Delta-hedging Correlation Risk?

Areski Cousin[1,4]*, Stéphane Crépey[2,4]†, Yu Hang Kan[3,†]

[1] Université de Lyon, Université Lyon 1, LSAF, France

[2] Équipe Analyse et Probabilité,
Université d'Évry Val d'Essonne, 91025 Évry Cedex, France

[3] IEOR Department
Columbia University, New York

[4] CRIS Consortium, France

July 16, 2010

Abstract

The Gaussian copula model is essentially a static quotation device, and its use for hedging is, in principle, questionable. The Gaussian copula delta thus assumes a constant tranche correlation, whereas in practice this correlation is dynamic, and correlated in particular to the credit index. It might therefore be expected that a dynamic model of credit risk, able to capture at least part of the dependence between implied correlation and index spreads, should have better hedging performances. In this paper, we compare two deltas which can be used for hedging a CDO tranche by its credit index: the market or Gaussian copula delta, and the local intensity delta, where the latter refers to the delta in a local intensity default model of portfolio credit risk, recalibrated to the market every day. A theoretical analysis is illustrated by data analysis and backtesting hedging experiments based on both pre-crisis and crisis market data. We observe that hedging performance are comparable for crisis period associated with CDX Series 9 and 10. However, the local intensity delta fails to outperform the market delta in pre-crisis period associated with CDX Series 5, even if the local intensity model is a sound, dynamic model of credit risk, fitting the market over the full set of CDO tranches, as opposed to a static model and a per tranche fit in the case of the Gaussian copula model.

*The research of this author benefited from the support of the DGE and the ANR project Ast&Risk.

†The research of this author benefited from the support of the 'Chaire Risque de crédit', Fédération Bancaire Française.

The authors warmly thank Rama Cont and Jean-Paul Laurent for stimulating discussions throughout the preparation of this work. We are also grateful to the participants at the Quantitative Methods in Finance (QMF) 2010 Conference, the 6th World Congress of the Bachelier Finance Society and the AFFI 2010 St Malo International Meeting for useful comments and discussions.

1 Introduction

A difficulty in financial modeling is the unavoidable gap between markets and their mathematics. There are at least two reasons for this. The first point is the complexity of financial markets, far beyond that of any tractable model. The second point is the scarceness of market data that can be used to determine the value of the models' parameters: Historical data for their statistical estimation or prices of liquid instruments for their calibration.

With portfolio credit derivatives these difficulties are exacerbated. Regarding the first point above it is enough to think of the complexity of the 'universe' underlying a CDO (not to mention a CDO of ABSs or CDO square). As for the second point, one must of course mention the rarity of default events and also the small number of liquid instruments (CDO tranches) quoted on a credit index at a given time. Given this uncertainty inherent to credit markets, a particularly important issue in the risk management of credit derivatives is of course that of the robustness of the models and of the hedging strategies.

Another perspective on the complexity of credit markets is that one clearly deals with fully incomplete markets, for which the definition of a price as the initial value of a replicating strategy does not make much sense. Even if a CDO tranche is essentially a vanilla option on the loss on a credit portfolio, one is with a CDO tranche at the complete opposite of a Black–Scholes world in which a trader actually *constructs* the payoff of a vanilla option by dynamic management of its underlying (which is not even traded in the case of a CDO tranche). Also considering the inherent computational complexity of dynamic models of portfolio credit risk, it is overall no surprise if static 'copula' models emerged as the market standard for dealing with portfolio credit derivatives.

1.1 Implied Black–Scholes Volatility Versus Base Correlation

Yet regarding hedging, the use of static copula models is, in principle, questionable. To try to see better why, let us denote by Π the market price process of a financial derivative, of two possible kinds: a vanilla equity index option, or a synthetic CDO tranche written on a credit index (protection bought, say). In the two cases, S refers to the market price of the underlying financial index. At this stage, jumps in the dynamics of S are neglected for simplicity. Let Θ be the first order partial derivative of the Black-Scholes or Gaussian copula pricing function u with respect to time t and Δ_x, $\Gamma_{x,y}$ be (resp.) the first and second order partial derivatives of u with respect to the variables x and y. One then has by application of the Itô formula:

- In the equity case,

$$\begin{aligned} d\Pi_t = du^{bs}(t, S_t, \sigma_t) &= \Theta^{bs}(t, S_t, \sigma_t)dt + \Delta_S^{bs}(t, S_t, \sigma_t)dS_t + \Delta_\sigma^{bs}(t, S_t, \sigma_t)d\sigma_t \\ &+ \frac{1}{2}\Gamma_S^{bs}(t, S_t, \sigma_t)d\langle S\rangle_t + \frac{1}{2}\Gamma_\sigma^{bs}(t, S_t, \sigma_t)d\langle \sigma\rangle_t + \Gamma_{S,\sigma}^{bs}(t, S_t, \sigma_t)d\langle S, \sigma\rangle_t \ , \end{aligned} \quad (1)$$

where σ_t is the option's Black–Scholes implied volatility at time t, and

• In the credit case,

$$
\begin{aligned}
d\Pi_t = du^{gc}(t, S_t, \rho_t) &= \Theta^{gc}(t, S_t, \rho_t)dt + \Delta_S^{gc}(t, S_t, \rho_t)dS_t + \Delta_\rho^{gc}(t, S_t, \rho_t)d\rho_t \\
&+ \frac{1}{2}\Gamma_S^{gc}(t, S_t, \rho_t)d\langle S\rangle_t + \frac{1}{2}\Gamma_\rho^{gc}(t, S_t, \rho_t)d\langle \rho\rangle_t + \Gamma_{S,\rho}^{gc}(t, S_t, \rho_t)d\langle S, \rho\rangle_t \ , \qquad (2)
\end{aligned}
$$

where ρ_t is the tranche's Gaussian copula implied correlation at time t.

