DYNAMIC MODELS OF PORTFOLIO CREDIT RISK: A SIMPLIFIED APPROACH

See: Journal of Derivatives, 15, 4 (Summer 2008), 9-28

John Hull and Alan White*

Joseph L. Rotman School of Management
University of Toronto

First Draft: December 2006
This Version: April 2008

ABSTRACT

We propose a simple dynamic model that is an attractive alternative to the (static) Gaussian copula model. The model assumes that the hazard rate of a company has a deterministic drift with periodic impulses. The impulse size plays a similar role to default correlation in the Gaussian copula model. The model is analytically tractable and can be represented as a binomial tree. It can be calibrated so that it exactly matches the term structure of CDS spreads and provides a good fit to CDO quotes of all maturities. Empirical research shows that as the default environment worsens default correlation increases. Consistent with this research we find that in order to fit market data it is necessary to assume that as the default environment worsens impulse size increases. We present both a homogeneous and heterogeneous version of the model and provide results on the use of the calibrated model to value forward CDOs, CDO options, and leveraged super senior transactions.

*We are grateful to Moody's Investors Services for providing financial support for this research. We are grateful to participants at the New York University Derivatives 2007 conference, the ICBI Global Derivatives Conference in May 2007, the Moody's/Copenhagen Business School conference in May 2007, and the University of Chicago conference on credit risk in October 2007. We would particularly like to thank Claudio Albanese, Leif Andersen, Peter Carr, Rama Cont, Pierre Collin-Dufresne, Darrell Duffie, Ruediger Frey, Michael Gordy, Igor Halperin, Jon Gregory, David Lando, Allan Mortensen, and Philipp Schönbucher for useful comments and suggestions. We would also like to thank the editor, Stephen Figlewski, for extremely perceptive and helpful comments.

Dynamic Models of Portfolio Credit Risk: A Simplified Approach

I. Introduction

The credit derivatives market has experienced meteoric growth since 1998. The most popular instruments are credit default swaps. These provide a payoff when a particular company defaults. However, in recent years portfolio credit derivatives have been attracting a lot of attention. These provide protection against the defaults experienced by a portfolio of companies. Statistics published by the Bank for International Settlements show that the outstanding notional principal for portfolio credit derivatives has grown from about $1.3 trillion in December 2004 (20% of the notional principal for all credit derivatives) to about $10.0 trillion in December 2006 (35% of the notional principal for all credit derivatives).

The most popular portfolio credit derivative is a collateralized debt obligation (CDO). In this a portfolio of obligors is defined and a number of tranches are specified. Each tranche is responsible for losses between U_1 % and U_2% of the total principal for some U_1 and U_2. As the market has developed standard portfolios and standard tranches have been specified to facilitate trading. One example is the iTraxx portfolio. This is a portfolio of 125 investment grade European companies with the notional principal (size of the credit exposure) being the same for each company. The equity tranche is responsible for losses in the range 0 to 3% of the total notional principal. The mezzanine tranche is responsible for losses in the range 3 to 6% of the total notional principal. Other tranches are responsible for losses in the ranges 6 to 9%, 9 to 12%, 12 to 22%, and 22 to 100% of total principal. The buyer of protection pays a predetermined annual premium (known as a spread) on the outstanding tranche principal and is compensated for losses that are in the relevant range. (In the case of the equity tranche the arrangement is slightly different: the buyer of protection pays a certain percentage of the tranche principal upfront and then 500 basis points on the outstanding tranche principal per year.)

Several other standard portfolios and associated tranches have been defined. For example, CDX NA IG is a portfolio of 125 investment grade North American credit exposures. (The tranches for this portfolio are 0 to 3%, 3 to 7%, 7 to 10%, 10 to 15%, 15 to 30%, and 30 to 100%) The most popular life of a CDO is five years. However, 7-year, 10-year, and to a lesser extent 3-year CDOs now trade fairly actively.

Default correlation is critical to the valuation of portfolio credit derivatives. Moody's statistics show that between 1970 and 2005 the default rate per year ranged from a low of 0.09% in 1979 to a high of 3.81% in 2001. This tendency of defaults to cluster has been studied by a number of researchers. One possible explanation is that default rates of all companies are influenced by one or more macroeconomic factors. Another is that defaults are "contagious" in the sense that a default by one company may induce other corporate failures. Das et al (2007) argue that contagion accounts for some part of the default clustering that is observed in practice.

The standard market model for valuing portfolio credit derivatives assumes a simple one-factor model for a company's time to default. This is referred to as the Gaussian copula model. Its origins can be found in Vasicek (1987), Li (2000), and Laurent and Gregory (2005). The Gaussian copula model is a static model. A single normally distributed variable determines the default environment for the whole life of the model. When the variable has a low value, the probability of each company defaulting during the life of the model is relatively high; when it has a high value, the probability of each company defaulting is relatively low. The model does not describe how the default environment evolves. Many alternatives to the Gaussian copula such as the *t*-copula, the double-*t* copula, the Clayton copula, the Archimedian copula, the Marshall Olkin copula, and the implied copula have been suggested. In some cases these models provide a much better fit to market data than the Gaussian copula model, but they are still static models.

The availability of CDO data for multiple time horizons presents researchers with an interesting and important challenge. This is to develop a dynamic model that fits market data and tracks the evolution of the credit risk of a portfolio. Dynamic models are important for the valuation of some structures. For example, options on tranches of CDOs cannot be valued in a satisfactory way without a dynamic model.

There are two types of dynamic credit risk models, which we will refer to as "specific" and "general" models. In a specific model the company or companies being modeled remain the same through time. In a general model they do not remain the same, but are defined to have certain properties. A model of the evolution of the credit spread for a particular company or the evolution of losses on a particular portfolio is a specific model. A model of the evolution of the average credit spread for A-rated companies or of the mezzanine spread for CDX NA IG is a general model. This paper is concerned with the development of a specific dynamic model for portfolios.

Extensions of the Merton (1974) structural model provide one approach for developing a specific dynamic model. Correlated processes for the values of the assets of the underlying companies are specified and a company defaults when the value of its assets reaches a barrier. The most basic version of the structural model is very similar to the

Gaussian copula model. Extensions of the basic model have been proposed by Albanese et al (2005), Baxter (2006), and Hull et al (2005). Structural models have the advantage that they have sound economic underpinnings. Their main disadvantage is that they are difficult to calibrate to market prices and usually have to be implemented with Monte Carlo simulation.

Reduced-form models provide an alternative to structural models. The most natural reduced-form approach to developing a dynamic model is to specify correlated diffusion processes for the hazard rates of the underlying companies. Our own experience and that of other researchers is that it is not possible to fit market data with this type of model. This is because there is a limit to how high the correlation between times to default can become. This has led researchers to include jumps in the processes for hazard rates. Duffie and Gârleanu (2001) for example assume that the hazard rate of a company is the sum of an idiosyncratic component, a component common to all companies, and a component common to all companies in the same sector. Each component follows a process with both a diffusion and a jump component. Other reduced form approaches are provided by Chapovsky et al (2006), Graziano and Rogers (2005), Hurd and Kuznetsov (2005), and Joshi and Stacey (2006).

Another approach to developing dynamic models involves the development of a model for the evolution of the losses on a portfolio. This is sometimes referred to as the "top down" approach. The behavior of individual companies in the portfolio is not considered. Sidenius et al (2004) use concepts from the Heath, Jarrow, and Morton (1992) interest-rate model to suggest a complex general no-arbitrage approach to modeling the probability that the loss at a future time will be less than some level.

Bennani (2005) proposes a model of the instantaneous loss as a percentage of the remaining principal. Schönbucher (2005) models the evolution of the loss distribution as a Markov chain. Errais et al (2006) suggest a model where the arrival rate of defaults experiences a jump when a default happens. In Longstaff and Rajan (2006) the loss follows a jump process where there are three types of jumps: firm specific, industry, and economy-wide. Putyatin et al (2005) suggest a model where the mechanism generating multiple defaults resembles the kinetics of certain chemical reactions. Walker (2007) uses a dynamic discrete-time multi-step Markov loss model.

Our objective in this paper is to develop a model that is easy to implement and easy to calibrate to market data. The model is developed as a reduced-form model, but can also be formulated as a top-down model. Under the model the hazard rate for a company follows a deterministic process that is subject to periodic impulses. This leads to a jump process for the cumulative hazard rate (or equivalently for the logarithm of the survival probability). CDOs, forward CDOs, and options on CDOs can easily be valued analytically using the model. For other instruments a binomial representation of the model can be used. We propose a calibration method where credit default swap spreads are exactly matched and CDO tranche quotes are matched as closely as possible.

The rest of the paper is organized as follows. Section II describes the model. Section III gives a number of progressively more complicated examples of the model and shows how they can be fitted to the market. In Section IV the model is applied to the valuation of a three different types of securities. Section V discusses extensions of the basic model. Conclusions are in Section VI.

II. THE MODEL

For any particular realization of the hazard rate, $h(t)$, between time zero and time t we can define

$$S(t) = e^{-\int_0^t h(\tau)d\tau}$$

The variable $S(t)$ is the cumulative probability of survival by time t conditional on a particular hazard rate path between time 0 and time t. A common approach to building a reduced form model is to define a process for $h(t)$ for each company over the life of the model. We choose instead to define a process for $S(t)$ for each company. The process for $S(t)$ provides the same information as the process for $h(t)$. However, it is much easier to work with because it leads to a much more straightforward way of valuing CDOs. (As we will see the process for $h(t)$ that corresponds to the process for $S(t)$ that we assume is one that has a deterministic drift and periodic impulses.) The default probability between time zero and time t as seen at time zero is the expected value of $1 - S(t)$.

It is important to specify what is known in this model. The underlying state variable is $S(t)$. While $S(t)$ is not directly observable it may be inferred from the prices of credit sensitive contracts. We also know how many defaults have occurred by time t. In developing and applying the model we will usually assume that at any given time, t, we know both $S(t)$ and the total number of defaults up to time t. This defines the filtration.[1]

Note that we do not assume that we know which particular companies have defaulted. In the homogeneous version of the model this additional information would of course be

---

[1] This approach to separating default probabilities and default events is considered in Ehlers and Schönbucher (2006) and is valid in our set up.

irrelevant. Unfortunately, the heterogeneous version of the model becomes computationally intractable if an attempt is made to incorporate this additional information.

