Rise and fall of synthetic CDO market: lessons learned

Juliusz Jabłecki*

Abstract

This paper uses a unique data set of more than 1,000 synthetic CDO deals to describe typical structures, their pricing and performance with the aim of identifying the factors behind the spectacular collapse of this important segment of structured credit market in late 2008. The data suggests that mark-to-market losses on many synthetic CDO tranches were much more significant than in case of simpler, lower-rated products despite the former experiencing little or no impairment of the notional. The losses were driven instead by the concentration of relatively limited number of defaults in a short period of time, suggesting that pre-crisis pricing must have seriously underestimated such risk of default clustering. In view of the post-crisis pick-up in synthetic CDO issuance, the paper attempts to heed this lesson and offer a simple factor model of default correlation in the spirit of Marshall-Olkin that is naturally suited to capturing the temporal dimension of default dependencies that have been crucial for CSO investors. The model allows building a rich dependence structure capable of consistently fitting standardized iTraxx and CDX index tranches, which makes it ideal for pricing bespoke CDOs.

1 Introduction

Synthetic CDOs – or collateralized synthetic obligations, CSOs – are derivative instruments which allocate the default risk of a pool of underlying credits to different tranches with different seniorities upon default. What makes synthetic CDOs "synthetic" is the fact that the referenced pool of assets is a basket of single name credit default swaps (CDS) rather than a portfolio of cash bonds of the relevant companies. According to an account popularized in the movie "Big Short" – winner of the 2016 Academy Award for the screenplay adapted

*Faculty of Economic Sciences, University of Warsaw and National Bank of Poland; e-mail: juliusz.jablecki@nbp.pl. I would like to thank the European Credit Strategy and Derivatives Team at JP Morgan for providing me the CSO data and insightful commentary without which this study could not have been completed. I also thank Stephen Figlewski for helpful comments on an earlier draft. That said, the responsibility for the views expressed herein – including potential errors – is entirely mine.

from Michael Lewis's book by the same title – "synthetic CDO was the atomic bomb with a drunk President holding his finger over the button." Such a colorful – if not very substantive – representation seems to capture the by now almost consensus view that write-downs on synthetic CDOs were at the center of "the Panic of 2007." Indeed, countless articles in popular press and the blogosphere have been written to demonstrate how synthetic CDOs apparently spread the contagion of toxic assets throughout the financial system, triggering the global financial crisis. And the belief that structured finance significantly underperformed due to outsize downgrades and credit losses found its way even to academic literature (see e.g. Antoniades and Tarashev, 2014). Yet, despite general stigma around structured credit products, rigorous empirical evidence on the performance of synthetic CDOs is surprisingly scant. This paper tries to fill this gap, and does so in two ways.

First, it seeks to enhance understanding of the performance of the synthetic CDO market during the recent crisis. To this end it employs a unique data set – graciously provided by JP Morgan[1] – with details of over 1,000 bespoke synthetic CDO deals issued between 2002 and 2011. Perhaps the most surprising finding revealed in the data is that the historical loss rates of bespoke portfolios have tended to be relatively small, with an average default rate of just 2.5% and no portfolios suffering default-related losses of more than 15%. Similarly, the average rating downgrade in each name in the database was merely 1.5 notches. Despite that, the downgrades of CSO tranches have typically been much more severe leading to considerable mark-to-market losses on most structures (although interestingly quite a few tranches investigated actually outperformed S&P 500). Hence, the major driver of mark-to-market losses was not the scale of actual defaults in a particular portfolio, but rather the perceived concentration of defaults in time, across multiple portfolios.

The second goal of the paper is to go beyond economic history and discuss the pricing methodology of synthetic CDO tranches. The market standard before the crisis was the so called Gaussian copula. However, as recently pointed out by Morini (2011), Brigo and Capponi (2010) and Gatarek and Jabłecki (2015), the copula model is by construction poorly suited to handling default concentration in time, which is crucial in explaining historical performance of synthetic CDO tranches as well as understanding pricing patterns and risk-management strategies in newly issued CSOs. While the market has moved to pricing in a significantly higher correlation for CSO tranches than before the crisis, the Gaussian copula model – still widely used in the industry – does not in general allow to control the risk of joint defaults in a short period of time using the correlation coefficient. In fact, in some cases the probability of joint defaults can decline as a function of correlation. Although these problems are generally known to practitioners and sometimes can be circumvented

[1] The data was initially presented in two notes to clients "Structured Credit After the Crisis" (January 28, 2010) and "CSO v2.0: The New Synthetic CDO Offering" (March 16, 2015).

by making more or less straightforward tweaks to the model, they indicate major source of model risk inherent in using the Gaussian copula. With this in mind, the present paper suggests a viable alternative to the Gaussian copula approach in the form of an intesity-based generalized Marshall-Olkin model, introduced previously in the context of estimating residual credit risk in a repo portfolio (Gatarek and Jabłecki, 2015). Unlike other alternatives to the Gaussian copula model, in particular the so called perfect copula of Hull and White (2005), Hull and White (2006) and Hull and White (2010), or the generalized Poisson loss model (Brigo and Capponi, 2010), the model proposed herein is a hybrid of top-down and bottom up (copula) models, consistent with equity, currency and interest rate models. Similarly as in Hull and White (2008), although in a completely different setting, the model allows for dynamic treatment of credit risk and in particular implies a dynamics for default correlation. As in classic intensity-based approaches (Giesecke, 2003; Morini, 2011), defaults of firms are driven in the model by firm-specific as well as economy-wide shocks. However, in contrast to previous approaches, the systematic factor is represented as a sequence of increasing random variables. This characterization has the interpretation of an invisible chain of dependencies running through the whole economy whereby a systematic default of i-th obligor causes defaults of all more systematically risky names. The model is arbitrage-free and allows to build an almost arbitrarily rich correlation structure, with a multiple systematic factors, each representing a different default pattern of names in the economy. A technique of calibrating the model to iTraxx and CDX tranches is discussed. The model produces a very good fit to market prices (well within bid-ask spreads), allowing for consistent valuation of bespoke CSO deals.

The rest of the paper proceeds as follows. Section 2 presents the empirical data on the historical performance of CSO tranches from the JP Morgan database. Section 3 discusses CDO pricing methodology along with the systematic factor model. Section 4 discusses calibration and Section 5 briefly concludes.

2 Historical performance of synthetic CDOs

A synthetic CDO – or CSO – is a contract that allocates the default risk of a portfolio of credits to different tranches.[2] A CSO involves a protection buyer, protection seller, a portfolio of equally weighted credit default swaps on an underlying pool of reference names $1, 2, 3, ..., d$, and finally an attachment point A and a detachment point B, determining the beginning and end of the portfolio loss tranche covered (with $0 \leq A \leq B \leq 1$). Tranched

[2]See Brigo and Capponi (2010) and O'Kane (2008) for an extensive overview.

loss at time t, $L_t^{A,B}$, can thus be formally expressed as

$$L_{A,B}(t) = \frac{1}{B-A}\left[(L(t)-A)^+ - (L(t)-B)^+\right], \tag{1}$$

where

$$L(t) = \sum_{i=1}^{d} \mathbb{1}_{\{\tau_i \leq t\}} \frac{(1-R_i)}{d} \tag{2}$$

denotes normalized portfolio cumulative loss, τ_i is the default time of name i and R_i is the recovery rate. For example, a 3-6% tranche will not suffer impairment on its notional as long as losses on the underlying portfolio stay below 3%; it will take a loss if portfolio losses are between 3-6% and will be entirely wiped out if portfolio losses exceed 6%. A tranche is called an equity tranche if its attachment point equals 0%, as a super-senior tranche if its detachment point equals 100%, and as a mezzanine tranche if it is neither of the above. It follows that any mezzanine tranche can be represented as a difference between two equity tranches: one with the detechment point of the original tranche and one detaching at the attachment point of the original tranche:

$$L_{A,B}(t) = \frac{1}{B-A}\left[(L(t)-A)^+ - (L(t)-B)^+\right] = \frac{1}{B-A}\left[BL_{0,B}(t) - AL_{0,A}(t)\right] \tag{3}$$

Both standardized and bespoke CSO are traded. The former reference the pool of credits making up US and European CDS indexes (CDX and iTraxx, respectively) with a standardized set of attachment and detachment points. The latter in turn are custom made whereby the composition of the underlying CDS portfolio as well as tranche subordination, width and other details can be chosen to fit investors' needs.

