Why CDOs Work

Jon Gregory[1]

Revised version, 3rd March 2014

1. Introduction

The growth of the structured credit market gave rise to many complex collateralised debt obligation (CDO) structures. An investment in a CDO can be broadly characterised paying a return as compensation for exposure to a certain range of losses on a static or managed portfolio. Precise quantification of the risks in a CDO is complex since one needs to assess the multidimensional loss distribution for the underlying portfolio.

Prior to the beginning of the global financial crisis in 2007, the CDO was a successful financial innovation. Investment banks made large profits from structuring CDOs and hedge fund and real money investors generally made good returns, at least when judged against the credit rating of the underlying investment. Banks also utilised CDOs so as to achieve "regulatory arbitrage" in relation to their capital requirements (a topic not dealt with in this article). Other players in the CDO market gained also, for example rating agencies earned good fees for rating the plethora of CDO structures and monoline insurers collected significant premiums for insuring or "wrapping" senior CDO tranches.

However, the global financial crisis was partly catalysed by an implosion of the CDO market. CDOs have been blamed for causing the crisis, pricing models for CDOs have been heavily criticised, litigation has been rife and investor demand has almost disappeared. All players in the CDO market, notably banks, rating agencies as well as investors have suffered as a result. Banks lost billions of dollars due to failed hedging and write-downs due to the counterparty risk in trades with failing monoline insurers. Investors have faced losses as either defaults hit or they were forced to sell investments at deeply discounted or fire sale prices. Ratings agencies have suffered due to structured credit revenues dropping dramatically and as a result of the reputational damage due to the perception that they had been giving high quality ratings to investments that were potentially quite the opposite. Only a small minority, such as the few hedge funds who anticipated the problems, have benefited from CDOs since 2007.

Recently, though, CDOs have attempted a comeback with collateralised loan obligation (CLO) issuance increasing substantially and reports that banks may again issue CDOs based on corporate portfolios[2]. An obvious and timely question to ask is whether the concept of a CDO is flawed and the market was doomed to eventual failure or if they, like many other financial investments, were simply a casualty of a largely unforeseen and completely unprecedented global financial crisis. In this article, we attempt to answer this fundamental question in a simple way. We note that our analysis will use corporate data and will use a

---

[1] Contact details: jon@solum-financial.com . The author is grateful for the comments of two anonymous referees.

[2] See, for example, "Frankenstein' CDOs twitch back to life", G. Tett, Financial Times 6th June 2013.

"full capital structure" synthetic CDO as an example. Nevertheless, it seems plausible that the general conclusions may apply more generally to other types of CDO.

2. Simple theory of a CDO

CDOs come in many forms such as cash or synthetic and cover various different assets from corporate to ABS. However, their basic principle is to take the risk on a given credit portfolio and redistribute it via tranches. A typical full capital structure CDO structure is represented in Table 1. A number of different classes of securities are issued to cover the full portfolio notional[3]. The riskiness of these securities changes from the bottom unrated equity tranche to the top so-called super senior tranche. Although this latter tranche has no rating, it is above the triple-A rated Class A notes and therefore is no worse than triple-A credit quality.

Table 1. *Illustration of the securities issued from a typical CDO.*

Class	**Amount**	**Tranching**	**Rating**	**Funding**
Super senior	850	[15-100%]	NR	Unfunded
Class A	50	[10-15%]	Aaa/AAA	Funded
Class B	30	[7-10%]	Aa2/AA	Funded
Class C	30	[4-7%]	Baa2/BBB	Funded
Equity	40	[0-4%]	NR	Funded

A CDO redistributes risk from 100% of an underlying portfolio to a series of tranches spanning the entire range of losses. Thinking in terms of the (real world) expected loss that each tranche of a CDO will suffer, we can write:

$$EL_p = \sum_{i=1}^{n} m_i EL_i, \qquad (1)$$

where EL_p represents the expected loss on the underlying unit size portfolio, EL_i represents the expected loss on the n underlying unit tranches with m_i the percentage tranche size assuming that these tranches span the entire range of losses (full capital structure) so that $\sum_{i=1}^{n} m_i = 1$.