Let $e = (e_t)$ with $e_0 = 0$ denote the cumulative tracking error, or profit-and-loss, of an index delta-hedged option position, where the market implied volatility or correlation parameter is used for computing the Black–Scholes or Gaussian copula implied delta. We assume nil interest rates for simplicity. In view of (1), in the equity case with $\sigma_t = \sigma$ constant, the profit-and-loss is identically equal to zero. Moreover, in the general case of a stochastic implied volatility σ_t, (1) immediately yields, using the Black–Scholes equation (for the level of volatility σ_t with zero interest rate),

$$
\begin{aligned}
de_t^{bs} = du^{bs}(t, S_t, \sigma_t) - \Delta_S^{bs}(t, S_t, \sigma_t)dS_t &= \frac{1}{2}\Gamma_S^{bs}(t, S_t, \sigma_t)S_t^2\left(\varsigma_t^2 - \sigma_t^2\right)dt + \Delta_\sigma^{bs}(t, S_t, \sigma_t)d\sigma_t \\
&+ \frac{1}{2}\Gamma_\sigma^{bs}(t, S_t, \sigma_t)d\langle\sigma\rangle_t + \Gamma_{S,\sigma}^{bs}(t, S_t, \sigma_t)d\langle S, \sigma\rangle_t \ , \quad (3)
\end{aligned}
$$

where $\varsigma_t^2 = \frac{d\langle S\rangle_t}{S_t^2 dt}$ is the index realized variance at time t. In the real-market case of a stochastic implied volatility σ_t, the delta-hedged profit-and-loss is thus of course non-zero, but equation (3) at least gives the trader some insights into the drivers of her P&L. Relying on (3), she can thus bet on either the *potential effect*, which is the dominant effect for high-Vega (ATM long-term) options, or the *carriage* effect, which is dominant for high-Gamma (ATM short-term) options. Or, alternatively to betting on the first order Gamma or Vega effect, she can hedge this effect by an auxiliary option hedging position, thus effectively taking bets on the second order Volga or Vanna effects which correspond to the last two terms in (3).
This explains how in spite of its obvious misspecifications and limitations, the Black–Scholes model can be effectively used for hedging, and not only as an implied volatility quotation device. This is reinforced by further robustness properties of the Black–Scholes model. It is thus well known that in case of a vanilla option, trading with a Black–Scholes model for a misspecified but conservative value of the volatility parameter yields a super-hedging strategy (see El Karoui et al. [15]).

In the credit case on the other hand, the situation is much less comfortable. Indeed, even in a virtual case where $\rho_t = \rho$ would be constant, one would have in view of (2) that

$$
de_t^{gc} = du^{gc}(t, S_t, \rho_t) - \Delta_S^{gc}(t, S_t, \rho_t)dS_t = \Theta^{gc}(t, S_t, \rho_t)dt + \frac{1}{2}\Gamma_S^{gc}(t, S_t, \rho_t)d\langle S\rangle_t \ , \qquad (4)
$$

which has no reason to and does typically not vanish (see Section 2.5). Moreover, in the real case of a stochastic implied correlation ρ_t, there is no way to rewrite (2) in such a useful form as (3). This is because the Gaussian copula model is not a consistent dynamic model of portfolio credit risk. In relation to this, there is no established analog in this framework of the before-mentioned Black–Scholes robustness property, as for instance the fact that trading and risk managing a CDO tranche with a conservative value of the Gaussian copula correlation would yield a super-hedging strategy.

1.2 Contents of the Paper

All this illustrates the fact that the Gaussian copula model is essentially a quotation device, and that its use for hedging is, in principle, questionable. Yet since this model has long been and probably still is the market standard for the risk management of credit derivatives (the basic model or variants with random recoveries), it would be nice to find ways to get insights into the delta-hedged profit-and-loss process (e_t^{gc}), beyond the not so helpful Itô formula (2). The reader is referred to, e.g., [21], [20], [27], [25], [8] or [4] for a review of market practices regarding risk management of index CDO tranches.

In the equity case, an alternative to (3) for getting some insight into the delta-hedged profit-and-loss process (e_t^{bs}) consists in introducing a notion of Markovian projection of the market, in the form of a local volatility model [12] recalibrated to the market at every trading time. It is then possible to gain some understanding into (e_t^{bs}) by comparing, on both theoretical and numerical grounds, the P&L of a delta-hedged option's position resulting from using either the implied Black–Scholes (as in (3)) or the local volatility delta (see Derman [13] and Crépey [11]).

In this paper we compare likewise two deltas with respect to the task of delta-hedging a CDO tranche with the related credit index and the savings account: the market or Gaussian copula implied delta Δ^{gc}, and the local intensity delta Δ^{lo}, where the latter stands for the delta of the tranche in a local intensity model recalibrated to the market every day.

The hedging of CDO tranches in local intensity models has been, among others, studied by Frey and Backhaus [17], [18], Laurent, Cousin and Fermanian [22], Cousin, Jeanblanc and Laurent [10]. As far as index CDO tranches are concerned, only few empirical papers analyze the performance of alternative quantitative methods for hedging. Cont and Kan [4] and further Cont, Deguest and Kan [3] perform a backtest study of hedging performances using different frameworks including the base correlation and the local intensity approaches. Ammann and Brommundt [1] investigates the ability of the one-factor Gaussian copula model to hedge iTraxx CDO tranches between June 2004 and September 2007. This empirical study compares the compound and the base correlation methods to hedge an iTraxx tranche with other tranches. They found that base correlation deltas outperforms compound correlation deltas and adjacent tranches give better hedge than other tranches. Cousin and Laurent [9] review the main theoretical and operational issues associated with hedging in the Gaussian copula and the local intensity approaches.

Our mainly empirical paper, in the line of Cont and Kan [4] (see also Cont et al. [3]), was in fact motivated by the recent revision of Basel II regulation on calculation of trading book capital requirement. Indeed, residual risks resulting from dynamic hedging strategies of credit derivatives must now be reflected in the capital charge, which of course makes the question of hedging performance a more than ever topical issue. With respect to the above references, our **contributions** are the following:

- In Cont and Kan [4], a wide variety of models were considered, but one of the conclusions is that essentially two concepts of credit delta emerge: the spread delta, and the jump-to-default delta. We thus focus in the present paper on two portfolio credit risk models, the one-factor Gaussian copula model[1] versus the local intensity model.

