

Arbitrage-Free Loss Surface Closest to Base Correlations

Andrei Greenberg*
Rabobank International

Original version: 28 April 2008

This version: 17 July 2008

Abstract

The drawbacks of base correlations are well-known to quantitative credit practitioners. The loss surface produced by any of its common implementations is arbitrageable either in the loss dimension, or the time dimension, or both. Yet the approach has been quite popular in the industry, especially with correlation traders, not least for its ability to fit the standard tranche market by definition, unlike any of the widely known “bottom-up” models. Consequently a large effort has been put into developing the base correlation framework into a workable pricing and risk management system, even though its fundamental problems were never resolved.

In the present work we start from a typical base correlation loss surface and seek to rectify it by relaxing as few conditions as possible. By focusing on the areas where arbitrage occurs most often, we are able to obtain an arbitrage-free loss surface with minimum modifications and still price standard tranches within the bid-offer and most of the times, very close to the mid. The new framework generates a series of loss distributions, and thus naturally offers a solution to such issues as pricing thin tranches and parts of the capital structure outside the quoted detachment points, previously dealt with by interpolation and extrapolation of base correlations.

1 Introduction

Base correlations were introduced in [9] as an extension to the Gaussian copula model with flat correlation (see e.g., [2, 7]). Default correlation was described

*Financial Engineering, Rabobank International, e-mail: andrei.greenberg@rabobank.com. The views and opinions expressed in this paper are those of the author alone acting in his personal capacity and do not necessarily reflect the views, opinions or policies of Rabobank International or of any of its affiliates. Neither Rabobank International nor any of its affiliates accepts any legal liability or responsibility for the accuracy, completeness or functionality of any information disclosed or described in this paper. The author is grateful to Hans-Jürgen Brasch and Saurav Sen for their comments.

by a single parameter ρ , a linear correlation between latent variables. With the advent of standardised tranche markets in 2003, it has become obvious that a single value of ρ was not sufficient to fit all tranche quotes simultaneously. Backing out correlation values from market prices of different tranches of the same portfolio (index) produced a “correlation smile”: junior mezzanine prices, where most of the demand was focused, implied a lower correlation than either equity or senior prices. This shape proved to be stable over time and was qualitatively the same across maturities and even different indices.

These pictures, as well as the characterisation of tranches as options on the portfolio loss, prompted an analogy between implied [Black-Scholes] volatility for equity options and implied [Gaussian copula] correlation for tranches. The analogy is not entirely correct for a variety of reasons. Without discussing the more fundamental ones we simply observe that, from the practical standpoint, while in the equity option world we set $\sigma = \sigma(K)$, in the correlation world $\rho = \rho(Tr_k) = \rho(A_k, D_k)$. Thus inter- and extrapolation can be used in the former to determine volatilities of non-quoted strikes, but a similar procedure is not possible for the latter.

This is precisely the drawback that the base correlation idea was designed to address. Additivity of [expected] tranche loss means that for a $[A, D]$ mezzanine tranche,

$$EL(A, D) = EL(0, D) - EL(0, A). \quad (1)$$

Therefore any mezzanine tranche can be decomposed into a difference of equity (or “base”) tranches, and prices of successive equity, mezzanine and senior tranches can be remapped to a series of equity tranche prices by bootstrapping. This means that unlike the original notion of implied (often called *compound*) correlation, we can define *base* correlation as a function of detachment point only. To this effect, given the prices for the equity $[0, A]$ and mezzanine $[A, D]$ tranches, $\rho(A)$ is such that it reprices $[0, A]$ and $\rho(D)$ is such that $EL[0, D; \rho(D)]$, combined with $EL[0, A; \rho(A)]$ via (1), will reprice $[A, D]$. A more detailed description can be found in the original publication [9] or in [10]. Now we are much closer with the implied volatility analogy, since for an arbitrary $[A', D']$ tranche we can write

$$EL(A', D') = EL(0, D'; \rho(D')) - EL(0, A'; \rho(A')), \quad (2)$$

which can be calculated if we find $\rho(A')$ and $\rho(D')$ from the quoted attachment-detachment points via a suitable interpolation scheme.

As practitioners sought to explore the implied volatility analogy, a number of important flaws in the base correlation framework were uncovered. Most importantly, tranche expected losses calculated using base correlations (with suitable interpolation, as required) are not consistent. It is possible within the framework to have expected loss of a more senior tranche to be larger than that of a more junior tranche (“capital structure arbitrage”); and for the same tranche, the expected loss at a later date can be smaller than at an earlier date (“maturity arbitrage”). While not arbitrages in the strict market sense (there

is no trading in expected loss), these inconsistencies mean that base correlation-implied loss surfaces (maturity by detachment point) *may* in general produce arbitrageable prices, whereby a riskier tranche will be assigned a smaller fair spread.

For a theoretically inconsistent construction, base correlations enjoyed remarkable popularity in the industry: for a long time it was considered market standard, and every player in the credit correlation trading business had an implementation of the framework. In our opinion, there are two main reasons behind this. First, base correlations is a modification of the one-factor Gaussian copula model, which was the approach of choice before 2003. Hence base correlations were “quickest to market” in terms of both model development and systems implementation. Second, base correlations fit liquid tranche markets exactly, for all maturities, indices and observation dates¹. This highly desirable feature could not be matched by the models trying to capture the “correlation skew”, from various extensions of the Gaussian copula to other copula function-based approaches and even dynamic models. Thus, while a lot of quantitative research effort was put into developing new skew-consistent models, none of them were able to completely depose base correlations.

This serves as the motivation for the current paper to take the base correlation loss surface as a starting point and explore the possibility of removing the internal inconsistencies (arbitrages) without spoiling the market fit too much. Our approach produces a sequence of full portfolio loss distributions at several horizons, from valuation date to maturity, each of which is consistent with base correlation-implied expected losses for liquid tranches to the largest possible degree. Since valid loss distributions always produce arbitrage-free prices, subject to certain time monotonicity conditions, the problem reduces to that of finding a set of loss probabilities subject to linear constraints given by base correlation-implied tranche expected losses. Similar approaches were taken in [6, 13, 17, 19], where formulations are very similar to the one we give in Section 3 below. While [13] only discusses feasibility, the other papers only use base correlations for liquid maturities and interpolate in time, which can introduce spurious maturity arbitrage and/or lead the resulting prices uncontrollably away from the market.

The rest of the paper is organised as follows. In Section 2 we give an account of the expected loss arbitrages occurring in the “classical” base correlations methodology. Section 3 gives a fixed-horizon formulation of the problem in terms of a loss distribution (density), including a procedure for removing arbitrage in the loss dimension and a recipe for generating a feasible initial guess. After discussing an important limitation in handling maturity arbitrage, we introduce an equivalent formulation in terms of cumulative loss probabilities in Section 4, which allows us to solve for the whole loss surface, arbitrage-free in both loss and maturity. Section 5 gives numerical examples, demonstrating the quality of fit of the modified loss surface to the original market prices of standard tranches,

¹Problems with CDX.IG 5Y and 7Y 15-30% tranches have emerged early in 2008, but as we argue later, they are not caused by a choice of correlation model per se.

as well as model-implied non-standard tranche prices and base correlations. Concluding remarks appear in Section 6.

We conclude the introduction with a brief discussion of CDO tranche pricing and loss surfaces. The “CDS analogy” for synthetic tranche pricing is presented, e.g., in [8, p. 40]. The idea is to build a term structure of tranche survival probabilities:

$$SP(A, D)(t) = 1 - \frac{\mathbb{E}[(L(t) - A)^+] - \mathbb{E}[(L(t) - D)^+]}{D - A}, \quad (3)$$

where $L(t)$ is the portfolio loss at time t . Then the price (e.g., PV or fair spread) of the tranche is the same as the corresponding price of a credit default swap of a fictitious name with the same survival probability term structure and zero recovery (as in, e.g., [11]).

We loosely define a “loss surface” as the profile of the equity tranche expected loss function $EL(0, K)(t)$ for various strikes (detachment points) $K \in [0, 100\%]$ and time horizons $t \in (0, T_{maturity}]$. Given a loss surface, we can use (1), (3) and the CDS analogy to price any tranche of the given portfolio.

2 Inconsistencies in Base Correlations

While looking like an extension to the Gaussian copula model, base correlations in fact introduce inconsistencies and generic flaws.