Investors in the underlying portfolio and tranches will naturally expected to be compensated for expected losses but also should demand risk premiums (for example for the default and liquidity risks they take). The success of a CDO structure will depend on the following relationship holding:

$$\alpha_p EL_p > \sum_{i=1}^{n} \alpha_i m_i EL_i, \qquad (2)$$

where α_i represents a "risk aversion coefficient" in relation to the relative return investors require for taking risk to a unit amount of the tranche i (or portfolio p). The $\alpha_i's$ will be effectively determined by the coupons demanded by the investors with risk aversion implying $\alpha_i > 1$. The left hand side of equation (2) represents the compensation paid for taking the credit risk of the underlying portfolio whilst the right hand side is the total compensation that

[3] Not all tranches will be sold as some may be held by the issuer either as investments or to be hedged as in the case of so-called "single tranche CDOs".

must be paid to the various CDO investors[4]. The focus on expected loss in this analysis is relevant since ratings methodologies typically used expected loss[5] as the primary (and often only) quantitative driver of the rating of a given CDO tranche.

We can see that the success of a CDO depends on the risk aversion that investors demand for taking the default risk of the various tranches compared to the same component for the underlying portfolio which we have expressed via the $\alpha_i's$. Since investors relied primarily on credit ratings which in turn are driven by expected losses, we assume that the coefficients depend directly on the expected losses of each tranche. We note that, especially in the early days of the CDO market, investors would demand a higher return compared to an equivalent rating in, for example, the secondary bond market. This component, together with other costs will be considered in a later example. We discuss the potential assessment of other aspects such as the systemic risk of a tranche at the end of this paper. The form assumed for the coefficients is:

$$\alpha_i = \left(\frac{a}{EL_i}\right)^b, \tag{3}$$

where $b = 0 => \alpha_i = 1$ would apply to risk-neutral investors and $a, b > 0$ corresponds to risk aversion. The risk aversion coefficient (α_i) tends to infinity as EL_i tends to zero (assuming the *relative* compensation for default risk will increase as the probability of loss reduces) and reduces as EL_i tends to unity (since here there is no uncertainty around default[6]). We will show below that the above relationship provides a good fit to empirical data.

With the assumptions regarding risk aversion, equation (2) becomes:

$$\left(EL_p\right)^{1-b} > \sum_{i=1}^{n} m_i (EL_i)^{1-b}, \tag{4}$$

The above expression, with the constraint defined by equation (1), is satisfied when $b < 1$. This is not surprising since when $b = 1$, investors would demand the same price irrespective of the risk of a tranche (i.e. they would ignore ratings).

3. Empirical test

There is a broad body of empirical analysis characterising risk premiums in corporate bonds, including Fons (1987), Altman (1989) and Giesecke (2010). We will use the results of Hull et al. (2005), who give the ratio corresponding to α_i in equation (3) directly via the ratio of risk-neutral to real world default losses as a function of credit rating. The $\alpha_i's$ are shown to be generally decreasing with higher expected loss (lower rating)[7]. Figure 1 shows the best fit of

[4] We note here that we consider the upfront present value of the total compensation for credit risk. In reality, this compensation is generally paid via a coupon stream on the various tranches. We will also show a real example later using the coupon stream.

[5] The exception to this being methodologies that considered the probability of default. This corresponds to the expected loss approach but with no recognition of recovery value in default.

[6] Note that this relationship can give a value of less than unity which would imply risk seeking behaviour. In the calibration shown, this occurs only at very large expected loss values. This effect can be corrected either by adding unity to the relationship in equation (3) or constraining the calibration. Neither of these changes the analysis or has a material impact on the conclusion.

[7] The values are 16.8, 13.0, 9.8, 5.1, 2.1, 1.2 and 1.3 for Aaa, Aa, A, Baa, Ba, B and Caa ratings respectively.

the parametric relationship defined by equation (3) using the expected losses and ratios defined in Hull et al. (2005). The fit is reasonably good and the resulting value of $b = 0.47$ is well below the critical boundary of unity where the CDOs would effectively not work.