As pointed out by Jarrow (2012), CSOs provide a transaction cost minimizing method of implementing a broad range of technical and fundamental credit views, and thus help to facilitate efficient allocation of capital in the economy. Although fundamentally a long risk tranche position resembles a long risk CDS index position, dependence of CDO tranches on the index spread is just the first-order ("delta") component of a more complex risk profile. Due to their unfunded nature and tranche structure – which makes them sensitive not only to average portfolio CDS spread but also the distribution and correlation of credit quality – CSOs allow investors to leverage their exposure to portfolio losses. To see this, note that an investor selling protection on the iTraxx Euro S.24 index suffers losses strictly proportional to the number of defaults. With 125 names in the index, a default of 5 names with 40% recovery rate (60% loss on a given credit) implies an impairment of 2.4% on the notional. In contrast, for a counterparty selling protection on the same index through a 0-3% equity tranche, the default of 5 names in the pool would imply a loss of 80% (i.e. the contractual

Exhibit 1. Synthetic CDO market: global outstanding and issuance

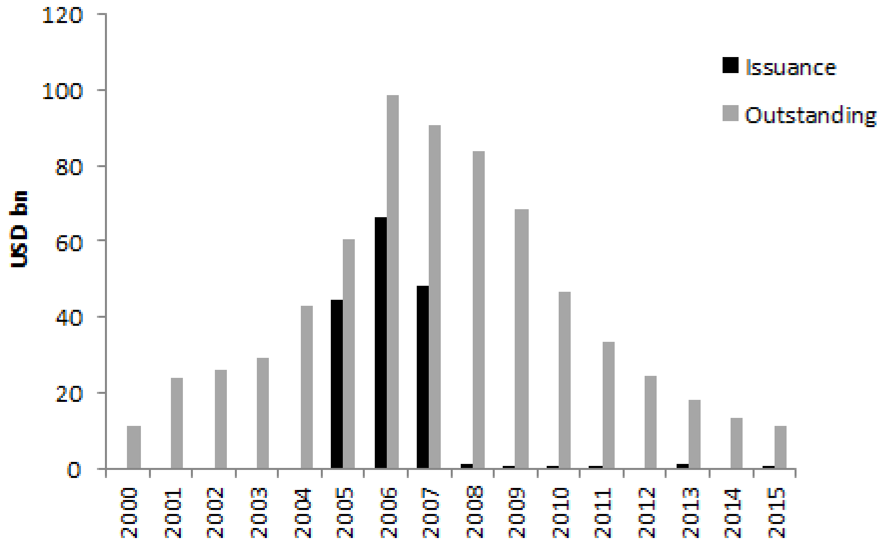

Source: SIFMA

spread would be paid on just 20% of the notional). This feature is obviously reflected in the relative pricing of the equity tranche and the CDS index. For example, on April 8, 2016, the spread on iTraxx Euro S.24 5Y 0-3% tranche was quoted at around 1318 bp[3] while the spread on the entire index was just 77 bp. This means that an investor with EUR 100 million capital looking to secure a return of about 10% could simply invest the entire capital in the iTraxx Euro equity tranche. Achieving the same return by investing in the index itself would require increasing leverage to 15:1 by borrowing EUR 1,500 million. Obviously such an expansion in the gross notional exposure would drastically increase investor's vulnerability, as the maximum potential loss would in that case be well above the EUR 100 million capital. In contrast, in case of the investment in the equity tranche, the maximum loss would be kept at EUR 100 million – even if the entire portfolio of index names were to default.

Given the OTC nature of synthetic CDOs determining the size of the market is challenging. Data compiled by Securities Industry and Financial Markets Association (SIFMA) based on input from a broad range of sources[4] suggests that the value of outstanding synthetic CDOs globally stood at just about USD 11 billion in 2015 down from USD 100 billion in 2006 which attests to the spectacular rise and subsequent collapse of the CSO market. A similar picture emerges from SIMFA's issuance data according to which over USD 66 billion worth of new deals was issued in 2006 compared to a mere USD 1.3 billion in 2013 and a mere USD 400 million in Q2 2015. One should bear in mind, though, that these statistics

[3]Using a more common quotation convention for equity tranches, the 1318 bp spread is equivalent to 43% upfront payment and a running spread of 100 bp.

[4]These include SIFMA members (securities firms, banks and asset managers) as well as Bloomberg, Dealogic, Thomson Reuters, and the main rating agencies.

probably grossly underestimate CSO market size, because they include only fully funded structures (i.e. those requiring the deposit of cash to an SPV at the inception of the deal) and exclude bespoke single-tranche notes. Indeed, data compiled by BNP Paribas indicates that as much as USD 20 billion synthetic CDO deals were issued in 2014 up from less than USD 5 billion in 2013 – a stark indication of the market coming back to life amid investors' hunt for yield in a zero-interest rate world.

To better understand the factors behind the rise and fall of synthetic CDOs between 2004-2009 it is essential to investigate the performance of standard CSO deals. Such historical information is virtually unavailable, though, as bespoke CSOs are OTC deals each with its own unique features. However, based on JP Morgan database of CSO deals – containing detailed information on more than 1,000 transactions – it is possible to identify typical underlying CDS portfolios and typical tranche attachment and detachment points. This exercise has been performed by JP Morgan credit strategists who have graciously allowed me to use their data in this publication. The analysis led to a formulation of 84 theoretical CSOs – based on 12 different CDS portfolios and 7 representative tranches most closely resembling those in the JP Morgan database. As for the database itself, it is difficult to say precisely how representative of the whole synthetic CDO market it actually is. However, in the absence of a market-wide database of such OTC deals, the sample provided by JP Morgan can be considered a useful proxy for describing the general state of the market for the following reasons: (i) the sample consists of over 1,000 different transactions, and so is fairly large; (ii) the database includes not just deals originated by JP Morgan, but also those inherited from Bear Stearns after the takeover, as well as any others purchased over time, thus covering different business models and practices; (iii) JP Morgan was one of the first banks to originate these transactions and as such had a variety of different clients and portfolios; (iv) CSO portfolio construction is to a large extent determined by the client, with JP Morgan typically hedging its own exposure, hence the sample should not be very restrictive on the names included and structures offered.

The construction is based on the frequency with which credits appear in the JP Morgan CSO database, taking into account their ratings, regions and sectors. Portfolios come in three different sizes (50, 100 and 200 names), have different risk profiles (high grade vs. high yield), as well as different geographic and sectoral concentration. Added as a benchmark is also a generic portfolio based on 100 most widely referenced corporate obligors in rated US CSOs, as reported by Standard & Poor's[5] (Exhibit 2). The selected portfolios are then tranched according to the most typical attachment and detachment points found in the JP Morgan database as also reflected in Creditflux data on CSO market activity in 2006-2008.

[5] See "The Most Widely Referenced Corporate Obligors In Rated U.S. Synthetic CDOs," Dec 16, 2008, Standard & Poor's.

Exhibit 2. Theoretical underlying CDS portfolios representative of the pre-crisis CSO market

Name	Description	# Credits	Portfolio loss
Generic	Top 100 names	100	1.70%
Small	Top 50 names	50	3.30%
Large	Top 200 names	200	1.00%
All High Grade	Top 100 IG names	100	1.70%
High Yield Biased	Top 85 IG and top 15 HG	100	2.60%
Top S&P	Top 100 names from S&P	100	1.50%
US Centric	Top 100 US names	100	1.70%
EU Centric	Top 100 EU names	100	0.30%
Asia Centric	Top 45 US, 40 EU and 15 Asia names	100	2.00%
Financials Biased	Top 65 Industrials and 35 Financials names	100	1.70%
Industrials Biased	Top 90 Industrials and 10 Financials names	100	0.30%
Sovereign Biased	Top 70 Industrials, 15 Financials and 15 Sovereign names	100	1.90%

Source: JP Morgan data

Exhibit 3. Attachment and detachment points for typical pre-crisis CSOs

Tranche name	Att-Det (%)
Equity	0-3
Thin Junior Mezzanine	3-4
Thick Junior Mezzanine	3-7
Thin Senior Mezzanine	7-8
Thick Senior Mezzanine	7-10
Senior	30-100
Portfolio	0-100

Source: JP Morgan data.