The parameter ρ represents correlation between defaults of individual names, and is therefore a characteristic of the whole portfolio. Using different correlations for different parts of the capital structure lacks theoretical foundation: since the order of defaults is not known, it is not clear which default correlation to use for any given name in, say, a times-to-default simulation. This breaks the mathematical basis of the Gaussian copula framework and must be one of the main reasons why base correlations were not well-received by the credit quant community.

Mathematical considerations aside, there are problems with base correlations from the financial point of view. A good systematic critique can be found in [20]. Below we focus on the most important problem which, in our opinion, is implied arbitrage opportunities. Absence of arbitrage for CDO tranches can be formulated as an imperative for the exposure to any band of portfolio losses to have a non-negative cost: the fair premium for buying protection on any tranche can never be negative, i.e., one cannot *be paid* to go *long* protection. Any of the following easily verifiable properties reflects this concept:

- fair spread for fixed-width mezzanine tranches is decreasing in subordination;
- fair spread for equity tranches is decreasing in thickness (or detachment point);
- expected loss of any mezzanine tranche is positive;

- etc.

(Strictly speaking, we cannot theoretically rule out the possibility of zero probability for some loss levels, so we should really be saying “non-decreasing” instead of “increasing”, etc.) Arbitrage opportunities in the markets indicate that one or both of the following is true for the loss surface:

- $EL(0, K_i)(t) < EL(0, K_j)(t)$ for $K_i \geq K_j$ (capital structure arbitrage)
- $EL(0, K)(t_i) < EL(0, K)(t_j)$ for $t_i \geq t_j$ (maturity arbitrage)

While these two conditions are not actually equivalent to arbitrage opportunities existing in the tranche market (expected losses are not tradable, they are only inputs to pricing algorithms), we adopt the convention to call them “arbitrages” for notational simplicity. (We recognise that “indications of arbitrage” or even “model inconsistencies” would be more accurate terminology.)

2.1 Capital structure arbitrage

Looking at the fundamental formula for mezzanine tranche expected loss with base correlations (2), we observe that, since equity tranche loss goes down with correlation, a sufficiently steep upward-sloping base correlation curve (i.e., very loosely speaking, $\rho(D) \gg \rho(A)$) may result in $EL[0, A; \rho(A)] > EL[0, D; \rho(D)]$, making $EL[A, D] < 0$.

While one could reasonably expect never to see this happen for liquidly traded mezzanine tranches (although see Table 2 below), controlling this behaviour for non-standard, especially thin tranches, is very hard. Interpolating in correlations is common practice; however, the resulting loss distribution (implied by the Breedan-Litzenberger methodology, along the lines of [14]) for any of the popular methods is not arbitrage-free. This is evident from Table 1 of mezzanine tranche spreads, where we have tried linear, cubic spline and monotonicity-preserving interpolation schemes to build the corresponding base correlation curve from the values for standard strikes. More sophisticated shape-preserving interpolation schemes used directly on tranche expected loss, rather than base correlations, give more promising results, but they tend to be more implementation-intensive (speed and/or storage) and in any case, fail to resolve the time dimension issues that we outline below.

2.2 Maturity arbitrage

So far we have only looked at the performance of base correlations across the capital structure. This makes sense since, just like a typical copula model, base correlations are built on fixed horizons: given the probabilities of default for individual names before a certain date, we can calculate expected losses of tranches for this date. However, as we pointed out at the end of Section 1, tranche pricing requires building a whole term structure of losses – an object which is commonly called the “tranche curve”, which is, roughly, a cross-section of the loss surface for fixed K .

Tranche		Fair spread, bp			
Attach	Detach	Linear	Spline	Monotonic	ExpLoss
13%	14%	76.72	83.05	87.05	93.75
14%	15%	60.98	75.81	91.62	88.86
15%	16%	118.19	70.07	88.12	75.01
16%	17%	104.46	64.85	74.07	62.17
17%	18%	93.50	61.80	63.54	57.05
18%	19%	85.20	61.49	56.91	55.17
19%	20%	77.85	61.97	52.43	54.41
20%	21%	69.74	60.69	47.71	54.06
21%	22%	60.38	56.71	41.56	53.88
22%	23%	50.67	51.46	35.11	53.79
23%	12%	42.00	47.55	33.89	53.74
24%	13%	36.04	48.24	37.18	53.70
25%	26%	32.41	50.49	42.50	53.67

Table 1: CDX.IG9 5Y thin tranche spreads, different interpolations (14-Jan-08).

The standard way of building these tranche curves is to use term structures of individual default probabilities to calculate tranche expected losses at a strip of dates. Unfortunately, there are even fewer liquid maturities than liquid strikes: typically, only 5Y, 7Y and 10Y index tranche markets can be found. Moreover, as in the Gaussian copula model, when pricing a 5Y tranche, say, we should use the same correlation parameters for all intermediate dates where the tranche expected losses are calculated, from the short end all the way to 5Y. This often introduces problems, best illustrated by the example below.

We calculated expected losses for equity tranches corresponding to standard detachment points of the CDX.IG index, using 5Y parameters, at several dates (Table 2). Although nothing stands out at maturity, for shorter horizons the 5Y base correlations produce arbitrageable expected losses – and the shorter the horizon, the more pronounced the effect. While it is hardly possible to realise these arbitrage opportunities (over five years these differences “even out” and there are no liquid markets in shorter maturities), this behaviour demonstrates yet another dimension of the fundamental problems about the base correlation framework.

Attempts to smoothen out this behaviour (e.g., cap tranche expected losses at 1 and floor them at 0) do not help, either. Table 3 shows that the expected loss of a 31-33% tranche has a dip around 20-Mar-2012. The framework seems to imply that holding this tranches for 5 years and 3 months is less risky than holding it for just 5 years, which is clearly nonsensical. Note that this behaviour could hardly have been predicted from the corresponding base correlation values.

It is important to note here that even though liquid tranche quotes are for a given maturity, we do not actually have information about tranche *expected losses* at this maturity. All we have from the market is a *price* (spread or upfront), though a term structure of tranche expected losses can be mapped to

Date	Tranche expected loss					
	0-3%	3-7%	7-10%	10-15%	15-30%	30-60%
20-Dec-07	2.0345%	-0.0017%	-0.0014%	-0.0021%	-0.0019%	0.0018%
20-Mar-08	5.5264%	-0.0051%	-0.0033%	-0.0062%	-0.0043%	0.0061%
20-Jun-08	8.9588%	0.0009%	-0.0015%	-0.0089%	-0.0058%	0.0113%
20-Sep-08	12.4436%	0.0232%	0.0051%	-0.0116%	-0.0077%	0.0203%
20-Dec-08	15.8583%	0.0721%	0.0198%	-0.0117%	-0.0083%	0.0317%
20-Mar-09	19.2266%	0.1562%	0.0458%	-0.0087%	-0.0062%	0.0472%
20-Jun-09	22.6186%	0.2879%	0.0871%	0.0000%	-0.0027%	0.0677%
20-Sep-09	25.9457%	0.4725%	0.1458%	0.0158%	0.0066%	0.0911%
20-Dec-09	29.1615%	0.7123%	0.2235%	0.0396%	0.0173%	0.1198%
20-Mar-10	32.6173%	1.0455%	0.3332%	0.0766%	0.0380%	0.1549%
20-Jun-10	36.0908%	1.4707%	0.4757%	0.1281%	0.0626%	0.1982%
20-Sep-10	39.4889%	1.9852%	0.6514%	0.1955%	0.0967%	0.2467%
20-Dec-10	42.7630%	2.5842%	0.8600%	0.2791%	0.1409%	0.2998%
20-Mar-11	45.9361%	3.2726%	1.1044%	0.3807%	0.1881%	0.3617%
20-Jun-11	49.0745%	4.0697%	1.3935%	0.5048%	0.2510%	0.4292%
20-Sep-11	52.0999%	4.9598%	1.7236%	0.6503%	0.3270%	0.5023%
20-Dec-11	54.9764%	5.9292%	2.0912%	0.8167%	0.4060%	0.5852%
20-Mar-12	57.9053%	7.0463%	2.5214%	1.0146%	0.5015%	0.6794%
20-Jun-12	60.7377%	8.2722%	3.0045%	1.2414%	0.6190%	0.7790%
20-Sep-12	63.4381%	9.5903%	3.5367%	1.4964%	0.7443%	0.8898%
20-Dec-12	65.9783%	10.9798%	4.1114%	1.7770%	0.8758%	1.0111%

Table 2: Expected losses of standard tranches (CDX 5Y, 29-Oct-07).