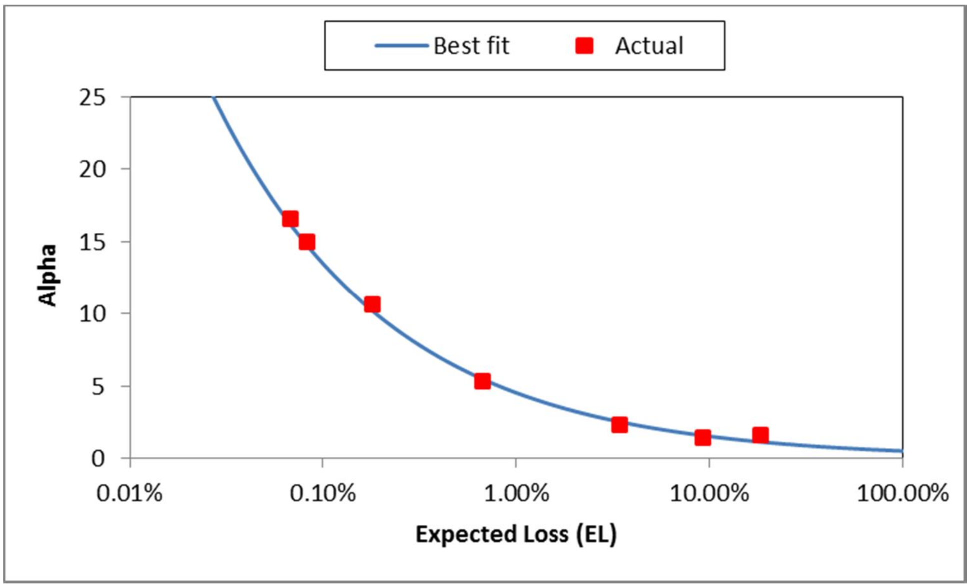

Figure 1. *Fit of the ratio of risk-neutral default intensity to real world default intensity as a function of the (real world) expected loss from Hull, Predescu and White (2005) together with the best fit of this data using equation (3) with coefficients given by a = 0.24 and b = 0.47.*

The above analysis shows a very important point which is that CDOs work due to the risk preferences of investors (where investments with more credit risk require a *relatively* lower multiplier as compensation for risk premiums) and rating agencies methodologies (which generally are based on expected losses of tranches). Indeed, a rating agency using an expected loss based rating approach cannot cause a CDO to "fail" no matter what assumptions they make, for example with respect to the correlation of defaults in the portfolio. This is significant because, in criticising rating agency modelling approaches, the choice of correlation may appear to some as the most obvious weakness. In addition, the use of the term "arbitrage CDO" seems to suggest that one party to a CDO transaction (e.g. an investor) must lose for the other (e.g. the issuer) to gain. On the basis of the analysis above, this is not the case.

We note that a CDO could in theory still fail due to the granularity of ratings. If a given tranche has "over subordination" in terms of achieving its given rating then this will impact the EL values in equation (4). This over subordination could be seen as arising from conservative rating agency approaches or issuers building extra subordination into tranches to provide rating stability. The fact that the calibrated b in Figure 1 is so significantly below the critical value of unity suggests that there is plenty of room for conservative assumptions although we will give a better insight into this in the example presented next.

4. Example

We can illustrate the key points with a very simple example of a CDO structure illustrated in Figure 2. A portfolio is divided into three tranches, equity, mezzanine and senior and has a 5-year fixed (bullet) maturity. Following Hull et al. (2005), we use the expected loss data based on Moody's ratings from Hamilton et al. (2004). The underlying credit portfolio is assumed to reference 100 bonds with triple-B ratings and the corresponding 5-year BBB default probability is 2.16%. Assuming a loss given default of 60%, this will give an expected loss of 2.06% × 60% = 1.296%. Finally, we know that a BBB portfolio has to compensate for a loss of more than this due to risk and liquidity premiums. The multiplier from Hull et al. (2005) is 5.1, which suggests that the overall compensation the investors would receive is in fact 1.296% × 5.1 = 6.610%. We assume the underlying portfolio will provide exactly this amount[8].