Exhibit 4. Risk-adjusted performance of the representative CSO tranches (December 2006-January 2010)

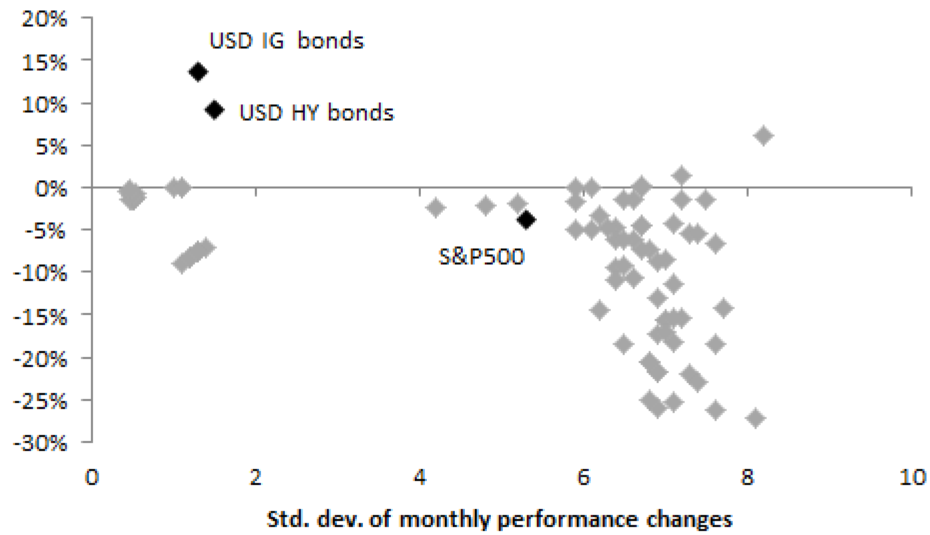

Note: calculations based on JP Morgan and Bloomberg data, assuming USD 100 notional on each portfolio; USD IG and HY bonds denote the iBoxx index of US liquid investment grade and high yield corporate bonds respectively.

To represent as closely as possible the diversity of products used, an equity, a supersenior and four types of mezzanine tranches are considered: junior thin and thick and senior thin and thick (Exhibit 3).[6] Different mezzanine tranche widths were popular before the crisis since – for a given subordination level – manipulating tranche width allowed to fine tune investors' leverage and thus also the spread received on a deal. Along with the trivial 0-100 tranche comprising the entire portfolio, this makes up 7 different tranching patterns allowing for the construction of 84 CSOs. Each of the 84 representative CSOs is thus formed by taking one of the 12 CDS portfolios and tranching it in one of the 7 specified ways, assuming issuance in December 2006 and 5-year maturity (to cover the crisis period). Risk-adjusted performance of all portfolios – i.e. average monthly return divided by standard deviation of monthly returns – is presented in Exhibit 4 which also features the performance of a long-only S&P 500 portfolio as well as standard bond portfolio benchmarks.

Although the performance metric chosen (average return scaled by standard deviation of returns) is a very simple one and in particular does not comprehensively account for CSO risk, it does nonetheless show that the overwhelming majority of CSO portfolios have produced negative risk-adjusted returns over 2006-2010. Even those few tranches that did deliver a positive return – e.g. the equity tranche of the EU centric CDS portfolio – have

[6]The 10-30 senior tranche is omitted, given that such tranches were not very popular before the crisis (with about 5% share by trade count in Creditflux market data) and because its behavior is not significantly different from the supersenior 30-100 tranche.

Exhibit 5. Average bespoke portfolio loss by year of issuance (left-hand panel) and the cumulative distribution of bespoke portfolio losses (right-hand panel)

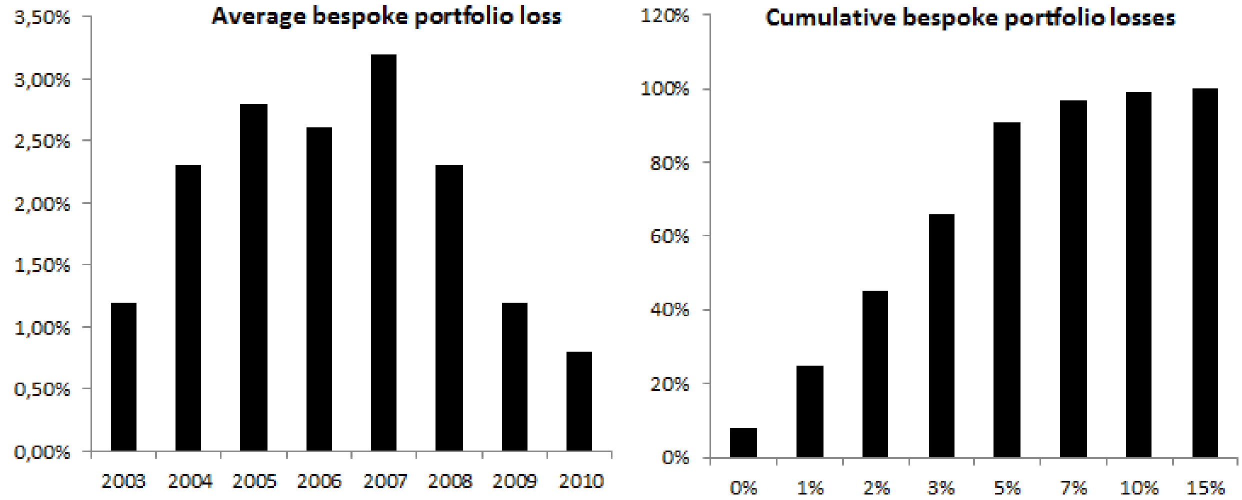

Source: JP Morgan CSO database.

done so with considerably greater risk than cash bonds. And this, of course, assuming CSO risk is captured only by standard deviation of returns – using a more sophisticated metric would produce even more depressing results. Interestingly, though, poor mark-to-market valuations have not been driven by actual defaults. The average credit loss (after recovery) on the CDS portfolios considered was just 1.62%, with a maximum loss of 3.30% on the small portfolio of 50 high yield corporates. Only one tranche in the sample was completely wiped out (0-3% on the Small portfolio) while 58 (of which 46 mezzanine) did not suffer any impairment.

A very similar pattern emerges from the analysis of credit losses on actual bespoke portfolios in the JP Morgan database. Exhibit 5 shows the average bespoke portfolio loss (by year of issuance) and the cumulative portfolio loss distribution. Consistently with the results for theoretical portfolios, losses on actual bespoke portfolios have been relatively modest, with the highest average loss of 3.2% on portfolios issued in 2007. In 66% portfolios losses were lower than 3% and only 3% of portfolios have suffered losses greater than 7%. Importantly, none of the 1,000 portfolios tracked by JP Morgan has experienced losses greater than 15%. By implication, none of the senior and upper mezzanine tranches (with attachment point above 7%) have suffered credit losses. Only 5% of junior mezzanine and 14% of equity tranches have been wiped out completely with the average loss on both kinds of tranches equal 11.8% and 48.5% respectively (Exhibit 6).

In conclusion, synthetic CDOs have underperformed plain vanilla credit products (even those initially lower rated ones) on a mark-to-market basis, delivering negative returns even though losses on CDS portfolios have been rather modest and many tranches have not ex-

Exhibit 6. Share of bespoke tranches in the JP Morgan database to have suffered losses

Tranche	Zero	Partial	Full	Average loss
Equity	12%	75%	14%	48.50%
Junior Mezzanine (Att<7%)	78%	16%	5%	11.80%
Upper Mezzanine (Att>7%)	99%	0%	0%	0.60%
Senior	100%	0%	0%	0.00%

Source: JP Morgan data.

perienced actual impairments. This suggests that, contrary to popular accounts, the drastic price erosion in the CSO market – which led to an effective cessation of issues by 2009 (Exhibit 1) – was not driven by widespread credit losses, but rather the concentration of relatively limited losses in time.[7] The initial pre-crisis pricing and risk assessment of synthetic CDOs must have seriously underestimated such risk of default clustering, a lesson that should be heeded in view of the post-crisis pick-up in CSO issuance. Thus, the following section looks at CSO pricing methodology, tries to pinpoint the source of past errors and hints at a potential solution.

3 CSO pricing

The valuation of synthetic CDO tranches is driven by two parameters: the average spread on the underlying CDS portfolio and default correlation among the referenced credits, or – more precisely – dependency among the defaults of respective names. Pricing CSO tranches starts from the assumption that the overall credit risk in the underlying portfolio – as reflected in the spread – is fixed. Tranches then divide that risk into components related to idiosyncratic default drivers and systematic default drivers. Since credit risk in the entire CSO structure must be equal to that of the underlying CDS portfolio (otherwise straightforward arbitrage would be possible), a change in the spread of one tranche must be reflected in a corresponding change in the spread of another tranche somewhere else in the capital structure in order to keep the average portfolio spread unchanged.

In turn, the pricing of portfolio tranches reflects the relative weight assigned by investors to idiosyncratic and systematic risk drivers, as captured in default correlation. A low level of default correlation indicates that companies are more independent and defaults are likely to be largely isolated, i.e. idiosyncratic. Higher default correlation implies greater likelihood of companies defaulting together, which translates into higher prices of credit protection on more senior tranches. Given an industry standard pricing model – the Gaussian copula –

[7] As pointed out by Brigo and Capponi (2010), in a time span of just one month – between September 7 and October 8, 2008 – seven credit events occurred involving major financial institutions, some of which happened on the very same day.

