1 Introduction: Credit Modelling Pre- and In-Crisis

This book aims to show the limits of popular models or pseudo-models (mostly quoting mechanisms with modelling semblance) that, in recent years, have been extensively used to mark to market and risk manage multi-name credit derivatives. We present a compendium of results that we first published in 2006 before the crisis pointing out the dangers in the modelling paradigms used at the time in the market, and showing how the situation has even worsened subsequently by analysing more recent data. We also point out that the current paradigm had been heavily criticized before the crisis, referring to our work and the works of other authors addressing the main limitations of the current market paradigm well before popular accounts such as Salmon (2009) appeared. The problems of the current paradigm include:

- An unrealistic Gaussian Copula assumption and the flattening of 7750 pairwise dependence parameters into one.
- Lack of consistency of the implied correlation market models with more than one tranche quote at the time.
- Occasional impossibility of calibration even of single tranches, or possibility to obtain negative expected tranched losses violating the arbitrage-free constraints.
- Lack of an implied loss distribution consistent with market CDO tranche quotes for a single maturity.
- Lack of a loss distribution dynamics consistent with CDO tranche quotes on several maturities.

- Lack of credit spread volatility, resulting for example in heavy consequences on the valuation of counterparty risk on CDS under wrong way risk.

In this respect we will introduce examples of models published before the crisis that partly remedy the above deficiencies. All the discussion is supported by examples based on market data, pre- and in-crisis. In addressing these issues we adopt the following path through the different methodologies.

1.1 BOTTOM-UP MODELS

A common way to introduce dependence in credit derivatives modelling is by means of copula functions. A typically Gaussian Copula is postulated on the exponential random variables triggering defaults of the pool names according to the first jumps of Poisson processes. In general, if one tries to model dependence by specifying dependence across single default times, one is in a so-called "bottom-up" framework, and the copula approach is typically within this framework. Such a procedure cannot easily be extended to a fully dynamical model in general. We cannot do justice to the huge copula literature in credit derivatives here; we only mention that there have been attempts to go beyond the Gaussian Copula introduced in the CDO world by Li (2000) and leading to the implied (base and compound) correlation framework, some important limits of which have been pointed out in Torresetti *et al.* (2006b). Li and Hong Liang (2005) also proposed a mixture approach in connection with CDO squared. For the results on sensitivities computed with the Gaussian Copula models, see for example Meng and Sengupta (2008).

An alternative to copulas in the bottom-up context is to insert dependence among the default intensities of single names – see, for example, the paper by Chapovsky *et al.* (2007). Joshi and Stacey (2006) resort to modelling business time to create default correlation in otherwise independent single-name defaults, resorting to an "intensity

P1: JYS
c01 JWBK474-Brigo March 19, 2010 18:14 Printer: Yet to come

gamma" framework. Similarly, but in a firm value-inspired context, Baxter (2007) introduces Levy firm value processes in a bottom-up framework for CDO calibration. Lopatin (2008) introduces a bottom-up framework that is also effective in the CDO context, having single-name default intensities that are deterministic functions of time and of the pool default counting process, then focusing on hedge ratios and analysing the framework from a numerical performances point of view, showing this model to be interesting even if lacking explicit modelling of single-name credit spread volatilities.

It is worth noticing that copula models are usually implemented with deterministic credit spreads. Credit spread volatility is assumed to be zero even if both historical (Hull and White, 2003) and implied (Brigo, 2005, 2006) CDS volatilities attain values above 50%. In first-to-default baskets, which are similar in a way to equity-tranches of CDOs (see, for example, Brigo and Mercurio, 2006), spread dispersion combined with lack of credit spread volatility may result in unintuitive features of the model prices with respect to the copula correlation parameter. This has been highlighted in the different context of counterparty risk for credit default swaps by Brigo and Chourdakis (2009), who show that credit spread volatility can be quite relevant, to the point that neglecting it changes the wrong way risk profile of CDS pricing under counterparty risk. This is further highlighted by Brigo and Capponi (2008), while Brigo and Pallavicini (2007, 2008), Brigo and Bakkar (2009) and Brigo *et al.* (2009) model credit spread volatility explicitly, showing that credit spread volatility has a relevant impact also on counterparty risk for other asset classes, including commodities and interest rate products. The stochastic credit spread models used in these works build on the work of Brigo and Alfonsi (2005), Brigo and Cousot (2006) and Brigo and El-Bachir (2008) on CDS options.