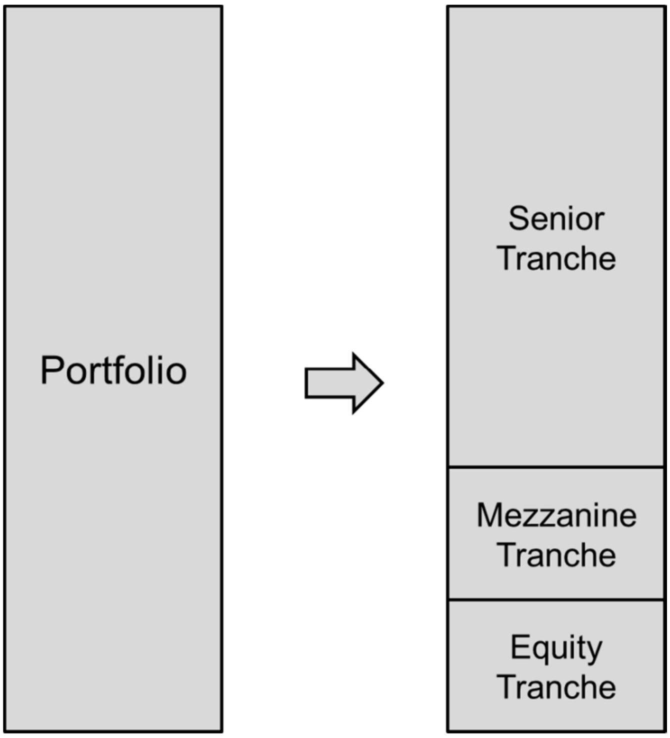

Figure 2. *Simple CDO structure used for the example.*

The approximate goal of a CDO is to sell the tranches for less than the return received on the underlying portfolio. In this simple example, this corresponds to paying investors an overall return of less than 6.610% for the equity, mezzanine and senior pieces. We assume arbitrarily that the tranches will be defined as [0, 4%], [4-8%] and [8-100%]. We furthermore assume that the ratings approach is expected loss based so that each tranche gains the best possible rating according to its modelled expected loss. The rating agency model is assumed to follow a Gaussian copula approach with a homogeneous correlation parameter of 20% which is broadly representative of approaches used throughout the development of the CDO market (for example, see Standard & Poor's 2002). We then calculate ratings of Caa, Ba and Aaa for the equity, mezzanine and senior respectively. Assuming investors will demand a return for these investments corresponding to the multipliers estimated by Hull et al. (2005) then the economics of the structure are shown in Table 2.

[8] It should be strongly emphasised that all the above numbers are based on empirical analysis over many years of data but given the results shown are so strong then we believe it is highly unlikely that other data would lead to different overall conclusions.

Table 2. Illustration of the economics of the base case 5-year CDO structure.

Rating	Tranche	5-year exp loss	Multiplier	Protection value	Size	Spread (bps)
Baa	[0-100%]	1.296%	5.1	6.610%	100%	144
Aaa	[8-100%]	0.072%	16.8	1.210%	92%	26
Ba	[4-8%]	6.702%	2.1	14.074%	4%	321
Caa	[0-4%]	36.498%	1.3	47.447%	4%	1376

The CDO works because most of the risk is sold in the equity tranche, which attracts a relatively low multiplier. It is *relatively* expensive to sell the AAA tranche as the multiplier assumes that for every unit of actual default risk passed on, 16.8 units of return must be paid. However, given the small amount of actual risk that is assessed as being in this tranche (under the modelling assumptions defined above), this does not affect the economics of the structure particularly adversely. The total value of protection bought via the CDO is 3.574% of notional[9] which is substantially higher than the value of protection sold to the market via the underlying portfolio (6.610%). Another way to see the value created is via the so-called excess spread, which is the spread paid in versus that paid out which is calculated as 52 bps[10]. This positive excess spread illustrates the approximate value within in the structure[11]. In Figure 3, we show the base correlation curve corresponding to the three tranches priced in Table 2. This shows that the well-known upwards sloping base correlation curve can be explained by the risk aversion preferences of investors.