Strike	Base corr	20-Dec-2011		20-Mar-2012		20-Jun-2012	
		Equity	Mezz	Equity	Mezz	Equity	Mezz
27%	84.59%	0.09465	9.465%	0.10140	10.140%	0.10836	10.136%
29%	87.42%	0.08836	0.344%	0.09466	0.371%	0.10123	0.498%
31%	90.08%	0.08286	0.312%	0.08876	0.315%	0.09495	0.387%
33%	92.85%	0.07786	0.034%	0.08338	0.003%	0.08921	0.021%
35%	95.49%	0.07336	-0.090%	0.07853	-0.160%	0.08402	-0.150%

Table 3: CDX.IG9 tranche expected losses using 10Y base correlations.

it. Therefore when we build a base correlation curve for a given maturity, we fix tranche expected losses not only for the maturity, but for all intermediate horizons as well. Thus we inherit the behaviour of this framework through time as well as at maturity. This will be important when we build our arbitrage-free loss surfaces in later sections.

2.3 Remark on scope

In this paper we do not deal with a few other sources of arbitrage, which have become more apparent in the wake of the senior and super-senior unwinds. First,

it is occasionally impossible to solve for the 5Y and 7Y CDX.IG base correlation at 30%, the highest quoted detachment point. Second, the super-senior tranche price implied by the index and all junior prices can come out as quite different from what is quoted in the market for this tranche. The reason why we do not devote specific attention to these problems is that we see them as coming from a different source altogether – namely, the price of systemic risk. As it was demonstrated in [13], a fixed 40% recovery assumption implies a tail of the distribution that is too “short”: it caps portfolio losses at 60%, while the market does not seem to agree with this. This can be dealt with by introducing either a systematic shock, in the flavour of [16], or stochastic recovery. In other words, this can be solved by adding an extra parameter set, which does not necessarily relate to correlation. Thus these problems do not appear to be specific to base correlations, and are therefore outside of the scope for this paper. We believe that a relatively simple stochastic recovery framework, such as introduction of several global recovery regimes, will serve as sufficient extension to handle these issues and still be able to apply our ideas below to the general case. Two recent papers [1, 5] describe the application of stochastic recovery in some detail.

3 Loss Distribution Formulation

Let $0 = K_0 < K_1 < K_2 < \dots < K_m < 100\%$ be an arbitrary sequence of detachment points (strikes). The following conditions for the corresponding equity tranche expected losses $EL_j = EL(0, K_j)$, must be satisfied for the resulting loss surface to be arbitrage-free:

- *positivity/boundedness* : $0 \leq EL_j \leq EL(0, 100\%) \leq 1$; $EL_j \leq U_j$ (U_j is the strike K_j rescaled in terms of loss units – see definition below);
- *monotonicity* : $EL_j \geq EL_{j-1}$;
- *concavity* : $\frac{EL_j - EL_{j-1}}{U_j - U_{j-1}} \geq \frac{EL_{j+1} - EL_j}{U_{j+1} - U_j}$;

Often a fourth condition of portfolio expected loss conservation is added; we discuss it in more detail below.

In practice, checking these conditions for all tranches (or even for a suitable “basis”) is too cumbersome a task. As we demonstrated in the previous section, base correlations do not, in general, produce a loss surface consistent with the above, even though it may be consistent for just the “standard” detachment points, or liquid tranches. Fortunately, for a sufficiently granular set of strikes K_j , these conditions are equivalent to the existence of a valid loss distribution, i.e., one with non-negative, unit-mass density. Therefore a suitable equivalent mathematical formulation in terms of the loss probabilities, guaranteeing the solution of the corresponding problem to satisfy the no-arbitrage conditions, is useful. While we note that these necessary conditions are only dealing with the loss dimension and ignore the time behaviour, understanding how a valid fixed-horizon loss distribution can be produced is nevertheless an important first step.

3.1 Fixed horizon

We assume that there are n obligors in the portfolio, so the portfolio loss at any time horizon t can be written as

$$L(t) = \sum_{k=1}^n N_k(1 - R_k) \mathbb{1}_{\{\tau_k \leq t\}} \quad (4)$$

Here N_k , R_k and τ_k denote, respectively, the notional, recovery rate and default time of name k , and $\mathbb{1}$ is the indicator function (in this case, of the default event for name k).

The common assumption of fixed recovery rates leads to a discrete loss distribution. (This will also be true for a class of stochastic recovery models with discrete support.) If all notionals and recoveries are the same (which is usually the case for standard indices), then each default results in an equal loss of $l = (1 - \bar{R})\bar{N}$ to the portfolio, so that the support of the loss distribution is $\{0, l, 2l, \dots, nl\}$, a finite set. For more complicated portfolio structures it is still possible, in most practical cases, to find a loss unit u , such that $(1 - R_k)N_k = u \cdot l_k$ for all k . In this case, the support of the loss distribution is comprised from various combinations of $\{0, l_1 u, l_2 u, \dots, l_n u\}$. Usually u is the greatest common divisor of all individual losses, possibly with some truncation. For a homogeneous portfolio the distribution size is $n+1$. Below we focus on the homogeneous case and discuss modifications for more general portfolios later in Section 4.4.

Consider the following formulation. Given a set of tranches and their expected losses (to a given time horizon), find a loss distribution at this horizon that is consistent with these expected losses and conserves the total expected loss on the whole portfolio of single names. The loss distribution is simply the set of all probabilities, i.e., the vector

$$\{P_0, P_1, \dots, P_n\} \quad : \quad P_k(t) = \text{Prob}\{L(t) = ku\}. \quad (5)$$

Denote as before the strikes of the given tranches by $0 = K_0 < K_1 < K_2 < \dots < K_m < 100\%$, so that there are m tranches. We note that there are two equivalent approaches: either take m equity tranches $[0, K_j]$, $j = 1, \dots, m$, or look at a contiguous sequence of equity and mezzanine tranches $[K_{j-1}, K_j]$, $j = 1, \dots, m$: in a loss-preserving model, both approaches produce the same loss surface. We choose the equity tranche sequence because it comes more naturally in the base correlation framework, where each equity tranche is priced with its own flat correlation. The senior-most detachment point is usually set to the maximum possible loss $K_{max} = u/\bar{N} = 1 - \bar{R}$.

Following the definition (4) of the portfolio loss, we define the expected loss of the j -th equity tranche as

$$EL_j(t) = \mathbf{E}[\min(L(t), K_j)],$$

and the expected loss on the whole portfolio as

$$\mu(t) = \sum_{k=1}^n N_k(1 - R_k) p_k(t).$$

(We suppress the time dependence for most of the paper, when we are dealing with fixed horizons only.)

We now reduce the problem to an exercise in finite optimisation. Note that there are no obvious restrictions on the shape of the distribution, other than fitting the expected losses. However it may be advisable to control its shape, for example, by forcing the distribution to be close to a given one or demanding it to be smooth (see, e.g., ideas in [6]). A simple measure of distribution “smoothness” is the sum of square differences between adjacent probabilities. In this case we can formulate the optimisation problem as follows: find a vector of $n + 1$ loss probabilities, such that the distribution is as smooth as possible; preserves the total expected loss on single names; and produces the given tranche expected losses. The formulation described in [6] is closest to ours, even though the entropy-based objective function the authors adopt precludes them from using quadratic programming. The time dimension and the issue of feasibility is also handled differently.

Before we write down the formulæ, we introduce rescaling of the tranche strikes in terms of loss units (recall the quantity u discussed earlier in this section). Denote

$$U_j = K_j \frac{n\bar{N}}{u}, \quad k_j = [U_j] = \max\{i \in \mathbb{Z} : i \leq U_j\}.$$

For any discrete loss distribution, the expected loss of an equity tranche detaching at $K\%$ or U loss units is

$$EL(0, U) = \sum_{j=1}^{[U]} j P_j + U \sum_{j=[U]+1}^n P_j. \quad (6)$$

We can thus write

$$f(P) = \frac{1}{2} \sum_{j=2}^n (P_j - P_{j-1})^2 \rightarrow \min \quad (7)$$

$$\text{s.t. :} \quad 0 \leq P_1, \dots, P_n \leq 1,$$

$$\sum_{j=1}^n j P_j = \tilde{\mu} \quad (8)$$

$$\sum_{j=1}^{k_1} j P_j + U_1 \sum_{j=k_1+1}^n P_j = EL_1,$$

$$\sum_{j=1}^{k_2} j P_j + U_2 \sum_{j=k_2+1}^n P_j = EL_2,$$

...

$$\sum_{j=1}^{k_m} j P_j + U_m \sum_{j=k_m+1}^n P_j = EL_m.$$

The $\tilde{\mu}$ in (8) is the portfolio expected loss rescaled in terms of loss units. Note that the probability of no loss, P_0 , does not appear – we recover it later from the condition that all probabilities must sum to unity.