Exhibit 7. Implied correlation for equity and mezzanine tranches on CDX Main portfolio

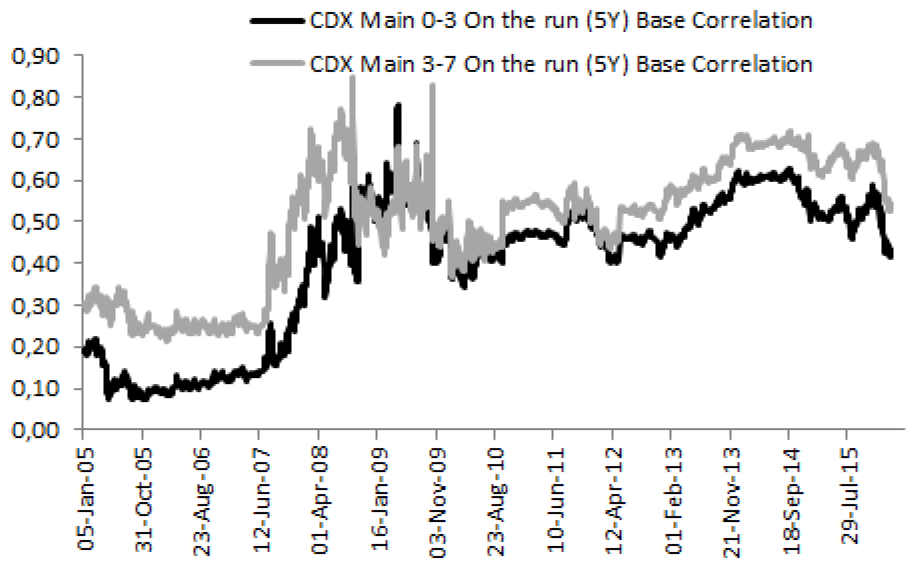

index spreads and tranche prices can be reverse-engineered to produce a correlation figure consistent with them. By analogy with the convention in option markets, this value is called "implied correlation" of CSO tranches.

Exhibit 7 shows implied correlations of equity and junior mezzanine tranches on the CDX Main portfolio. The crisis has triggered a clear regime shift in implied correlations, which – in case of equity tranches – have increased since mid-2007 by roughly a factor of five. It is this marked increase in CSO implied correlations that has triggered massive mark-to-market valuation losses on most structures, despite – as has been argued above – the overall sound credit performance of the referenced obligors. The increase in correlation was supposed to reflect market participants' concern about the previously underpriced risk of default clustering. However, within the Gaussian copula framework – applied before the crisis and unfortunately still widely used as the default valuation model[8] – correlation does not have a clear and intuitive relationship with the probability of losses concentrated in time, i.e. the very risk that the market has been trying to price in.

3.1 The market standard: Gaussian copula approach

In the Gaussian copula model[9] each obligor i is assigned a standard normal variable A_i and a time-dependent default threshold $z_i(T)$. It is assumed that default occurs before time T

[8] A variant of the Gaussian copula model is – as of this writing – implemented e.g. in Bloomberg CDO pricer used by practitioners.

[9] The discussion will be purposefully brief as the model has been described in detail elsewhere; see e.g. Andersen and Sidenius (2004); Hull and White (2004).

if the variable A_i (which itself is not observable and has no dynamics) finds itself below the threshold $z_i(T)$. Formally,

$$\mathbb{P}(\tau_i \leq T) = \mathbb{P}(A_i \leq z_i(T)) = PD(T), \tag{4}$$

where τ_i is the default time of credit i and $PD(T)$ is the T-year probability of default, calibrated to an obligor's CDS curve.

Correlation structure in the model is introduced by assuming that all A_i are driven by a common systematic factor Z and a name-specific idiosyncratic factor Y:

$$A_i = \rho_i Z + \sqrt{1-\rho_i^2} Y_i \tag{5}$$

where Z and Y_i are standard normal with $\text{cov}(Y_i, Y_j) = 0$ and $\text{cov}(Y_i, Z) = 0$. Now the vector $A = [A_1, A_2, ..., A_n]$ is multivariate Gaussian with correlation given by[10]

$$Corr(A_i, A_j) = \rho_i \rho_j \tag{6}$$

Denote the distribution of A by

$$\mathbb{P}\left(A_1 \leq a_1, ..., A_n \leq a_n\right) = \Phi(a_1, ..., a_n) \tag{7}$$

then the copula is defined as

$$\begin{aligned} C_{A_1,...A_n}(u_1, ..., u_n) &= \Phi\left(\Phi^{-1}(u_1), ..., \Phi^{-1}(u_n)\right) = \\ &= \mathbb{E}\left(\mathbb{P}\left(A_1 \leq \Phi^{-1}(u_1), ..., A_n \leq \Phi^{-1}(u_n)\right) | Z\right) \end{aligned} \tag{8}$$

Since for a given Z the A_i are independent, the copula can be conveniently reexpressed as

$$\begin{aligned} C_{A_1,...,A_n} &= \mathbb{E}\left(\prod_{i=1}^{n} \mathbb{P}\left(A_i \leq \Phi^{-1}(u_i)\right) | Z\right) = \mathbb{E}\left(\prod_{i=1}^{n} \mathbb{P}\left(\rho_i Z + \sqrt{1-\rho_i^2} Y_i \leq \Phi^{-1}(u_i)\right) | Z\right) = \\ &= \mathbb{E}\left(\prod_{i=1}^{n} \Phi\left(\frac{\Phi^{-1}(u_i) - \rho_i Z}{\sqrt{1-\rho_i^2} Y_i}\right)\right) = \int_{-\infty}^{\infty} \prod_{i=1}^{n} \Phi\left(\frac{\Phi^{-1}(u_i) - \rho_i Z}{\sqrt{1-\rho_i^2} Y_i}\right) \phi(Z) dZ \end{aligned} \tag{9}$$

where $\phi(Z)$ is the standard normal density. The probability of joint default is now given by

$$\mathbb{P}(\tau_1 \leq T_1, ...\tau_n \leq T_n) = \int_{-\infty}^{\infty} \prod_{i=1}^{n} \Phi\left(\frac{\Phi^{-1}(PD(T_i)) - \rho_i Z}{\sqrt{1-\rho_i^2}}\right) \phi(Z) dZ \tag{10}$$

[10]This is the so called asset correlation, not to be confused with default time correlation $\rho(\tau_i, \tau_j)$ and default indicator correlation $\rho(\mathbb{1}_{\{i\}} \mathbb{1}_{\{j\}})$. In general, conditional on default probabilities, default correlation is an increasing function of asset correlation (Hanson, Pesaran, and Schuermann, 2008).

Although (10) is very convenient numerically (it involves only one-dimensional integration), this simplicity comes at a cost of reducing the entire dependence structure of a portfolio of n credits to linear correlations determined by $\frac{n(n-1)}{2}$ parameters. In practice, to avoid technical difficulties, popular implementations of the model consisted in setting all ρ_i equal to one parameter ρ. Thus, for example, to capture the entire dependence structure in a CDS index with 125 names (such as iTraxx or CDX), $\frac{125\times 124}{2}$, i.e. 7750 parameters, would be collapsed to a single number ρ. Given that the average spread and spread dispersion in the underlying CDS portfolio are taken as given for the purpose of pricing tranches – and therefore cannot be changed to reflect the supply-demand imbalance for particular tranches – correlation has become the main method for quoting prices of standardized and bespoke synthetic CDOs. Specifically, using the one-factor Gaussian copula, a correlation number is backed out from traded tranche spreads and increased or decreased as needed to reflect demand or supply for a particular tranche. This is done either by assuming that each single $[A, B]$ tranche has its own correlation parameter (compound correlation approach), or – more commonly – by assigning a different unique correlation to equity tranches $[0, A]$ (base correlation approach). However, such practice of implying correlation out of market spreads leads to several important problems.[11]

First, there is an inherent inconsistency in implying correlations (especially in the more common base correlation approach), as two components of the same trade – say tranche $[A, B]$ – are essentially valued with two different models, i.e. Gaussian copulas with two different correlation parameter values. One practical implication of this is that interpolation or extrapolation of these mutually inconsistent values – needed e.g. to price bespoke CSO tranches – can easily produce tranche spreads that are not arbitrage free or may not represent feasible market scenarios.

Second, collapsing the whole complex portfolio dependence structure – as parameterized by 7750 parameters for iTraxx or CDX – to a single number leaves out a host of potentially relevant information. Some names in the portfolio can be very closely related – e.g. those operating in the same sector – whereas other pairs might not have much in common. Even assuming that an average between those high- and low-dependence states can be reliably estimated it is unclear how meaningful it really is.