The model assumes that most of the time the default probabilities of companies are predictable and defaults are independent of one another. Periodically there are economy-wide shocks to the default environment. When a shock occurs each company has a non-zero probability of default. As a result there are liable to be one or more defaults at that time. It is these shocks and their size that create the default correlation.

The model is of course a simplification of reality. In practice shocks to the credit environment do not cause several companies to default at exactly the same time. The defaults arising from the shocks are usually spread over several months. However, the model's assumption is reasonable because it is the total number of defaults rather than their precise timing that is important in the valuation of most portfolio credit derivatives. Another simplification is that shocks to the credit environment affect all companies. In practice they are liable to affect just a subset of companies in the portfolio. However, as an approximation we can think of a shock affecting a subset of companies as being equivalent to a smaller shock affecting all companies as far as its effect on the number of defaults is concerned.

For ease of exposition in explaining the model we assume homogeneity so that all companies have the same default probabilities. We also assume that the recovery rate is constant. Later we explain how these assumptions can be relaxed.

Consider a portfolio of obligors with total notional principal L. The protection seller for a tranche of a CDO provides protection against losses on the portfolio that are in the range

$a_L L$ to $a_H L$ for the life of the instrument. The parameter a_L is known as the attachment point and the parameter a_H is known as the detachment point. The protection buyer pays a certain number of basis points on the outstanding notional principal of the tranche. This principal equals $a_H L - a_L L$ initially and declines as losses in the range a_L to a_H are experienced.

As explained in Hull and White (2006) the key to valuing a CDO tranche at time zero is the calculation of the expected tranche principal on payment dates. The expected payment by the buyer of protection on the CDO tranche on a payment date equals the expected tranche principal on the payment date multiplied by the spread. The expected payoff by the seller of protection between two payments dates equals the reduction in the expected tranche principal between those dates.[2] The expected accrual payments required in the event of a default between two payment dates can be calculated from the reduction in the expected tranche principal between the dates and an assumption about when the reduction occurs. As we will see the expected tranche principal can be easily calculated from $S(t)$. This is why a model of the behavior of $S(t)$ is much easier to work with than a model of the behavior of $h(t)$.

The properties of $S(t)$ are that $S(0) = 1$, S is non-increasing in time, and $S(t) \geq 0$ for all t. A convenient transformation of S is to define a variable $X = -\ln(S)$ that has the properties

[2] Our terminology reflects the viewpoint of the buyer of protection. We will refer to the payments made by the buyer of protection as 'payments', and the payment in the event of default by the seller of protection as a 'payoff'.

that $X(0)=0$ and X is non-decreasing in time with no upper bound. We assume that in a risk-neutral world X follows a jump process with intensity λ and jump size H:[3]

$$dX = \mu dt + dq \tag{1}$$

In any short interval of time, Δt, $dq = H$ with probability $\lambda\Delta t$, and $dq = 0$ with probability $1 - \lambda\Delta t$. The non-decreasing nature of X requires that $\mu \geq 0$, and $H \geq 0$. We will assume that μ and λ are functions only of time and H is a function only of the number of jumps so far.[4]

For the purpose of using the model we do not need to explicitly consider the hazard rate process. However it is interesting to note that the process followed by the hazard rate, $h(t)$, is more extreme than the jump process considered by other researchers. It is

$$dh = \mu'(t)dt + dI$$

The term dI is an impulse that has intensity λ. The impulse takes the form of a Dirac delta function. The effect of an impulse at time t is to cause the hazard rate to become infinite in such a way that the integral of the hazard rate over any short interval around time t is finite.

The jumps in the model can lead to several companies defaulting at the same time. For example, suppose that S decreases from 1 to $1 - q$ as a result of a jump at time t. There

---

[3] Processes of the form $dX = \mu dt + \sigma dz$ where dz is a Wiener process and $\sigma > 0$ are inappropriate because they allow X to decrease. We could assume a process where $\sigma = 0$ and μ follows a positive diffusion process. However, this would be difficult to handle and would not calibrate well to market data. In our experience large jumps in X are necessary to fit market data.

[4] The default intensity, λ, can be allowed to depend on X. The model is then less analytically tractable, but can still be represented as a binomial tree.

was no chance of default before the jump, but each company has a probability q of defaulting by time t. If there are N companies in the portfolio the probability that n of them will default at time t is then

$$\frac{N!}{n!(N-n)!}q^n(1-q)^{N-n}$$

Another model which allows multiple defaults at the same time is the generalized Poisson loss model of Brigo *et al* (2007). In this model there are several independent Poisson processes. An integer z_i is associated with the ith process. When an event occurs in the ith process, z_i defaults occur simultaneously. The model provides a good fit to data on CDOs with several maturities but does not have the dynamic structure of our model.

A. Notation

We will use the following notation:

$p(J, t_1, t_2)$ The probability of exactly J jumps between times t_1 and t_2 ($t_2 > t_1$)

$P(J,t)$ The probability of exactly J jumps between time zero and time t (= $p(J, 0, t)$)

H_J The size of the Jth jump in X

$\phi(n,t_1,t_2)$ The probability of exactly n defaults between times t_1 and t_2

$\Phi(n,t)$ The probability of exactly n defaults between time zero and time t ($=\phi(n,0,t)$)

$S(J, t)$ — The cumulative probability of survival by time t when there have been J jumps between time zero and time t

$\Lambda(t)$ — $\int_0^t \lambda(\tau)\, d\tau$

$M(t)$ — $\int_0^t \mu(\tau)\, d\tau$

$E(t)$ — Expected principal on a tranche at time t when the initial tranche principal is $1.

$v(t)$ — Present value of $1 received at time t

$W(n, t)$ — The remaining CDO tranche principal at time t when there have been n defaults. The initial tranche principal is $1.

B. Valuation of a CDO

CDOs can be valued analytically using the model. Since the payments and payoffs associated with a CDO do not depend on decisions by the buyers and sellers of protection the value of the CDO is the same under all filtrations. As before, let a_L and a_H be the tranche attachment and detachment points, respectively. Define

$$n_L = \frac{a_L N}{1-R}$$

$$n_H = \frac{a_H N}{1-R}$$

where N is the number of companies in the portfolio and R is the fixed recovery rate. Denote the smallest integer greater than x by $g(x)$. For a tranche with an initial principal of \$1 the tranche principal at time t when there have been n defaults is

$$W(n,t)=\begin{cases} 1 & \text{when} \quad n<g(n_L) \\ \dfrac{a_H-n(1-R)/N}{a_H-a_L} & \text{when} \quad g(n_L)\le n<g(n_H) \\ 0 & \text{when} \quad n\ge g(n_H) \end{cases} \tag{2}$$

The probability of J jumps between time zero and time t is

$$P(J,t)=\frac{\Lambda(t)^J e^{-\Lambda(t)}}{J!} \tag{3}$$

The value of S at time t if there have been J jumps is

$$S(J,t)=\exp\left[-M(t)-\sum_{j=1}^{J}H_j\right] \tag{4}$$

The probability of n defaults in the portfolio by time t conditional on J jumps is

$$\Phi(n,t|J)=b(n,N,1-S(J,t)) \tag{5}$$

where b is the binomial probability function:

$$b(n,N,q)=\frac{N!}{n!(N-n)!}q^n(1-q)^{N-n}$$

The expected principal on the CDO tranche at time t conditional on J jumps is

$$E(t|J)=\sum_{n=0}^{N}\Phi(n,t|J)W(n,t) \tag{6}$$

The unconditional expected principal at time t is therefore

$$E(t)=\sum_J P(J,t)E(t|J) \quad (7)$$

Let the payment times be $t_1, t_2 \ldots t_m$, define $t_0 = 0$, and assume that defaults always happen half way through the period between payments. If the initial principal is \$1 the present value of the regular payments that are made on payment dates is sA where s is the spread and

$$A=\sum_{k=1}^{m}(t_k-t_{k-1})E(t_k)v(t_k) \quad (8)$$

The present value of the accrual payments made in the event of a default are sB where

$$B=0.5\sum_{k=1}^{m}(t_k-t_{k-1})\left[E(t_{k-1})-E(t_k)\right]v(t_k^*) \quad (9)$$

where $t_k^*=0.5(t_k+t_{k-1})$. The present value of the payoffs arising from defaults is

$$C=\sum_{k=1}^{m}\left[E(t_{k-1})-E(t_k)\right]v(t_k^*) \quad (10)$$

The total value of the contract to the seller of protection is $sA + sB - C$. The breakeven spread is $C/(A + B)$.

C. The Loss Distribution

The model has been presented as a reduced form model. However, it can be converted to a top down model where the process for the loss is modeled. Because we are assuming a constant recovery rate it is sufficient to model the number of defaults n. In this version of the model the filtration is different from that given at the beginning of this section. It is assumed that at time t we know the number of defaults, but we do not know survival probabilities. The process for survival probabilities becomes nothing more than a convenient tool for generating the model.

The proportion of the original portfolio lost by time t when the number of defaults is n is

$$\frac{n}{N}(1-R)$$

First consider the unconditional distribution for n at time t. This can be calculated from equations (3) and (5)

$$\Phi(n,t)=\sum_{J}\Phi(n,t|J)P(J,t) \tag{11}$$

This is a mixture of Poisson distributions. The loss distribution given by the model has thin tails when jumps are small and fat tails when they are large. This means that the model is a flexible tool for handling a variety of default correlation environments.