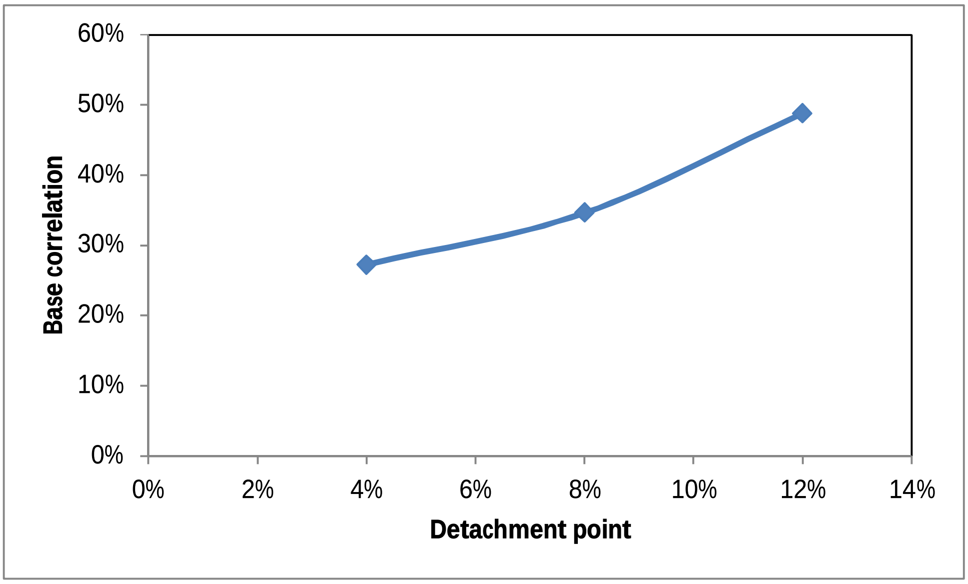

***Figure 3.** Base correlation corresponding to the CDO tranche prices in Table 2.*

The results above are consistent with other approaches. For example, Brennan et al. (2009) show that there is a "marketing gain" can be achieved from tranching corporate debt when ratings are based on expected loss methodologies. They also find this marketing gain is greater when there is increased systematic risk within the underlying portfolio.

[9] 1.210% × 92% + 14.074% × 4% + 47.447%.× 4%.

[10] 144 - 26 × 92% - 321 × 4% - 1376 × 4%.

[11] The excess spread is not a perfect guide to the profit since it changes over the lifetime of the CDO as defaults occur. Indeed, due to the equity tranche having the shortest duration, the excess spread will overestimate the value in the structure. However, it is a reasonable guide to the approximation economics.

Note that based on equation (4) and the empirical test, as long as the rating methodology is based on expected losses, then the correlation assumptions of rating agencies cannot possibly cause a CDO to fail. However, we should also discuss the granularity inherent in the rating process. Due to the arbitrary choices regarding the tranching and the granularity of ratings, there is already “over-subordination” in the above structure since the same ratings could be achieved by reducing the attachment point of each tranche. For example, in the above example, the same size mezzanine tranche covering losses defined by [3-7%] would still achieve a Ba rating. However, despite this implicit over-subordination, the structure is still significantly profitable and therefore even if CDO investors demand a higher relative return for each rating (which was certainly true in the early days of the CDO market), then there is seemingly ample value in the structure for all parties to benefit.

In order to illustrate the granularity effect, we consider the impact on the excess spread of changing correlation assumptions for the structure under both fixed and optimised tranching. In the former case, the tranches are always held fixed at their assumed values in Table 2. In the latter case, we assume a relatively simple optimisation[12] process where the equity tranche is sized so as to just achieve a CCC rating and then the detachment point of the mezzanine tranche is chosen so as to achieve the highest excess spread[13]. The results of this comparison are shown in Figure 4, with the optimised tranching depicted in Figure 5. Whereas the base case CDO “fails” at very high correlation values, the optimised structure maintains a materially positive excess spread (the smaller excess spread at high correlation is due to the granularity introduced by using ratings).

