A Dynamic Correlation Modelling Framework with Consistent Stochastic Recovery

Li, Yadong

Barclays Capital

26 February 2009

A Dynamic Correlation Modelling Framework with Consistent Stochastic Recovery

Yadong Li*
Qauntitative Analytics, Barclays Capital

Initial Version: Feb 26, 2009
Current Revision: April 2, 2009

Abstract

This paper describes a flexible and tractable bottom-up dynamic correlation modelling framework with a consistent stochastic recovery specification. In this modelling framework, only the joint distributions of default indicators are determined from the calibration to the index tranches; and the joint distribution of default time and spread dynamics can be changed independently from the CDO tranche pricing by applying one of the existing top-down methods to the common factor process. Numerical results showed that the proposed modelling method achieved good calibration to the index tranches across multiple maturities under the current market conditions. This modelling framework offers a practical approach to price and risk manage the exotic correlation products.

Keywords: Credit, Correlation, CDO, Dynamic, Copula, Stochastic Recovery, Bottom-up, Top-down

1 Introduction

The base correlation model remains the most common method to price and risk manage synthetic CDOs (O'Kane & Livesey, 2004). It is well known that the base correlation model is not arbitrage free, and it cannot produce a consistent joint default time distribution; therefore the base correlation model cannot be used to price and risk manage any default path-dependent or spread-dependent products. Not too long ago, the deterministic recovery assumption was the common practice within the base correlation framework. However, in the recent market environments, models with the deterministic recovery often fail to calibrate to the index tranche market because it forces the senior most tranches to be risk free, leaving too much risk in the junior part of the capital structure. (Andersen & Sidenius, 2004) first proposed the stochastic recovery for Gaussian Copula. Recently, a number of stochastic recovery specifications were suggested for the base correlation framework, e.g. (Amraoui & Hitier, 2008) and (Krekel, 2008). With these stochastic recovery specifications, the senior most tranches become risky, allowing the base correlation model to calibrate. However, most of the existing stochastic recovery specifications are not internally consistent, i.e., they can't be used to drive a Monte Carlo simulation and match the underlying CDS curves' expected recovery across time. The stochastic recovery specifications therefore introduced another source of inconsistency to the already inconsistent base correlation framework.

*Email: yadong.li@gmail.com. The views expressed in this paper are the author's own and may not necessarily reflect those of Barclays Capital. The author is very grateful to Hongwei Cheng, Bjorn Flesaker, Sebastien Hitier, Ken Kyoung-Kuk Kim, Sam Morgan, Marco Naldi and Ariye Shater for many helpful comments and discussions. The remaining errors belong to the author. Part of the research described in this paper was done in former Lehman Brothers.

There have been a lot of efforts in developing alternative models to the base correlation model in order to better price and risk manage the exotic correlation products whose payoff may depend on the default paths and tranche spreads. One alternative modelling approach is to find a consistent static copula, which can produce the joint default time distribution in order to price default path-dependent instruments. Random Factor Loading (Andersen & Sidenius, 2004) and the Implied Copula (Hull & White, 2006) (Skarke, 2005) are examples of the alternative static copulas. Another alternative modelling approach is to develop dynamic correlation models, which can price the spread-dependent correlation instruments, e.g., tranche options.

There are two main categories of dynamic correlation models: the top-down approach and the bottom-up approach. The top-down approach directly models the dynamics of the portfolio loss distribution and ignores all the single name identities. The advantages of the top-down models include: 1) it is relatively easy to implement and calibrate and 2) it offers very rich spread dynamics. The main disadvantages of the top-down models include: 1) it lacks the single name risk and sensitivity 2) it can't be used to price a bespoke CDO from the index tranches because the spread dispersion, which is a critical factor in CDO pricing, is not captured by the top-down models. (Shonbucher, 2006), (Sidenius *et al.*, 2006), (Bennani, 2005), (Errais *et al.*, 2009) and (Arnsdorf & Halperin, 2007) are some representative examples of the top-down models.

The bottom-up approach, on the other hand, starts with the single name spread dynamics and a correlation structure; and then computes the portfolio and tranche spread dynamics as functions of the single name spread dynamics and the correlation structure. The advantage of the bottom-up approach is that it preserves the single name identities and the spread dispersion, and offers the single name sensitivity. A bottom-up model can produce the joint distribution of default times and spreads; therefore, it can cover a wider range of exotic correlation products than a top-down model. For example, any exotic contract whose payoff depends on the identity of an underlying issuer[1] cannot be easily handled with a top-down model. However, a bottom-up model is much more difficult to implement and calibrate. Often, the model parameters that control the spread dynamics also affect the tranche prices; therefore the calibration to the index tranche prices can put severe restrictions on the resulting spread dynamics, making it difficult to produce the desired spread dynamics and the goodness of fit to the index tranches simultaneously. Due to these difficulties, there is no known bottom-up model that can produce good index tranche calibration and flexible spread dynamics to the best knowledge of the author. (Mortensen, 2006), (Chapovsky *et al.*, 2006) and (Kogan, 2008) are some representative bottom-up dynamic correlation models.

Under the current market conditions, the stochastic recovery is required for a bottom-up dynamic correlation model to achieve good calibration to the index tranche prices. Most of the existing stochastic recovery specifications cannot be directly used by a bottom-up dynamic correlation model because of their intrinsic inconsistencies. Defining a consistent and tractable stochastic recovery specification remains a challenge.

This paper proposes a very flexible bottom-up dynamic correlation modelling framework, along with a consistent and tractable stochastic recovery specification. The proposed dynamic modelling framework can be easily calibrated to the index tranche prices across multiple maturities; and it also allows the spread dynamics to change independently from the tranche pricing, thus being able to produce very rich spread dynamics.

2 Consistent Stochastic Recovery

This section first describes the generic properties of recovery rates; then proposes a tractable and consistent stochastic recovery specification.

Define τ as the default time of an issuer, and $\mathbf{1}_{\tau<t}$ as the indicator that the issuer defaults before time t. The recovery rate $r(t_1, t_2)$ is a conditional random variable that represents the recovery rate conditioned on the issuer defaults between time t_1 and t_2, i.e. $\tau \in (t_1, t_2)$. $r(t, t)$ is used to denote the instantaneous

[1] For example, a vanilla bespoke CDO traded against a risky counterparty who does not post the full collateral. In this case, the identity of the counterparty is important.

recovery rate when the issuer defaults exactly at time t, i.e., $\tau \in (t, t+dt)$.

Definition 2.1 *The following terms are defined to simplify the explanation:*

1. *instantaneous mean:* $\mu(t,t) = \mathbb{E}[r(t,t)]$
2. *instantaneous variance:* $\sigma^2(t,t) = \mathrm{Var}[r(t,t)] = \mathbb{E}[r^2(t,t)] - \mu^2(t,t)$
3. *cumulative mean:* $\mu(0,t) = \mathbb{E}[r(0,t)]$
4. *cumulative variance:* $\sigma^2(0,t) = \mathrm{Var}[r(0,t)] = \mathbb{E}[r^2(0,t)] - \mu^2(0,t)$

The cumulative mean and variance of recovery rate are important for building the loss distribution at a given time horizon t using the semi-analytical method (Andersen *et al.*, 2003). The instantaneous mean and variance are useful inside a Monte Carlo simulation.

Proposition 2.2 *The recovery rate has the following properties:*

1. *The recovery rate is range bounded:* $r(t_1,t_2) \in [0,1]$, $\mu(t_1,t_2) \in [0,1]$
2. *The variance of the recovery rate is range bounded:* $\sigma^2(t_1,t_2) \in [0, \mu(t_1,t_2)(1-\mu(t_1,t_2))]$

The lower bound 0 of the recovery rate variance corresponds to the deterministic recovery. The upper bound of the recovery variance corresponds to a two point distribution with recovery rate values of $\{0,1\}$ and probabilities of $\{1-\mu, \mu\}$, whose variance is the largest among all recovery distributions with mean μ.

