index Company, 2004.

Solnik, B., and J. Roulet. "Dispersion as Cross-sectional Correlation." *Financial Analysts Journal*, 56 (2000), pp. 54–61.

Standard & Poor's. "CDO Evaluator Applies Correlation and Monte Carlo Simulation to the Art of Determining Portfolio Quality." Commentary, Standard & Poor's, November 2001.

van der Voort, M. "Double Default Correlation." Working Paper, Econometric Institute, Erasmus University, Rotterdam, 2004.