The probability of n defaults between times t_1 and t_2 conditional on J_1 jumps by time t_1, J_2 jumps by time t_2, and n_1 defaults between times zero and t_1 is

$$\phi(n,t_1,t_2|J_1,J_2,n_1)=b\Big(n,N-n_1,\big(S(J_1,t_1)-S(J_2,t_2)\big)\big/S(J_1,t_1)\Big) \tag{12}$$

The probability of J jumps between times t_1 and t_2 is

$$p(J,t_1,t_2)=\frac{\left[\Lambda(t_2)-\Lambda(t_1)\right]^J e^{\Lambda(t_1)-\Lambda(t_2)}}{J\,!} \tag{13}$$

and from Bayes rule the probability of J jumps by time t conditional on n defaults by time t is

$$P(J,t|n)=\frac{\phi(n,t|J)P(J,t)}{\sum_j \phi(n,t|j)P(j,t)} \tag{14}$$

Equations (12) and (13) can be used to calculate the probability that there will be n_2 defaults at time t_2 conditional on n_1 defaults and J_1 jumps by time t_1:

$$\Phi\left(n_2,t_2\,|n_1,J_1,t_1\right)=\sum_J \phi\left(n_2-n_1,t_1,t_2\,|J_1,J_1+J,n_1\right)p\left(J,t_1,t_2\right) \tag{15}$$

Using equation (14) we obtain the transition probability from n_1 defaults at time t_1 to n_2 defaults at time t_2

$$\Phi\left(n_2,t_2\,|n_1,t_1\right)=\sum_{J_1} \Phi\left(n_2,t_2\,|n_1,J_1,t_1\right)P\left(J_1,t_1\,|n_1\right) \tag{16}$$

This equation defines the process for the number of defaults (or equivalently the loss) for the portfolio.

III. ALTERNATIVE VERSIONS OF THE MODEL AND THEIR CALIBRATION

In this section we first present a version of the model where (similar to Black-Scholes) there is just one free parameter. We then discuss extensions to this simple model and present a three-parameter version of the model that is designed to provide a good fit to all CDO spreads of all maturities.

The calibration of the model to the market will be illustrated using the data in Exhibit 1 for iTraxx and CDX NA IG on January 30, 2007.[5] We assume a recovery rate of 40%. The term structure of CDS spreads is assumed to be a piece-wise linear function. It equals the three year spread for maturities up to three years; between years three and five the spread is interpolated between the three and five year CDS spreads; and so on. Payments on CDOs and CDSs are assumed to be quarterly in arrears.

A: Zero Drift; Constant Jumps; Time-Dependent Intensity

A particularly simple version of the model is the case in which $\mu(t) = 0$ and the jump size is constant. For any given value of the jump size, H, the jump intensity $\lambda(t)$ is chosen to match the term structure of CDS spreads. We assume that the value of $\lambda(t)$ is constant between CDO/CDS payment dates and work forward in time choosing λ's so that the CDS term structure is matched.

[5] Following the usual conventions the quotes in Exhibit 1 are the rate of payment in basis points per year to purchase protection from defaults in the indicated range. The exception is the 0 to 3% tranche where the quote is the up front payment as a percentage of the notional that is paid in addition to 500 basis points per year.

This version of the model has just one free parameter: the jump size. The jump size can be implied from a CDO market quote or vice versa. The procedure for calculating the value of a CDO tranche is described in Section IIB. The jump sizes implied from the iTraxx market quotes in Exhibit 1 are shown in Exhibit 2. The numbers in this table are analogous to the numbers in a volatility surface that is determined from market quotes for option prices using the Black-Scholes formula. Just as option traders monitor volatility surfaces, credit derivatives traders can monitor the jump size surface.

The implied jump size is a measure of default correlation. As the jump size approaches zero the default correlation approaches zero. As the jump size becomes large the default correlation approaches one. Exhibit 3 compares the implied jump sizes reported in Exhibit 2 with the tranche (or compound) correlations implied from the tranche quotes using a Gaussian copula model. It can be seen that the two exhibit very similar patterns. Results are similar for CDX NA IG.

The advantage of calculating an implied jump size rather than an implied copula correlation is that the jump size is associated with a dynamic model whereas the copula correlation is associated with a static model.

B. Extensions of the Constant Jump Model

There are a number of ways the model we have just considered can be extended. The most obvious extension is to allow a non-zero drift. In this case the jump size and intensity can be chosen to fit CDO quotes while the drift is selected to fit the CDS term structure. In practice this does not provide material improvement over the zero-drift case.

We can obtain a better fit to CDO prices for a particular maturity by having multiple jump processes each with its own size and each with its own intensity or by having a single jump process in which the jump size is a random variable. This would be in the spirit of Longstaff and Rajan (2006) or Brigo *et al* (2006). These approaches work reasonably well in fitting panels of CDO quotes but pose computational problems in the convolution of the survival probabilities. When either of these approaches is implemented the best fit to market CDO quotes arises when there is a high probability of small jump sizes and a low probability of large jump sizes. These approaches do not have the property that large jump sizes are associated with adverse states of the world.

Rather than developing the model along these lines we have chosen to consider a version of the model that is particularly easy to implement and involves relatively few parameters.

C. A Three-Parameter Model

Empirical evidence suggests that default correlations are stochastic and increase in adverse credit conditions. For example, Servigny and Renault (2002) who look at historical data on defaults and ratings transitions to estimate default correlations, find that the correlations are higher in recessions than in expansion periods. Das *et al* (2006) employ a reduced form approach and compute the correlation between hazard rates. They conclude that correlations increase when hazard rates are high. Hull *et al* (2005) show that a structural model fits CDO data well when asset correlations are positively related to the default rate. Their work is consistent with that of Ang and Chen (2002) who find that the correlation between equity returns is higher during a market downturn.

This research suggests that a model where the jump sizes in X are larger in adverse market conditions might fit market data better than a constant jump size model. To test this we relax the constant jump size assumption and explicitly build in the property that large jump sizes (high correlation) are associated with low survival probability states. The size of the Jth jump is given by

$$H_J = H_0 e^{\beta J}$$

where H_0 and β are positive constants. The intensity of the process, λ, is a constant and the drift, $\mu(t)$, is determined to match CDS spreads.

In calibrating the model the objective is to find values of H_0, β, and λ that minimize the sum of squared differences between market tranche spreads and model tranche spreads. The procedure involves repeatedly a) choosing trial values of H_0, β, and λ, b) calculating the $\mu(t)$ function so that the term structure of index spreads is matched for the trial parameters, and c) calculating the sum of squared differences between model spreads and market spreads for all tranches of all maturities (15 spreads in total). An iterative procedure is used to find the values of H_0, β, and λ that lead to the sum of squared differences being minimized when this three-step procedure is used.

For the iTraxx data in Exhibit 1 the best fit parameter values are H_0=0.00223, β=0.9329, and λ=0.1486. The corresponding values for the CDX NA IG data are H_0=0.00147, β=1.2813, and λ=0.1310. The pricing errors are shown in Exhibit 4. The model fits market data well − much better than versions of the model where the jump size is constant.

Exhibit 5 shows the loss distribution 3-, 5- , 7- and 10-years in the future for iTraxx as seen at time zero based on the calibrated model. As the time horizon increases the probability mass of the distribution shifts to the right and becomes more spread out. There is also some fine structure in the distribution. To make this visible all probabilities for losses greater than 9% are scaled up by a factor of 100. This reveals a complex shape to the right tail of the loss distribution. Results for CDX NA IG are similar.

The values of H_J are initially fairly small, but increase fast. For example in the case of iTraxx on January 30, 2007 $H_3 = 0.057$, $H_5 = 0.386$, and $H_7 = 2.513$. There is a small probability of low values of S being reached. For example, the probability that S is less than 0.90 at the end of 5 years is about 0.007, and the probability that it is less than 0.75 is about 0.0001. Some of these results values may seem extreme. However, they are consistent with the results in papers such as Hull and White (2006) which show that it is necessary to assign a very low, but non-zero, probability to a very high hazard rate in a static model in order to fit market quotes.

It should be recalled that the results shown here are risk-neutral probabilities that are inferred from tranche prices. Consider the 2 basis point cost of protection for the 5-year 12 to 22% tranche of iTraxx. Given the recovery rate that is assumed, the buyer of protection only receives payoffs when more than 20% of the entire portfolio has defaulted and payments for protection only stop when the tranche is wiped out after 36.7% of the portfolio has defaulted. Suppose there are only two possible outcomes: the tranche is wiped out with probability p or is untouched with probability $1–p$. If the tranche is wiped out the buyer of protection receives a payoff of $1; if the tranche is untouched he pays 2 basis points per year for 5 years on a notional of $1. Ignoring discounting, the value of *p* that makes this a fair contract is about 0.001 which is roughly consistent with the probabilities implied by the calibrated model.

This calibration procedure was repeated for all the iTraxx tranche data that was available from Reuters between July 4, 2006 and January 10, 2007. This data includes the spreads on 5-, 7- and 10-year CDO tranches as well as 3- to 10- year index spreads. All 15 CDO tranche spreads were available on 51% of the days, 14 spreads were available on 34% of the days, 13 spreads were available on 13% of the days and on 2% of the days only 12 spreads were quoted. The quality of the fit on each day was similar to that reported for the January 30, 2007 market data. Exhibit 6 shows ten-day moving averages for the parameter values.[6]

[6]A ten-day moving average provides a better indication of parameter values than the daily results because of the impact of noise in the quotes and missing quotes.

IV. APPLICATIONS OF THE MODEL

In this section we show how the model we calibrated to iTraxx market data in Section IIIC can be used to value a number of different instruments.

A. Forward CDOs

Consider first forward CDOs.[7] As discussed earlier the model in this paper is a specific dynamic model. We are modeling defaults for a portfolio of companies that is defined at time zero and remains unchanged. The forward CDO spreads that we calculate using the model are the spreads for this portfolio. For example, the 3×2 forward spread for the mezzanine (3% to 6%) tranche is the spread that must be paid in years 4 and 5 on the remaining amount of the original tranche principal in order to provide protection. The protection is against those default losses that occur in years 4 and 5 and are between 3% and 6% of the principal of the original underlying portfolio. The calculated spread is not that for a forward start *de novo* iTraxx contract. The latter contract would be based on a portfolio that will be selected to be investment grade at the start of the protection period covered by the forward contract.