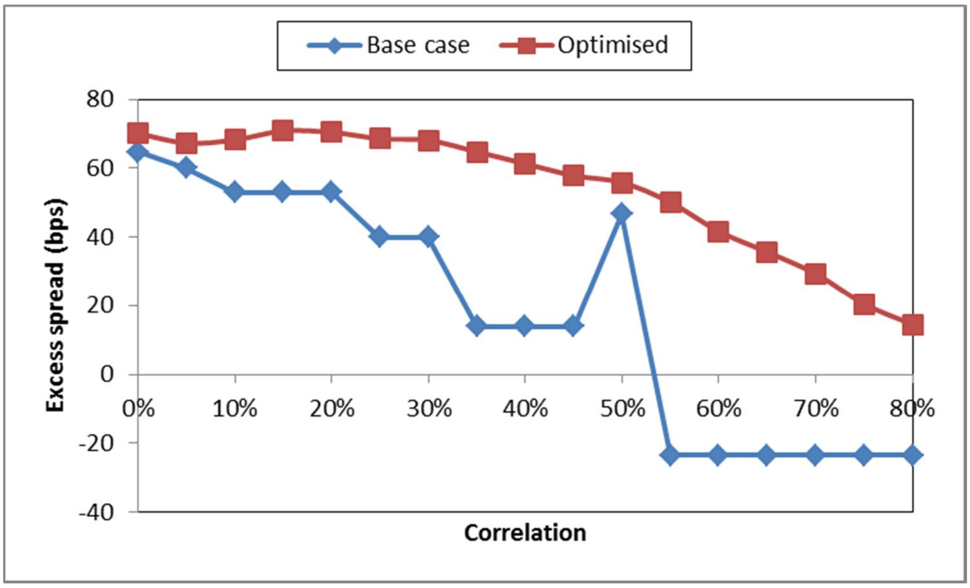

Figure 4. *Excess spread for the base case and optimised CDO portfolio.*

[12] Given the granularity of the rating process leading to a discontinuous objective function, a more complex optimisation is difficult to achieve but leads to similar results as the optimal tranching is usually to start with the relatively cheap CCC tranche and then there is only one variable across which to solve.
[13] We have repeated this process using a B tranche as the equity and the results are quite similar. An optimisation was also repeated with a more granular 18-state ratings scale and the results were qualitatively similar.

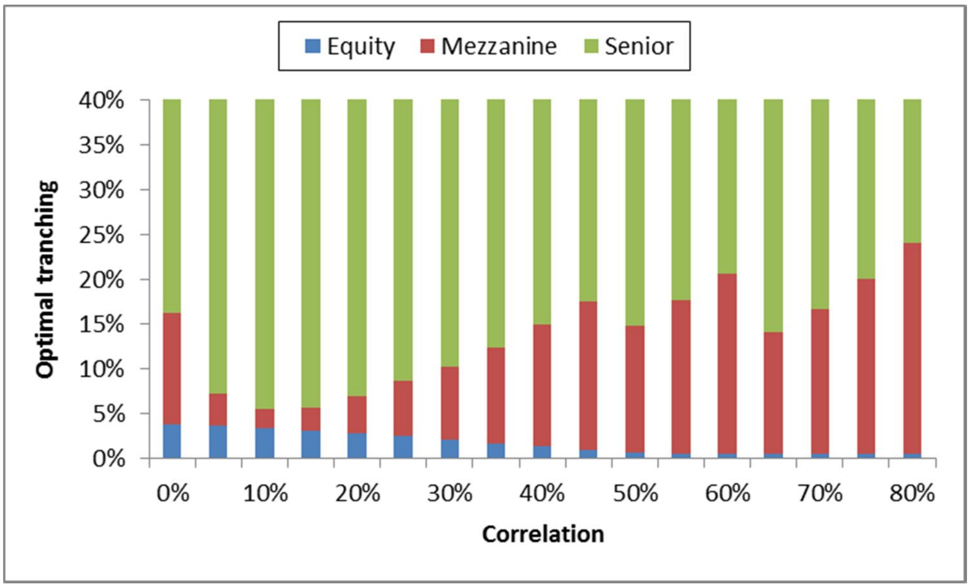

Figure 5. *Tranching corresponding to the optimised portfolio shown in Figure 4.*