The problem above is a bounded quadratic programme, which is guaranteed to have a unique global solution if only the Hessian in (7) is positive definite and the constraints are feasible. The Hessian has the following structure:

$$H(f) = \frac{\partial f}{\partial P_i \partial P_j} = \begin{pmatrix} 1 & -1 & 0 & \cdots & 0 & 0 & 0 \\ -1 & 2 & -1 & \cdots & 0 & 0 & 0 \\ & \vdots & & \ddots & & \vdots & \\ 0 & 0 & 0 & \cdots & -1 & 2 & -1 \\ 0 & 0 & 0 & \cdots & 0 & -1 & 1 \end{pmatrix},$$

and its positive definiteness is straightforward to check (all principal minors are equal to 1). We therefore turn to the issue of feasibility.

3.2 Constructing a feasible initial guess

For a typical portfolio and tranche structure, we will have $m \ll n$, i.e., much fewer constraints (tranches) than variables (probabilities). However the importance of getting a good initial guess should not be underestimated. An arbitrary distribution, such as one coming from a different model, can be so far from satisfying the expected loss constraints that the optimiser will fail to find a valid search direction. For larger, more computationally intensive problems, such as very heterogeneous portfolios, a good initial guess may reduce the running time considerably. Therefore we devote some time to finding a feasible initial guess².

The simplest idea is to use the leanest possible distribution, i.e., the one that has only the absolute minimum number of non-zero probabilities. Since we have $m + 1$ constraints, we need at least $m + 1$ probabilities to fit them. We use the maximum loss probability, P_n , to tie up the total portfolio expected loss. Away from that point, we assume that $P_j = 0$ for $j \neq k_i$, $i = 1, \dots, m$. In other words, we try to fit the tranche expected losses by putting probability mass at the node just below the tranche strike. This simplifies the constraint equations in (7):

$$\begin{aligned} k_1 P_{k_1} + U_1 \sum_{i=2}^m P_{k_i} + U_1 P_n &= EL_1, \\ k_1 P_{k_1} + k_2 P_{k_2} + U_2 \sum_{i=3}^m P_{k_i} + U_2 P_n &= EL_2, \\ &\dots \\ k_1 P_{k_1} + \dots + k_m P_{k_m} + U_m P_n &= EL_m, \\ \sum_{i=1}^m k_i P_{k_i} + n P_n &= \tilde{\mu}. \end{aligned} \tag{9}$$

²We thank Hans-Jürgen Brasch for useful suggestions and discussions on the subject of this section.

Solving this system does not present any difficulty. By taking the difference of each of the equations with the previous one (except for the very first equation, which is left intact) we obtain a system in upper triangular form:

$$\begin{aligned}
k_1 P_{k_1} + U_1 \left(\sum_{i=2}^m P_{k_i} + P_n \right) &= EL_1, \\
(k_2 - U_1) P_{k_2} + (U_2 - U_1) \left(\sum_{i=3}^m P_{k_i} + P_n \right) &= EL_2 - EL_1, \\
&\dots \\
(k_m - U_{m-1}) P_{k_m} + (U_m - U_{m-1}) P_n &= EL_m - EL_{m-1}, \\
(n - U_m) P_n &= \tilde{\mu} - EL_m.
\end{aligned}$$

The solution can be represented by the following recursive formulæ:

$$P_n = q_{m+1}; \quad P_{k_i} = \frac{q_i - \sum_{j=i+1}^m P_{k_j} - P_n}{d_i}, \quad 1 \leq i \leq m, \quad (10)$$

where we have defined

$$d_j = \frac{k_j - U_{j-1}}{U_j - U_{j-1}}; \quad q_j = \frac{EL_j - EL_{j-1}}{U_j - U_{j-1}} = EL(K_{j-1}, K_j), \quad 1 \leq j \leq m,$$

and

$$q_{m+1} = \frac{\tilde{\mu} - U_m EL_m}{n - U_m} = EL(K_m, K_{max}).$$

Note the interpretation of the right-hand sides as the expected losses on the corresponding mezzanine tranches.

We note the special case $U_1 = k_1$, i.e., when the first liquid detachment point falls on a loss distribution node. The first equation in (9) then reduces to

$$U_1 \sum_{i=1}^m P_{k_i} + U_1 P_n = U_1 (1 - P_0) = EL_1,$$

so that P_{k_1} drops out and P_0 comes in. We will then need to recover P_{k_1} from the condition $\sum P_j = 1$, as we do with P_0 in the main case below.

The set of probabilities (10), together with

$$P_0 = 1 - \sum_{i=1}^m P_{k_i} - P_n,$$

will define a feasible initial distribution if we can show that $0 \leq P_{k_i} \leq 1$ for all $i = 1, \dots, m$ and the same is true for P_n and P_0 . This follows by construction: we simply take the probability mass between any two adjacent k_i 's and concentrate it at the right endpoint.

Indeed, the first two no-arbitrage conditions (boundedness and monotonicity) guarantee that the percentage expected losses of all mezzanine tranches are between 0 and 1. This immediately implies that $0 \leq P_n \leq 1$. We then keep moving to the left, looking at each P_{k_i} in backward succession. Recall that by definition of the expected loss of mezzanine tranche

$$\begin{aligned} q_i &= \sum_{j=k_{i-1}+1}^{k_i} \frac{j - U_{i-1}}{U_i - U_{i-1}} P_j + \mathbf{P}\{L > k_i\} \\ &\leq \frac{k_i - U_{i-1}}{U_i - U_{i-1}} \mathbf{P}\{k_{i-1} < L \leq k_i\} + \mathbf{P}\{L \geq k_i\}, \end{aligned}$$

so that, for example, for the senior-most detachment point K_m we have

$$P_{k_m} \leq \mathbf{P}\{k_{m-1} < L \leq k_m\} + \frac{\mathbf{P}\{L \geq k_m\} - P_n}{d_m} = \mathbf{P}\{k_{m-1} < L \leq k_m\}$$

since in the given construction, the tail of the distribution is concentrated at n , as we have established at the previous step. Thus P_{k_m} is bounded by the probability mass contained between the points K_{m-1} and K_m . We can follow the same argument for all subsequent K_i , $i < m$, to establish the correspondence between P_{k_i} and $\mathbf{P}\{K_{i-1} < L \leq K_i\}$. This guarantees the validity of the constructed loss distribution, as long as the originally given tranche expected losses satisfy the no-arbitrage conditions. We demonstrate the connection in more detail below.

Using (6), we take the difference between expected losses of equity tranches detaching at adjacent loss grid points – say, k and $k + 1$ loss units:

$$EL(0, k+1) - EL(0, k) = (k+1)P_{k+1} + (k+1-k) \sum_{j=k+2}^n P_j - kP_{k+1} = \sum_{j=k+1}^n P_j.$$

This is just the tail probability beyond the upper strike. If we take the second difference now, we obtain

$$\begin{aligned} EL(0, k+1) - 2EL(0, k) + EL(0, k-1) &= (EL(0, k+1) - EL(0, k)) \\ &\quad - (EL(0, k) - EL(0, k-1)) \\ &= \sum_{j=k+1}^n P_j - \sum_{j=k}^n P_j = -P_k. \end{aligned}$$

This makes the no-arbitrage conditions clear: they simply guarantee the positivity of both the cumulative loss distribution function and the loss probabilities themselves. It also demonstrates why these conditions are sufficient for the specific loss distribution that we have just constructed (i.e., the feasible initial guess) to be valid.

Our final note is to observe that the given interpretation of the first and second differences of the expected loss function gives an intuitive geometric

visualisation of the initial guess. From the above formulæ it is evident that the probability of a given loss can be zero only when the second difference of the expected loss function at this gridpoint vanishes – in other words, if the expected loss function is linear around this point. Therefore in our construction, the expected loss is piecewise linear, switching at the points k_1, \dots, k_m, n . Therefore the initial construction corresponds to a piecewise linear interpolation of tranche expected loss, with the subtle point that the nodes are not the given strikes K_1, \dots, K_m , but the closest loss gridpoints just below these strikes.