Forward contracts can be valued analytically using the model in a similar way to CDOs. As with the CDO, since the payments and payoffs associated with the forward contract do not depend on decisions by the buyers and sellers the value of the forward contract is the same under all filtrations. Suppose that the forward contract lasts between payment times t_u and t_m. Define $A(t_u)$, $B(t_u)$, and $C(t_u)$ similarly to A, B, and C in equations (8), (9)

[7] Forward CDOs provide a simple application of the model and lead into the calculation of options on CDOs in the next section. As explained in Hull and White (2007) a dynamic model is not necessary to determine forward CDO spreads.

and (10) except that they accumulate expected payoffs and payments between times t_u and t_m rather than between time zero and time t_m:

$$A(t_u) = \sum_{k=u+1}^{m} (t_k - t_{k-1}) E(t_k) v(t_k)$$

$$B(t_u) = 0.5 \sum_{k=u+1}^{m} (t_k - t_{k-1}) \left[E(t_{k-1}) - E(t_k) \right] v(t_k^*)$$

$$C(t_u) = \sum_{k=u+1}^{m} \left[E(t_{k-1}) - E(t_k) \right] v(t_k^*)$$

The total value of the forward CDO contract to the seller of protection when the tranche spread is s is $sA(t_u) + sB(t_u) - C(t_u)$. The breakeven forward spread is $C(t_u)/[A(t_u) + B(t_u)]$.

The breakeven spreads for forward contracts on CDO tranches that mature in five years are shown in Exhibit 7. The tranche spreads rise fairly fast. One reason for this is that the CDS spreads are upward sloping so that default probabilities tend to increase as time passes. Another reason is that, in the case of all tranches except the equity tranche, there is very little probability of loss in the first one or two years. It follows that when these years are excluded the spread for the remaining years increases.

B. Options on CDO Tranches

A European option on a CDO tranche is an option to buy or sell protection for a particular tranche at a particular strike spread. The option expires at time t_u and if exercised the protection lasts between times t_u and t_m. As with the forward CDO contract, the underlying portfolio is defined at time zero and remains unchanged. Upon option exercise the spread is applied to the remaining principal (if any) of the tranche. Like forwards European options can be valued analytically using the model.

Since the buyer of the option must decide when to exercise the option the information available, the filtration, will affect the value of the option. It will be recalled that we are working in a filtration where both $S(t)$ and the number of defaults between time zero and time t are known at time t. This information will therefore be used to determine exercise decisions. The number of defaults is directly observable by the holder of an option and we assume that the prices of CDSs and CDO tranches reflect the state of the economy, $S(t)$. These prices are observable and can be used to formulate an exercise decision. An option to buy protection will be exercised at time t_u if the cost of protection on the tranche at time t_u is higher than the strike spread.

Given the structure of our model, knowing the value of S at time t_u is equivalent to knowing the number of jumps J_u before time t_u. Conditional on n_u defaults and J_u jumps by time t_u the expected tranche principal at time t_k ($k > u$) can be calculated from equations (2) and (15) as

$$E\left(t_k \middle| n_u, J_u, t_u\right) = \sum_{n=n_u}^{N} W\left(n, t_k\right) \Phi\left(n, t_k \middle| n_u, J_u, t_u\right)$$

From this the present value of payments, accrual payments, and payoffs condition on n_u and J_u can be calculated:

$$A(t_u | n_u, J_u) = \sum_{k=u+1}^{m} (t_k - t_{k-1}) E(t_k | n_u, J_u, t_u) v(t_k)$$

$$B(t_u | n_u, J_u) = 0.5 \sum_{k=u+1}^{m} (t_k - t_{k-1}) [E(t_{k-1} | n_u, J_u, t_u) - E(t_k | n_u, J_u, t_u) v(t_k^*)]$$

$$C(t_u | n_u, J_u) = \sum_{k=u+1}^{m} [E(t_{k-1} | n_u, J_u, t_u) - E(t_k | n_u, J_u, t_u) v(t_k^*)]$$

Suppose the strike spread for a European call option is s_K. The present value of the option when there have been n_u defaults and J_u jumps by time t_u is

$$\max\left[C\left(t_u|n_u,J_u\right)-s_K A\left(t_u|n_u,J_u\right)-s_K B\left(t_u|n_u,J_u\right),0\right]$$

The value of the call option can therefore be calculated from equations (3), (4), and (5) as

$$\sum_{J_u}\sum_{n_u=0}^{N}\Phi\left(n_u,t_u|J_u\right)P\left(J_u,t_u\right)\max\left[C\left(t_u|n_u,J_u\right)-s_K A\left(t_u|n_u,J_u\right)-s_K B\left(t_u|n_u,J_u\right),0\right]$$

Similarly the value of a European put option is

$$\sum_{J_u}\sum_{n_u=0}^{N}\Phi\left(n_u,t_u|J_u\right)P\left(J_u,t_u\right)\max\left[s_K A\left(t_u|n_u,J_u\right)+s_K B\left(t_u|n_u,J_u\right)-C\left(t_u|n_u,J_u\right),0\right]$$

When the strike spread is set equal to the forward spread an at-the-money option is created and the put and call have the same price. Exhibit 8 shows the prices in basis points of at-the-money options with varying maturities on a 5-year CDO tranche. For example, a one-year at-the-money option to buy at-the-money protection on the 6% to 9% tranche for the period from one year to five years costs 23.2 basis points or 0.232% of the initial (time zero) tranche principal. The option strike spreads are the corresponding forward spreads given in Exhibit 7.[8]

Tranches with higher spreads (lower attachment points) have higher option values. The value of the option initially increases as we increase the option maturity and then starts to

[8] We can also value options using the top down version of the model in Section IIC. Only n_u is then used in generating the exercise decision. The results in Section IIC can be used to value options analytically in this case. Conditioning exercise on n_u results in substantially lower option prices than conditioning on both n_u and S. The differences vary by tranche and are largest for the mezzanine tranche. The average price reduction for the options in Exhibit 8 is about 13%.

decline to reflect the diminishing maturity of the underlying. This is similar to what we observe for options on bonds and swaps.

As discussed in Hull & White (2007) if we are willing to assume that CDO tranche spreads are lognormally distributed it is possible to derive an analytic expression for the prices of European put and call options on CDO tranches. Exhibit 9 uses the result in Hull and White (2007) to calculate the implied spread volatilities from the prices in Exhibit 8. For the equity tranche the implied volatility decreases with option maturity while for all other tranches the implied volatility increases with option maturity.

Exhibit 10 shows the implied volatilities for two-year options on 5-year CDO tranches for strike prices between 75% and 125% of the forward tranche spread. Again the results are based on Hull and White (2007). If the assumption of lognormality held exactly the implied volatility for an option would be the same for all strike prices. The table shows that the variation of implied volatility with the strike price is quite small. This suggests that the lognormality assumption in Hull and White (2007) is (at least for January 30, 2007 data) approximately consistent with the model in this paper.

C. Leveraged Super Senior Transaction

Our final application of the model is to a leveraged super senior (LSS) tranche with a loss trigger. This is a CDO tranche which is automatically cancelled when losses reach some level. On cancellation the tranche is marked to market and the seller of protection must pay the buyer an amount equal to the value of the tranche. However, the total amount paid by the seller (including any losses for which the seller is responsible prior to the cancellation date) is capped at a fraction x of the tranche notional. Since the payments and payoffs associated with the LSS do not depend on decisions by the buyers and sellers the value of the contract is the same under all filtrations.

Because we are assuming a constant recovery rate of 40% the loss that triggers cancellation can be translated to a number of defaults. We suppose that cancellation is triggered when the number of defaults reaches n^*. The parameters defining the LSS on a particular portfolio are therefore the tranche attachment point a_L, the tranche detachment point a_H, the spread s, the maturity T, the loss cap x, and the number of losses triggering cancellation n^*.[9] To simplify matters we only consider cases in which the number of defaults that triggers a termination, n^*, is less than $125a_L/(1-R)$. This ensures that the seller is not responsible for any losses prior to the termination date.[10]

The tranche in a LSS is generally a senior one. The seller of protection considers it highly unlikely that the tranche will experience losses, but does not find selling protection on the tranche to be appealing because the spread is small. The LSS provides a way of leveraging the spread. Suppose that the tranche principal is $10 million and the spread is 15 basis points. Selling protection on the tranche in the usual way requires $10 million to be deposited by the seller of protection at an interest rate of LIBOR so that, if the tranche does not experience defaults, the interest earned is LIBOR plus 15 basis points. Consider an LSS with $x = 0.1$. Only $1 million needs to be deposited at the beginning of the transaction because that is all that the seller of the protection is risking. As a result the

[9] For ease of exposition we assume n^* is constant. The methodology we present can easily be extended to the situation (often encountered in practice) where n^* changes with the passage of time.

[10] Our methodology can be extended to cases where this condition is not satisfied.

spread earned on this principal, assuming the tranche does not experience defaults and the deal is not cancelled is LIBOR plus 150 basis points.

The model for S that we developed in Section IIIC can be represented as a binomial tree. To construct the tree the life of the model is divided into a number of short time intervals. Denote the time corresponding to the end of the ith interval by τ_i and let $\tau_0 = 0$. During each time interval it is assumed that there is either zero or one jump in S. This leads to a tree with the geometry shown in Exhibit 11. In this figure H_J is the size of the Jth jump and $M_i = M(\tau_i)$. The probability on the upper and lower branches emanating from a node at time τ_i are $\lambda_i\Delta_i$ and $1-\lambda_i\Delta_i$, respectively, where $\lambda_i = \lambda(\tau_i)$ and $\Delta_i = \tau_{i+1}-\tau_i$. The τ_i are chosen so that there are nodes on each payment date (i.e., for each k, $\tau_i = t_k$ for some i.) In practice this is achieved by creating ν equal time steps between each payment date for some integer ν.

Note that the binomial tree as it is presented in Exhibit 11 only reveals $S(t)$. As a result it would produce different values for options on a CDO tranche than those discussed in section IV B.[11] To value an option where the exercise decision depends on both $S(t)$ and the number of defaults to date it would be necessary to construct a binomial tree for S where there are N+1 states at each S-node corresponding to the N+1 alternative values for the number of defaults at the node. The use of this type of tree was proposed by Hull and

[11] Similar to the development of the top down version of the model in which only the number of defaults is known, it is possible to develop a version of the model in which only $S(t)$ is known. The resulting prices for European options on CDO tranches are usually about 1% lower than the prices discussed in section IV B where the exercise decision is based on both $S(t)$ and the number of defaults to date.

White (1993). It is now sometimes referred to as a "binomial forest." As one rolls back through the tree calculations are carried out for all N+1 states at each node.

As it happens the binomial tree can be used to value a LSS without becoming a forest. The key difference between an LSS and an option on a CDO is that there are no decisions to be made by either party at any node in an LSS.