Consider two consecutive time periods of $(0,t_1)$ and (t_1,t_2), the following equation holds because both sides are the recovery amount between time $(0,t_2)$:

$$r(0,t_2)\mathbf{1}_{\tau\in(0,t_2)} = r(0,t_1)\mathbf{1}_{\tau\in(0,t_1)} + r(t_1,t_2)\mathbf{1}_{\tau\in(t_1,t_2)} \tag{1}$$

Take the expectation on the previous equation:

$$p(t_2)\mu(0,t_2) = p(t_1)\mu(0,t_1) + (p(t_2)-p(t_1))\mu(t_1,t_2) \tag{2}$$

where $p(t) = \mathbb{E}[\mathbf{1}_{\tau<t}]$ is the default probability over time. Squaring both sides of (1), the cross term disappears because the two periods do not overlap, also note $\mathbf{1}^2 = \mathbf{1}$:

$$r^2(0,t_2)\mathbf{1}_{\tau\in(0,t_2)} = r^2(0,t_1)\mathbf{1}_{\tau\in(0,t_1)} + r^2(t_1,t_2)\mathbf{1}_{\tau\in(t_1,t_2)} \tag{3}$$

Then taking the expectation yields:

$$p(t_2)\mathbb{E}[r^2(0,t_2)] = p(t_1)\mathbb{E}[r^2(0,t_1)] + (p(t_2)-p(t_1))\mathbb{E}[r^2(t_1,t_2)] \tag{4}$$

Dividing the period between $(0,t)$ into infinitesimal time intervals, (2) and (4) can be written in the following continuous form:

Proposition 2.3 *Suppose the default probability of the issuer* $p(t) = \mathbb{E}[\mathbf{1}_{\tau<t}]$ *is continuous and differentiable with* t. *The following relationship exists between the instantaneous mean recovery* $\mu(t,t)$ *and the cumulative mean recovery* $\mu(0,t)$*:*

$$\mu(0,t) = \frac{1}{p(t)}\int_0^t \mu(s,s)p'(s)ds = \frac{1}{p(t)}\int_0^{p(t)} \mu(p,p)dp \tag{5}$$

It is always possible to write the $\mu(t,t)$ *as* $\mu(p,p)$ *because the inverse function* $t^{-1}(p)$ *always exists since the* $p(t)$ *is monotonic and continuous. Similarly:*

$$\mu^2(0,t) + \sigma^2(0,t) = \frac{1}{p(t)}\int_0^t [\mu^2(s,s)+\sigma^2(s,s)]p'(s)ds = \frac{1}{p(t)}\int_0^{p(t)} [\mu^2(p,p)+\sigma^2(p,p)]dp \tag{6}$$

Note that the $\sigma^2(0,t)$ is not just an integration of the $\sigma^2(p,p)$, it also includes the contribution from changes in the $\mu(p,p)$. An observation that immediately follows the Proposition 2.3 is that if the $\mu(p,p)$ and $\sigma^2(p,p)$ are chosen to be analytical functions of the default probability p, the $\mu(0,t)$ and $\sigma^2(0,t)$ can be computed just from the value of $p(t)$ at time t using (5) and (6) regardless of the detailed shape of the $p(t)$ over time. This property is critical in developing the dynamic correlation modelling framework in the next section of this paper.

Considering a basket of n credits indexed by the subscript $i = 1...n$, the notional amount of each credit is w_i. The portfolio loss at time t is the sum of all the individual losses $L(t) = \sum_{i=1}^n w_i l_i$, where $l_i = \mathbf{1}_{\tau_i<t}(1 - r_i(0,t))$ is the loss for a unit notional amount of name i. The mean and variance of the l_i is easy to compute:

$$\begin{aligned}\mathbb{E}[l_i] &= p_i(t)(1-\mu_i(0,t)) &(7)\\ \mathrm{Var}[l_i] &= p_i(t)\sigma_i^2(0,t) + p_i(t)(1-p_i(t))(1-\mu_i(0,t))^2 &(8)\end{aligned}$$

If the $\mathbf{1}_{\tau_i<t}$ and $r_i(0,t)$ are independent between names, it is well known that the portfolio loss distribution at time t can be approximated by a normal distribution according to the central limit theorem (Shelton, 2004). The normal approximation to the loss distribution is fully characterized by the mean and variance of the portfolio loss $L(t)$, which can be computed as:

$$\begin{aligned}\mathbb{E}[L(t)] &= \mathbb{E}[\sum_1^n w_i l_i] = \sum_1^n w_i \mathbb{E}[l_i] = \sum_1^n w_i p_i(t)[1-\mu_i(0,t)] &(9)\\ \mathrm{Var}[L(t)] &= \mathrm{Var}[\sum_1^n w_i l_i] = \sum_1^n w_i^2 \mathrm{Var}[l_i] = \sum_1^n w_i^2 p_i(t)[\sigma_i^2(0,t) + (1-p_i(t))(1-\mu_i(0,t))^2] &(10)\end{aligned}$$

Therefore, the only recovery rate measures that are required to compute the loss distribution with the independent defaults and recovery rates are the $\mu_i(0,t)$ and $\sigma_i^2(0,t)$. The fine details of the recovery rate distribution other than the first two moments do not affect the portfolio loss distribution if n is reasonably large so that the normal approximation is sufficiently accurate. The same argument can be made for any conditional independent correlation models, e.g., Gaussian Copula.

Proposition 2.4 *Given a conditional independent correlation model, the loss distribution at time t is only sensitive to the first two moments of the cumulative recovery distribution, i.e., $\mu_i(0,t), \sigma_i^2(0,t)$. The higher moments of the recovery rate distribution's have very limited impact on the portfolio loss distribution.*

The effects of the higher moments of the stochastic recovery distribution are quantified in section 5.3 of this paper. Since the $\sigma_i^2(0,t)$ enters the variance of the portfolio loss in (10), a stochastic recovery model has to specify both the mean and variance of the recovery rate in order to correctly reproduce the portfolio loss distributions over time. Any stochastic recovery specification that does not capture the variance of recovery is inconsistent by construction. Also, the stochastic recovery models that directly specify the cumulative $\mu_i(0,t)$ and $\sigma_i^2(0,t)$, or the distribution of $r_i(0,t)$ are generally not consistent because their implied instantaneous recovery $r_i(t,t)$ is not guaranteed to satisfy the constraints in the Proposition 2.2. Most of the popular stochastic recovery specifications for the base correlation model, such as (Amraoui & Hitier, 2008) and (Krekel, 2008), are not internally consistent for the reasons above.

In conclusion, a consistent and tractable stochastic recovery specification can be easily constructed by defining the analytical functions for the $\mu_i(p,p)$ and $\sigma_i^2(p,p)$. In a conditional independent model, the μ_i and σ_i^2 can be defined as functions of the conditional default probability. It is natural to choose the $\mu(p,p)$ to be a decreasing function of conditional default probability p, since it forces the recovery rates to be lower in the bad states of the economy when a lot of names default. In a conditional independent model, the overall unconditional recovery rate is a weighted average of the conditional recovery rates over all possible states of the market factor.

One implication of this stochastic recovery specification is that the expected recovery does not remain constant over time unless the $\mu(p,p)$ is a constant. Today, the CDS curves are usually quoted with a constant expected recovery term structure, which is largely a result of the trading and quoting convention rather than any economic reasons. Therefore in practice, the requirement of a constant expected recovery rate over time can be relaxed to allow the $\mu(p,p)$ to be a non-constant function of p. The function $\mu(p,p)$ can be chosen so that the $\mu(0,t)$ matches the CDS curve recovery rate at the most liquid tenor, e.g., 5Y for the high grade names and 1Y for the distressed names. With this approach, the default probabilities at the less liquid tenors have to be adjusted so that the underlying names' expected losses are preserved. Given the lack of justification for a constant expected recovery term structure, allowing $\mu(p,p)$ to be non-constant is a good compromise since it allows more flexible term structure of the recovery rate, and it also helps the model calibration to the index tranches.