3.3 Tackling arbitrage

The previous section presented a recipe for constructing a valid fixed-horizon loss distribution from an arbitrage-free tranche expected loss sequence for a fixed horizon. We showed how this sequence can be used to construct a feasible initial guess, to be plugged into the quadratic programme (7) to obtain an arbitrage-free loss distribution that fits the given expected losses.

However examples in Section 2.2 suggest that even for a sequence of liquid tranche detachment points, the corresponding tranche expected losses may not be consistent, i.e., present arbitrage, especially for shorter horizons.

In search of a solution in these cases, we look at the expected loss profiles obtained in practice and try to single out the points where arbitrage-free conditions get violated most frequently. We observe that most of the time arbitrages emerge at the senior end of the capital structure. Often the total expected loss of all portfolio constituents puts too much loss into the part of the loss distribution above the senior-most quoted detachment point. This can even happen at maturity and is consistent with the findings in [13]. For shorter horizons, it can also be the case that senior and even mezzanine losses are mis-aligned. In all cases we have observed, however, it was always possible to fit the equity price and usually also the junior mezzanine, without introducing arbitrage.

Our recipe is therefore to “filter” the given tranche expected losses and only include, as the constraints in (7), the ones that do not violate no-arbitrage conditions. Thus at any given horizon we may end up with fewer than $m + 1$ equality constraints in the optimisation problem. This results in divergence of model-implied liquid tranche prices from the original market (usually at the senior end), since the loss surface is no longer consistent with base correlations at *all* intermediate dates. In practice the discrepancies in tranche losses turn out to be relatively small. This is due to the fact that they tend to occur on just a handful of dates close to valuation date, while the tranche market maturities are considerably longer (5 years at the least). Consequently the resulting tranche curves – and prices – are fairly close to the original values used to back out the base correlation curve. We demonstrate this with several examples later.

3.4 Time dimension revisited

So far we have been able to produce a sequence of fixed-horizon loss distributions that are consistent with as many liquid tranche expected losses as possible at

each horizon. However putting this sequence together will only produce tranche curves if their behaviour in time is correct (i.e., arbitrage-free) as well. In other words, we need to ensure that expected losses of all possible tranches of the portfolio increase in time.

The bad news is that in the current formulation given by (7) this is impossible to check. Individual loss probabilities P_k (i.e., the values of the loss distribution at any node) do not have to be monotonic in time for tranche losses to keep going up – the corresponding condition, though sufficient for tranche loss monotonicity, is clearly not necessary and in fact, too restrictive³. Thus the formulation in terms of loss probabilities P_k is not easily extendable in the time dimension. A modification is required.

4 Cumulative Loss Formulation

Fortunately, while a no-arbitrage condition in time is problematic for the loss *density*, it is much more obvious in *cumulative* loss distribution, where we can write

$$\mathbf{P}\{L(t_i) \leq x\} \geq \mathbf{P}\{L(t_j) \leq x\} \quad \text{for } t_j > t_i. \quad (11)$$

In other words, cumulative loss to any point is non-increasing in time. Thus if we are able to recast the formulation (7) in terms of cumulative probabilities $Q_k = \mathbf{P}\{L \leq ku\}$, an extra monotonicity constraint between adjacent horizon will ensure that arbitrage in time cannot emerge.

Since the loss distribution is discrete, we have

$$Q_0 = P_0, \quad Q_k = \sum_{j=0}^k P_j, \quad 1 \leq k < n, \quad Q_n = 1,$$

so that

$$P_0 = Q_0, \quad P_k = Q_k - Q_{k-1} \text{ for } k > 0,$$

and therefore for an equity tranche expected loss (6) we have

$$EL(0, U) = \sum_{j=1}^{[U]} j(Q_j - Q_{j-1}) + U(1 - Q_{[U]}). \quad (12)$$

The first term can be rearranged more conveniently as

$$\sum_{j=1}^{[U]} j(Q_j - Q_{j-1}) = [U] Q_{[U]} - \sum_{j=0}^{[U]-1} Q_j.$$

³It is most likely to be violated in pairs of adjacent dates, where we have dropped or reinstated an extra tranche expected loss condition.

This enables us to rewrite any of the equity tranche expected loss equations as follows

$$\begin{aligned} k_i Q_{k_i} - \sum_{j=0}^{k_i-1} Q_j + U_i(1 - Q_{k_i}) &= EL_i \iff \\ \sum_{j=0}^{k_i-1} Q_j + (U_i - k_i)Q_{k_i} &= U_i - EL_i, \end{aligned}$$

for any i between 1 and m .

We also need to add extra constraints to the new formulation: in addition to the usual conditions of probabilities being bounded between 0 and 1, the cumulative loss probabilities should also be monotonic: $0 \leq Q_j - Q_{j-1} \leq 1$ (so that $0 \leq P_j \leq 1$). The dimensionality of the problem stays the same at n , as it is now the last probability, Q_n , that drops out (as it must be equal to 1), while Q_0 figures in the equations. The objective functions also change from

$$f(P) = \frac{1}{2} \sum_{j=1}^n (P_j - P_{j-1})^2$$

to

$$F(Q) = \frac{1}{2} (Q_1 - 2Q_0)^2 + \frac{1}{2} \sum_{j=1}^{n-2} (Q_{j-1} - 2Q_j + Q_{j+1})^2 + \frac{1}{2} (Q_{n-2} - 2Q_{n-1} + 1)^2$$

making the Hessian

$$H(F) = \frac{\partial F}{\partial Q_i \partial Q_j} = \begin{pmatrix} 5 & -4 & 1 & 0 & 0 & \cdots & 0 & 0 & 0 & 0 & 0 \\ -4 & 6 & -4 & 1 & 0 & \cdots & 0 & 0 & 0 & 0 & 0 \\ 1 & -4 & 6 & -4 & 1 & \cdots & 0 & 0 & 0 & 0 & 0 \\ & \vdots & & \vdots & & \ddots & & \vdots & & \vdots & \\ 0 & 0 & 0 & 0 & 0 & \cdots & 1 & -4 & 6 & -4 & 1 \\ 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 1 & -4 & 5 & -2 \\ 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 1 & -2 & 1 \end{pmatrix}.$$

We thus have the following formulation in cumulative loss probabilities:

$$\begin{aligned} F(Q) &\rightarrow \min & (13) \\ \text{s.t. : } & 0 \leq Q_0, Q_j \leq 1, \quad 0 \leq Q_j - Q_{j-1} \leq 1, \quad j = 1, \dots, n-1, \\ & \sum_{j=0}^{k_1-1} Q_j + (U_1 - k_1)Q_{k_1} = U_1 - EL_1, \\ & \dots \\ & \sum_{j=0}^{k_m-1} Q_j + (U_m - k_m)Q_{k_m} = U_m - EL_m \\ & \sum_{j=0}^{n-1} Q_j = n - \tilde{\mu}. \end{aligned}$$

A solution to the problem (13) gives a vector of cumulative loss probabilities $\{Q_0, \dots, Q_{n-1}\}$, such that the corresponding probabilities

$$P_0 = Q_0, \quad P_1 = Q_1 - Q_0, \quad \dots, \quad P_{n-1} = Q_{n-1} - Q_{n-2}, \quad P_n = 1 - Q_{n-1}$$

form a valid loss distribution, which is consistent with the tranche expected losses at the given time horizon. In other words, this is an exact analogue of the problem (7) in the cumulative loss space.

To complete the presentation of the framework, we need to address several technical points.