Denote the jth node at time τ_i by (i, j). Let S_{ij} and E_{ij} be the cumulative survival probability and expected tranche principal at node (i, j). The value of S_{ij} is given by equation (4) and E_{ij} can be calculated from equations (2), (5) and (6).

We first roll back through the tree calculating V_{ij}, the value of the tranche to the protection buyer per dollar of principal at node (i, j) assuming no cancellation and no limit on the liability of the protection seller.[12] Some times τ_i correspond to payment dates and others do not. Define δ_i as follows. When τ_i is a payment date so that $\tau_i = t_k$, δ_i equals the accrual fraction $t_k - t_{k-1}$. When τ_i is not a payment date $\delta_i = 0$. Variables A_{ij}, B_{ij}, and C_{ij} can be defined analogously to A, B, and C in Section IIB. At the final nodes $A_{ij} = \delta_i E_{ij}$, $B_{ij} = 0$, and $C_{ij} = 0$ and at earlier nodes they can calculated by working backward through the tree using

[12] This could be calculated analytically, but since the tree is used for incorporating the impact of the cancellation it makes sense to use the tree for this as well. The results are slightly different from those previously reported because now defaults are no longer restricted to occur in the middle of an accrual period.

$$
\begin{aligned}
A_{ij} &= \left[\lambda_i \Delta_i A_{i+1,j+1} + (1-\lambda_i \Delta_i) A_{i+1,j}\right] v(\tau_{i+1}) / v(\tau_i) + \delta_i E_{ij} \\
B_{ij} &= \left[\lambda_i \Delta_i B_{i+1,j+1} + (1-\lambda_i \Delta_i) B_{i+1,j}\right] v(\tau_{i+1}) / v(\tau_i) \\
&\quad + 0.5\left[\lambda_i \Delta_i (E_{ij} - E_{i+1,j+1}) + (1-\lambda_i \Delta_i)(E_{ij} - E_{i+1,j})\right] v(\tau_i^*) / v(\tau_i) \qquad (17) \\
C_{ij} &= \left[\lambda_i \Delta_i C_{i+1,j+1} + (1-\lambda_i \Delta_i) C_{i+1,j}\right] v(\tau_{i+1}) / v(\tau_i) \\
&\quad + \left[\lambda_i \Delta_i (E_{ij} - E_{i+1,j+1}) + (1-\lambda_i \Delta_i)(E_{ij} - E_{i+1,j})\right] v(\tau_i^*) / v(\tau_i)
\end{aligned}
$$

where $\tau_i^* = 0.5(\tau_i + \tau_{i-1})$. The V_{ij} are then calculated as

$$V_{ij} = C_{ij} - s\left(A_{ij} + B_{ij}\right)$$

To value the LSS it is necessary to roll back through the tree again. At node (i, j) we know the probability that a company has not defaulted, S_{ij}. The probability of n defaults by node (i, j), is given by equation (5) as $b(n, N, 1 - S_{ij})$. The probability that termination has not been triggered at a node (i, j) so that the deal is alive is therefore

$$w_{ij} = \sum_{i=0}^{n^*-1} b\left(i, N, 1 - S_{ij}\right)$$

As in the case of a CDO we assume without loss of generality that the principal is \$1. Variables A_{ij}^{LSS} and C_{ij}^{LSS} are defined analogously to A_{ij} and C_{ij} so that sA_{ij}^{LSS} is the present value of payments on the LSS from node (i, j) to the end of the life of the LSS and C_{ij}^{LSS} as the present value of payoffs from node (i, j) to the end of the life of the transaction. (We do not need B_{ij}^{LSS} because the protection seller is never responsible for any defaults prior to cancellation in the LSS).

At node (i, j) we are in one of three situations:

1. The structure is still alive.

2. Cancellation was triggered at an earlier node.

3. Cancellation is triggered at node (i, j).

The probability that the structure is still alive at node (i, j) is w_{ij}. In this case the backwards induction equations for A_{ij}^{LSS} and C_{ij}^{LSS} are the same as those for A_{ij} and C_{ij} in equation (17).

The probability that the structure is not alive at node (i, j) is $1-w_{ij}$. We are then in situations 2 or 3. Assume that situation 3 applies. The protection buyer must make a final payment equal to $\varepsilon_i E_{ij}$ where ε_i is the time since the previous payment date. There are two components to C_{ij}^{LSS}. The seller of protection must pay $-V_{ij}$, the present value of future net payments on a regular CDO tranche as well as the payment that would be due on a regular CDO tranche at time $i\Delta t$. The latter is $1-E_{ij}$ because we know that the protection seller is not responsible for defaults prior to cancellation. Payoffs are capped at x so that C_{ij}^{LSS} equal the minimum of x and $-V_{ij}+1-E_{ij}$.

Assuming situation 1 or situation 3 applies leads to the following formulas for calculating A_{ij}^{LSS} and C_{ij}^{LSS}

$$A_{ij}^{LSS} = w_{ij}\left\{\left[\lambda_i\Delta_i A_{i+1,j+1}^{LSS} + (1-\lambda_i\Delta_i)A_{i+1,j}^{LSS}\right]v(\tau_{i+1})/v(\tau_i) + \delta_i E_{ij}\right\} + (1-w_{ij})\varepsilon_{ij}E_{ij}$$

$$C_{ij}^{LSS} = w_{ij}\left[\lambda_i\Delta_i C_{i+1,j+1}^{LSS} + (1-\lambda_i\Delta_i)C_{i+1,j}^{LSS}\right]v(\tau_{i+1})/v(\tau_i) + (1-w_{ij})\min\left[-V_{ij} + (1-E_{ij}), x\right]$$

Because of the way the values of A_{ij}^{LSS} and C_{ij}^{LSS} get overwritten as we work back through the tree the values calculated at the first node take into account the possibility of situation 2 applying at nodes (i,j). The breakeven spread is therefore $C_{00}^{LSS}/A_{00}^{LSS}$.

To illustrate the properties of LSS transactions we calculated the breakeven spread for a 5-year CDO iTraxx tranche on January 30, 2007 using the model in Section IIIC with a_L = 0.12, a_H = 0.22. We considered leverage levels between 5% (x=0.05) and 20% (x=0.2) and for trigger levels from 1 to 16 defaults.

The results are shown in Exhibit 12. It should be noted that while the market spread for the 12% to 22% iTraxx tranche is 2 basis points, the spread for this tranche in the calibrated model is 1.53 basis points. All the results we present (for forward contracts, options, and LSSs) are of course based on the calibrated model. In particular the spread for the LSS converges to the calibrated 1.53 basis points. If our sole objective had been to price the LSS considered here, when carrying out the calibration we would have given the squared error in the spread for the 12 to 22% tranche a high weight compared with that for other tranches. This would have resulted in that tranche having a spread very close to 2 basis points in the calibrated model.

As the degree of leverage and the number of defaults required to trigger termination are increased the breakeven spread declines. Higher leverage means that the buyer of protection gets a smaller payoff in the event of default and so should be required to pay less for the protection. When the number of defaults required to trigger termination is increased it is less likely that the deal will be terminated prematurely and more likely that the seller of protection will benefit from the loss cap. Again this reduces the amount the buyer of protection should be required to pay.

V. EXTENSION OF MODEL

We now consider a number of ways in which the basic homogeneous constant recovery rate model can be extended.

A. Stochastic Recovery Rate

Let the cumulative default probability, 1–*S*, be denoted as *Q*. A negative relationship between the recovery rate and the cumulative default rate can be incorporated into the model by assuming that the average recovery rate $\bar{R}$ applying to all the defaults that have occurred up to time *t* is a function only of the cumulative default probability, *Q*. One possibility is the relationship

$$\bar{R} = R_0 \frac{1}{aQ}\left(1 - e^{-a(Q-Q_E)}\right)$$

where *a* and R_0 are positive constants and Q_E is the expected value of *Q* at the particular time being considered. Under this model the marginal recovery rate when the cumulative default probability is *Q* is

$$\frac{\partial(\bar{R}Q)}{\partial Q} = R_0 e^{-a(Q-Q_E)}$$

This is greater than R_0 when *Q* is lower than its expected value and less than R_0 when it is greater than its expected value. A constraint is that *a* must be chosen so that $R_0 e^{-a(Q-Q_E)}$ is always less than one.

Define $\bar{R}(J,t)$ as the average recovery rate between times 0 and *t* when there have been *J* jumps in *X*. In equation (2) *W*(*n*,*t*) is replaced by $W(n,t|J)$, *R* is replaced by $\bar{R}(J,t)$, and

n_L and n_H are defined in terms of $\overline{R}$ so that they become dependent on *J* and *t*. Equation (6) becomes

$$E(t|J) = \sum_{n=0}^{N} \Phi(n,t|J)W(n,t|J)$$

Apart from these changes the valuation of a CDO tranche is the same as before. Similar changes enable forward CDOs, European options on CDOs, and leveraged super seniors to be valued. In the case of leveraged super senior it is necessary to calculate the critical number of defaults at each node that lead to the loss threshold be exceeded.

B. Heterogeneous Model

The model in Section IIIC can be extended so that it becomes a heterogeneous model where each company has a different CDS spread.[13] The model that has been presented is then assumed to represent the evolution of the cumulative default probability for a representative company in the portfolio. For any particular company in the portfolio the jump size and jump intensity are assumed to be the same as that for the representative company. However the deterministic drift, $\mu(t)$, is adjusted to match the CDS spread. The binomial model for determining the probability of *n* defaults by time *t* conditional on *J* jumps in equation (5) must be replaced by an iterative procedure such as that in Andersen et al (2003) and Hull and White (2004). Once this has been done the model structure and calculations for valuing a CDO are much the same as we have presented them.

[13] Recall that we are working a filtration where, in addition to $S(t)$, we know the number of defaults by time *t*. We do not know the particular companies that have defaulted. Incorporating the latter into our model, or any other model, would be prohibitively time consuming.