Going back to bottom-up models in the context of CDOs, Albanese *et al.* (2006) introduce a bottom-up approach based on structural model ideas that can not only be made consistent with several inputs under historical and pricing measures, but also manages to calibrate CDO tranches.

1.2 COMPOUND CORRELATION

Building on Torresetti *et al.* (2006b), in the context of bottom-up models, we start with the net present value (NPV) of synthetic Collateralized Debt Obligations (CDOs) tranches on pools of corporate credit references in its original layout: the compound correlation framework.

We highlight two of the major weaknesses of the compound correlation:

- Lack of robustness of the compound correlation framework in view of the non-invertibility of mainly the 10-year maturity DJ-iTraxx 6–9% and CDX 7–10% tranches and, more recently, the non-invertibility of mainly the 10-year maturity DJ-iTraxx 12–22% and CDX 10–15% tranches.
- More importantly from a practical standpoint, we highlight the typical non-smooth behaviour of the compound correlation and the resulting difficulties in pricing bespoke CDO tranches.

1.3 BASE CORRELATION

We then introduce the next step the industry took (see, for example, McGinty and Ahluwalia, 2004), namely the introduction of base correlation, as a solution to both problems, given the fact that:

- the resulting map is much smoother, thus facilitating the pricing of bespoke tranche spreads from liquid index tranches;
- until early 2008 the heterogeneous pool one-factor Gaussian Copula base correlation had been consistently invertible from index market tranche spreads.

Nevertheless we will expose some of the known remaining weaknesses of the base correlation framework:

(1) Depending on the interpolation technique being used, tranche spreads could not be arbitrage free. In fact, for senior tranches it may well be that the expected tranche loss plotted versus time is initially decreasing.

(2) The impossibility of inverting correlation for senior AAA and super-senior tranches. This problem arose only recently following the eagerness of market participants to buy protection from systemic risk exposure almost at any cost.
(3) Inconsistency at single tranche valuation level, as two components of the same trade are valued with models having two different parameter values.
(4) Last, but not least, flattening information on 7750 pairwise correlation parameters into a single one for each equity tranche trade.

As an explanation for weakness (1) we point to the fact related to item (3), namely that this arises because the NPV of each tranche is obtained by computing the expected tranche loss and outstanding notional under two different distributions (the distribution corresponding to the attachment base correlation and that corresponding to the detachment base correlation) so that base correlation is already an inconsistent notion at the single tranche level.

As an explanation for weakness (2), we point to the fact that the deterministic recovery assumption, while being computationally very convenient, does not allow us to capture the more recent market conditions. This has been addressed in the implied correlation framework by Amraoui and Hitier (2008) and Krekel (2008). However, even with this update, the base correlation remains flawed and may still lead to negative loss distributions.

Base correlation, with updates and variants, remains to this day the main pricing method for synthetic corporate CDOs, regardless of the body of research criticizing it that we hint at above and below.

1.4 IMPLIED COPULA

We next summarize the concept of Implied Copula (introduced by Hull and White (2006) as the "perfect copula") as a non-parametric model by which to deduce, from a set of market CDO spreads spanning the entire capital structure, the shape of the risk-neutral pool loss distribution. The general use of flexible systemic factors was later generalized and vastly improved by Rosen and Saunders (2009), who also discuss the dynamic implications of the systemic factor framework. Factors and dynamics are also discussed in Inglis *et al.* (2008).

Our calibration results, based on the Implied Copula – already seen in Torresetti *et al.* (2006c) – point out that a consistent loss distribution across tranches for a single maturity features modes in the tail of the loss distribution. These probability masses on the far right tail imply default possibilities for large clusters (possibly sectors) of names of the economy. These results had been published originally in 2006 on ssrn.com. We will report such features here and will find the same features again by following a completely different approach below.