4.1 Feasible initial guess

An example initial guess that satisfies all constraints can be constructed using the same idea as in Section 3.2, where we set all but $m + 1$ loss probabilities to zero. The equivalent structure here is piecewise flat cumulative loss with $m + 1$ “steps”. To this end, we seek an initial guess in the following form:

$$Q_0 = Q_1 = \dots = Q_{k_1-1}, \quad Q_{k_1} = \dots = Q_{k_2-1}, \quad \dots, \quad Q_{k_m} = \dots = Q_{n-1},$$

so that the loss distribution is completely determined by setting $m + 1$ probabilities $Q_0, Q_{k_1}, \dots, Q_{k_m}$. In this case, the equality (expected loss) constraints are reduced to

$$\begin{aligned} k_1 Q_0 + (U_1 - k_1) Q_{k_1} &= U_1 - EL_1, \\ k_1 Q_0 + (k_2 - k_1) Q_{k_1} + (U_2 - k_2) Q_{k_2} &= U_2 - EL_2, \\ &\dots \\ k_1 Q_0 + \sum_{j=1}^{m-1} (k_{j+1} - k_j) Q_{k_j} + (U_m - k_m) Q_{k_m} &= U_m - EL_m, \\ k_1 Q_0 + \sum_{j=1}^{m-1} (k_{j+1} - k_j) Q_{k_j} + (n - k_m) Q_{k_m} &= n - \tilde{\mu}. \end{aligned}$$

To solve the above, we subtract from each equation the previous one, leaving the first equation intact. Some rearrangement of terms yields the following system

$$\begin{aligned} U_1 Q_0 + (U_1 - k_1) \Delta Q_{k_1} &= U_1 - EL_1, \\ \Delta U_2 Q_{k_1} + (U_2 - k_2) \Delta Q_{k_2} &= \Delta U_2 - (EL_2 - EL_1), \\ &\dots \\ \Delta U_m Q_{k_{m-1}} + (U_m - k_m) \Delta Q_{k_m} &= \Delta U_m - (EL_m - EL_{m-1}), \\ (n - U_m) Q_{k_m} &= (n - U_m) - (\tilde{\mu} - EL_m), \end{aligned}$$

where we have denoted $\Delta U_j = U_j - U_{j-1}$ and $\Delta Q_{k_j} = Q_{k_j} - Q_{k_{j-1}}$. The solution can be found recursively:

$$Q_{k_m} = SP(K_m, K_{\max}), \quad Q_{k_j} = SP(K_j, K_{j+1}) - (1 - d_{j+1}) \Delta Q_{k_{j+1}}. \quad (14)$$

In the above we recall the definition of the survival probability of a tranche attaching at A and detaching at D percent as

$$SP(A, D) = 1 - EL\%(A, D) = 1 - \frac{EL(0, D) - EL(0, A)}{D - A}$$

and, as in (10)

$$d_{j+1} = \frac{k_{j+1} - U_j}{U_{j+1} - U_j}, \quad j \geq 0,$$

with $K_0 = 0\%$.

Since by definition, $k_j = [U_j] \leq U_j$, we have $d_{j+1} \leq 1$; also, if the standard tranches are at least one loss thick (a constraint which is comfortably satisfied by any of the current liquid tranche markets), then $d_{j+1} \geq 0$. In this case we have the following intuitive bounds on the cumulative loss probabilities:

$$SP(K_{j-1}, K_j) \leq Q_{k_j} \leq SP(K_j, K_{j+1}), \quad j = 1, \dots, m,$$

where we set $K_0 = 0\%$ and $K_{m+1} = K_{\max}$.

The second equation in (14) implies that

$$\Delta Q_{k_{j+1}} = \frac{1}{d_{j+1}} (Q_{k_{j+1}} - SP(K_j, K_{j+1})),$$

which gives the desired recurrence relation.

The obtained probabilities satisfy the expected loss constraints by definition. They will also be non-negative and monotonically increasing (non-decreasing), as long as the same is true for the tranche survival probabilities. The latter, in its turn, is guaranteed if there is no arbitrage at the given horizon. The next section summarises the recipes for handling arbitrage, some of which were described in Section 3.3.

4.2 Arbitrage

Capital structure arbitrages arise when the tranche survival probabilities decrease with seniority at a given fixed horizon, or are negative. In accordance with the scheme proposed in Section 3.3, we reduce the number of expected loss constraints until all of the remaining ones are arbitrage-free. As before, this means that the corresponding liquid tranches will no longer be exactly repriced to market, but since the occurrences of such arbitrage are usually few and limited to short horizons, the practical impact tends to be quite low. Another possibility is that the arbitrage only occurs in the super-senior part of the capital structure, so that all the standard mezzanine tranches are repriced, but the portfolio expected loss is not conserved. In this case we will not be able to start the recurrence (14) with the super-senior tranche survival probability. Instead, we can use any value between the senior-most mezzanine tranche survival probability and 1 as Q_{k_m} (as long as it satisfies any additional constraints, such as the ones we mention below, in connection with arbitrage in time). In any

case fixed-horizon arbitrage is handled in the current formulation in the same manner as in the loss distribution formulation of the Section 3.

However, unlike the previous approach, the current setup also gives a way of removing arbitrage in time. For any two successive time horizons $T < T'$, the fundamental relationship (11) implies that

$$Q_j(T) \geq Q_j(T').$$

Therefore to enforce no-arbitrage in maturity, it is sufficient to add the cumulative loss distribution at the preceding horizon as the upper bound for the optimisation problem at the given horizon. In other words, the constraint $0 \leq Q_j \leq 1$ should be replaced with $0 \leq Q_j \leq Q_j^-$, where Q_j^- is the cumulative loss probability at the same node at the previous horizon. We can thus solve a sequence of problems given by (13) for all dates on the portfolio grid in succession, updating the upper bound with the solution obtained at the previous step.

4.3 Different objective functions

As we have mentioned in Section 3.1, the choice of the objective function in the optimisation problems (7), (13) is more or less arbitrary. By this we mean that consistency with the tranche market is guaranteed by the constraints, while the objective function embeds the “target shape” of the loss distribution, which is an internal construction that is neither set, nor even observable in the market. Different objective functions may produce different loss distributions, ultimately affecting prices of non-standard tranches. Choosing a particular shape over another one is hard to formalise: since either will fit the liquid tranche market, there are no “right” or “wrong” shapes. Our choice of the smoothest distribution may be thought of as a natural first choice; other choices include entropy-minimising (cf. [6]) distributions and distributions close to a given one.

For the latter, if we would like the shape of the resulting loss distributions not to deviate too far from a given law, the corresponding objective function will be

$$\tilde{F}(Q) = \frac{1}{2} \sum_{j=0}^{n-1} \left(Q_j - Q_j^{(0)} \right)^2. \quad (15)$$

Here $Q^{(0)}$ represents the target distribution – for example, $Q_j^{(0)} = j/n$ for the uniform target or $Q_j^{(0)} = LCDF_{\Phi, \rho}(j/n)$ for the loss CDF obtained from the Gaussian copula model with correlation ρ . The dependence of non-standard tranche prices on the choice of target distributions will be evident from the examples we give in the next section.

When a target distribution is given, it may be beneficial to modify the procedure for constructing an initial guess. Instead of a piecewise flat distribution, one can consider a piecewise flat *increment* to the target distribution:

$$Q_j = Q_j^{(0)} + \theta_j,$$

where θ_j has the $m+1$ -step structure described in Section 4.1. The i -th expected loss constraint then reads

$$k_1 Q_0 + \sum_{j=1}^{i-1} (k_{j+1} - k_j) Q_{k_j} + (U_i - k_i) Q_{k_i} = EL_i^{(0)} - EL_i$$

and the total expected loss constraints becomes

$$k_1 Q_0 + \sum_{j=1}^{m-1} (k_{j+1} - k_j) Q_{k_j} + (n - k_m) Q_{k_m} = \tilde{\mu}^{(0)} - \tilde{\mu},$$

where we use the superscript “(0)” for the quantities calculated from the target distribution. We can solve for the increments to obtain the following equivalent of (14):

$$\begin{aligned} Q_j &= Q_j^{(0)} + \theta_{k_m}, & j \geq k_m \\ Q_j &= Q_j^{(0)} + \theta_{k_i}, & k_i \leq j \leq k_{i+1} - 1, \quad i = 1, \dots, m-1 \end{aligned}$$

with

$$\begin{aligned} \theta_{k_m} &= SP(K_m, K_{\max}) - SP^{(0)}(K_m, K_{\max}), \\ \theta_{k_{i+1}} - \theta_{k_i} &= \frac{1}{d_{i+1}} \left(\theta_{k_{i+1}} - SP(K_i, K_{i+1}) + SP^{(0)}(K_i, K_{i+1}) \right), \end{aligned}$$

for $i = 1, \dots, m-1$.

Note that while these probabilities satisfy the expected loss constraints by construction, verifying the inequality constraints (positivity and monotonicity in loss and maturity) will either impose restrictions on the target distributions (via admissible values for mezzanine tranches survival probabilities), or lead to exclusion of more expected loss constraints during the arbitrage “filtering” stage. Therefore care must be taken in applying this procedure for constructing a feasible initial guess.

4.4 Non-homogeneous portfolios

The formulæ presented above are, strictly speaking, only valid for homogeneous portfolios, such as liquid indices under standard assumptions, where exposures and recovery rates of all constituents are the same. In mathematical terms, we have only considered the case when $l_i = 1$ for all $i = 1, \dots, n$, and consequently $n = L_{\max} \equiv \sum l_i$: the dimension of the optimisation problem (13) is equal to the number of portfolio constituents. In other words, the support of the loss distribution consists of all consecutive integers from 0 up to n , the number of names in the portfolio.