Third, over and above any problems with inferring correlations or interpolating them in a no-arbitrage way, Gaussian correlation parameter has no stable and predictable relation with the probability of joint defaults driving tranche prices. In particular, depending on the time frame under consideration and obligors' credit spreads, the probability of joint defaults can be an increasing, decreasing or non-monotonic, hump shaped, function of asset correlation (Gatarek and Jabłecki, 2015). The problem with such a non-monotonic pattern is not only

[11] See e.g. Brigo and Morini (2010) or Morini (2011) for a more detailed critique of implied correlation.

that it is difficult to understand from an economic point of view, but also that it makes it in general impossible to know a priori whether – with a given set of PDs – an increase in correlation will increase or decrease risk exposure.

Finally, we have seen above that the mark-to-market losses on CSO portfolios that have driven many investors out of the market, leading to its virtual collapse, were related to investors' fear of default concentration in time. Gatarek and Jabłecki (2015) show using a set of simple examples that the Gaussian copula model is incapable of handling such questions. In fact, for some parameterizations even perfect asset correlation does not produce default clustering.

3.2 A simple factor model of joint defaults

As a remedy for the flaws inherent in the Gaussian copula approach, Gatarek and Jabłecki (2015) have recently suggested a simple factor model for correlated defaults in the Marshall-Olkin tradition, demonstrating its application to residual credit risk measurement of a repo portfolio. This paper shows how to extend that analysis and use the model for pricing CDOs. To avoid repetition, the key ideas behind the approach are summarized below only in the simplest single-factor, deterministic hazard rate setting which will be applied below to pricing CDOs. However, interested readers should consult Gatarek and Jabłecki (2015), which presents the general version of the model with stochastic hazard rates and multiple systematic factors.

Consider d obligors with default times $\tau_1, ..., \tau_d$. Assume that default times are exponentially distributed with parameters $\lambda_1, ..., \lambda_d$, which admit natural interpretation as hazard rates or conditional default probabilities.[12] As in the Gaussian model dependence between default times is introduced by stating that each default can result from the materialization of either an idiosyncratic factor or a systematic factor – whichever hits sooner. Being hit by either factor has the mathematical interpretation of the first jump of a specific Poisson process. Hence, for each obligor i the time until the arrival of the idiosyncratic factor is represented simply by an exponential variable Y_i with parameter λ_i^{idio}. However, unlike in (5) where the systematic factor was a single random variable – in fact one lacking any dynamics – here a systematic factor is an increasing sequence of exponential variables $Z_1 \leq Z_2 \leq ... \leq Z_d$, representing essentially the domino-like sequence in which names are likely to default for systematic reasons. Under such assumptions, individual obligors' default times can be represented as:

$$\tau_i = \min \{Y_i, Z_i\}, \tag{11}$$

[12] Both here and in model calibration hazard rates are time-dependent, although this is suppressed above to simplify notation.

where $Y_1, ..., Y_d$ and $Z_1, ..., Z_d$ are independent exponential variables. Obviously, default times of all obligors, τ_i, are also exponentially distributed with parameters $\lambda_i = \lambda_i^{idio} + \lambda_i^{sys}$ and survival probabilities

$$\mathbb{P}(\tau_i > T) = e^{-\lambda_i T}. \tag{12}$$

The construction of a systematic factor $(Z_1, ..., Z_d)$ proceeds as follows. Suppose, for example, that credit "1" is the most sensitive to the systematic factor, credit "2" is less so but still highly sensitive, and so on, while credit d is the least exposed. Hence, the systematic factor should first trigger the default of name "1", then "2" etc. before ultimately hitting d. To reflect this, assign to each name i a Poisson process $\widetilde{Z}_i$, with intensity λ_i^{sys}, whose arrival triggers the default of credit i, but also – due to the ordering relation – the default of all the more systematically risky names $i-1$, $i-2, ..., 1$ (we assume that Poisson processes $\widetilde{Z}_i$ are independent). Thus, the systematic intensity of each obligor i will be a sum of its own intensity λ_i^{sys} and the intensities of the Poisson processes triggering defaults of more senior names, i.e. $\sum_{j=i}^{d} \lambda_j^{sys}$. This can be formalized by setting

$$Z_i = \min\left\{\widetilde{Z}_i : i \geq j\right\}, \tag{13}$$

where Z_i is the Poisson process representing the total systematic exposure of obligor i. Note that $Z_i \leq Z_{i+1}$ for $i = 1, ..., d-1$, so indeed the family $Z_1, ..., Z_d$ is a systematic factor. Since each obligor is also affected by an idiosyncratic shock Y_i with intensity λ_i^{idio}, default times τ_i are exponentially distributed with parameters $\lambda_i = \lambda_i^{idio} + \sum_{j=i}^{d} \lambda_j^{sys}$ and survival probabilities

$$\begin{aligned}
\mathbb{P}(\tau_i > T) &= \mathbb{P}\left(\min\left\{Y_i, \min\left\{\widetilde{Z}_i, \widetilde{Z}_{i+1}, ..., \widetilde{Z}_d\right\}\right\} > T\right) = \\
&= \mathbb{P}\left(\min\left\{Y_i, \widetilde{Z}_i, \widetilde{Z}_{i+1}, ..., \widetilde{Z}_d\right\} > T\right) = \\
&= \mathbb{P}(Y_i > T)\mathbb{P}(\widetilde{Z}_i > T)\mathbb{P}(\widetilde{Z}_{i+1} > T)...\mathbb{P}(\widetilde{Z}_d > T) = e^{-\lambda_i T}.
\end{aligned} \tag{14}$$

The model has several properties that underscore its usefulness in pricing CSO tranches. First, in line with economic intuition, in the model idiosyncratic defaults tend to be more frequent than systematic defaults. Second, the definition and construction of a systematic factor as an increasing sequence of random variables allows capturing the phenomenon of default clustering in time (in fact in the model only systematic defaults can be multiple). We have seen above that fear of defaults clustered in time is the most likely factor behind the drastic re-pricing of the CSO market throughout the crisis and that Gaussian copula does not capture such phenomena in a reliable way. Third, unlike the Gaussian copula approach, the model does not rely on any asset correlation input which critically determines tranche prices but, as we have seen above, can be related to the probability of joint defaults in unintuitive ways. Instead, our approach starts from a qualitative analysis of systematic risk which then determines the breakdown of market-implied hazard rates of each corporate into idiosyncratic and systematic components, making default correlation the output, not input, of the model. Importantly, default correlation as implied by the model is not a flat number but has time dynamics of its own.

4 Calibration examples

This section shows how to calibrate the simple factor model outlined above to market prices of the most liquid standardized tranches of iTraxx Euro Main, CDX.NA.IG and CDX.NA.HY CDS indices, referencing European and US investment grade and high yield corporate issuers respectively.

4.1 Allocating idiosyncratic and systematic risk drivers

The starting point of model calibration is the allocation of market-implied corporate hazard rates into idiosyncratic and systematic risk components. For example, Gatarek and Jabłecki (2015) suggest introducing a uniform dependence parameter $\rho \in [0, 1]$ which determines the extent to which default times in the economy are systematic and propose to determine ρ on the basis of principal components analysis on the CDS spreads of the referenced corporates. An alternative approach, more convenient in CSO pricing, consists in splitting corporate obligors into groups, or clusters, that are homogeneous in terms of their exposure to systematic risk. Naturally, there is more than one way of doing that, just as there may be different sources of systematic risk in the financial system. Thus, names can be grouped e.g. on the basis of their sector, their sensitivity to the general state of the economy, vulnerability to economic conditions abroad, etc. The model can accommodate that by including different systematic factors, whereby each factor would be characterized by the order in which it triggers the default of a sequence of risky names. In what follows – consistent with the narrative in section 3.2 – only a single systematic factor is considered, which however appears enough to fit current market prices. Although the precise construction of the factor can be optimized to improve calibration (as will be done in section 4.2), it is useful to first consider a step-by-step approach, limiting the exercise for presentation purposes to the subset of largest financial companies included in the iTraxx Euro index (Exhibit 8). As evidenced by the recent global crisis, financials are clearly exposed to a common risk driver. Such systematic dependence can be represented in the model by splitting the 10 selected names into 3 groups similar in terms of their exposure to systematic risk as follows:

- Cluster 1: comprises the three best credits among the euro area financials, i.e. Allianz

Exhibit 8. Selected iTraxx Euro Financials

Name	5Y CDS spread (bp)	S&P rating
Allianz SE	48	AA
Zurich Insurance Co Ltd	48	AA-
Hannover Rueck SE	47	AA-
ING Bank NV	70	A
Aegon NV	85	A-
Aviva PLC	85	A-
Intesa Sanpaolo SpA	133	BBB-
Mediobanca SpA	184	BBB-
UniCredit SpA	185	BBB-
Deutsche Bank AG	198	BBB+

(AA), Zurich Insurance and Hannover Rueck (both AA-). All three insurers have generally robust business and financial profiles and the main downward pressure on their credit profile is likely to originate from truly systemic sources such as a prolonged period of low interest rates (constraining profitability of all financials) or from risks such as natural catastrophes and pandemic;

- Cluster 2: comprises two insurers (Aviva, Aegon), a notch riskier than their highest rated peers, along with a well-diversified European universal bank (ING); this is a mezzanine risk group;

- Cluster 3: comprises the most risky yet still investment grade European banks, three of which (Intesa, Unicredit, Mediobanka) have large direct exposure to Italian government which makes them vulnerable to any reemergence of tensions in euro area sovereign debt markets, and one is Germany's mega-bank (Deutschebank) which has been struggling to maintain profitability over the past several years, and is currently particularly sensitive to investors' risk perception changes; this is the most risky group.