One advantage of this stochastic recovery specification is that it gives user control of the recovery variance through the parameter $\mu(p,p)$ and $\sigma_i^2(p,p)$. The recovery variance is very important to the CDO tranche pricing and risk especially when a name is very close to default.

3 Dynamic Correlation Modelling Framework

In this section, a bottom-up dynamic correlation modelling framework is described. In this framework, the spread dynamics can be changed independently from the CDO tranche pricing. The decoupling of CDO tranche pricing and spread dynamics offers great flexibility in constructing the desired spread dynamics, and it also greatly reduces the complexity of the model calibration. The decoupling of CDO tranche pricing and spread dynamics is achieved by defining a set of copula functions on the default indicators:

Definition 3.1 *Denote the unconditional default probability of the i-th name in the portfolio as $p_i(t)$. A set of copula functions on default indicators can be defined by the following three components:*

1. *A non-negative and increasing stochastic process $\tilde{X}(t)$ that represents the common market factor. The cumulative distribution function of the $\tilde{X}(t)$ is denoted as $F(x,t) = \mathbb{P}\{\tilde{X}(t) < x\}$ and its density function is denoted as $f(x,t) = \partial F(x,t)/\partial x$. The $f(x,t)$ is also referred as the marginal distribution of $\tilde{X}(t)$. A increasing $\tilde{X}(t)$ implies that its cumulative distribution function $F(x,t)$ has to be decreasing over time:*
$$\frac{\partial F(x,t)}{\partial t} \leq 0 \tag{11}$$
2. *A conditional default probability function $p_i(x,t) = \mathbb{E}[\mathbf{1}_{\tau_i<t}|\tilde{X}(t) = x]$ that satisfies the following constraints:*
$$\begin{aligned} p_i(x,t) &\in [0,1] &(12)\\ p_i(t) &= \mathbb{E}[p_i(x,t)] = \int p_i(x,t)f(x,t)dx &(13)\\ \frac{\partial p_i(x,t)}{\partial x} &\geq 0 &(14)\\ \frac{\partial p_i(x,t)}{\partial t} &\geq 0 &(15)\end{aligned}$$
The $p_i(x,t)$ function needs to have some name specific parameters so that it can be calibrated to the individual names' unconditional default probabilities according to (13). Constraints (14) and (15) ensure that the conditional default probability $p_i(x,t)$ are increasing for any possible path of $\tilde{X}(t)$ given that the $\tilde{X}(t)$ itself is increasing.
3. *conditional independence: individual names' default indicators $\mathbf{1}_{\tau_i<t}$ are independent conditioned on $\tilde{X}(t) = x$. The conditional independence ensures the uniqueness of the joint distribution of default indicators.*

The copula functions of default indicators in Definition 3.1 uniquely determines the joint distribution of default indicators at any time horizon from the unconditional single name default probabilities $p_i(t)$. In

Figure 1: $JDDI(t)$ vs $JDDT$

A Set of $JDDI(t)$

$\mathbf{1}_{\tau_1<t}$	$\mathbf{1}_{\tau_2<t}$	$JDDI(t=1)$	$JDDI(t=2)$
0	0	40%	20%
0	1	30%	30%
1	0	20%	20%
1	1	10%	30%

Two $JDDT$s

τ_1 **Range**	τ_2 **Range**	$JDDT_1$	$JDDT_2$
(2,∞)	(2,∞)	20%	20%
(2,∞)	(1,2)	0%	10%
(1,2)	(2,∞)	0%	10%
(1,2)	(1,2)	20%	0%
(2,∞)	(0,1)	30%	20%
(1,2)	(0,1)	0%	10%
(0,1)	(2,∞)	20%	10%
(0,1)	(1,2)	0%	10%
(0,1)	(0,1)	10%	10%

practice, any $p_i(x,t)$ function that satisfies the constraints (12) to (15) can be used to construct the copula functions of default indicators. A sample specification of $p_i(x,t)$ is given in section 4.

There is an important distinction between the copula function of default indicators, as in Definition 3.1, and a typical copula function of default time, as in the Gaussian Copula, Random Factor Loading or Implied Copula. A copula function of default time fully specifies the joint distribution of default time (abbreviate as $JDDT$), whereas a copula function on default indicators only specifies the joint distribution of default indicators (abbreviate as $JDDI(t)$) at a given tenor. The argument t in $JDDI(t)$ indicates that it is specific to a given time, while $JDDT$ does not take any time argument as it is globally applicable to all time horizon. A set of $JDDI(t)$ over time is always less restrictive than the $JDDT$ in the sense that there can be infinitely many $JDDT$s that are compatible with a set of $JDDI(t)$.

To illustrate the difference between the $JDDI(t)$ and the $JDDT$, Figure 1 showed a sample set of $JDDI(t)$ for a portfolio with two names over two time periods, as well as two different $JDDT$s which are compatible with the same set of $JDDI(t)$. Please note that both the $JDDT$ and $JDDI(t)$ are based on the time zero information only; and the $JDDT$ always have more information than the $JDDI(t)$. For example in Figure 1, considering an instrument that pays \$1 only if both name default within the time period (1, 2), this instrument cannot be priced by the $JDDI(t)$, but it can be priced by the $JDDT$s. It is also interesting to note that the two $JDDT$s in Figure 1 produce different prices for this instrument. The fact that the $JDDT$ contains more information than the $JDDI(t)$ is not a result of the discrete sampling in time. Even if the $JDDI(t)$ were specified in the continuous time, it still couldn't price the instrument that pays \$1 only if both name default within a specific future time period.

Definition 3.1 is very general, and it can produce the $JDDI(t)$ of typical one-factor default time copulas. For example, Gaussian Copula's $JDDI(t)$ can be produced by choosing:

$$f(x,t) = \phi(x) \tag{16}$$

$$p_i(x,t) = \Phi\left(\frac{\Phi^{-1}(p_i(t)) - \sqrt{\rho}x}{\sqrt{1-\rho}}\right) \tag{17}$$

where $\rho \in [0,1)$ is the correlation, $\phi(x)$ is the normal distribution density function and $\Phi(x)$ is the normal cumulative distribution function. Since the common factor x in Gaussian Copula can take negative values, a change of variable on the common factor such as $z = e^{-x}$ is required for the common factor to be non-negative as required by Definition 3.1. In general, the non-negative and increasing constraints on $\tilde{X}(t)$ in Definition 3.1 can be relaxed as long as $p_i(x,t)$ is increasing for any possible path of $\tilde{X}(t)$. In the case of Gaussian Copula, the common factor $\tilde{X}(t)$ remains constant, so that (11), (14) and (15) are trivially satisfied.

The transition from 1 to 0 is forbidden for the default indicators since a default event is irreversible. The irreversibility of default events create certain restrictions on the $JDDI(t)$ across time, for example, the probability of zero default in the portfolio has to be decreasing over time. If the irreversibility constraint is violated, the set of $JDDI(t)$ would not qualify for the distributions of default indicators, and it does not admit any valid $JDDT$. The conditional independence of the $\mathbf{1}_{\tau_i<t}$ and the monotonicity of the $p_i(x,t)$ ensures that the resulting set of the $JDDI(t)$ over time from Definition 3.1 is always consistent and conforming to the irreversibility constraint:

Proposition 3.2 *The copula functions of default indicators in Definition 3.1 uniquely determines a valid and consistent set of $JDDI(t)$ over time, which is fully consistent with the unconditional single name default probability term structures $p_i(t)$ and it admits infinitely many $JDDT$ for practical purposes*[2].