This relationship is violated when we allow different recovery rates and/or notional for the constituents. For example, a simple portfolio of two names, one with notional \$10 million and recovery rate 40% and the other with notional \$20

million and recovery 20%, will have very different properties from a homogenous two-name portfolio. The support of the loss distribution is $\{0, 3, 8, 11\}$, since the loss unit size is \$2 million, so that $l_1 = 3$ and $l_2 = 8$. Thus the size of the distribution is $N = 4$, but the number of names is $n = 2$, while the maximum number of loss units $L_{\max} = 11$, which is clearly very different from the homogeneous case, and we would not be able to use the formulæ above directly.

To extend the framework to include these cases, which are especially useful in bespoke portfolios, we introduce a new quantity, $\lambda(j)$ for $j = 0, \dots, N$, which is simply the loss amount (number of loss units) at the j -th node of the loss distribution. For the homogeneous case, $\lambda(j) = j$ and $\lambda(N) = n = N = L_{\max}$; however, in the general case, $\lambda(j) \geq j$ and $\lambda(N) = L_{\max} \geq N \geq n$. We can then rewrite the expected loss formula (12) as

$$EL(0, U) = \sum_{j=1}^{[U]} \lambda(j)(Q_j - Q_{j-1}) + U(1 - Q_{[U]})$$

and rearrange the first term:

$$\sum_{j=1}^{[U]} \lambda(j)(Q_j - Q_{j-1}) = \lambda([U])Q_{[U]} - \sum_{j=0}^{[U]-1} (\lambda(j+1) - \lambda(j))Q_j.$$

For consistency, we also need to note that the probabilities Q_j in the above are now understood as $Q_j = \mathbf{P}\{L \leq \lambda(j)u\}$.

Consequently, for any $i \in [1, m]$, the corresponding tranche expected loss constraint will have the following form in the general case

$$\sum_{j=0}^{k_i-1} [\lambda(j+1) - \lambda(j)]Q_j + [U_i - n(k_i)] = U_i - EL_i,$$

and the portfolio expected loss constraint is

$$\sum_{j=0}^{N-1} [\lambda(j+1) - \lambda(j)]Q_j = \lambda(N) - \tilde{\mu}.$$

The rest of the analysis can be carried out as before.

5 Numerical results

We present various examples illustrating the performance of the current framework. In the following sections we demonstrate calibration quality, thin index tranche prices and results for a bespoke portfolio.

5.1 Market fit

Tables 4-5 show the quality of fit for our framework to the liquid tranche markets on various dates over the life of CDX.IG Series 9 and iTRAXX Main Series 8. These series are interesting because they have seen both relatively tight and quite wide markets. We can see that model-implied spreads (or upfront premia with 500 bp running spread for equity tranches) are usually very close to the corresponding mid-market values. The only time when we do not get inside the bid-offer is for the CDX 10Y 15-30% tranche on 7 December, but even there the error is quite small (0.1 bp above the offer quote, with 1.5bp bid-offer spread and level of the order of 50 bp). We also observe that the equity tranche always fits exactly, as expected from construction: a pricing error is introduced for a given tranche only if we have to remove the corresponding expected loss constraint, as part of arbitrage “filtering” on some dates, but we should get to keep at least the first constraints on all dates.

Tranche	29-Oct-2007 (5Y)				4-Mar-2008 (5Y)			
	Bid	Offer	Mid	Model	Bid	Offer	Mid	Model
0-3%	16.625	17	16.8125	16.8125	35.25	36.25	35.75	35.75
3-6%	104.5	106.5	105.5	105.53	445	455	450	450.1
6-9%	44.5	46	45.25	45.31	303	318	310.5	310.48
9-12%	26.5	29	27.75	27.82	225	235	230	230
12-22%	16.5	18.5	17.5	17.56	125	135	130	130.01
Tranche	29-Oct-2007 (7Y)				4-Mar-2008 (7Y)			
	Bid	Offer	Mid	Model	Bid	Offer	Mid	Model
0-3%	25.625	26	25.8125	25.8125	40.5	42	41.25	41.25
3-6%	160	164	162	162.03	495	510	502.5	502.56
6-9%	77	79	78	78.03	330	345	337.5	337.53
9-12%	45	47.5	46.25	46.32	260	275	267.5	267.49
12-22%	27	30	28.5	28.57	145	155	150	150
Tranche	29-Oct-2007 (10Y)				4-Mar-2008 (10Y)			
	Bid	Offer	Mid	Model	Bid	Offer	Mid	Model
0-3%	34.125	34.625	34.375	34.375	46	47.5	46.75	46.75
3-6%	329	334	331.5	331.5	590	610	600	600
6-9%	132	135	133.5	133.522	375	395	385	385
9-12%	74	77	75.5	75.58	310	325	317.5	317.4
12-22%	37	39	38	38.36	155	170	162.5	164.32

Table 4: iTRAXX Series 8 tranche spreads and fit.

5.2 Thin tranches

We now revisit Table 1 and append it with thin tranche prices given by the new framework. In Table 6 we provide results for two types of objective functions

Tranche	7-Dec-2007 (5Y)				14-Jan-2008 (5Y)			
	Bid	Offer	Mid	Model	Bid	Offer	Mid	Model
0-3%	48.375	49.375	48.875	48.875	52.812	53.812	53.312	53.31
3-7%	233.28	247.72	240.5	240.55	373.45	396.55	385	385.1
7-10%	108.64	115.36	112	112.03	170.72	181.28	176	176.04
10-15%	56.26	59.74	58	58.16	96.03	101.97	99	99.07
15-30%	37.83	40.17	39	39.13	54.32	57.68	56	57.33
Tranche	7-Dec-2007 (7Y)				14-Jan-2008 (7Y)			
	Bid	Offer	Mid	Model	Bid	Offer	Mid	Model
0-3%	55.25	57.25	56.25	56.25	58.062	60.062	59.062	59.062
3-7%	301.95	369.05	335.50	335.52	447.75	547.25	497.5	497.55
7-10%	142.2	173.8	158	158	201.15	245.85	223.5	223.5
10-15%	78.3	95.7	87	87.01	112.5	137.5	125	125
15-30%	42.3	51.7	479	47.01	61.2	74.8	68	68
Tranche	7-Dec-2007 (10Y)				14-Jan-2008 (10Y)			
	Bid	Offer	Mid	Model	Bid	Offer	Mid	Model
0-3%	60.125	61.125	60.625	60.625	63.25	64.25	63.75	63.75
3-7%	554.19	576.81	565.5	565.5	637.98	664.02	651	651
7-10%	200.3	212.7	206.5	206.52	287.12	304.88	296	296.1
10-15%	106.7	113.3	110	110.06	152.77	162.23	157.5	157.53
15-30%	49.47	52.53	51	52.64	75.418	80.082	77.75	78.16

Table 5: CDX.IG Series 9 tranche spreads and fit.

(smoothness maximising and closest to the one-factor Gaussian copula loss distribution). We can observe that one seems to be somewhat more effective at diffusing the initial flat portions of the loss distribution than the other.

Tranche		Fair spread, bp					
Attach	Detach	Base corr interp			ExpLoss	Loss dist	
		Linear	Spline	Monotonic	Interp	Smooth	Gauss
13%	14%	76.72	83.05	87.05	93.75	91.71	83.79
14%	15%	60.98	75.81	91.62	88.86	75.33	72.22
15%	16%	118.19	70.07	88.12	75.01	61.51	72.09
16%	17%	104.46	64.85	74.07	62.17	57.07	72.06
17%	18%	93.50	61.80	63.54	57.05	56.53	71.47
18%	19%	85.20	61.49	56.91	55.17	56.52	64.75
19%	20%	77.85	61.97	52.43	54.41	56.51	55.97
20%	21%	69.74	60.69	47.71	54.06	56.50	52.76
21%	22%	60.38	56.71	41.56	53.88	56.49	52.76
22%	23%	50.67	51.46	35.11	53.79	56.49	52.76
23%	12%	42.00	47.55	33.89	53.74	56.49	52.76
24%	13%	36.04	48.24	37.18	53.70	56.49	52.76
25%	26%	32.41	50.49	42.50	53.67	56.48	52.75

Table 6: CDX.IG9 5Y thin tranche spreads, different interpolations (14-Jan-08).