Valuing a forward CDO is similar to valuing a regular CDO. To value a leveraged super senior it is necessary to calculate at each node of the tree the probability the critical loss level has been exceeded. One application of the iterative procedure just mentioned is therefore necessary for each node of the tree. For options on CDOs calculations are more time consuming. It is necessary to calculate the probability of n-n_u defaults between times t_u and t_k conditional on n_u defaults by time t_u and J_u jumps by this time. For this we require, for each of the N companies, the probability of survival by time t_u conditional on n_u defaults and J_u jumps by time t_u. This is $\gamma_1\gamma_2/\gamma_3$ where γ_1 is the unconditional probability of the company surviving to time t_u, γ_2 is the probability of n_u out of the remaining N–1 companies defaulting, and γ_3 is the probability of n_u out of N companies defaulting. A large number of applications of the iterative procedure are required.[14]

C. Modeling Two Portfolios Simultaneously

The model can be extended so that iTraxx and CDX NA IG are modeled simultaneously. One way of doing this is to have three independent jump processes. The first process leads only to jumps in the cumulative hazard rate for iTraxx companies; the second jump process leads only to jumps in the cumulative hazard rate for CDX NA IG companies; the third jump process leads to jumps in the cumulative hazard rate for both iTraxx and CDX NA IG companies.

D. Bespoke Portfolios

In practice it is often the case that derivatives dependent on bespoke portfolios have to be valued with dynamic models. The model we have presented must first be calibrated to

[14] It may well be possible to speed up the calculations by developing a robust approximation to the probability that a particular company will survive by time t_u conditional on n_u defaults.

iTraxx or CDX NA IG (or both). The drifts for the cumulative hazard rates of individual names comprising the bespoke portfolio can then chosen to match their CDS spreads.

VI. CONCLUSIONS

We have presented a simple one-factor model for the evolution of defaults on a portfolio. The model has two advantages over the Gaussian copula model. First, it is simpler and easier to implement. Second, it is a dynamic model that allows a wider range of products to be valued. The model is an alternative to the more complex dynamic models suggested by other researchers. To our knowledge this is the first paper to use a dynamic credit model for pricing a variety of different types of portfolio credit derivatives.

The model has the attractive feature that it has many analytic properties and can be represented in the form of a binomial tree. The variable modeled on the tree is the cumulative survival probability for a representative company. The model is easy to use and appears to have the property that future CDO spreads are approximately lognormal.

We have shown how the model can be calibrated to market data. Our results indicate that in a risk neutral world there is a small chance that the default probability for a representative company during the life of the model will be very high. This is consistent with the results from static copula models. To fit market data it is necessary for jump sizes to become increasingly large. This is consistent with empirical data showing that default correlations are higher in recessionary periods.

REFERENCES

Albanese, C.; O. Chen; A. Dalessandro; and A. Vidler. "Dynamic Credit Correlation Modeling," Working paper, Imperial College (2005).

Andersen, L.; J. Sidenius; and S. Basu. "All Your Hedges in One Basket," *Risk* , 16 (2003), November, 67-72.

Ang, A., and J. Chen. "Asymmetric Correlations of Equity Portfolios," *Journal of Financial Economics*, 63 (2002), 443-94.

Baxter, M. "Levy Process Dynamic Modeling of Single Name Credits and CDO Tranches," Working paper, Nomura International (2006).

Bennani, N. "The Forward Loss Model: A Dynamic Term Structure Approach for the Pricing of Portfolio Credit Risk," Working Paper, The Royal Bank of Scotland (2006).

Brigo, D.; A. Pallavicini; and R. Torresetti. "Calibration of CDO Tranches with the Dynamical Generalized-Poisson Loss Model," Risk, 20 (2007), May, 70-75.

Chapovsky, A.; A. Rennie; and P.A.C. Tavares. "Stochastic Modeling for Structured Credit Exotics," Working paper, Merrill Lynch International (2006).

Das, S. R.; D. Duffie; N. Kapadia; and L. Saita. "Common Failings: How Corporate Defaults are Correlated," *Journal of Finance*, 62 (2007), 1, 93-117.

Das, S. R.; L. Freed; G. Geng; and N. Kapadia. "Correlated Default Risk," *Journal of Fixed Income*, 16 (2006), 2, 7-32.

Duffie, D., and N. Gârleanu. "Risk and the Valuation of Collateralized Debt Obligations," *Financial Analysts Journal*, 57 (2001), January-February, 41-59.

Ehlers, P., and P. Schönbucher. "Background Filtrations and Canonical Loss Processes for Top-down Models of Portfolio Credit Risk," Working paper, ETH Zurich (2006).

Errais, E.; K. Gieseke; and L. Goldberg. "Pricing Credit from the Top Down Using Affine Point Processes," Working paper, Department and Management Science and Engineering, Stanford University (2006).

Graziano, G., and L.C.G. Rogers. "A Dynamic Approach to the Modeling of Correlation Credit Derivatives Using Markov Chains," Working paper, University of Cambridge (2005)

Gregory, J., and J-P. Laurent. "Basket Default Swaps, CDOs, and Factor Copulas," *Journal of Risk*, 7 (2005), 8-23.

Heath, D., R. Jarrow, and A. Morton. "Bond Pricing and the Term Structure of Interest Rates: A New Methodology," *Econometrica*, 60 (1992), 77-105.

Hull, J., and A. White. "Efficient Procedures for Valuing European and American Path-dependent Options," *Journal of Derivatives*, 1 (1993), 1-13.

Hull, J., and A. White. "Valuation of a CDO and nth to Default CDS without Monte Carlo Simulation," *Journal of Derivatives*, 12 (2004), 8-23.

Hull, J., and A. White. "The Valuation of Credit Default Swap Options" *Journal of Derivatives*, 10 (2003), 40-50.

Hull, J., and A. White, "Valuing Credit Derivatives Using an Implied Copula Approach," *Journal of Derivatives*, 14 (2006), 8-28.

Hull, J. and A. White. "Forward and European Options on CDO Tranches," *Journal of Credit Risk*, 3 (2007), 63-73.

Hull, J.; M. Predescu; and A. White. "The Valuation of Correlation-dependent Credit Derivatives Using a Structural Model," Working Paper, University of Toronto (2005).

Hurd, T.R., and A. Kuznetsov. "Fast CDO Computations in Affine Markov Models," Working Paper, McMaster University (2005)

Joshi, M., and A. Stacey. "Intensity Gamma," *Risk*, 19 (2006) , July, 78-83.

Li, D.X. "On Default Correlation: A Copula Approach," *Journal of Fixed Income*, 9 (2000), 43-54.

Longstaff, F., and A. Rajan. "An Empirical Analysis of the Pricing of Collateralized Debt Obligations," *Journal of Finance*, 63, 2 (April 2008), 529-563.

Merton, R.C. "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance*, 29 (1974), 449-470.

Putyatin, V.; D. Prieul; and S. Maslova. "A Markovian Approach to Modelling Correlated Defaults," *Risk*, 18 (2005), May, 65-69.

Schönbucher, P. "Portfolio Losses and the Term Structure of Loss Transition Rates: A New Methodology for Pricing Portfolio Credit Derivatives," Working Paper, ETH Zurich (2005).

Servigny, A., and O. Renault. "Default Correlation: Empirical Evidence," Working Paper, Standard and Poors (2002).

Sidenius, J.; V. Piterbarg; and L. Andersen. "A New Framework for Dynamic Credit Portfolio Loss Modeling," Working paper, Bank of America (2004).

Vasicek, O. "Probability of Loss on a Loan Portfolio," Working Paper, KMV (1987). (Published in *Risk*, December 2002 with the title "Loan Portfolio Value")

Walker, M. "Simultaneous Calibration to a Range of Portfolio Credit Derivatives with a Dynamic Discrete-Time Multi-Step Loss Model," Working paper, University of Toronto.

Exhibit 1

iTraxx CDO tranche quotes January 30, 2007.

a_L	a_H	3 yr	5 yr	7 yr	10 yr
0	0.03	n/a	10.25	24.25	39.30
0.03	0.06	n/a	42.00	106.00	316.00
0.06	0.09	n/a	12.00	31.50	82.00
0.09	0.12	n/a	5.50	14.50	38.25
0.12	0.22	n/a	2.00	5.00	13.75
Index		15.00	23.00	31.00	42.00

CDX NA IG CDO tranche quotes January 30, 2007.

a_L	a_H	3 yr	5 yr	7 yr	10 yr
0	0.03	n/a	19.63	38.28	50.53
0.03	0.07	n/a	63.00	172.25	427.00
0.07	0.10	n/a	12.00	33.75	96.00
0.10	0.15	n/a	4.50	14.50	43.25
0.15	0.30	n/a	2.00	6.00	13.75
Index		19.00	31.00	43.00	56.00

Exhibit 2

Implied jump sizes for iTraxx on January 30, 2007 for the one-parameter model in Section IIIA where the drift is zero and the jump size is constant

a_L	a_H	3 yr	5 yr	7 yr	10 yr
0	0.03	n/a	0.0247	0.0221	0.0221
0.03	0.06	n/a	0.0120	0.0054	0.2378
0.06	0.09	n/a	0.0336	0.0268	0.0112
0.09	0.12	n/a	0.0578	0.0501	0.0340
0.12	0.22	n/a	0.0981	0.0900	0.0749

Exhibit 3: Implied jump size using the one-parameter model in Section IIIA compared with the implied tranche (i.e., compound) correlation from the Gaussian copula model for 5-, 7-, and 10-year tranches of iTraxx on January 30, 2007.