Because the individual name's default indicators $\mathbf{1}_{\tau_i<t}$ are independent conditioned on $\tilde{X}(t) = x$, the semi-analytical method can be applied to build the portfolio loss distribution at time t. In Definition 3.1, the $p_i(x,t)$ is only a function of the value $x = \tilde{X}(t)$ but not the path of the common factor $\tilde{X}(t)$ over time. In the stochastic recovery specification described in section 2, the cumulative mean and variance of the recovery rate, i.e. the $\mu_i(0, p_i(x,t))$ and $\sigma_i^2(0, p_i(x,t))$, are not functions of the $\tilde{X}(t)$ path either. Therefore the conditional portfolio loss distribution at time t can be well approximated by a normal distribution with the mean and variance in (9) and (10) regardless of the path of the common factor $\tilde{X}(t)$.

Property 3.3 *Given the conditional probability function $p_i(x,t)$ and $\mu(p,p), \sigma_i^2(p,p)$ of recovery rates, the conditional loss distribution at time t is only a function of the value $x = \tilde{X}(t)$, but not the path of the common factor $\tilde{X}(t)$ over time. Since the unconditional loss distribution is an integration of the conditional loss distribution over all the possible values of $x = \tilde{X}(t)$, the unconditional loss distribution is fully determined by the marginal distribution function $f(x,t)$; and it does not depend on any other property of the process $\tilde{X}(t)$.*

The Property 3.3 is very convenient because it simplifies the model calibration across multiple maturities. At a given maturity t, the model can be calibrated to the expected tranche loss (ETL) by changing the marginal distribution $f(x,t)$. The $p_i(x,t)$ has to be recalibrated using (13) when changing the $f(x,t)$ in order to maintain the consistency with the input single name default probability $p_i(t)$. The distribution $f(x,t)$ can be represented either as a non-parametric distribution or a parametric distribution such as a mixture of Gamma processes[3]. Numerical optimization routines work quite well for a one-dimensional distribution function, and it is not difficult to find a suitable marginal distribution $f(x,t)$ that matches the index ETL at time t using either the parametric or the non-parametric representation. A bootstrap procedure can be readily applied to calibrate the model sequentially across multiple maturities. The only constraints on the calibration to later maturities are the monotonic constraint from (14), (15) and (11), which are technical in nature and normally do not pose a serious limitation. Without the Property 3.3, the loss distribution would depend on the path of the $\tilde{X}(t)$, and the calibration to multiple maturities becomes difficult because the calibration procedure has to consider all the possible paths in $\tilde{X}(t)$ over time instead of just its value at t. A bootstrap calibration without the Property 3.3 is therefore very difficult to implement because the state space of all the possible paths of $\tilde{X}(t)$ often becomes too large for numerical optimization routines to be effective at later maturities.

Proposition 3.4 *With deterministic discount factors*[4]*, the $JDDI(t)$ and the stochastic recovery specification fully determine the value of instruments whose payoffs are not default path-dependent or spread-dependent, such as vanilla bespoke CDO, long short CDO, CDO^n, long short CDO^n or NTD basket. Therefore, no information beyond the $JDDI(t)$ and the recovery can be obtained by observing the market prices of these instruments.*

[2]There can be certain extreme situations where a set of $JDDI(t)$ admit only one valid $JDDT$, for example, when all names' default time are deterministic and are known at time zero. These extreme situations are clearly irrelevant in practice

[3]The Gamma process is convenient because it is increasing.

[4]It is possible to incorporate stochastic interest rates and correlate them with the $\tilde{X}(t)$ in this framework, but that is of very limited practical interests.

Figure 2: Progressive Calibration of the Model

Steps	Model Info.	Model Parameters	Market Input	Products Covered
1	$JDDI(t)$	$f(x,t)$	Single name CDS and index tranches, very liquid	Bespoke CDOs, NTD Basket, long/short CDO and CDO^n
2	$JDDT$	static Markov chain on $\tilde{X}(t)$	Some market observables on default path dependent instruments, illiquid	All default path-dependent instruments, such as waterfall synthetics, forward starting or step-up tranches, loss triggered LSS
3	$JDDT$ + systemic spread dynamics	Full dynamics of $\tilde{X}(t)$	Very few market observables on tranche options, almost no liquidity	Products that depend on systemic spread dynamics: such as senior tranche options, spread triggered LSS etc
4	$JDDT$ + systemic and idiosyncratic spread dynamics	Dynamics in $\tilde{X}(t)$ + idiosyncratic dynamics compatible with $p_i(x,t)$	Some market observables on single name swaption, some liquidity	Products that depend on both systemic and idiosyncratic spread dynamics, such as junior tranche options, etc.

The Proposition 3.4 suggests that it is a dangerous practice for any correlation model to determine the $JDDT$ purely from the calibration to the index tranche prices because the index tranche prices have no information beyond the $JDDI(t)$. All default time copulas have this problem because a default time copula by definition fully specifies the $JDDT$ from the index tranche calibration. Therefore, the $JDDT$ and the spread dynamics from default time copulas are not based on the relevant market information, and great caution is required when pricing default path-dependent instruments with the static copulas. Most bottom-up dynamic correlation models also have this problem of specifying the $JDDT$ prematurely from the index tranche calibration, e.g. (Mortensen, 2006), (Chapovsky *et al.*, 2006) and (Kogan, 2008). Therefore, their $JDDT$ and spread dynamics are severely limited by the index tranche calibration. The modelling framework proposed in this paper circumvents this limitation by only constructing the $JDDI(t)$ and $f(x,t)$ from the calibration to the index tranche and thus lends great freedom to the choice of $JDDT$ and spread dynamics.

After obtaining the $f(x,t)$ from the index tranche calibration, any $\tilde{X}(t)$ process can be used to construct an arbitrage free dynamic bottom-up model that reproduce the index tranche and underlying single name CDS prices as long as the following two conditions are met:

1. The $\tilde{X}(t)$ is non-negative and increasing.
2. The $\tilde{X}(t)$ reproduces the calibrated marginal distribution $f(x,t)$.

These requirements are exactly the same as those of a typical top-down model on the portfolio loss process, where the portfolio loss is non-negative and increasing, and the marginal distributions of the portfolio loss have to be preserved. Therefore, top-down models that were developed for the portfolio loss process can be applied to the $\tilde{X}(t)$ process, and producing a full bottom-up dynamic correlation model. In this approach, the $JDDT$ and the systemic spread dynamics can be freely adjusted by the top-down method without affecting the underlying single name and tranche prices. The top-down process on the $\tilde{X}(t)$ should be constructed using market observables that contain additional information beyond the $JDDI(t)$, e.g., forward-starting tranche or tranche options.

Figure 2 suggests a progressive calibration procedure within this modelling framework. In this progressive calibration procedure, each step adds only the necessary restrictions to the model in order to accommodate additional market information. The earlier steps do not limit the generality of the later steps; and the later steps always preserve all the model parameters and properties from the earlier steps. In step 2, the static Markov chain on $\tilde{X}(t)$ fully specifies the $JDDT$. In general, the $\tilde{X}(t)$ process is not Markovian. For example,

$\tilde{X}(t)$ could be a path-dependent function (e.g. integral) of an underlying process $y(t)$; in which case, the static Markov chain in step 2 only gives the transition probability of $\tilde{X}(t)$ based on time zero information; the step 3 then specifies the full dynamics of $(\tilde{X}(t), y(t))$, which determines the systemic spread dynamics in addition to the $JDDT$.