[1]We assume a homogenous specification of the Gaussian copula and base correlation calibration, whereas [4] used an inhomogeneous one with compound correlation calibration.

The former approach can be understood as a static version of a multivariate structural credit risk model where default risk is driven by spread risk. Conversely, in the local intensity model, spread risk is governed by default risk and the hedging issue can be fully analyzed in dynamic way.

- We use *crisis* and *pre-crisis* data, whereas [4] only used crisis data, and we identify and characterize two market regimes, *normal or steady* as opposed to *crisis or systemic.*
- We provide a *theoretical explanation* of the relative position of the Gaussian copula delta Δ^{gc} with respect to the local intensity Δ^{lo} depending on the market regime.
- Data analysis and backtesting hedging experiments on one-day (as in [4]) *and five-days* time-lag and hedging rebalancing time are provided. Indeed the time scale of trading and risk managing credit derivatives is typically of the order of one week (5 open days) rather than one day. This led some people to argue that the results of [4] were not reliable as dominated by 'noise' and short-term volatility. Therefore we systematically present all the numerical results for two time horizons, one day (as would be standard for equity derivatives) and 5 days.

The paper is organized as follows. Section 2 presents the set-up and the data through the market prism of the Gaussian copula. The model underlying the dynamic local intensity delta (as opposed to the market Gaussian copula delta) is reviewed in Section 3. In Section 4 we analyze the hedging problem from a theoretical point of view. Backtesting hedging experiments are conducted and discussed in Section 5.

2 Data analysis

Our numerical analysis of hedging performance is conducted on three (on-the-run) series of the 5-years CDX North America Investment Grade index (CDX.NA.IG), namely series 5, 9 and 10. These three series have been chosen to facilitate comparison before and during the credit crisis.

2.1 Standardized CDO Tranches

Let us recall that synthetic CDOs are structured products based on an underlying portfolio of reference entities subject to credit risk. It allows investors to sell protection on specific risky portions or tranches of the underlying credit portfolio depending on their desired risk-profile. We concentrate our numerical investigation of hedging performance on the most liquid segment of the market, namely CDO tranches written on standard CDS indexes such as the CDX.NA.IG index. As illustrated in Figure 1, the CDX.NA.IG index contains 125 investment grade CDS, written on large European corporations. Market-makers of this index have also agreed to quote standard tranches on these portfolios from equity or first loss tranches to the most senior tranches. Each tranche is defined by its attachment point which is the level of subordination and its detachment point which is the maximum loss of the underlying portfolio that would result in a full loss of tranche notional. The first-loss 0-3% equity tranche is exposed to the first several defaults in the underlying portfolio. This tranche is the riskiest as there is no benefit of subordination but it also offers high returns if no default occurs. The junior mezzanine 3-7% and the senior mezzanine 7-10% tranches are less immediately exposed to the portfolio defaults but the premium received by the protection seller is smaller than for the equity tranche. The 10-15% tranche is the senior tranche, while the 15-30% tranche is the low-risk super senior piece.

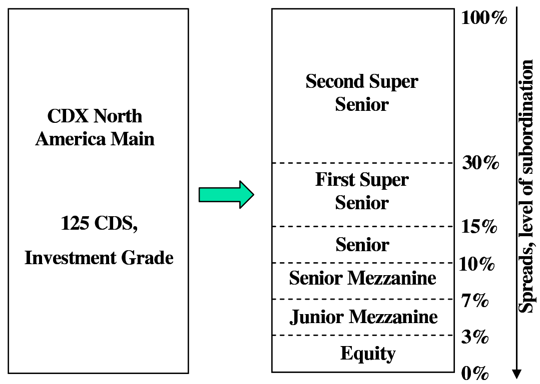

Figure 1: Standardized CDO tranches on CDX.NA.IG.

Considering a pool (portfolio) of n credit names, we denote by τ_i the default time corresponding to the i-th name, by $H_t^i = \mathbb{1}_{\tau_i \leq t}$, $i = 1, \ldots, n$, the related default indicator processes, and by R an homogeneous and constant recovery at default. We define the *cumulative default process* N and the *cumulative loss process* L by $N_t = \sum_{i=1}^n H_t^i$ and $L_t = \frac{1}{n}(1-R)N_t$. Note that L is expressed per unit of nominal exposure (in percentage). The cash-flows associated with a synthetic CDO tranche with attachment point a and detachment point b (a and b in percentage) are driven by losses that affect the tranche, i.e.,

$$L_t^{(a,b)} = (L_t - a)^+ - (L_t - b)^+ .$$

In few words, a CDO tranche is a swapped product with two legs, a default protection leg and a fee leg, and a notion of fair spread Σ_t at time t defined much as in the case of interest-rate swaps, so that the two legs of the contract would have equal values at time t, if the contractual spread was equal to Σ_t. Of course the contractual spread Σ_0 of the contract is fixed once for all at time 0 (starting time of the swap), and at time t the contract has therefore a value which can be positive or negative depending on Σ_t.

We denote by $\mathbb{F} = (\mathcal{F}_t)_{t \in [0,T]}$ a filtration that represents flow of information we use for pricing, and by $\mathbb{E}$ expectation relative to a risk-neutral pricing measure $\mathbb{P}$ on a primary market of hedging instruments. We assume that the short-term interest rate r is constant. Let us recall that the time-t price of a CDO tranche $[a, b]$ maturing at time T can be expressed as a function of conditional discounted loss expectations $\mathbb{E}(e^{-r(s-t)} L_s^{(a,b)} | \mathcal{F}_t)$, for $t \leq s \leq T$. We refer the reader to, for instance, Cousin and Laurent [7] or Cont and Kan [4] for more details on synthetic CDO tranches and related products.

Unless otherwise stated, the short-term interest rate and the recovery rate used to computed numerical results are fixed to (resp.) $r = 3\%$ and $R = 40\%$.

2.2 Data Sets

For numerical illustration throughout this paper, we consider the 5-year CDX.NA.IG index and tranches data in three periods:

- Series 5: 20 September 2005 - 20 March 2006,
- Series 9: 20 September 2007 - 20 March 2008,
- Series 10: 25 March 2008 - 25 September 2008.