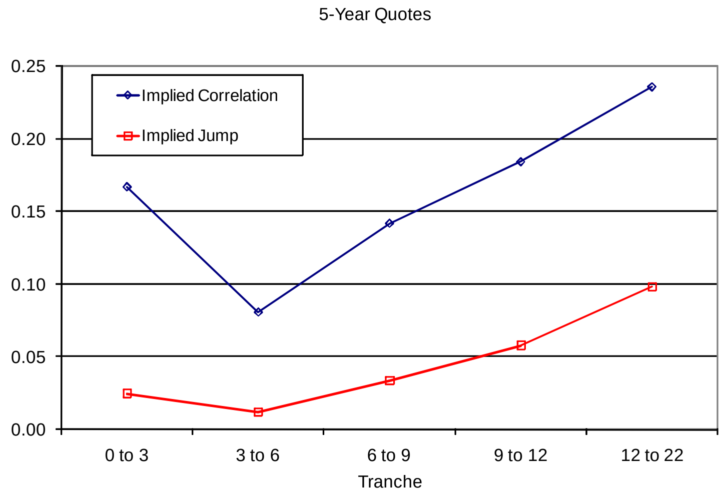

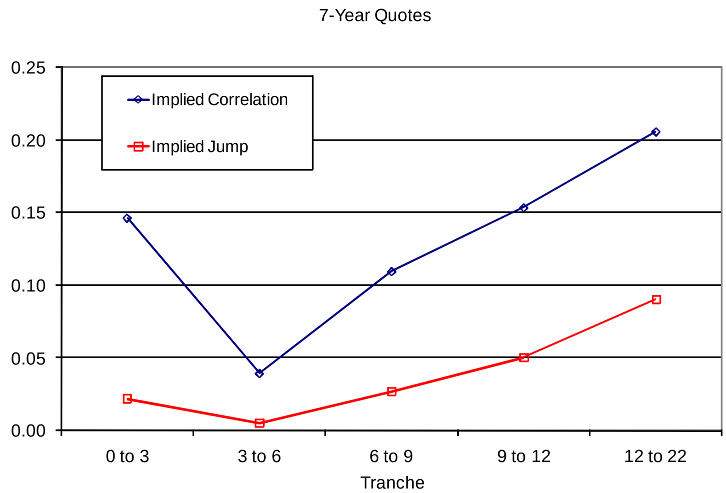

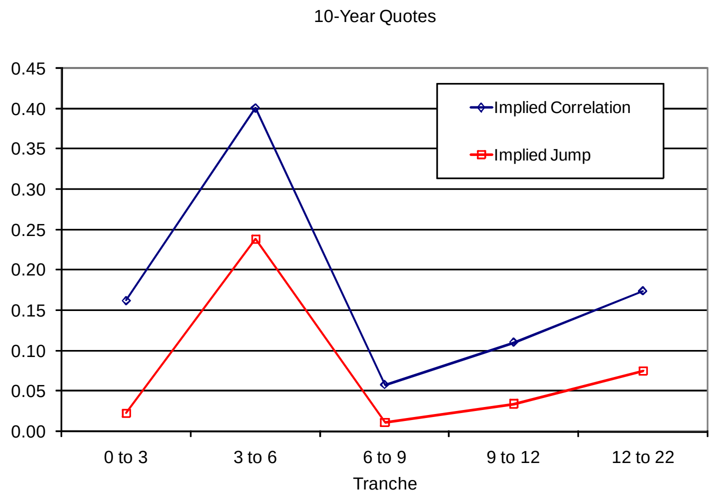

Exhibit 4

Errors resulting from calibration of three-parameter model in Section IIIC to the iTraxx data in Exhibit 1 for January 30, 2007. (For example, the quote for the 3% to 6% 10-year tranche is 316 and the best fit spread is 314.63.)

a_L	a_H	3 yr	5 yr	7 yr	10 yr
0	0.03	n/a	1.34	2.75	4.32
0.03	0.06	n/a	0.37	3.12	–1.37
0.06	0.09	n/a	–0.54	–2.69	–1.92
0.09	0.12	n/a	–1.01	–1.55	–0.12
0.12	0.22	n/a	–0.47	–0.21	1.28
Index		0.00	0.00	0.00	0.00

Errors resulting from calibration of the three-parameter model in Section IIIC to the CDX NA IG data in Exhibit 1 for January 30, 2007.

a_L	a_H	3 yr	5 yr	7 yr	10 yr
0	0.03	n/a	1.63	3.20	2.85
0.03	0.07	n/a	–4.01	–2.16	1.99
0.07	0.10	n/a	2.30	3.39	2.51
0.10	0.15	n/a	4.00	4.79	1.44
0.15	0.30	n/a	0.69	1.28	5.55
Index		0.00	0.00	0.00	0.00

Exhibit 5: Unconditional loss distribution for iTraxx on January 30, 2007 at four maturities. Results are based on the three-parameter model in Section IIIC calibrated to the market data in Exhibit 1.

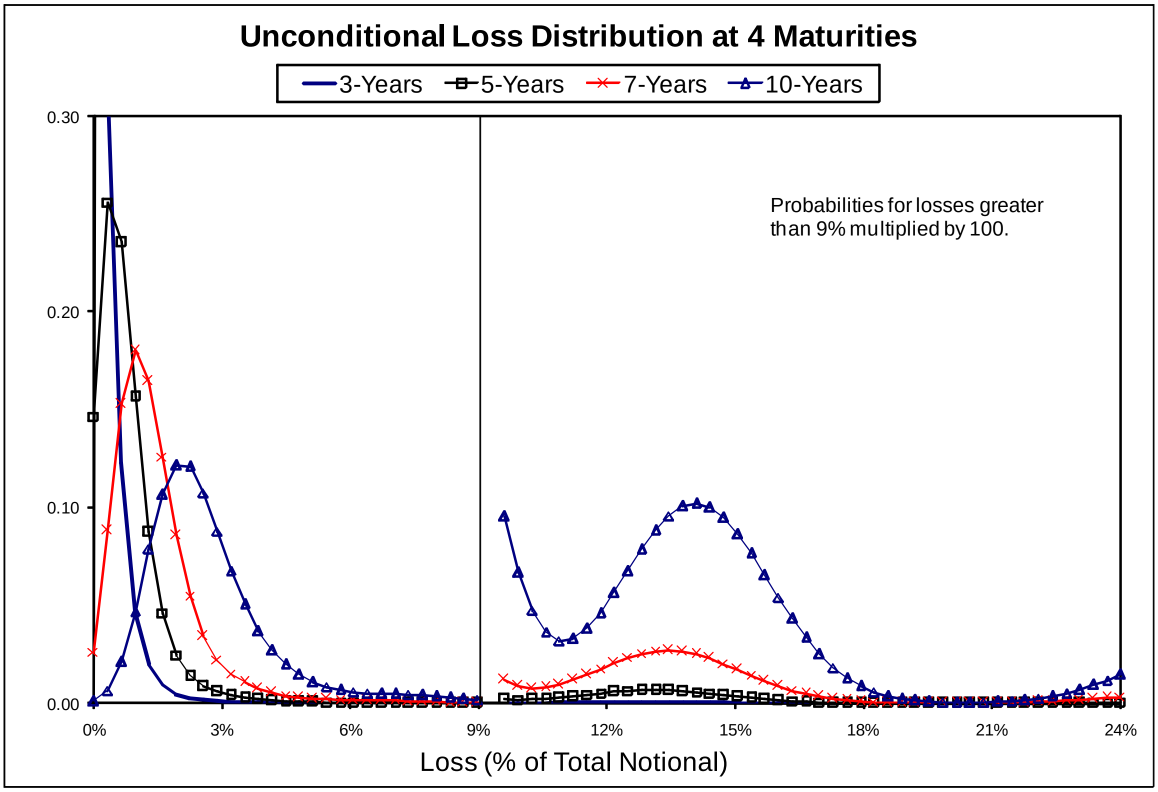

Exhibit 6: Parameters H_0, β and λ estimated for iTraxx for the model in Section III C between July 4, 2006 and January 10, 2007.