Here we highlight the persistence of the modes (bumps) in the right tail of the implied loss distribution:

- through time, via historical calibrations;
- through regions, comparing the results of the historical calibration to the DJ-iTraxx and the CDX; and
- through the term structure comparing the results of the calibration to different maturities.

The Implied Copula can calibrate consistently across the capital structure, but not across maturities, as it is a model that is inherently static. The next step thus consists in introducing a dynamic loss model. This moves us into the so-called top-down framework (although dynamic approaches are also possible in the bottom-up context, as we have seen in some of the above references). But before analysing the top-down framework in detail, we will make a quick diversion for a model-independent approach to CDO tranches pricing and interpolation.

P1: JYS
c01 JWBK474-Brigo March 19, 2010 18:14 Printer: Yet to come

1.5 EXPECTED TRANCHE LOSS SURFACE

The expected tranche loss (ETL) for different detachment points and maturities can be viewed as the basic brick on which the components of synthetic CDO formulas are built with linear operations (but under some non-linear constraints). We explain in detail how the payoffs of credit indices and tranches are valued in terms of expected tranched losses (ETLs). This methodology, first illustrated pre-crisis in Torresetti *et al.* (2006a), reminds us of Walker's earlier work (2006) and of the formal analysis of the properties of expected tranche loss in connection with no arbitrage in Livesey and Schlögl (2006).

ETLs are natural quantities that can be implied from market data, and no-arbitrage constraints on ETLs as attachment points and maturities change are briefly introduced. As an alternative to the inconsistent notion of implied correlation illustrated earlier, we consider the ETL surface, built directly from market quotes given minimal interpolation assumptions. We check that the type of interpolation does not interfere excessively with the results. Instruments bid/ask spreads enter our analysis, contrary to Walker's (2006) earlier work on the ETL implied surface. By doing so we find less violations of the no-arbitrage conditions.

We also mention some further references that appeared later and dealt with evolutions of this technique: Parcell and Wood (2007), again pre-crisis, consider carefully the impact of different kinds of interpolation, whereas Garcia and Goossens (2007) compare the ETL between the Gaussian Copula and Lévy models.

In general the ETL implied surface can be used to value tranches with non-standard attachments and maturities as an alternative to implied correlation. However, deriving hedge ratios as well as extrapolation may prove difficult. Also, the ETL is not really a model but rather a model-independent stripping algorithm, although the particular choice of interpolation may be viewed as a modelling choice. Eventually the ETL is not helpful for pricing more advanced derivatives such as tranche options or cancellable tranches, because the ETL does not specify an explicit dynamics for the loss of the pool. To that we turn now, by looking at the top-down dynamic loss models.

1.6 TOP (DOWN) FRAMEWORK

One could completely give up single-name default modelling and focus on the pool loss and default-counting processes, thus considering a dynamical model at the aggregate loss level, associated with the loss itself or to some suitably defined loss rates. This is the "top-down" approach – see, for example, Bennani (2005, 2006); Giesecke and Goldberg (2005); Schönbucher (2005); Di Graziano and Rogers (2005); Brigo *et al.* (2006a, 2006b); Errais *et al.* (2006); Lopatin and Misirpashaev (2007); among others. The first joint calibration results of a dynamic loss model across indices, tranches attachments and maturities (available in Brigo *et al.*, 2006a), show that even a relatively simple loss dynamics, like a capped generalized Poisson process, suffices to account for the loss distribution dynamical features embedded in market quotes. This work also confirms the Implied Copula findings of Torresetti *et al.* (2006c), showing that the loss distribution tail features a structured multi-modal behaviour, implying non-negligible default probabilities for large fractions of the pool of credit references, showing the potential for high losses implied by CDO quotes before the beginning of the crisis. Cont and Minca (2008) use a non-parametric algorithm for the calibration of top models, constructing a risk-neutral default intensity process for the portfolio underlying the CDO, looking for the risk-neutral loss process "closest" to a prior loss process using relative entropy techniques. See also Cont and Savescu (2008).