Series 5 will be considered in this paper as a representative example of 'normal' or 'idiosyncratic' data, as opposed to the 'systemic crisis data' of Series 9 and 10. On can see on Figure 2 that spreads are low and little volatile in the case of the pre-crisis Series 5 data, increasing and volatile in the Series 9, and high and volatile in the Series 10.

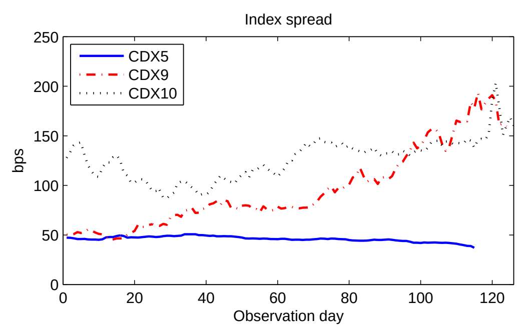

Figure 2: Index spread of 5-year CDX.NA.IG Series 5 from 20 September 2005 to 20 March 2006, Series 9 from 20 September 2007 to 20 March 2008 and Series 10 from 21 March 2008 to 20 September 2008.

2.3 Gaussian Copula Model and Implied Correlations

The one factor Gaussian copula model, that was first introduced in Li [23], is the financial industry quotation standard for multi-name credit derivatives. Without entering into details, let us only mention that at the current time t, the Gaussian copula model parameters are t, a correlation parameter $\rho \in [0, 1]$, and a family $F = (F^i)_{1 \le i \le n}$ of marginal time-to-default cumulative distribution functions over $[t, +\infty)$.

In the rest of the paper, we consider a homogeneous specification of the Gaussian copula model in which the F_t − Gaussian copula model parameter at time t is determined by assuming that all the marginal CDS spread curves at time t are constant and equal to the index spread S_t. More specifically, all marginal default distribution functions at time t are equal to $F_t(x) = 1 - \exp(-xS_t/(1 - R))$. This restriction is motivated by two principal reasons. First, we consider the hedging of standardized CDO tranches by the corresponding CDS index, so marginal effects are not of primary importance although this is the case when single-name CDS are considered for hedging. Second, our aim is to compare this model with the local intensity model in terms of hedging, the latter being a top-down model.

Let us denote by $u^{gc}(T, a, b; t, S_t, \rho)$ and $\Sigma^{gc}(T, a, b; t, S_t, \rho)$ the price (protection bought) and the fair spread of the (T, a, b)-tranche at time t computed in the one-factor Gaussian copula model. Note that these two quantities are described as a function of the index market spread S_t at time t and a correlation parameter ρ. Similarly, we denote by $v^{gc}(T, 0, 100\%; t, S_t)$ the price (protection bought) of the corresponding CDS index at time t. Observe that in the case of the index, the Gaussian copula price does not depend on the correlation parameter.

As the Black(–Scholes) formula on volatility markets, the Gaussian copula model is usually used in the reverse-engineering mode for quoting CDO tranches in terms of their *Gaussian copula implied correlations*. More precisely, if $\Sigma_t^{ma}(T, a, b)$ refers to the market spread of standard tranche $[a, b]$ at time t, the *base correlation* of that tranche is the value of the correlation ρ_t in a Gaussian copula model such that

$$\Sigma^{gc}(T, 0, b; t, F_t, \rho_t) = \Sigma_t^{ma}(T, 0, b)\,, \tag{5}$$

where $\Sigma_t^{ma}(T, 0, b)$ denotes a synthetic market spread computed from the observed market spreads of tranches with detachment point $\leq b$ (see, e.g., [26]).

Figure 3 shows the base correlation at 3% strike during the three sample periods. Observe in particular the increase of the correlation during the sub-prime crisis (Series 9 and 10).

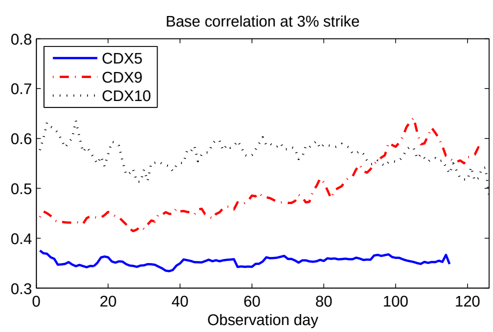

Figure 3: Base correlation at 3% strike of 5-year CDX.NA.IG Series 5 from 20 September 2005 to 20 March 2006, Series 9 from 20 September 2007 to 20 March 2008 and Series 10 from 21 March 2008 to 20 September 2008.

We define the steepness of a base correlation skew at time t as

$$\nabla\rho_t = \rho_t^{b_{\max}} - \rho_t^{b_{\min}}$$

where $b_{\min}$ and $b_{\max}$ are the minimum and maximum strikes to compute the base correlations. For CDX series, $b_{\min} = 3\%$ and $b_{\max} = 30\%$. In other words, the steepnees of a base correlation skew is the length of the interval that contains all implied base correlations. In a sense, this quantity measures to what extent the model is unable to reproduce simultaneously the price of all standard tranches (using a single correlation parameter). Figure 4 shows the base correlation steepness for CDX series 5, 9 and 10. Note that the steepness of the base correlation skew is even higher in the pre-crisis period of Series 5 compared to Series 9 and 10.

2.4 Correlation Between Index Spreads and Base Correlations

As we will see later on, the usual method for hedging CDO tranches is based on (index) spread sensitivities computed in a one factor Gaussian copula model where the correlation

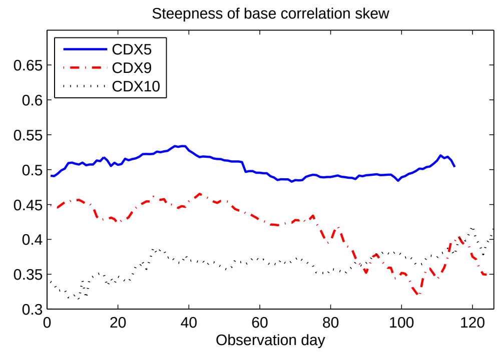

Figure 4: Steepness of the base correlation skew.

parameter has been fixed or arbitrarily updated depending on market trends. As far as one wants to be cover against changes in index spreads, an important question in terms of risk management is whether this risk is correlated with other risks that may significantly affect the price of the product to be hedged. In particular, a fixed or misspecified correlation parameter used when spreads are bumped may have dramatic consequences on hedging performance. We illustrate in this section that, index spreads and base correlations move in the same direction during distressed periods of Series 9 and Series 10 although this tends to be the opposite for the pre-crisis period associated with Series 5.