The copula functions in Definition 3.1 also admit the idiosyncratic spread dynamics that are independent of the systemic process $\tilde{X}(t)$. The $p_i(0,t)$ term structure defines the default probability from the idiosyncratic spread dynamics because there is no systemic default risk if $\tilde{X}(t)$ remains constant at 0. The idiosyncratic spread dynamics could have a strong impact on the pricing of certain exotic instruments, e.g., junior tranche options. The calibration of the idiosyncratic spread dynamics is also quite difficult because the CDS swaption is only liquid for very few names, and they have only very short maturities.

The progressive calibration procedure in Figure 2 is very attractive in practice because it allows instruments to be priced from the most necessary and the most reliable market information. For example, if the model is calibrated to the step 2 and is used to risk manage a book containing bespoke CDOs and loss triggered LSS, it is certain that the bespoke CDO prices are fully determined by the liquid index tranches and underlying CDS curves; and they are not affected by the views or observations on the forward losses which may be used to calibrate the step 2. Suppose there is new market information on the forward losses, then only step 2 of the model calibration needs to be updated, which affects only the pricing of the loss triggered LSS.

The step 3 and 4 of the calibration procedure are not yet feasible under the current market conditions because of the lack of reliable market observations. Therefore, exotic correlation instruments whose payoff depends on the spread dynamics, e.g. tranche options, cannot be exactly priced due to the incompleteness of the market. However, if the model can be calibrated to the step 2, the model implied $JDDT$ often imposes range bounds on the valuation of these spread dependent instruments. The range bounds of spread dependent instruments can be very useful in practice given the inability to obtain the exact pricing of these instruments.

4 Model Implementation

This section describes the details of a non-parametric implementation of this modelling framework, where the $p_i(x,t)$ function in Definition 3.1 is chosen to follow that of (Chapovsky *et al.*, 2006):

$$p_i(x,t) = 1 - c_i(t)e^{-\beta_i(t)x} \tag{18}$$

The $\beta_i(t) \geq 0$ is a loading factor on the systemic process. For simplicity, $\beta_i(t)$ is chosen so that the systemic process contributes a constant fraction to the cumulative hazard:

$$\log(\mathbb{E}[e^{-\beta_i(t)x}]) = \gamma_i \log(1 - p_i(t)) \tag{19}$$

The $\gamma_i \in [0,1]$ denotes the constant systemic fraction, which directly affects the correlation between individual name's spread movements. $1 - c_i(t)$ is the default probability from the idiosyncratic dynamics, which has to make up the rest of the cumulative hazard according to (13):

$$\log(c_i(t)) = (1 - \gamma_i)\log(1 - p_i(t)) \tag{20}$$

This $p_i(x,t)$ specification is convenient because (12) to (14) are automatically satisfied. (15) is satisfied as long as the $\beta_i(t)$ is increasing in t. A constant γ_i in (19) implies that the $\beta_i(t)$ is not guaranteed to be increasing for all possible $f(x,t)$. Therefore, the choice of either $f(x,t)$ or γ_i has to be constrained in order to maintain the monotonicity of the $\beta_i(t)$.

Consider two time periods $t_1 < t_2$ and suppose $f(x,t_1)$ and $\beta_i(t_1)$ are already calibrated to market prices at time t_1. With a constant γ_i, a $\beta_i(t_2) \geq \beta_i(t_1)$ can always be found when the $f(x,t_2)$ is very close to the $f(x,t_1)$ since in the limiting case of $f(x,t_2) = f(x,t_1)$, the $\beta_i(t_2)$ cannot be less than the $\beta_i(t_1)$ given the

Figure 3: Stochastic Recovery

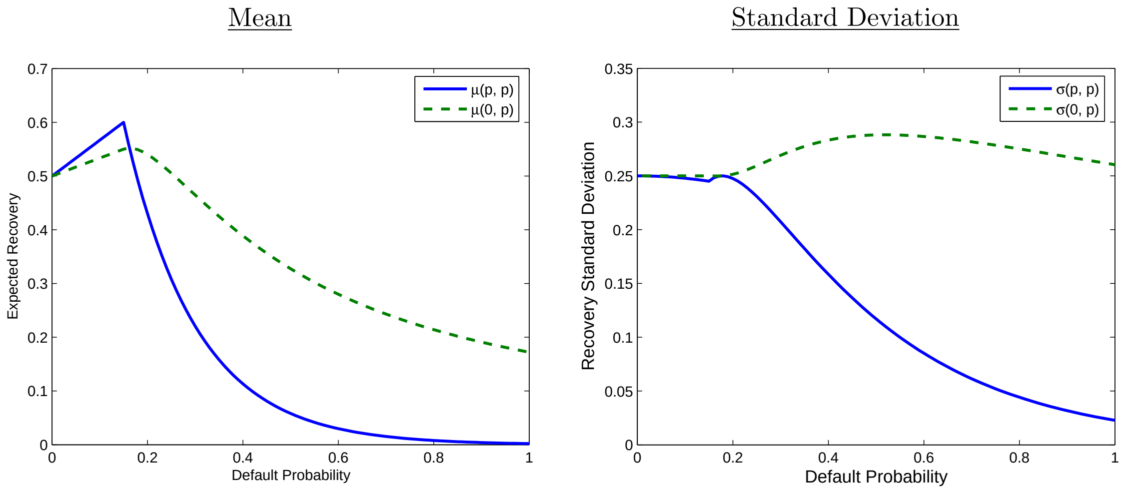

default probability in (19) is increasing: $p_i(t_2) \geq p_i(t_1)$. Therefore, the monotonicity of $\beta_i(t)$ can always be enforced by making the $f(x, t_2)$ close to the $f(x, t_1)$.

In a diverse portfolio, the distressed names usually impose more constraints on the choice of $f(x, t_2)$ since their default risk are concentrated in the front end before time t_1, and their $p_i(t_2)$ can be very close to $p_i(t_1)$. A constant γ_i may force $f(x, t_2)$ to be very close to $f(x, t_1)$ in order to satisfy the monotonicity constraint of $\beta_i(t)$ for the most distressed names in the portfolio, which could undermine the model's ability to calibrate to the index tranches. Therefore for distressed names, it is better to have a time dependent $\gamma_i(t)$ which starts with a low value and increases over time, thus leaving more freedom in the choice of $f(x, t_2)$. It also makes economic sense for very distressed names to have lower systemic dependencies in the short time horizon.

The γ_i factors have to be high ($> 80\%$) for the majority of the names in order to obtain good calibration to the index tranches, which suggests that the main risk factor in current market is the systemic risk. For simplicity, γ_i is chosen to be 90% for all names except for very distressed names in this implementation.

As discussed in section 2, only the $\mu_i(p, p)$ and $\sigma_i^2(p, p)$ of the recovery rate need to be specified in order to price CDO tranches consistently. For simplicity, all credits are assumed to have the same functional form of $\mu(p, p)$ and $\sigma^2(p, p)$. Figure 3 shows the mean and standard deviation of the recovery function used in the non-parametric model implementation. The choice of $\mu(p, p)$ function is somewhat arbitrary, its overall trend is chosen to be decreasing in p because it is desirable for the recovery to be lower in the bad states of the market factor. A peak is created in $\mu(p, p)$ at 15% default probability just to show the ability to create an arbitrary shape of the recovery term structure. The $\sigma^2(p, p)$ is assumed to be a fixed percentage of the maximum variance for the given $\mu(p, p)$: $\sigma^2(p, p) = \alpha\mu(p, p)(1 - \mu(p, p))$, where the α is chosen to be 25% somewhat arbitrarily. If there are observations or views about the variance of a name's recovery rate, the α parameter can be changed to match those.