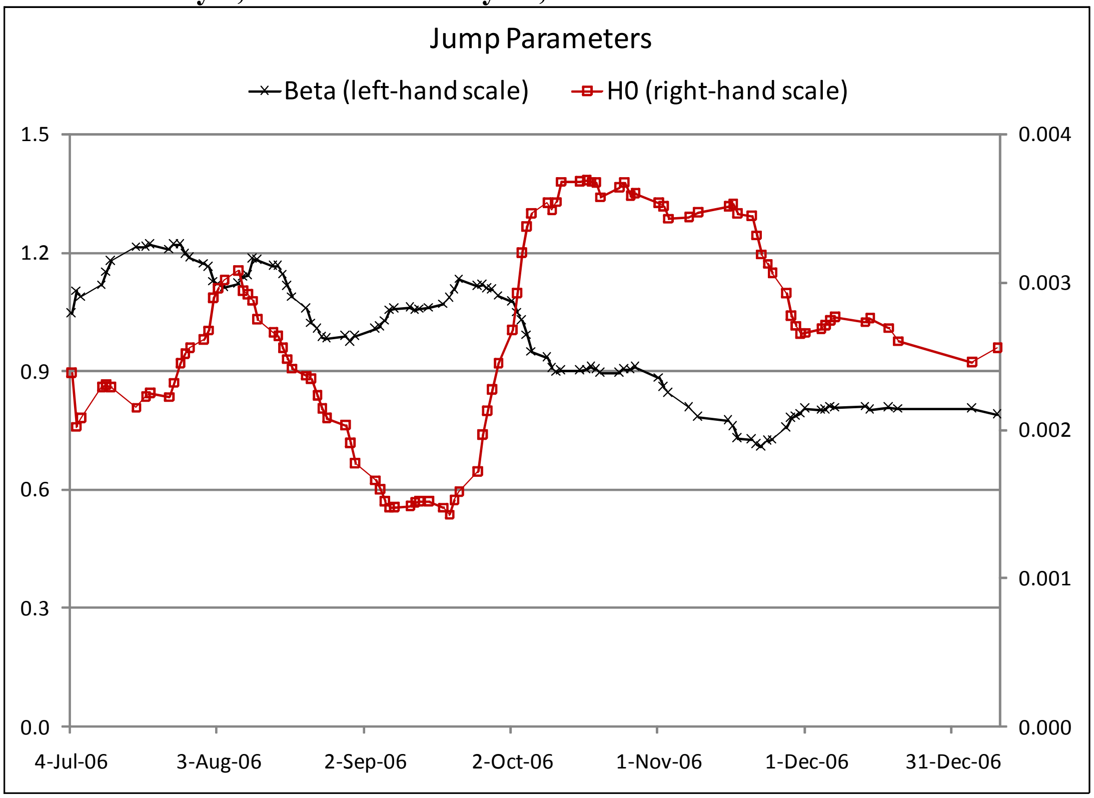

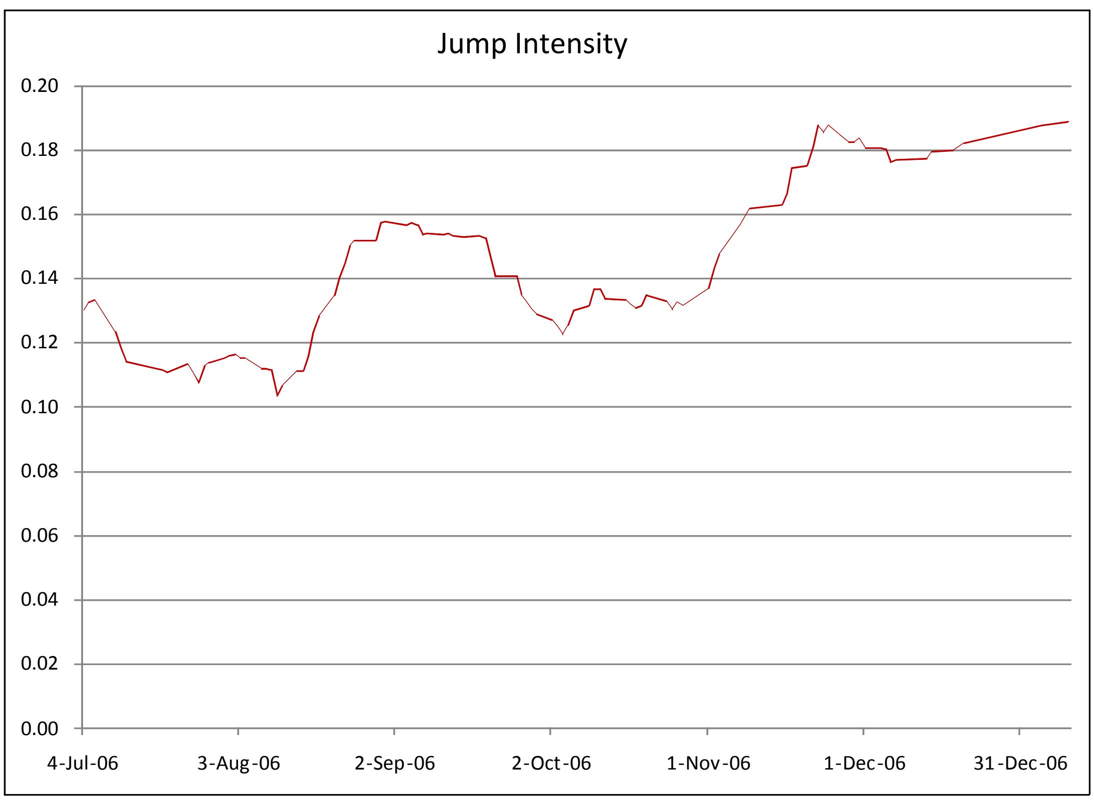

Exhibit 7

Breakeven tranche spread for forward start CDO tranches on iTraxx on January 30, 2007. The tranches mature in five years. Results are based on the three-parameter model in Section IIIC calibrated to the market data in Exhibit 1.

Tranche Start		1.0	2.0	3.0	4.0	4.5
a_L	a_H	Breakeven Tranche Spreads				
0	0.03	11.5	11.3	11.1	6.4	3.4
0.03	0.06	54.0	70.1	93.2	124.4	144.2
0.06	0.09	14.7	19.4	26.1	35.2	40.7
0.09	0.12	5.8	7.7	10.6	14.8	17.5
0.12	0.22	2.0	2.6	3.7	5.3	6.3
Index		25.3	29.1	36.7	41.4	43.7

Exhibit 8

Prices in basis points of at-the-money European options on iTraxx CDO tranches on January 30, 2007. The tranches mature in five years. Results are based on the three-parameter model in Section IIIC calibrated to the market data in Exhibit 1.

		Option Expiry in Years				
a_L	a_H	1.0	2.0	3.0	4.0	4.5
0.03	0.06	67.8	91.3	89.7	68.3	41.4
0.06	0.09	23.2	29.8	30.7	23.1	13.5
0.09	0.12	9.7	12.2	13.3	10.0	6.1
0.12	0.22	3.7	4.4	5.0	3.8	2.4

Exhibit 9

Implied volatilities of at-the-money European style options on iTraxx CDO tranches on January 30, 2007. The tranches mature in five years. Results are based on the three-parameter model in Section IIIC calibrated to the market data in Exhibit 1.

		Option Expiry in Years				
a_L	a_H	1.0	2.0	3.0	4.0	4.5
0.03	0.06	96.1%	100.9%	96.8%	104.3%	107.5%
0.06	0.09	123.2%	122.6%	125.6%	137.2%	135.3%
0.09	0.12	133.0%	128.6%	139.4%	144.8%	148.9%
0.12	0.22	149.3%	137.3%	160.5%	160.6%	181.8%

Exhibit 10

Implied volatilities for 2-year European options on iTraxx CDO tranches for strike prices between 75% and 125% of the forward spread on January 30, 2007. The tranches mature in five years. Results are based on the three-parameter model in Section IIIC calibrated to the market data in Exhibit 1.

	CDO Tranche			
K / *F*	3 to 6%	6 to 9%	9 to 12%	12 to 22%
0.75	100.1%	125.6%	132.9%	143.5%
0.80	100.6%	125.2%	132.1%	142.3%
0.85	100.9%	124.7%	131.3%	141.1%
0.90	101.1%	124.1%	130.5%	139.9%
0.95	101.0%	123.4%	129.5%	138.6%
1.00	100.9%	122.6%	128.6%	137.3%
1.05	100.6%	121.7%	127.5%	136.0%
1.10	100.3%	120.9%	126.5%	135.2%
1.15	99.8%	119.9%	125.4%	135.3%
1.20	99.3%	119.0%	125.0%	136.0%
1.25	98.8%	118.0%	124.8%	137.1%

Exhibit 11: Four-step binomial tree for the variable X, which is minus the log of the cumulative survival probability. M_i is the value of X at time τ_i when there are no jumps; H_J is the size of the Jth jump. The probability on the upper and lower branches emanating from a node at time τ_i are $\lambda_i\Delta_i$ and $1- \lambda_i\Delta_i$, respectively, where $\lambda_i = \lambda(\tau_i)$ and $\Delta_i = \tau_{i+1} - \tau_i$. The value of the cumulative survival probability, S, at any node is $\exp(-X)$.

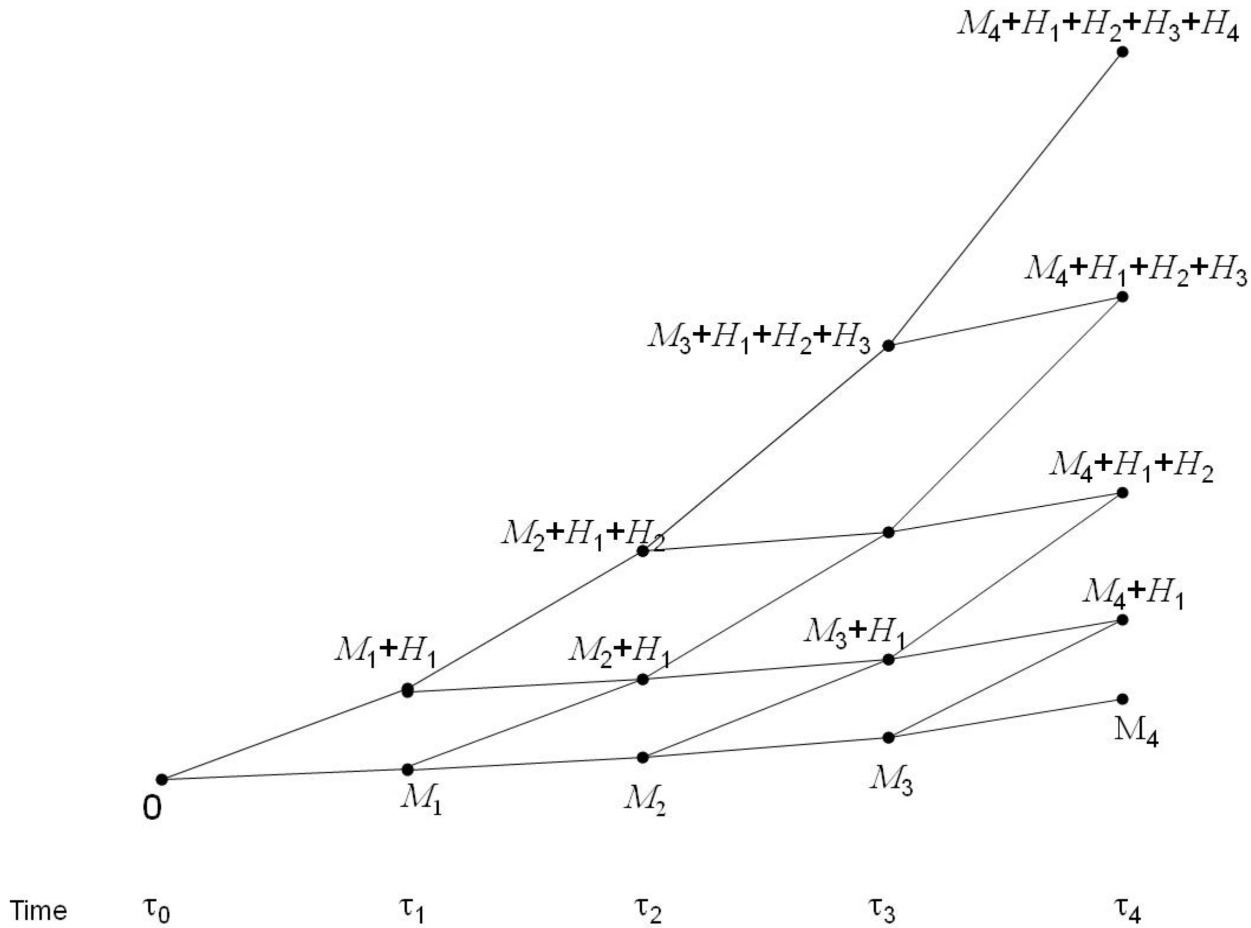

Exhibit 12: Breakeven spread in a leveraged super senior transaction in basis points as a function of the number of defaults required to trigger termination for the iTraxx 12% to 22% tranche on January 30, 2007. Results are based on the three-parameter model in Section IIIC calibrated to the market data in Exhibit 1. The maximum loss borne by the seller of protection is *x*.