5. Conclusion

In this paper, we have shown that a corporate CDO seems to work due to the risk aversion characteristics of credit investors and the rating process for CDO tranches. Under these conditions, everyone can gain from CDOs (issuers, investors and third parties). An issuer profiting from a CDO structure does not necessarily imply that the investors are getting a bad deal or vice versa.

Suppose one took the view that a CDO cannot possibly create value and that an issuer would only gain at the investor's expenses or vice versa. It follows that there is only one explanation as to what fuelled the growth of the CDO market and led partially to the global financial crisis. This is the reliance of expected loss as driving the ratings process for tranches. Banks did not need complex financial engineering to profit from CDOs as value was already there inherently due to the risk aversion of investors and rating agency methodologies.

One obvious rationalisation of the results shown is that investors in the tranches of a CDO should demand different returns compared to the same ratings for the underlying portfolio. This is consistent with Gibson (2004), who argues that highly leveraged CDO tranches could be much more risky than comparable corporate bonds. Coval et al. (2009) have also argued that senior tranches of CDOs contain significant systemic risk and should therefore command higher risk premia since they are "economic catastrophe bonds". The aforementioned paper by Brennan et al. (2009) shows a "marketing gain" in a CDO which is greater when there is increased systemic risk.

Such effects would dilute the value within a CDO although not obviously to a point when they no longer serve any useful economic function. Interestingly, the recent revival of the synthetic CDO seems to have been dampened by the unwillingness of investors to accept the prices offered on the super senior piece[14]. It could be therefore that a better appreciation of the systemic risk in senior tranches has made CDOs unprofitable.

[14] See "Bid to relaunch synthetic CDO unravels", T. Alloway, T. Braithwaite, D. McCrum, Financial Times.

With hindsight bias, it is possible to criticise rating agencies for making incorrect assumptions about default probabilities and correlations. However, when CDOs are rated via expected loss based quantitative assumptions then, due to the risk aversion characteristics of investors, it seems likely that a CDO will always “work” when considering tranches across the entire capital structure. CDOs were clearly a casualty of the crisis but there is not obviously any arbitrage in an arbitrage CDO.

References

Altman, E.I., 1989, “Measuring Corporate Bond Mortality and Performance”, Journal of Finance, Vol. 44, No. 4, pp. 909-922.

Brennan, M. J., J. Hein, and S-H. Poon, 2009, “Tranching and rating”, European Financial Management, Vol. 15, No. 5, pp 891-922.

Coval, J., J. Jurek., and E. Stafford, 2009, “Economic catastrophe bonds,” American Economic Review, 99(3), 628—66.

Fons, J.S., 1987, “The Default Premium and Corporate Bond Experience” Journal of Finance, Vol. 42, No. 1, pp. 81-97.

Gibson, M., 2004, “Understanding the risk of synthetic CDOs”, Finance and Economics Discussion Paper, 2004–36, Federal Reserve Board, Washington DC.

Giesecke, K., F.A. Longstaff, S. Schaefer, and I. Strebulaev, 2010, “Corporate Bond Default Risk: A 150−Year Perspective.” NBER Working Paper No. 15848.

Hamilton, D., P. Varma, S. Ou., and R. Cantor, “Default & Recovery Rates of Corporate Bond Issuers, A Statistical Review of Moody's Ratings Performance, 1920-2003”, 2003, Moody’s Investor Research, January.

Hull, J., M. Predescu and A. White, 2005, “Bond Prices, Default Probabilities and Risk Premiums” Journal of Credit Risk, Vol. 1, No. 2, pp. 53-60.

Standard & Poor’s, 2002, “Global Cash Flow and Synthetic CDO Criteria“, S&P Structured Finance, March.