However, in general to justify the "down" in "top-down" one needs to show that from the aggregate loss model one can recover *a posteriori* consistency with single-name default processes when they are not modelled explicitly. Errais *et al.* (2006) advocate the use of random thinning techniques for their approach; see also Halperin and Tomecek (2008), who delve into more practical issues related to random thinning of general loss models, and Bielecki *et al.* (2008) who also build semi-static hedging examples and consider cases where the portfolio loss process may not provide sufficient statistics.

Still, it is not often clear for specific models whether a fully consistent single-name default formulation is possible, given an aggregate model as the starting point.

There is a special "bottom-up" approach that can lead to distinct and rich loss dynamics. This approach is based on the common Poisson shock (CPS) framework, reviewed in Lindskog and McNeil (2003). This approach allows for more than one defaulting name in small time intervals, contrary to some of the above-mentioned "top-down" approaches. In the "bottom-up" language, one sees that this approach leads to a Marshall-Olkin Copula linking the first jump (default) times of single names. In the "top-down" language, this model looks very similar to the GPL model in Brigo *et al.* (2006a) when one does not cap the number of defaults.

The problem of the CPS framework is that it allows for repeated defaults, which is clearly wrong as any name could default more than once.

In the credit derivatives literature, the CPS framework has been used, for example, in Elouerkhaoui (2006) – see also references therein. Balakrishna (2006) introduces a semi-analytical approach allowing again for more than one default in small time intervals and hints at its relationship with the CPS framework, and also shows some interesting calibration results. Balakrishna (2007) then generalizes this earlier paper to include delayed default dependence and contagion.

1.7 GPL AND GPCL MODELS

Brigo *et al.* (2007) address the repeated default issue in CPS by controlling the clusters default dynamics to avoid repetition. They calibrate the obtained model satisfactorily to CDO quotes across attachments and maturities, but the combinatorial complexity for a non-homogeneous version of the model is forbidding, so that the resulting

GPCL approach is hard to use successfully in practice when taking single names into account.

Still, in the context of the present book, the Generalized-Poisson Loss (GPL) and the Generalized-Poisson Cluster Loss (GPCL) models will be useful to show how a loss distribution dynamics consistent with CDO market quotes should evolve.

In this book we summarize the GPL model, leaving aside the GPCL model. As explained above, the GPL is a dynamical model for the loss, and is able to reprice all tranches and all maturities at the same time. We employ here a variant that models the loss directly rather than the default-counting process plus recovery. The loss is modelled as the sum of independent Poisson processes, each associated with the default of a different number of entities, and capped at the pool size to avoid infinite defaults. The intuition of these driving Poisson processes is that of defaults of sectors, although the amplitudes of the sectors vary in our formulation of the model pre- and in-crisis. In the new model implementation in-crisis for this book, we fix the amplitude of the loss triggered by each cluster of defaults *a priori*, without calibrating it, as we did in our earlier GPL work. This makes the calibration more transparent and the calibrated intensities of the default of sectors easier to interpret. We point out, however, that the precise default of sectors is made rigorous only in the GPCL.

We highlight how the GPL model is able to reproduce the tail multimodal feature that the Implied Copula proved was indispensable for accurately repricing the market spreads of CDO tranches on a single maturity.

We also refer to the later related results of Longstaff and Rajan (2008), that point in the same direction but add a principal component analysis on a panel of CDS spread changes, with some more comments on the economic interpretation of the default clusters being sectors.

Incidentally, we draw the reader's attention to the default history, pointing to default clusters being concentrated in a relatively short time period (a few months) like the thrifts in the early 1990s at the height of the loan and deposit crisis, the airliners after 2001 and the autos and financials more recently. In particular, from 7 September 2008 to 8 October 2008 – a time window of one month – we witnessed seven credit events occurring to major financial entities: Fannie Mae, Freddie Mac, Lehman Brothers, Washington Mutual, Landsbanki, Glitnir and Kaupthing. The Fannie Mae and Freddie Mac conservatorships were announced on the same date (7 September 2008) and the appointment of a "receivership committee" for the three Icelandic banks (Landsbanki, Glitnir and Kaupthing) was announced between 7 and 8 October.