Building on the general idea presented in section 3.2, each cluster will now be associated with its own systematic risk driver Z_1, Z_2, Z_3. As in (13), each risk driver is represented as a Poisson process whose arrival triggers the default of all cluster members and those in the more junior clusters. The default intensity of systematic risk drivers is linked to the CDS spread of the least risky name in each cluster. For example, Z_1 represents the systematic risk driver of Cluster 1 and its default intensity is equal to that of Allianz. The first jump of Z_1 – associated e.g. with a Europe-wide natural disaster or a military conflict – triggers the default of the three highest rated insurers, but in a domino-like way also hits entities in Clusters 2 and 3. Similarly, a shock affecting Z_2 members – e.g. a general drop in European interest rates deep into the negative territory – causes defaults of the riskier insurers (Aviva,

Exhibit 9. Stripping of market-implied 5Y hazard rates into various shock components for selected iTraxx Euro financials

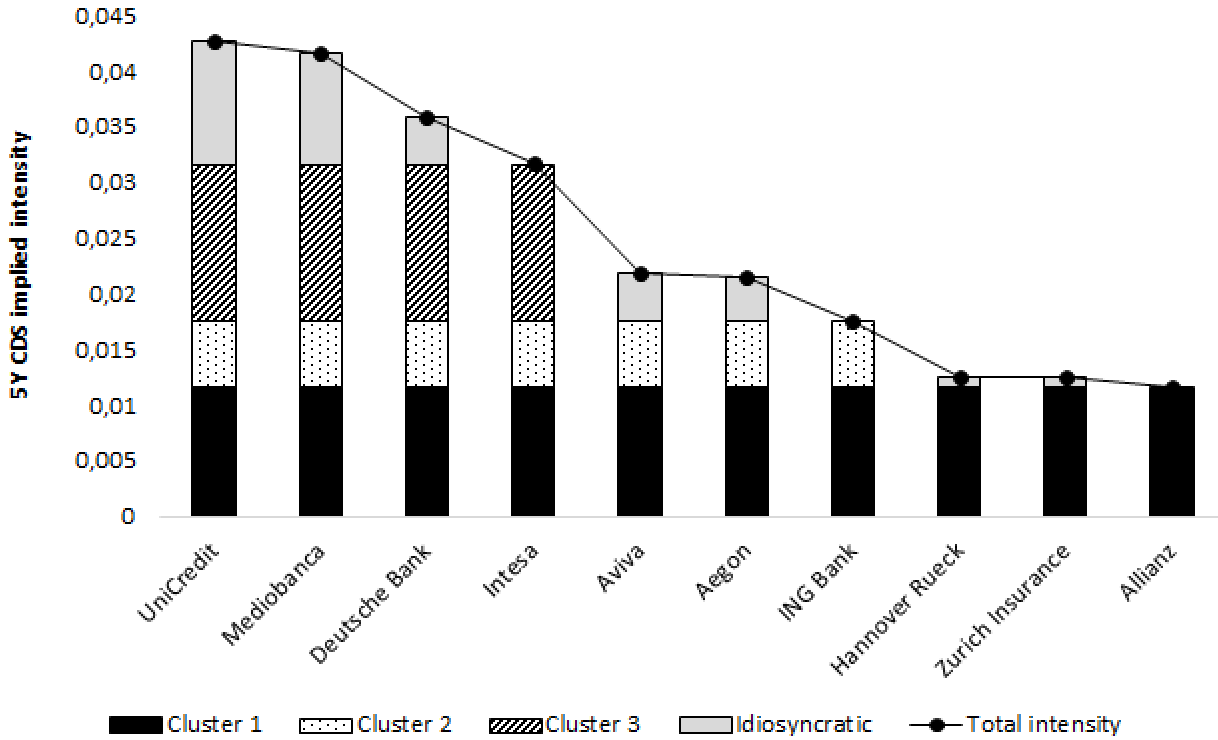

Aegon) as well as all the banks.

$$\begin{aligned} \lambda_{Z1} &= \lambda_{ALV} \\ \lambda_{Z2} &= \max(\lambda_{ING} - \lambda_{Z1}, 0) \\ \lambda_{Z3} &= \max(\lambda_{ISP} - \lambda_{Z2} - \lambda_{Z1}, 0) \end{aligned} \tag{15}$$

Once the systematic part of credit spreads is allocated, the residual spread component can be attributed to idiosyncratic risk.[13] Such a decomposition ensures consistency, i.e. guarantees that the intensity of systematic and idiosyncratic risk factors for each name will add up to that name's overall default intensity as stripped from the CDS curve. Exhibit 9 shows the result of this process using corporate CDS curves as of March 31, 2016.

Note that, as already hinted above, the proposed construction does not require making any assumptions about default correlation. A full representation of the default structure among the selected companies requires instead postulating the order in which the names default if hit by a systematic shock. This means that default correlation between any two names over any time horizon – as implied by their CDS curves and the chosen default pattern – can be simulated and becomes model output, rather than input. The results for our selected

[13]This construction by default assumes that cluster leaders (ALV, ING, ISP) have no idiosyncratic risk, but this can be altered if necessary.

Exhibit 10. Simulated default correlations for selected iTraxx Euro financials (5 years, 10,000 Monte Carlo runs)

	UCG	MB	DB	ISP	AV	AGN	ING	HNR	ZUR	ALV
UCG	100%	63%	67%	77%	42%	42%	45%	35%	35%	35%
MB	63%	100%	69%	78%	42%	44%	47%	36%	36%	36%
DB	67%	69%	100%	87%	46%	48%	52%	40%	41%	41%
ISP	77%	78%	87%	100%	53%	56%	60%	46%	47%	47%
AV	42%	42%	46%	53%	100%	74%	85%	65%	67%	67%
AGN	42%	44%	48%	56%	74%	100%	87%	67%	67%	68%
ING	45%	47%	52%	60%	85%	87%	100%	77%	77%	79%
HNR	35%	36%	40%	46%	65%	67%	77%	100%	95%	98%
ZUR	35%	36%	41%	47%	67%	67%	77%	95%	100%	97%
ALV	35%	36%	41%	47%	67%	68%	79%	98%	97%	100%

iTraxx Euro financials – based on 10,000 Monte Carlo runs – are shown in Exhibit 10. In line with intuition, average correlations in Cluster 1 are clearly much higher than in Cluster 3. Also, correlations between the highest rated names are close to 100% and by some 30 percentage points higher than between the most risky names in the sample. This reflects the relatively higher share of idiosyncratic risk in the respective hazard rates (on average 20% in Cluster 3 vs. 8% in Cluster 1). Even this relatively simple exercise demonstrates how complex default correlation can be and that collapsing the matrix in Exhibit 10 to a single number by taking the average between the "highs" and the "lows" – as would be required in the Gaussian copula approach – would drastically misrepresent the otherwise complex reality.

The method outlined above of simulating jumps of the Poisson processes associated with each factor can be used naturally to price tranches of the selected basket of iTraxx financials. Although tranching a 10-name CDS basket would probably not be practical, size was selected here for presentation purposes only. More realistic extensions are straightforward and described in greater detail in the following section.