Remark 2.1 The time scale of trading and risk managing credit derivatives is typically of the order of one week (5 open days) rather than one day. At a one day horizon a credit derivatives business is likely to be dominated by 'noise' and short-term volatility. Therefore we systematically present all the numerical results for two time horizons, one day (as would be standard for equity derivatives) and 5 days.

Table 1 shows the correlation between 1-day/5-day returns of the index spread and the base correlations in each of the sample period. Observe that in the case of the pre-crisis data of 2005 (Series 5), the correlations between 1-day return of the spread and base correlations are close to zero on the 1-day scale, and even negative on the 5-days scale. On the other hand, during the systemic credit crisis of 2007-08 (Series 9 and 10), there is a significant positive correlation between the two. When we increase the observation period from one day to five days, the correlation decreases across all periods and strikes, becoming 'less positive' or 'more negative' than on the 1-day scale.

Figure 5 shows the correlation between the index spread return and the 3% base correlation return based on a 20-day rolling window. Observe that for CDX series 5 and 10, the correlation between the 1-day returns is steady along time. On the other hand, the correlation becomes more volatile when we consider the 5-day returns. Note that for CDX series 9, which is a period of increasing spread, the correlation varies significantly, for either 1-day or 5-day returns.

Market data exhibits two correlation regimes. The pre-crisis period is associated with rather negative correlation levels between changes in index spreads and changes in base correlations although this is the opposite for the more recent crisis periods. In a sense, the topic of this paper is to see whether we can benefit from the use of a dynamic model of credit risk that may capture (albeit partially) the latter observed feature to 'delta-hedge correlation risk',

	1-Day			5-Day		
Strike	CDX5	CDX9	CDX10	CDX5	CDX9	CDX10
3%	-0.03	0.30	0.55	-0.30	0.07	0.40
7%	0.03	0.50	0.60	-0.22	0.41	0.48
10%	0.05	0.55	0.61	-0.18	0.45	0.50
15%	0.07	0.62	0.63	-0.15	0.52	0.51
30%	0.10	0.65	0.61	-0.11	0.62	0.50

Table 1: Correlations between returns of the index spread and the base correlations.

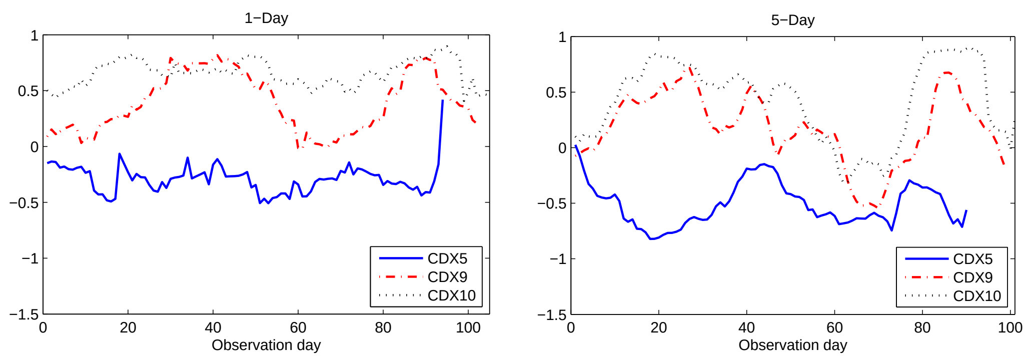

Figure 5: Correlation between index spread return and 3% base correlation return based on 20-day rolling window. Left: 1-day returns; Right: 5-day returns.

or at least, the component of this risk which is correlated to the index.

2.5 Greeks in the Gaussian Copula Model

Let us recall that the Itô decomposition (2) of tranche prices involves partial derivatives (or Greeks) of the Gaussian copula pricing function with respect to its parameters:

$$\begin{aligned} d\Pi_t = du^{gc}(t, S_t, \rho_t) &= \Theta^{gc}(t, S_t, \rho_t)dt + \Delta_S^{gc}(t, S_t, \rho_t)dS_t + \Delta_\rho^{gc}(t, S_t, \rho_t)d\rho_t \\ &+ \frac{1}{2}\Gamma_S^{gc}(t, S_t, \rho_t)d\langle S\rangle_t + \frac{1}{2}\Gamma_\rho^{gc}(t, S_t, \rho_t)d\langle\rho\rangle_t + \Gamma_{S,\rho}^{gc}(t, S_t, \rho_t)d\langle S, \rho\rangle_t \ . \end{aligned} \quad (6)$$

In what follows, we consider the price of a CDO tranche from the point of view of the protection buyer. We aim at investigating the empirical validity of the latter decomposition over a 6-month period corresponding to CDX series 9. In particular, we check that changes in tranche prices estimated using a discrete-time version of (6) fairly corresponds to observed changes in price. We also give some insight on the contribution of each term to the increment of tranche price.