The $\mu(p, p)$ function in Figure 3 is multiplied by a name specific scaling factor to match the individual credits' CDS curve recovery at the 5Y tenor. Since the μ, σ^2 are functions of the conditional default probability, the unconditional cumulative recovery rate at time t for name i can be computed by integrating over all the possible market factor values:

$$R_i(0, t) = \frac{1}{p_i(t)} \int \mu_i(0, p_i(x, t)) p_i(x, t) f(x, t) dx \tag{21}$$

Even though $\mu(0, p)$ has a strong trend over p as shown in the Figure 3, the unconditional recovery rate

Figure 4: Model Calibration to CDX-IG9 on Jan. 15, 2009

Market Input ETL

Att	Det	5Y	7Y	10Y
0.0%	2.6%	83.51%	87.23%	91.12%
2.6%	6.7%	57.22%	64.36%	71.28%
6.7%	9.8%	30.05%	41.47%	54.94%
9.8%	14.9%	18.02%	26.07%	36.49%
14.9%	30.3%	4.87%	7.20%	10.57%
30.3%	61.0%	4.05%	6.24%	8.54%
0.0%	100.0%	8.72%	10.96%	13.47%

ETL from Model Calibration

Att	Det	5Y	7Y	10Y
0.0%	2.6%	83.54%	87.11%	90.62%
2.6%	6.7%	57.27%	64.10%	70.43%
6.7%	9.8%	30.16%	41.42%	54.52%
9.8%	14.9%	18.03%	25.83%	35.57%
14.9%	30.3%	5.02%	7.73%	11.01%
30.3%	61.0%	3.98%	5.74%	7.42%
0.0%	100.0%	8.72%	10.96%	13.47%

Figure 5: Calibrated $F(x,t)$

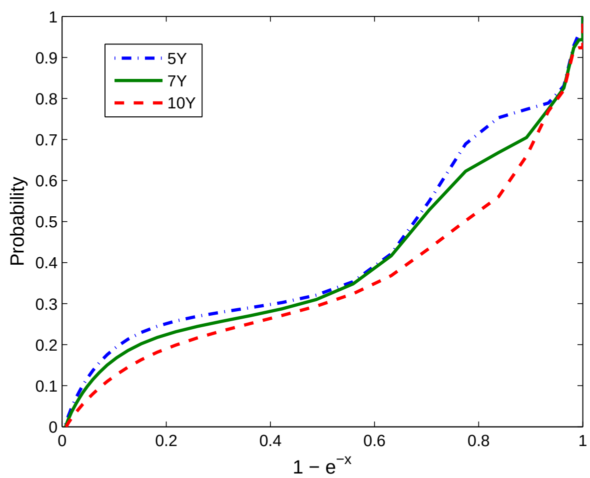

$R_i(0,t)$ would exhibit a much milder trend over time due to the averaging effects through the integral in (21). More results about the unconditional recovery rates are shown in the following sections.

The $p_i(x,t)$, $\mu(p,p)$ and $\sigma^2(p,p)$ given in this section are just one example of possible model specifications. There could be many different specifications which are equally valid and effective under the general principles of Definition 3.1.

5 Model Results

In this section, some numerical results are presented from the non-parametric implementation described in section 4.

5.1 Calibration to Index Tranches

The non-parametric implementation of this model is calibrated to the expected tranche loss (ETL) of CDX-IG9 index as of the close of Jan. 15, 2009. Figure 4 shows the input ETL from the tranche market and the model calibration results. The model calibration is quite close to the input ETL across the three maturities. Figure 5 showed the calibrated cumulative distribution function $F(x,t)$ at 5Y, 7Y and 10Y. The constraint

Figure 6: Expected Recovery Change

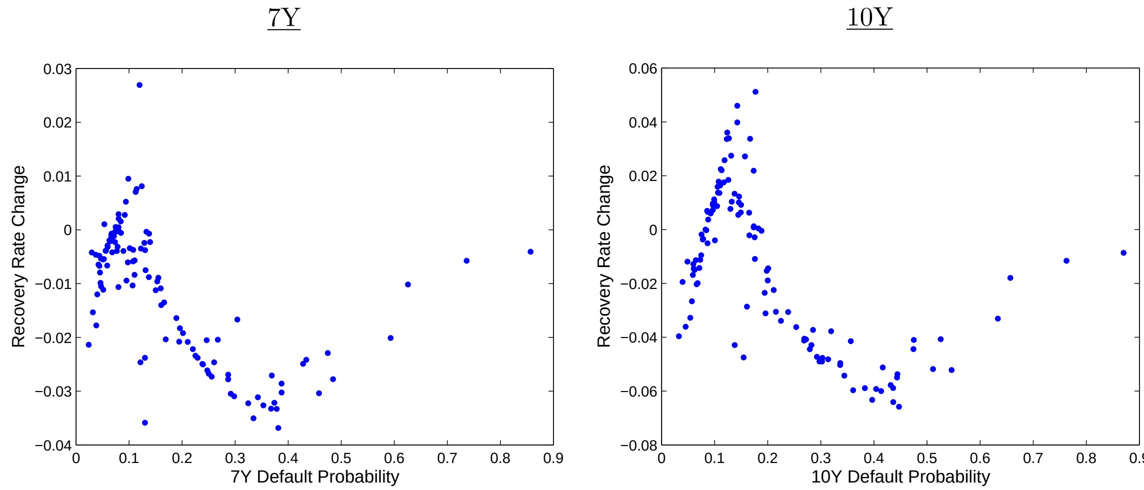

(11) is built into the bootstrap process so that the resulting marginal distributions are compatible with a increasing process. It is clear that the calibrated $F(x,t)$ does satisfy (11) since the three CDF curves never cross each other.

Since the model expected recovery matches only the CDS curve recovery at the 5Y tenor, the single name default probabilities at the 7Y and 10Y tenors are adjusted to preserve the expected loss of the individual CDS curve. The calibration results showed that the total portfolio losses of the 0-100% tranche are exactly preserved at all the maturities.

In this example, the calibration inputs are the ETL at the 5Y, 7Y and 10Y maturities. It is easy to extend the calibration to all other quarterly dates if their ETLs are available. However in reality, the only market observables are tranche prices at 5Y, 7Y and 10Y, the ETLs are not directly observable. This issue can be addressed with one of the following two methods:

1. use another model, e.g., base correlation, to extract the ETL surface at each quarterly date and calibrate the model to the full ETL surface. This ensures the maximum consistency to the existing base correlation framework. This method works better if the $f(x,t)$ is specified as a non-parametric distribution.
2. create an interpolation method on the distributions of the $f(x,t)$ so that the distributions at all quarterly dates can be interpolated from the distributions at the market maturities of 5Y, 7Y and 10Y. This allows the model to be directly calibrated to the market tranche prices instead of the ETLs. This approach is most suited if the $f(x,t)$ is specified as a parametric distribution.

Method 1 is ideal if it is required to maintain maximum consistency with the existing base correlation model. Method 2 is preferred if it is desirable to preserve the possibility of pricing a bespoke CDO from the liquid index tranche prices.

5.2 Implied Recovery Rate Term Structure

The calibrated model matches the CDS curve recovery exactly at the 5Y tenor. The expected recoveries from the model are different from the CDS curve recoveries at other tenors because the $\mu(p,p)$ function is not a constant. The default probabilities at 7Y and 10Y are adjusted according to the model implied recovery rate so that the expected losses of individual curves are preserved. Figure 6 showed the scatter plots of the difference between the 7Y and 10Y model implied recoveries from their CDS curve recoveries for all the 122

Figure 7: Monte Carlo Simulation of Tranche Loss

Co-monotonic Markov Chain

Att	**Det**	**5Y**	**7Y**	**10Y**
0.0%	2.6%	83.59%	87.15%	90.64%
2.6%	6.7%	57.18%	64.10%	70.46%
6.7%	9.8%	30.12%	41.35%	54.48%
9.8%	14.9%	18.02%	25.81%	35.49%
14.9%	30.3%	5.02%	7.73%	11.01%
30.3%	61.0%	3.97%	5.73%	7.40%
0.0%	100.0%	8.71%	10.95%	13.46%

Maximum Entropy Markov Chain

Att	**Det**	**5Y**	**7Y**	**10Y**
0.0%	2.6%	83.54%	87.08%	90.62%
2.6%	6.7%	57.18%	64.08%	70.46%
6.7%	9.8%	30.16%	41.36%	54.47%
9.8%	14.9%	18.00%	25.84%	35.52%
14.9%	30.3%	4.97%	7.70%	11.00%
30.3%	61.0%	3.93%	5.72%	7.39%
0.0%	100.0%	8.68%	10.94%	13.45%

names in the CDX-IG9 portfolio. The horizontal axis is the default probability at the corresponding tenor.