4.2 Calibration to iTraxx Euro and CDX indices

The pricing of CSOs using the model presented in section 3.2 is relatively straightforward and can be greatly facilitated by using a proper approximation for the process counting defaults in the underlying CDS pool. To fix ideas, consider a CDO on a CDS portfolio referencing d names (in the case of Itraxx Europe $d = 125$) with a uniform deterministic hazard rate R and loss-given-default $LGD = 1 - R$. To facilitate presentation assume also that the notional of each obligor is LGD^{-1}. As explained above, for any tranche $[A, B]$ with attachment point A and detachment point B, the protection seller in a CDO pays tranched loss increments at defaults. In exchange, the protection buyer pays a periodic premium K (spread) on the survival amount on dates $T_1, T_2, ..., T_N$. For equity tranches (and recently also mezzanine) part of the premium can also be paid at inception as an upfront $U_{A,B}$ to limit counterparty credit risk faced by the protection seller. The survival amount – outstanding notional – associated with a given tranche equals

$$1 - L_{A,B}(t) = \left(B\frac{d}{LGD} - N(t)\right)^{+} - \left(A\frac{d}{LGD} - N(t)\right)^{+} \tag{16}$$

where $N(t)$ counts the number of defaults, i.e.

$$N(t) = \sum_{i=1}^{d} \mathbb{1}_{\{\tau_i < t\}} \tag{17}$$

The discounted protection leg payoff at time $t = 0$ can be written as

$$\int_0^{T_N} D(0,t) dL_{A,B}(t) \tag{18}$$

where $D(s,t)$ is the discount factor evaluated at s for time t. Similarly, the discounted premium leg, at $t = 0$ is

$$K_{A,B} \sum_{i=1}^{N} \delta_i D(0,T_i) \left(1 - L_{A,B}(T_i)\right) + U_{A,B} \tag{19}$$

where $\delta_i = T_i - T_{i-1}$. CSO tranches are quoted in the market by the breakeven value of $K_{A,B}$ that sets the risk-neutral price of a given tranche to zero. Since the tranche value is calculated by taking the risk-neutral expectation of the difference between the premium and protection legs, the break-even spread is given by

$$S_{A,B}(0) = \frac{\mathbb{E}\left(\int_0^{T_N} D(0,t) dL_{A,B}(t)\right) - U_{A,B}}{\mathbb{E}\left(\sum_{i=1}^{N} \delta_i D(0,T_i) \left(1 - L_{A,B}(T_i)\right)\right)} \tag{20}$$

which under the assumption of deterministic default-free interest rates typically taken in the market simplifies to

$$S_{A,B}(0) = \frac{\int_0^{T_N} D(0,t) d\mathbb{E}\left(L_{A,B}(t)\right) - U_{A,B}}{\sum_{i=1}^{N} \delta_i D(0,T_i) \mathbb{E}\left(1 - L_{A,B}(T_i)\right)} \tag{21}$$

In view of (16), the risk-neutral expectation of the survival amount can be expressed using the default counting process $N(t)$

$$\mathbb{E}\left(1 - L_{A,B}(t)\right) = \mathbb{E}\left(\frac{Bd}{LGD} - N(t)\right)^{+} - \mathbb{E}\left(\frac{Ad}{LGD} - N(t)\right)^{+} =$$
$$= \sum_{j=0}^{n(B)} \mathbb{P}\left(N(t) = j\right)\left(\frac{Bd}{LGD} - j\right) - \sum_{j=0}^{n(A)} \mathbb{P}\left(N(t) = j\right)\left(\frac{Ad}{LGD} - j\right) \quad (22)$$

where $n(k) = \max\left\{n \in \mathbb{N} : \, n < \frac{kd}{LGD}\right\}$.

It is clear from (21) that the price of CSO tranches and break-even spreads are determined essentially by the default-counting process $N(t)$. Consider first a very rough closed-form approximation of $N(t)$, useful in Monte Carlo pricing, which relies crucially on the breakdown of hazard rates into systematic and idiosyncratic components, with the systematic factor Z represented by an ordered family $Z_1 \leq Z_2 \leq ... \leq Z_d$. Let

$$M(t) = \sum_{i=1}^{d} \mathbb{1}_{\{Z_i < t\}} \quad (23)$$

and

$$N_j(t) = \sum_{i=j+1}^{d} \mathbb{1}_{\{Y_i < t\}} \quad (24)$$

denote the number systematic defaults and idiosyncratic defaults in the set of names $\{j + 1, ..., d\}$. Using the property that Z_i are ordered we easily get

$$\mathbb{P}(M(t) = j) = \mathbb{P}(Z_j < t < Z_{j+1}) = \left(1 - e^{-t\lambda_j}\right) - \left(1 - e^{-t\lambda_{j+1}}\right) = e^{-t\lambda_{j+1}} - e^{-t\lambda_j}. \quad (25)$$

Consequently,

$$\mathbb{P}(N(t) = m) = \sum_{j=0}^{m} \mathbb{P}(M(t) = j)\mathbb{P}(N_j(t) = m - j) =$$
$$= \sum_{j=0}^{m} \left(e^{-t\lambda_{j+1}} - e^{-t\lambda_j}\right) \mathbb{P}(N_j(t) = m - j). \quad (26)$$

Assuming that defaults of individual obligors can repeat themselves,[14] the point process

[14]Since in practice defaults of individual obligors cannot repeat themselves – which is equivalent to randomly drawing defaulting names *without* replacement – this assumption slightly overstates the total number of defaults. A more rigorous approach to estimating such probabilities is based on the Bernoulli triangle as suggested originally by Hull and White (2004).

$N_j(t)$ is a Poisson process with intensity $\sum_{i=j+1}^{d} \lambda_i^{idio}$. Hence,

$$\mathbb{P}(N_j(t) = m) \approx \frac{1}{m!}\left(t \sum_{i=j+1}^{d} \lambda_i^{idio}\right)^m \exp\left(-t \sum_{i=j+1}^{d} \lambda_i^{idio}\right). \tag{27}$$

Such an approximation can be particularly useful in Monte Carlo pricing. A more sophisticated approximation – one that is used in the calibration below – consists in assuming that the process counting all systematic defaults is Markovian, so that with $d \geq n > i$

$$\begin{aligned}\mathbb{P}\left(M(T_{j+1}) = n | M(T_j) = i\right) &= \mathbb{P}\left(M(T_{j+1}) \geq n\right) - \mathbb{P}\left(M(T_{j+1}) \geq n+1\right) = \\ &= \mathbb{P}\left(Z_n < T_{j+1} < Z_{n+1}\right) = \\ &= \exp\left(\int_0^{T_{j+1}} \lambda_{n+1}^{sys}(s)ds\right) - \exp\left(\int_0^{T_{j+1}} \lambda_n^{sys}(s)ds\right) \end{aligned} \tag{28}$$

where Z_i are constructed as in (13). Assuming furtheremore that all idiosyncratic hazard rates are equal[15], i.e. that $\mathbb{P}(Y_i \leq t) = \mathbb{P}(Y_1 \leq t) = 1 - \exp\left(\int_0^t \lambda_1^{idio}(s)ds\right)$ for $i = 2, ..., d$, leads to an analogous Markovian approximation for the process counting all idiosyncratic defaults $N_0(t)$:

$$\mathbb{P}\left(N_0(T_{j+1}) = n | N_0(T_j) = i\right) \approx \begin{cases} (d-i)\left(\mathbb{P}(Y_1 \leq T_{j+1}) - \mathbb{P}(Y_1 \leq T_j)\right) & \text{for } n = i+1 \\ 0 & \text{for } n < i \text{ and } n > i+1 \\ 1 - (d-i)\left(\mathbb{P}(Y_1 \leq T_{j+1}) - \mathbb{P}(Y_1 \leq T_j)\right) & \text{for } n = i \end{cases} \tag{29}$$

This methodology, relying crucially on closed-form formulas for transition probabilities of processes $M(t)$ and $N_0(t)$, will be used below to calibrate the model to three main CDS indices: iTraxx Euro series 24, CDX.NA.IG and CDX.NA.HY (Exhibit 11). The general idea behind calibration mimics that described in section 4.1, however to make use of Markovian approximations idiosyncratic hazard rates are assumed to be uniform.

As explained above, to reduce dimensionality the names making up each index (125 in the case of iTraxx and CDX.NA.IG and 100 for CDX.NA.HY) are grouped into clusters. Then, each cluster is assigned its own systematic shock the arrival of which triggers the default of all names in that cluster as well as the ones in more junior clusters. The number of clusters and their respective sizes are determined so as to ensure best fit of tranches prices to the market. It turns out that fitting the entire tranche structure for iTraxx Euro simultaneously for three maturities (3Y, 5Y, 7Y) requires five clusters consisting of 9, 10, 16, 23, and 125 names, respectively. Four clusters of 10, 16, 23, and 125 names were needed to fit the

[15] This assumption is validated by the fact that the degree of heterogeneity among idiosyncratic hazard rates of individual names does not materially affect CDO prices. What matters instead is the total sum of idiosyncratic hazard rates.

Exhibit 11. iTraxx.EUR.24 tranche mid quotes in basis points and bid-ask spreads in parentheses (March 21, 2016)

Att-Det (%)	3Y	5Y	7Y
0-3	2125 (225)	4050 (236)	5387.50 (500)
3-6	42.90 (200)	737.50 (149)	1550 (400)
6-12	45.90 (14)	107.99 (15)	165.86 (38)
12-100	12.89 (3)	33.83 (3)	50.43 (6)

Note: The 0-3 and 3-6 tranches are quoted as an up-front with a running spread of 100 bp. The remaining tranches are quoted as periodic premium.