The Greeks for the Gaussian copula model are computed by finite differencing:

$$
\begin{aligned}
\widehat{\Theta}^{gc}(t_j, S_{t_j}, \rho_{t_j}) &= \left[u^{gc}(t_j + \delta t, S_{t_j}, \rho_{t_j}) - u^{gc}(t_j, S_{t_j}, \rho_{t_j})\right]/\delta t,\\
\widehat{\Delta}^{gc}_S(t_j, S_{t_j}, \rho_{t_j}) &= \left[u^{gc}(t_j, S_{t_j} + \delta S, \rho_{t_j}) - u^{gc}(t_j, S_{t_j}, \rho_{t_j})\right]/\delta S,\\
\widehat{\Delta}^{gc}_\rho(t_j, S_{t_j}, \rho_{t_j}) &= \left[u^{gc}(t_j, S_{t_j}, \rho_{t_j} + \delta\rho) - u^{gc}(t_j, S_{t_j}, \rho_{t_j})\right]/\delta \rho,\\
\widehat{\Gamma}^{gc}_S(t_j, S_{t_j}, \rho_{t_j}) &= \left[u^{gc}(t_j, S_{t_j} + \delta S, \rho_{t_j}) - 2u^{gc}(t_j, S_{t_j}, \rho_{t_j}) + u^{gc}(t_j, S_{t_j} - \delta S, \rho_{t_j})\right]/(\delta S)^2,\\
\widehat{\Gamma}^{gc}_\rho(t_j, S_{t_j}, \rho_{t_j}) &= \left[u^{gc}(t_j, S_{t_j}, \rho_{t_j} + \delta\rho) - 2u^{gc}(t_j, S_{t_j}, \rho_{t_j}) + u^{gc}(t_j, S_{t_j}, \rho_{t_j} - \delta\rho)\right]/(\delta \rho)^2,\\
\widehat{\Gamma}^{gc}_{S,\rho}(t_j, S_{t_j}, \rho_{t_j}) &= \left[u^{gc}(t_j, S_{t_j} + \delta S, \rho_{t_j} + \delta\rho) - u^{gc}(t_j, S_{t_j}, \rho_{t_j} + \delta\rho) - u^{gc}(t_j, S_{t_j} + \delta S, \rho_{t_j})\right.\\
&\quad \left. + \, u^{gc}(t_j, S_{t_j}, \rho_{t_j})\right]/(\delta S \delta \rho)
\end{aligned}
$$

where $\delta t = 1$ day, $\delta S = 1$bp and $\delta \rho = 0.1\%$. Then, we approximate the variations of the spread and correlation by

$$
\begin{aligned}
dt &\approx t_{j+k} - t_j,\\
dS_t &\approx S_{t_{j+k}} - S_{t_j},\\
d\langle S \rangle_t &\approx (S_{t_{j+k}} - S_{t_j})^2,\\
d\langle \rho \rangle_t &\approx (\rho_{t_{j+k}} - \rho_{t_j})^2,\\
d\langle S, \rho \rangle_t &\approx (S_{t_{j+k}} - S_{t_j})(\rho_{t_{j+k}} - \rho_{t_j})
\end{aligned}
$$

where the time lag k (hedging period) will be set to 1 day or 5 days.

Figure 6 shows the decomposition of the changes in the equity tranche value of CDX series 9 due to first and second order changes in index spread and base correlation. As visible in Figure 7, actual changes in the equity tranche and changes in the equity tranche estimated by summation of all terms of this Itô formula are very similar. This suggests that if index spreads and base correlations would correspond to market prices of some tradable assets, it would be then possible to design a perfect hedge using the one-factor Gaussian copula model and Itô decomposition (6).

Observe on Figure 6 that the most influential terms are the first order sensitivities with respect to changes in spread and base correlation. Note that these two sensitivities are about of the same order of magnitude, so that hedging spread risk only in the Gaussian copula model leaves the trader with a significant exposure to correlation risk (see Section 2.4). The second order changes of equity tranche value with respect to changes in spread also contribute a substantial amount of volatility, even if it may be partly due to index jumps, that effectively contribute to this term in our decomposition. Moreover, unlike the finding by Cont and Kan [4], the second order changes with respect to spread appear to be consistently negative (for a buy protection position). Let us note that this feature has been proved formally by [24] for an equity tranche. The difference may be due to our specification of a homogeneous Gaussian copula model with base correlation calibration, whereas a heteregeneous specification of the Gaussian copula model was used in Cont and Kan [4] with compound correlation calibration.

Remark 2.2 This backtesting investigation also illustrates an assertion previously made in the Introduction (see (4)), i.e., the fact that, contrary to the Black-Scholes setting, the term $\Theta^{gc}(t, S_t, \rho_t)dt + \frac{1}{2}\Gamma^{gc}_S(t, S_t, \rho_t)d\langle S \rangle_t$ does not vanish in the Gaussian copula model. Indeed, as can be seen from Figure 6, the previous quantity corresponds (after discretization) to the summation of two terms that are significantly negative.