As shown in Figure 6, the model expected recoveries at 7Y and 10Y do not deviate too much from the CDS curve recoveries. The scatter plots in Figure 6 roughly follow the shape of the $\mu(0, p)$ in Figure 3, but at a much milder pace because of the averaging effects in the (3). The slight difference between the model expected recoveries and the CDS curve recoveries should not be a problem in practice as long as individual name's expected losses are preserved, since there is very little reliable market information on the term structure of recoveries.

5.3 Monte Carlo Simulation

A simple Monte Carlo simulator is implemented in this section to verify the consistency and correctness of the proposed modelling framework. In order to drive the default time simulation, the model has to be calibrated at least to the step 2 in Figure 2, i.e., a static Markov chain has to be built from the calibrated marginal distribution $f(x, t)$. Two different methods of building the static Markov chain on $\tilde{X}(t)$ are implemented: co-monotonic and maximum entropy. A detailed description of these two methods can be found in (Epple *et al.*, 2006), where both of these methods were applied to the portfolio loss process following the typical top-down approach. The $f(x, t)$ distributions are taken from the calibration in section 5.1.

Figure 7 showed the simulated ETLs at the three maturities from drawing 1,000,000 independent default time and recovery paths from both of the Markov chains. The default time and recovery paths are drawn using the following steps:

1. Draw a full path of $\tilde{X}(t)$ over time from the Markov chain.
2. Use the $p_i(x, t)$ function to compute the conditional default probability term structures of all the underlying names for the given path of $\tilde{X}(t)$.
3. For each name, draw an independent uniform random number d_i which represents the conditional default probability. d_i is then used to determine the default period of the corresponding name according to the conditional default probability term structure.
4. For each name defaulted before the final maturity (10Y), compute its instantaneous recovery mean and variance $\mu_i(d_i, d_i)$, $\sigma_i^2(d_i, d_i)$.
5. Draw an independent recovery rate for any defaulted name from a two point distribution whose mean and variance are given by the $\mu_i(d_i, d_i)$, $\sigma_i^2(d_i, d_i)$.

After drawing the default time and recovery path, the tranche losses at all tenors are computed from the same default time and recovery path to ensure full consistency across all maturities. Then the tranche losses from these independent default time and recovery paths are averaged to produce the ETL.

The simulated ETLs from the two Markov chains are very close to each other, which is expected since they have identical $JDDI(t)$ by construction. Both of the simulated ETLs are very close to the semi-analytical

Figure 8: Temporal Correlation of Incremental Portfolio Losses

Co-monotonic Markov Chain

-	0-5Y	5Y-7Y	7Y-10Y
0-5Y	1	.5027	.4887
5-7Y	.5027	1	.2109
7-10Y	.4887	.2109	1

Maximum Entropy Markov Chain

-	0-5Y	5Y-7Y	7Y-10Y
0-5Y	1	.4199	.3936
5-7Y	.4199	1	.1227
7-10Y	.3936	.1227	1

Figure 9: Simulated Single Name Expected Loss

Co-monotonic Markov Chain

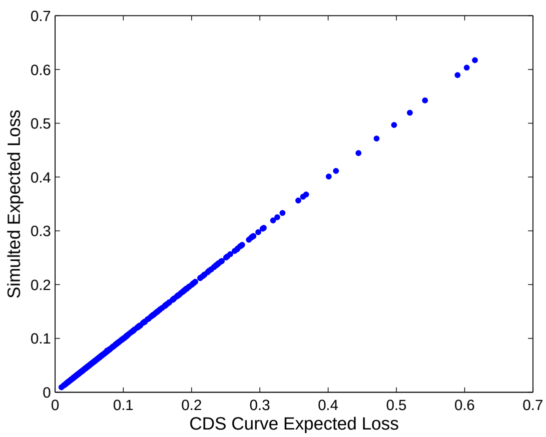

Maximum Entropy Markov Chain

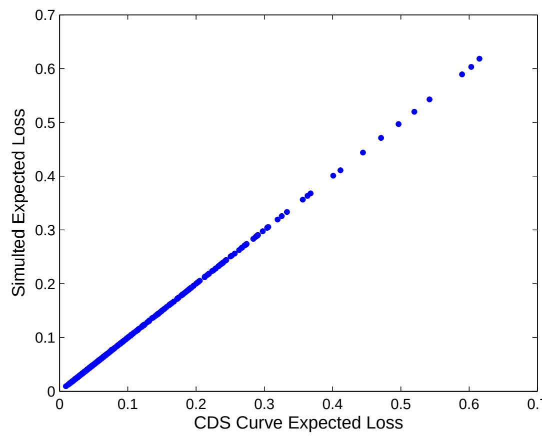

calibration results shown in Figure 4, where the normal approximation is used to build the conditional loss distribution. The maximum difference in the ETL between the Monte Carlo simulation and the semi-analytical pricing with normal approximation is less than 0.1%. The ETL difference of this magnitude is clearly negligible for practical purposes. It is also verified that a different instantaneous recovery rate distribution, such as the beta distribution, produces very similar results to those in Figure 7, as long as the $\mu(p, p)$ and $\sigma^2(p, p)$ of the recovery rate are matched.

However, the two Markov chains lead to very different $JDDT$s. Figure 8 showed the correlation matrix between the simulated incremental portfolio losses in the three periods (0-5Y, 5Y-7Y and 7Y-10Y), conditioned on the portfolio loss before 5Y is less than 10%. It is evident that the temporal loss correlation from the co-monotonic Markov chain is much stronger than that of the maximum entropy Markov chain. The temporal loss correlation is a critical factor in pricing exotic correlation instruments such as forward-starting tranche and loss-triggered LSS.

Figure 9 showed two scatter plots of the simulated vs. the CDS curve expected losses for all 122 underlying names over the three maturities. All the dots in Figure 9 are perfectly aligned along the diagonal line in this scatter plot, which showed that the Monte Carlo simulation correctly preserves all the single names' expected losses across all three maturities. Because the recovery rates can be different from the CDS curve recovery at tenors other than 5Y, the default probabilities are not preserved at 7Y and 10Y. However, keeping the expected losses consistent with the underlying CDS curves is often good enough in practice since the PV01 change caused by the recovery rate shift is usually small. This simulation exercise confirms the following claims made in this paper:

1. The stochastic recovery and dynamic correlation modelling framework proposed in this paper is fully

consistent and arbitrage free[5].

2. A full bottom-up dynamic correlation model can be obtained by applying a top-down method to the $\tilde{X}(t)$ process.
3. Different ways of building the $\tilde{X}(t)$ process can lead to very different $JDDT$s and temporal loss correlations even if their $JDDI(t)$s are identical over time.
4. Single name expected loss term structures are correctly preserved in this modelling framework.
5. The details of the recovery distribution do not affect CDO pricing as long as their first two moments are preserved.
6. The normal approximation to the conditional loss distribution is very accurate under the current market conditions.

6 Conclusion

This paper described a tractable and consistent stochastic recovery specification, and a very generic dynamic correlation modelling framework that could combine the best features of the top-down and bottom-up approaches. The modelling framework is equipped with the important Property 3.3, which allows easy calibration to the index tranche prices across multiple maturities. Close calibration to the index tranches across multiple maturities in a consistent model is a very difficult problem in its own right.