Exhibit 12. CDX tranche mid quotes in basis points and bid-ask spreads in parentheses (March 21, 2016)

Att-Det (%)	CDX.NA.IG.25 5Y	Att-Det (%)	CDX.NA.HY.25 5Y
0-3	5628.3 (200)	0-15	6175 (102)
3-7	1504.4 (125)	15-25	961 (88)
7-15	43.8 (75)	25-35	-850 (50)
15-100	-364.1 (30)	35-100	-1916 (30)

Note: CDX.NA.IG.25 tranches are quoted as an up-front with a running spread of 100 bp, whereas the CDX.NA.HY.25 with a spread of 500 bp.

standardized tranches referencing investment grade members of CDX index while tranches on CDX.NA.HY required introducing five clusters (with 16, 21, 23, 80, and 100 names). In each case the recovery rate was taken to be standard 40%.

The calibration errors in upfront payment (for tranches 0-3% and 3-6%) and quoted spread (for tranches 6-12% and 12-100%) are shown in Figure 13, which reveals a very good fit, well within bid-ask spreads.

An additional virtue of the pricing approach presented above is that once the disaggregation of hazard rates into idiosyncratic and systematic components is performed, it becomes possible to gauge the extent to which expected losses are driven by systematic and idiosyncratic factors. To see this, note that by (21) the credit spread for the entire index (tranche 0-100%) effectively represents the risk-neutral expectation of credit loss on the CDS basket:

$$S_{0-100}(0)\sum_{i=1}^{N}\delta_i D(0,T_i)\left(125-\mathbb{E}(N(T_i))\right)=\int_0^{T_N} D(0,t)d\mathbb{E}\left(L_{0,100}(t)\right) \tag{30}$$

The discounted expected loss can be rewritten as

$$ELoss=\int_0^{T_N} D(0,t)\mathbb{E}(N(dt))=\sum_{i=1}^{125}\int_0^{T_N} D(0,t)\lambda_i(t)(1-F_i(t))dt \tag{31}$$

with $F_i(t)=\mathbb{P}(\tau_i<t)$ being the distribution function describing the default times $\tau_1,...,\tau_{125}$

Exhibit 13. Calibration errors expressed in relation to bid-ask spreads

Att-Det (%)	iTraxx.EUR.24		
	3Y	5Y	7Y
0-3	0.25	0.23	-0.07
3-6	0.02	-0.24	-0.17
6-12	-0.11	-0.03	0.00
12-100	-0.25	-0.23	-0.09

Att-Det (%)	CDX.NA.IG.25 5Y
0-3	0.03
3-7	0.50
7-15	-0.01
15-100	0.45

Att-Det (%)	CDX.NA.HY.25 5Y
0-15	-0.00
15-25	-0.00
25-35	-0.46
35-100	2.01

Note: Relative calibration error defined as $\frac{x_{mid}-x_{model}}{x_{ask}-x_{bid}}$, for different tranche spreads (upfronts) x.

of the respective index names. In the model $\lambda_i(t) = \lambda_i^{idio}(t) + \lambda_i^{sys}(t)$ which leads to the following decomposition of expected losses into idiosyncratic and systematic parts:

$$
\begin{aligned}
&ELoss = ELoss^{sys} + ELoss^{idio} = \\
&\quad = \sum_{i=1}^{125} \int_0^{T_N} D(0,t)\lambda_i^{sys}(t)(1-F_i(t))dt + \sum_{i=1}^{125} \int_0^{T_N} D(0,t)\lambda_i^{idio}(t)(1-F_i(t))dt \qquad (32)
\end{aligned}
$$

Hence, assuming that index spreads reflect approximately expected losses, the model-based decomposition of hazard rates can be used to gain insight into the relative weight of idiosyncratic and systematic loss drivers. The results of such decomposition for the iTraxx Europe index are presented in Exhibit 14. Such representation reveals that the longer the maturity of a CSO, the lower the contribution of idiosyncratic factors to expected losses and the greater the contribution of the fatal shock scenario. For the 7Y tenor, almost 60% of portfolio risk is attributable to the largest cluster, reflecting fear of defaults concentrated in time. As argued in section 2, this is also the precise factor that has driven mark-to-market losses on bespoke CSOs throughout the crisis. The only way to express such fears in the language of the Gaussian copula model is by increasing default correlation parameter (cf. Exhibit 7). However, given the unstable and opaque relation of default correlation to probability of joint defaults, the approach suggested in (32) appears much more suitable.

Exhibit 14. Model-based attribution of expected losses to idiosyncratic factors and clusters associated with systematic risk (iTraxx Europe as of March 21, 2016)

Factors	3Y	5Y	7Y
Idiosyncratic	28%	26%	25%
Cluster 1 shock (9 names)	14%	11%	12%
Cluster 2 shock (10 names)	31%	25%	23%
Cluster 3 shock (16 names)	0%	2%	4%
Cluster 4 shock (23 names)	3%	2%	2%
Cluster 5 shock (125 names)	19%	38%	57%

5 Conclusions

The goal of this paper was twofold. First, using a unique data set of more than 1,000 CSO deals, it has tried to shed light on the performance of synthetic CDOs, and in particular the extent and source of losses that have led to a virtual collapse of that important segment of structured credit market in 2008. The empirical evidence presented suggests that significant mark-to-market losses on a vast majority of tranches were not related to widespread credit events, but rather reflected the concentration of relatively limited credit losses in a short period of time, that have not been properly reflected in pre-crisis pricing. More generally, the standard market pricing model – still relying explicitly or implicitly on Gaussian copula – is not suited very well to handling such concentration of defaults in time. Thus, in view of the green shoots of recovery in post-crisis CSO issuance, the second goal of the paper was to present a viable alternative to copula models in the form of an intuitive, analytically tractable and flexible intensity-based model of default correlation. The model relies on a redefinition of the systematic factor as a sequence of increasing random variables characterizing the chain of dependencies running through the financial system. As such, the model is naturally suited to capturing the temporal dimension of default dependencies that have been crucial for CSO investors. The calibration demonstrates that the model can consistently fit the standardized index tranches of iTraxx and CDX and provides a useful decomposition of expected losses across idiosyncratic and systematic risk drivers.

References

Andersen, L., and J. Sidenius (2004): “Extensions to the Gaussian copula: Random recovery and random factor loadings,” *Journal of Credit Risk Volume*, 1(1), 05.

Antoniades, A., and N. A. Tarashev (2014): “Securitisations: tranching concentrates uncertainty,” *BIS Quarterly Review December*.

Brigo, D., and A. Capponi (2010): “Bilateral counterparty risk with application to CDSs,” *Risk*, 23(3), 85.

Brigo, D., and M. Morini (2010): “Dangers of Bilateral Counterparty Risk: the fundamental impact of closeout conventions,” Papers 1011.3355, arXiv.org.

Gatarek, D., and J. Jabłecki (2015): “A Simple Model of Correlated Defaults with Application to Repo Portfolios,” *The Journal of Derivatives*.

Giesecke, K. (2003): “A simple exponential model for dependent defaults,” *Journal of Fixed Income*, 13(3), 74–83.

Hanson, S. G., M. H. Pesaran, and T. Schuermann (2008): “Firm heterogeneity and credit risk diversification,” *Journal of Empirical Finance*, 15(4), 583–612.

Hull, J., and A. White (2005): “The Perfect Copula: The Valuation of Correlation-Dependent Derivatives Using the Hazard Rate Path Approach,” .

——— (2010): “An improved implied copula model and its application to the valuation of bespoke CDO tranches,” *Journal of Investment Management*, 8(3), 11.

Hull, J. C., and A. D. White (2004): “Valuation of a CDO and an n-th to default CDS without Monte Carlo simulation,” *The Journal of Derivatives*, 12(2), 8–23.

——— (2006): “Valuing credit derivatives using an implied copula approach,” *The Journal of Derivatives*, 14(2), 8–28.

——— (2008): “Dynamic models of portfolio credit risk: A simplified approach,” *The Journal of Derivatives*, 15(4), 9–28.

Jarrow, R. A. (2012): *Rethinking the Financial Crisis*chap. The Role of ABS, CDS and CDOs in the Credit Crisis and the Economy, pp. 210–234. Russel Sage Foundation.

Morini, M. (2011): *Understanding and Managing Model Risk: A practical guide for quants, traders and validators*. Wiley.

O’Kane, D. (2008): *Modelling Single-name and Multi-name Credit Derivatives*. John Wiley and Sons.