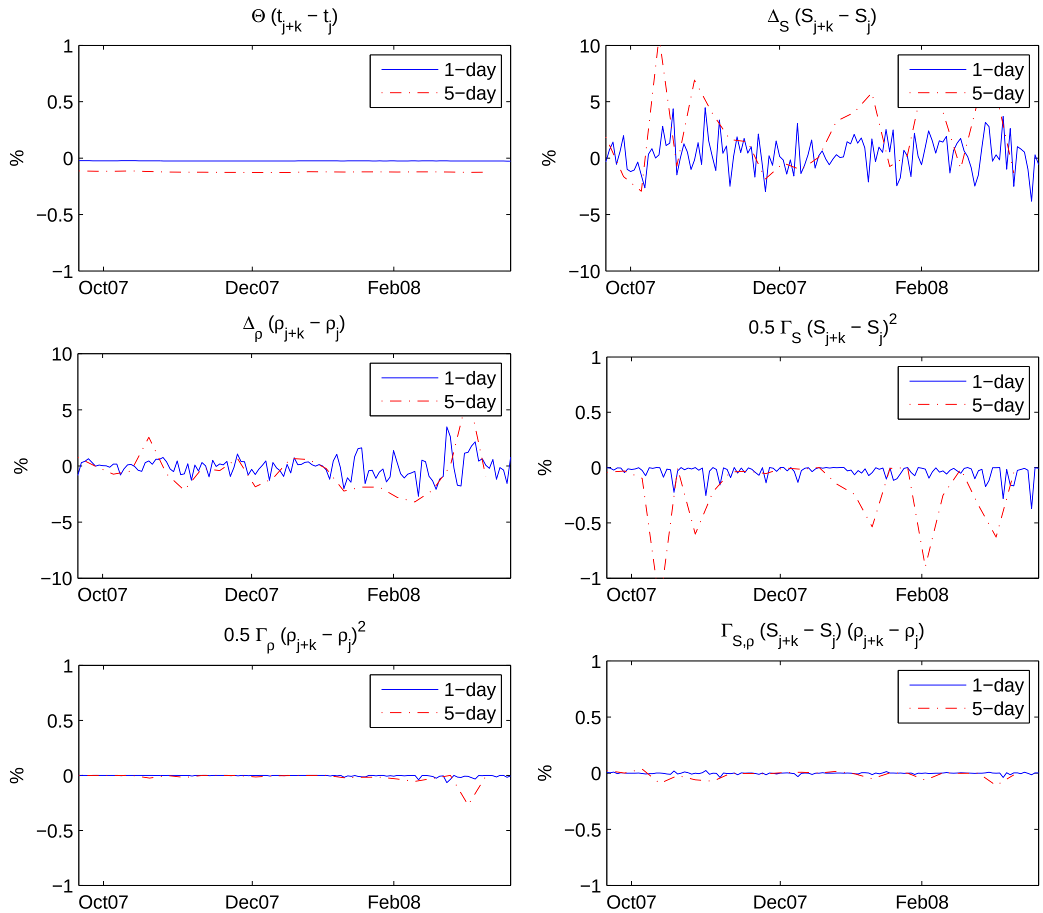

Figure 6: Changes of equity tranche value (100% notional) of CDX series 9 with respect to first and second order changes in time, index spread and base correlation.

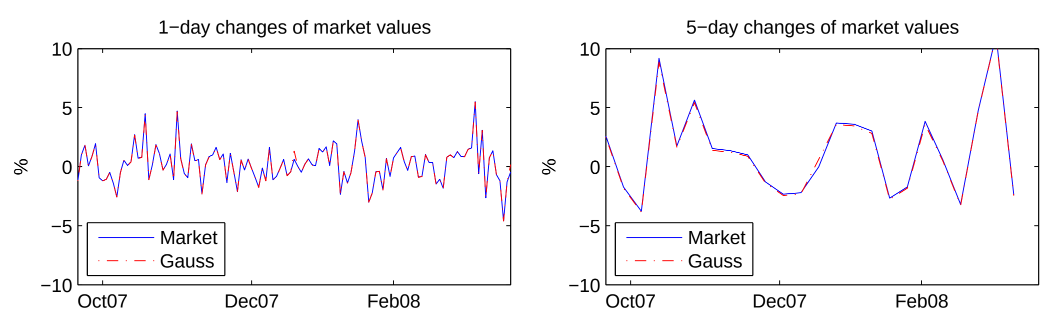

Figure 7: Actual changes in equity tranche (100% notional) of CDX series 9 and the estimation of changes in equity tranche values based on discrete-time approximation of Itô formula under the Gaussian copula model.

3 Local Intensity Model

In the Local Intensity Model (see, e.g., Laurent, Cousin and Fermanian [22], Cont and Minca [5] or Cont, Deguest and Kan [3]), the cumulative number of defaults $N = \{N_t; t \geq 0\}$ of a credit portfolio of n names is modeled as a Markov point process. More specifically, we assume that $N_0 = 0$ and that N is a pure birth process with local intensity $\lambda(t, N_t)$ given by a deterministic function $\{\lambda(t,i)\}_{t\geq 0, i\geq 0}$ satisfying $\lambda(t,i) = 0$ for $i \geq n$. This last condition guarantees that the process N is actually stopped at the level n, as there are n names in the pool. Conditionally on the information $\mathcal{F}_t = \mathcal{F}_t^N$ available at time t, the probability of a jump in the infinitesimal time interval $(t, t+dt)$ is given by $\lambda(t, N_t)dt$. The infinitesimal generator of the process is thus given by the $(n+1) \times (n+1)$ matrix

$$\Lambda_t = \begin{pmatrix} -\lambda(t,0) & \lambda(t,0) & 0 & 0 & 0 \\ 0 & -\lambda(t,1) & \lambda(t,1) & 0 & 0 \\ & & \dots & & \\ 0 & 0 & 0 & -\lambda(t,n-1) & \lambda(t,n-1) \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} .$$

For the sake of notational simplicity, let us consider the default leg of a stylized 'European-type' CDO tranche $[a, b]$ that provides the payoff $\phi(N_T) = (L_T - a)^+ - (L_T - b)^+$ at time T. Let us recall that the aggregate loss is simply proportional to the number of defaults in this top-down approach, i.e., $L_T = (1-R)\frac{N_T}{n}$, where R is the recovery rate. If Π_t denotes the value at time t of this product, we obtain that, for $t \in [0, T]$

$$\Pi_t = \mathbb{E}(\phi(N_T) \,|\, \mathcal{F}_t) = u^{lo}(t, N_t) \; , \tag{7}$$

where $u^{lo}(t, i)$, for $0 \leq t \leq T$, $i = 0, \dots, n$, is the *pricing function*, solution to the system of backward Kolmogorov differential equations

$$(\partial_t + \Lambda_t)u^{lo}(t, \cdot) = 0 \text{ on } [0, T) \, , \quad u^{lo}(T, \cdot) = \phi(\cdot) \; . \tag{8}$$

Moreover we have the following martingale representation, for $t \in [0, T]$:

$$\Pi_t = u^{lo}(t, N_t) = \mathbb{E}(\phi(N_T)) + \int_0^t \Big(u^{lo}(s, N_{s-} + 1) - u^{lo}(s, N_{s-}) \Big) \, dM_s \tag{9}$$

where M is the compensated jump martingale of N, i.e., $dM_t = dN_t - \lambda(t, N_{t-})dt$. Using the analogous martingale representation for the price process $P_t = v^{lo}(t, N_t)$ of the credit index with payoff $p(N_T) = L_T$, it follows that in this model one can replicate the tranche by the index (and the riskless asset) by using the delta in the index defined by, for $t \in [0, T]$,

$$\Delta_t = \Delta_t^{lo} = \Delta^{lo}(t, N_{t-}) \; , \tag{10}$$

with, for $i = 0, \dots, n$,

$$\Delta^{lo}(t, i) = \frac{u^{lo}(t, i+1) - u^{lo}(t, i)}{v^{lo}(t, i+1) - v^{lo}(t, i)} \; . \tag{11}$$

Let us also denote by $S_t = S^{lo}(t, i)$ the CDS index spread computed at time t in the local intensity model where $N_t = i$.