Moreover, Standard and Poors issued a request for comments related to changes in the rating criteria of corporate CDOs.[1] Thus far agencies have been adopting a multifactor Gaussian Copula approach to simulate the portfolio loss in the objective measure. S&P proposed changing the criteria so that tranches rated AAA should be able to withstand the default of the largest single industry in the asset pool with zero recoveries. We believe that this goes in the direction of modelling the loss in the risk-neutral measure via GPL-like processes, given that the proposed changes to S&P's rating criteria imply admitting – as a stressed but plausible scenario – the possibility that a cluster of defaults in the objective measure exists. See also Torresetti and Pallavicini (2007) for the specific case of Constant Proportion Debt Obligations (CPDOs).

We finally comment more generally on the dynamical aggregate models and on their difficulties to lead to single-name hedge ratios when trying to avoid complex combinatorics. The framework thus remains incomplete to this day, because obtaining jointly tractable dynamics and consistent single-name hedges that can be realistically applied in a trading floor remains a problem. We provided above some references for the latest research in this field. We emphasize, however, that even a simple dynamical model like our GPL or the single-maturity Implied Copula is enough to appreciate that the market quotes were implying the presence of large default clusters with

[1] See "Request for Comment: Update to Global Methodologies and Assumptions for Corporate Cash Flow CDO and Synthetic CDO Ratings", 18 March 2009, Standard & Poors.

P1: JYS
c01 JWBK474-Brigo March 19, 2010 18:14 Printer: Yet to come

non-negligible probabilities well in advance of the credit crisis, as we documented in 2006 and early 2007.

Finally, it is important to point out that most of the above discussion and references (with very few exceptions) centre on corporate CDOs, and mostly on synthetic CDOs, and little literature is available for valuation of CDOs on other asset classes, with possibly complex waterfalls and prepayment risk that are cash rather than synthetic, including Collateralized Loan Obligations (CLOs), Residential Mortgage-Backed Securities (RMBSs) and CDOs of RMBSs, which are more related to the asset class that triggered the crisis. For many such deals the main problem is often the data. The few works on this area include Jaeckel (2008) and Papadopoulos and Tan (2007).

1.8 STRUCTURE OF THE BOOK

The book is structured as follows: Chapter 2 introduces the market quotes we are looking at and provides general discounted payoffs, arbitrage-free prices and spread formulas for CDS indices and CDO tranches.

Chapter 3 introduces the Gaussian Copula model, in its different formulations concerning homogeneity and finiteness, and then illustrates the notions of implied correlation from CDO tranche quotes. The two paradigms of base correlation and compound correlation are explained in detail. Existence and uniqueness of implied correlation are discussed on a number of market examples, highlighting the pros and cons of compound and base correlations, and the limitations inherent in these concepts. In particular, Section 3.9 summarizes issues with implied correlations, pointing out the danger for arbitrage when negative expected tranche losses appear, and the lack of consistency across capital structure and maturity. The first inconsistency is then addressed in Chapters 4 and 5, with the Implied Copula, illustrated with a number of studies throughout a long period, whereas both inconsistencies are addressed in Chapter 6, where our fully-fledged GPL dynamic loss model is illustrated pre-crisis. *En passant*, in

P1: JYS
c01 JWBK474-Brigo March 19, 2010 18:14 Printer: Yet to come

Chapter 5 we introduce and analyse the notion of Expected Tranche Loss (ETL), a quantity that can be implied in a model-independent way from index and CDO tranche quotes and can be used for the interpolation of CDO quotes across attachments and maturities.

All these paradigms are then analysed in-crisis in Chapter 7, while the final discussion, including the reasons why implied correlation is still used despite all its important shortcomings, are given in Chapter 8. In particular, the need for hedge ratios with respect to single names, random recovery modelling and speed of calibration remain issues that are hard to address jointly outside the base correlation framework.