In Figure 1 we look at the debated issue of the shape of the base correlation curve below the 3% strike. We recall that in [18] the authors used the concept of local correlation to argue that it should curve upwards, forming a “smile in the skew”, contrary to the natural extrapolation of the usual base correlation interpolation methods, which produce a downwards-sloping curve below 3%. For the other methods, correlations can be bootstrapped from a sequence of model-implied mezzanine tranches spreads. We observe that one of the popular non-Gaussian models – the gamma model [4] – also implies a monotonic base correlation curve. At the same time, extrapolation in equity tranche expected loss (rather than base correlations) induces an upward slope at the short end, as does our framework. Even if we ignore extrapolation-based methods, results from the two arbitrage-free constructions are in disagreement. While interpretations of this discrepancy may vary – e.g., one may recall problems with global fit of the gamma model or the fact that the loss distribution framework is influenced by the initial guess – it is important to keep in mind that multiple arbitrage-free models fitting the market prices may indeed exist, and they are not restricted to producing similar results away from the liquid strikes. Our own framework was built constructively, there was no evidence that it was the *unique* solution to the loss distribution generation problem.




Figure 1: iTRAXX8 5Y base correlations around and below 3% (4-Mar-08).

5.3 Bespoke portfolios

To demonstrate the performance of our framework outside of the liquid indices, we present results for pricing a series of tranches of a bespoke portfolio. Equivalent correlations for the bespoke are determined by means of a portfolio expected loss-based mapping of standard strikes (cf. [3]). The results obtained from different liquid index base correlations are combined, via an averaging procedure, based on portfolio composition and maturity, to form a tranche curve for a given bespoke tranche.

The example portfolio consists of 100 names, of which 43 are European and 57 are American. Maturity is roughly half-way between 7Y and 10Y liquid index tenors. We take real-life CDS curves as of February, 2008, end-of-month; the average spread is 135 bp, with names ranging from 35 bp (Johnson & Johnson) to 755 bp (First Data Corp); we include senior, as well as subordinated debt, with 40% and 20% assumed recovery, respectively. To explore the framework a bit further, we use three different types of objective function, adding to the customary smoothness maximisation and closest to the one-factor Gaussian copula distribution, a closest to uniform distribution. The piecewise flat loss CDF is still used as the initial guess.

Table 7 presents results for 2%-thick tranches from 0 all the way up to 62%. We can see traces of the initial distribution displaying themselves in the form of tranche sequences with almost identical spreads, but tending the distribution to the one coming out of a one-factor Gaussian copula model of correlated defaults looks the most effective in producing a more spread-out distribution.

6 Conclusion

We examined a typical base correlation loss surface and explored the possibility of modifying it to get rid of arbitrages. We looked for minimal changes to the original surface, which would restore tranche expected loss consistency without deteriorating the fit to market prices too much. Our findings indicate that base correlation loss surfaces are not actually that far from being arbitrage-free. While violations can be quite blatant, a simple procedure allows to filter them out, while at the same time containing liquid tranche prices within the bid-offer on the absolute majority of the observed dates. This certainly serves as an encouragement for the development of arbitrage-free market-fitting models.

In the process of our investigations, we have developed and presented an algorithm for generating a strip of portfolio loss distributions at a pre-specified set of dates. Together with a choice of initial guess and target distribution shape, this algorithm can be used as an arbitrage-free CDO tranche pricing method. We demonstrated how this method produces consistent prices for arbitrary index tranches and, coupled with a suitable index-to-bespoke correlation mapping method, for bespoke tranches as well.

The current approach is perhaps the most direct way of getting an arbitrage-free pricing tool from base correlations. The presented pricer is not unique, since

Tranche		Fair spread, bp		
Attach	Detach	Smooth	Uniform	Gauss
0%	2%	2395.0	2282.84	2352.6
2%	4%	1116.0	1226.7	1198.6
4%	6%	664.2	647.9	655.0
6%	8%	436.3	451.3	443.6
8%	10%	336.1	332.5	330.3
10%	12%	235.9	241.9	241.1
12%	14%	206.3	197.9	199.0
14%	16%	135.0	148.0	151.6
16%	18%	122.3	135.4	135.8
18%	20%	122.2	126.4	126.2
20%	22%	122.2	123.9	122.4
22%	24%	119.3	115.4	113.8
24%	26%	109.9	110.1	108.2
26%	28%	102.8	108.4	106.4
28%	30%	97.4	103.3	103.0
30%	32%	95.0	98.7	99.0
32%	34%	94.0	98.5	97.5
34%	36%	93.8	98.5	97.4
36%	38%	93.7	98.5	97.3
38%	40%	93.7	98.5	97.2
40%	42%	93.7	98.5	96.9
42%	44%	93.7	98.5	96.6
44%	46%	93.7	98.4	95.1
46%	48%	93.7	98.4	92.0
48%	50%	93.7	98.4	89.0
50%	52%	93.7	98.3	87.2
52%	54%	93.7	98.1	86.6
54%	56%	93.7	97.5	86.3
56%	58%	93.7	90.3	86.1
58%	60%	93.6	63.4	86.0
60%	62%	64.9	34.7	59.8

Table 7: Breakeven spreads for tranches of a bespoke portfolio.

it depends on the arbitrage “filtering” algorithm, the initial guess and the choice of the target distribution. While we have chosen the simplest construction, the principle and the resulting framework is generic. It can be used for any products amenable to modelling with structural copula-like methods (notably, those without path-dependency and dynamic features). It is likely to find use in practical situations when base correlation consistency and market fit are key.

References

- [1] Amraoui S. and Hitier S., “Optimal Stochastic Recovery for Base Correlations”, *BNP Paribas*, June 2008.
- [2] Andresen L., Sidenius J. and Basu S., “All Your Hedges in One Basket”, *Risk*, November 2003.
- [3] Baheti P. and Morgan S., “Base Correlation Mapping”, *Lehman Brothers Quantitative Credit Research Quarterly*, March 2007.
- [4] Baxter M., “Lévy process dynamic modelling of single-name credits and CDO tranches”, *Nomura Fixed Income Quant Group*, April 2006.
- [5] Krekel M., “Pricing distressed CDOs with Base Correlation and Stochastic Recovery”, *UniCredit Markets & Investment Banking*, May 2008.
- [6] Krekel M. and Partenheimer J., “The Implied Loss Surface of CDOs”, *HVB*, December 2006.
- [7] Laurent J.P. and Gregory J., “Basket Default Swaps, CDO’s and Factor Copulas”, September 2003.
- [8] “The Lehman Brothers Guide to Exotic Credit Derivatives”, *Risk books*, October 2003.
- [9] McGinty L., Beinstein E., Ahluwalia R. and Watts M., “Introducing Base Correlations”, *JP Morgan Credit Derivatives Strategy*, March 2004.
- [10] O’Kane D. and Livesey M., “Base Correlations Explained”, *Lehman Brothers Quantitative Credit Research Quarterly*, November 2004.
- [11] O’Kane D. and Turnbull S., “Valuation of Credit Default Swaps”, *Lehman Brothers Quantitative Research Quarterly*, April 2003.
- [12] Pain A., Renault O. and Shelton D., “Base Correlation: the Term Structure Dimension,” *Citigroup Structured Credit Products*, December 2005.
- [13] Schlögl L., “Stochastic Recovery Rates and No-Arbitrage in the Tranche Markets”, *Presentation at KCL Seminar*, January 2006.
- [14] Schlögl L. and Greenberg A., “The Shape of Implied Loss Distributions”, *Lehman Brothers Quantitative Credit Research Quarterly*, March 2005.

- [15] Sidenius J., “On the Term Structure of Loss Distributions – a Forward Model Approach”, *JP Morgan*, September 2006.
- [16] Tavares P.A.C., Nguyen T.-U., Chapovsky A. and Vaysburd I., “Composite Basket Model”, *Merrill Lynch*, July 2004.
- [17] Torresetti R., Brigo D. and Pallavicini A., “Implied Expected Tranched Loss Surface from CDO Data”, *Banco IMI*, May 2007.
- [18] Turc J., Very P. and Benhamou D., “Princing CDOs with a Smile”, *Société Générale Quantitative Strategy*, February 2005.
- [19] Walker M.B., “CDO Models – Towards the Next Generation: Incomplete Markets and Term Structure”, *University of Toronto*, May 2006.
- [20] Willemann S., “An Evaluation of the Base Correlation Framework for Synthetic CDOs”, *Aarhus School of Business*, December 2004.