This modelling framework also circumvents a serious problem found in all default time copulas and a number of existing bottom-up dynamic spread models, where the $JDDT$ is determined prematurely from the calibration to the index tranche prices which contain only information on $JDDI(t)$. The proposed modelling framework specifies only the $JDDI(t)$ from the calibration to the index tranches; thus allowing all the compatible $JDDT$ to be considered in the later calibration steps in Figure 2. Therefore, the model can produce a very rich set of $JDDT$s and spread dynamics by applying a top-down method on the common factor process $\tilde{X}(t)$. Changing the dynamics of $\tilde{X}(t)$ does not change the calibrated CDO tranche prices as long as the marginal distribution $f(x,t)$ remains intact.

The affine jump diffusion (AJD) process is a popular choice in building the bottom-up dynamic correlation models. Models based on AJD processes also suffer the same problem of determining $JDDT$ prematurely from the index tranche calibration, therefore their $JDDT$ and spread dynamics are not generic enough. For example, the jump in the AJD intensity process is usually modelled as an independent Poisson jump process with a deterministic hazard rate and random jump sizes for tractability, as in (Chapovsky *et al.*, 2006). Under an AJD model, the senior tranches only suffer loss if a large jump in the intensity arrives. Since the jump arrivals in a Poisson process are totally unexpected and stateless, the probability of large jump arrivals does not change with the filtration of the common process and default events. Therefore the senior tranche's expected loss and spread often exhibit very low volatility in an AJD dynamic model. In the proposed modelling framework, the $\tilde{X}(t)$ process can have large jumps, and the probability of large jump arrivals can be a function of the states of the common process, therefore this modelling framework can produce very high senior tranche volatility as observed in the current market.

(Giesecke & Goldberg, 2005) and (Halperin & Tomecek, 2008) suggested random thinning as a possible method to incorporate the single name dynamics in a top-down model. In the random thinning approach, the dynamics of the portfolio loss process are determined first, then the random thinning technique is applied to retrofit the single name identity and spread dynamics to the portfolio loss process. The random thinning technique often requires the Monte Carlo simulation, which is numerically expensive. One notable drawback with the random thinning approach is that the portfolio loss process has to be fully specified before the random thinning. However, under the current market conditions, there is very little reliable observations to calibrate the dynamics of the portfolio loss process. Therefore, the portfolio loss dynamics

[5]In theory, the difference between the model expected recovery and the underlying CDS curve recovery at tenors other than 5Y is an arbitrage violation, but that is almost impossible to exploit in practice. The $\mu(p,p)$ can be set to be a constant if matching the constant term structure of the curve recovery is a requirement. The $\sigma(p,p)$ has to increase over p to push the risk to the senior most tranche if the $\mu(p,p)$ is kept constant.

are often determined by ad hoc model assumptions or unreliable market observations or views. With random thinning, even the most basic multi-name instruments, such as vanilla bespoke CDO or NTD basket, can be numerically expensive to price and risk manage via a mapping method to the index tranches; the resulting prices and single name risks are also susceptible to the ad hoc model assumptions or unreliable market inputs on the portfolio loss dynamics.

In the modelling framework proposed in this paper, only the most reliable market information like the single name CDS spreads and the index tranche prices are used to calibrate the model in the first step in Figure 2. The valuation and single name risks of vanilla instruments such as bespoke CDOs and NTD baskets do not depend on any model assumptions, observations or views about the portfolio loss dynamics. Also, the semi-analytical method is supported because of the conditional independence in Definition 3.1. Therefore the pricing and single name risk of vanilla instruments are very straight forward and numerically efficient in the proposed modelling framework.

Another notable advantage of this modelling framework is its ease of implementation. There are no requirements for complicated numerical techniques other than the simple one-dimensional optimization of the marginal distribution of $f(x,t)$, where standard optimization techniques work very well. After finding $f(x,t)$, existing top-down methodologies can be applied to build the $JDDT$ and spread dynamics. In comparison, existing bottom-up models based on AJD often require solving PDEs with finite difference methods.

The proposed modelling framework is also very flexible, the functional form of the $\mu_i(p,p)$ and $\sigma_i^2(p,p)$ can be changed to match the observations or views on the underlying credits' recovery distributions. The copula functions on the default indicators are also very flexible because any function $p_i(x,t)$ can be used as long as the technical constraints in (14), (15) are met. The model can also be implemented either parametrically or non-parametrically. A non-parametric approach makes it easy to reproduce the full ETL surface from a base correlation model, thus keeping the exotic correlation products consistent with the vanilla CDOs which are normally managed under the base correlation modelling framework. On the other hand, a parametric approach makes it possible to price a bespoke CDO from the index tranches, and easier to compute risk measures to the model parameters.

The specification and calibration of the full spread dynamics are not attempted in this paper due to the lack of reliable market observables on spread dynamics to calibrate the model beyond the initial steps in Figure 2. However, just calibrating the model to initial steps would be very useful in practice since reliable range bounds in valuation can often be derived for spread-dependent exotic instruments once the $JDDT$ is specified. This could be a practical approach to manage very exotic correlation instruments in an incomplete market.

References

Amraoui, S., & Hitier, S. 2008. Optimal stochastic recovery for base correlation. *defaultrisk.com*, Jun.

Andersen, L., & Sidenius, J. 2004. Extensions to the gaussian copula: Random recovery and random factor loadings. *Journal of credit risk*, Jun.

Andersen, L., Sidenius, J., & Basu, S. 2003. All your hedges in one basket. *Risk*, Nov.

Arnsdorf, M., & Halperin, I. 2007. Bslp: Markovian bivariate spread-loss model for portfolio credit derivatives. *defaultrisk.com*, Mar.

Bennani, N. 2005. The forward loss model: A dynamic term structure approach for the pricing of portfolio credit derivatives. *defaultrisk.com*, Nov.

Chapovsky, A., Rennie, A., & Tavares, P. 2006. Stochastic intensity modelling for structured credit exotics. *defaultrisk.com*, Oct.

Epple, F., Morgan, S., & Schloegl, L. 2006. Joint distribution of portfolio losses and exotic portfolio products. *Lehman brothers: Quantitative credit research quarterly*, Oct, 44–55.

Errais, E., Giesecke, K., & Goldberg, L. 2009. Affine point processes and portfolio credit risk. *defaultrisk.com*, Mar.

Giesecke, K., & Goldberg, L. 2005. A top-down approach to multi-name credit. *defaultrisk.com*, Feb.

Halperin, I., & Tomecek, P. 2008. Climbing down from the top: Single name dynamics in credit top down models. *defaultrisk.com*, Oct.

Hull, J., & White, A. 2006. Valuing credit derivatives using an impplied copula approach. *Journal of derivatives*, Nov.

Kogan, L. 2008. A dynamic default correlation model. *Lehman brothers: Quantitative credit research quarterly*, Jul.

Krekel, M. 2008. Pricing distressed cdos with base correlation and stochastic recovery. *defaultrisk.com*, May.

Mortensen, A. 2006. Semi-analytical valuation of basket credit derivatives in intensity-based models. *defaultrisk.com*, Jan.

O'Kane, D., & Livesey, M. 2004. Base correlation explained. *Lehman brothers: Quantitative credit research quarterly*, Nov, 3–20.

Shelton, D. 2004. Back to normal. *Citigroup global structured credit research.*

Shonbucher, P. 2006. Portfolio losses and the term structure of loss transition rates: A new methodology for the pricing of portfolio derivatives. *defaultrisk.com*, Jun.

Sidenius, J., Piterbarg, V., & Andersen, L. 2006. A new framework for dynamic credit portfolio loss modeling. *defaultrisk.com*, Jun.

Skarke, H. 2005. Remarks on pricing correlation products. *defaultrisk.com*, Jul.