Note that in all the numerics, prices computed in the local intensity model take into account the real cash-flows of traded tranches as opposed to the above stylized presentation. Moreover, in order to be consistent with market conventions, all prices, and also later the deltas, are computed from unit notional value (or 100% notional value). Therefore, both the tranche values and deltas are 'scaled' by the tranche width. It turns out that the price for all tranches and the CDS index is between 0 and 100% and the delta of a $[a, b]$ tranche is between 0 and $\frac{100\%}{b-a}$. For instance, the delta of the $[0, 3\%]$ equity tranche belongs to the interval $[0, 33.33]$.

3.1 Model Calibration

Various methods have been proposed to recover the local intensity function $\lambda(t, i)$ (see for instance Chapter 2 of [10] for a review of such approaches). However, Cont, Deguest and Kan [3] show that even if the local intensity function is calibrated to the same set of market data, model dependent quantities such as the local intensity delta, can be significantly different across the calibration methods. Therefore, we study the local intensity based on two calibration approaches:

- Parametric: A time-homogeneous specification of local intensities extending the one proposed by Herbertsson [19]. Here, the dependence on number of defaults is piecewise linear with grid points corresponding to the attachment points. More specificaly, $\lambda(t, k) = \lambda(k)$, $k = 0, \ldots, 125$ is the linear interpolation at points $k = 0, \ldots, 125$ associated with the constraints $\lambda(125) = 0$ and $\lambda(x_i) = \lambda_i$ for real numbers x_i, $i = 1, \ldots, 6$, such that $\frac{1}{n}(1 - R)x_i = a_i$ where $a_i \in \{0\%, 3\%, 7\%, 10\%, 15\%, 30\%\}$ are the standard CDX.NA.IG attachment points. Note that, at a given quotation date, there is a one-to-one correspondence between model parameters $\lambda_1, \ldots, \lambda_6$ and spreads of tranches and index.

- Non-parametric: Entropy minimization calibration introduced by Cont and Minca [5].

One advantage of the parametric model is that since the local intensity function is assumed to be time-homogeneous, the dependence on defaults can be easily compared by looking at the function at any fixed time. On the other hand, the non-parametric approach, as shown by Cont et al. [3], usually produces an irregular local intensity function which is difficult to compare. However, Cont et al. [3] shows that the non-parametric approach is significantly more stable with respect to shifting in the market spreads.

Table 2 shows the root mean squared errors of the calibrated spreads across quotation dates of each Series. The Gaussian copula model, by design, is well calibrated to the data (one correlation parameter per standard tranche). For the local intensity models, the errors are about 2% for the tranches and about 5% for the index.

More specifically, the first panel of Table 3 shows the CDX series 9 spreads of 20 September 2007, as well as the spreads calibrated to these data in the Gaussian copula model and in the two specifications of the local intensity model. The calibrated spreads are nearly identical for both specifications of the local intensity.

Figure 10 shows further the local intensity functions calibrated by the two approaches on the CDX series 9 data of 20 September 2007, with different levels of 'zoom' on the left tail of the distributions. On a global scale (lower right graph), the two calibrated local intensity functions look completely different. The left tails of the distributions are closer (upper left graph), but are still clearly distinct, and this is in fact most likely this divergence between the left tails which is responsible for the difference between the related hedge ratios to be commented upon below.

	CDX5			CDX9			CDX10		
Tranche	Gauss	Para	EM	Gauss	Para	EM	Gauss	Para	EM
Index	0.04	5.15	5.14	0.03	4.40	4.81	0.02	6.73	6.77
0%-3%	0.01	2.35	2.36	0.00	1.31	1.32	0.01	1.69	1.68
3%-7%	0.00	0.51	0.69	0.00	0.61	0.86	0.00	1.04	1.03
7%-10%	0.00	0.08	1.32	0.00	0.24	0.91	0.00	0.43	0.39
10%-15%	0.00	0.06	1.77	0.00	0.24	1.15	0.00	0.40	0.36
15%-30%	0.00	0.29	1.97	0.01	1.19	1.74	0.01	1.80	1.68

Table 2: Root mean squared errors (in percentage) of calibrated spreads. Gauss: Gaussian copula model; Para: Parametric local intensity model; EM: Local intensity model with entropy minimization calibration.

3.2 Data Revisited

In Section 2, we analyzed the data set in the market scale of the Gaussian copula. Let us now examine the data through the prism of the local intensity model. In order to do so, we need to introduce a concept of correlation as implied by prices computed in the local intensity model.

We then denote by $\rho^{lo}(t, N_t)$ the base correlation implied from the equity tranche price computed in a local intensity model pre-calibrated at time t on market spreads.

Figure 8 shows the time series of the differences between the index spread $S^{lo}(t, N_t)$ (resp. market base correlation $\rho^{lo}(t, N_t)$) and the values implied by the local intensity model with one additional default $S^{lo}(t, N_t + 1)$ (resp. $\rho^{lo}(t, N_t + 1)$). Note that N_t represents the number of defaulted obligors in the underlying portfolio, which is not necessarily equal to zero. In particular, after Fannie Mae and Freddie Mac defaults in 2008, $N_t = 2$ for CDX series 10.

Observe that index spreads implied by the local intensity model with one more default are always higher than the market spreads, so the index spread is increasing in the (at least, first) number of defaults across all the data sets. Now, rather consistently with Table 1 and Figure 5, in the CDX5 sample period, the base correlation implied by the local intensity model with one more default, and hence a greater index spread, is very close to (first part of the time series) or lower than (end of the time series) the market value. On the opposite, in the CDX9 and CDX10 sample periods, the base correlation implied by the local intensity model with one more default, and hence a greater index spread, is always greater than the market value.

3.3 Stability of Hedge Ratios

Figure 9 shows time series of equity tranche deltas for each period. Observe that the deltas computed from entropy minimization (‘entropic deltas’ in an abusive short-hand henceforth)

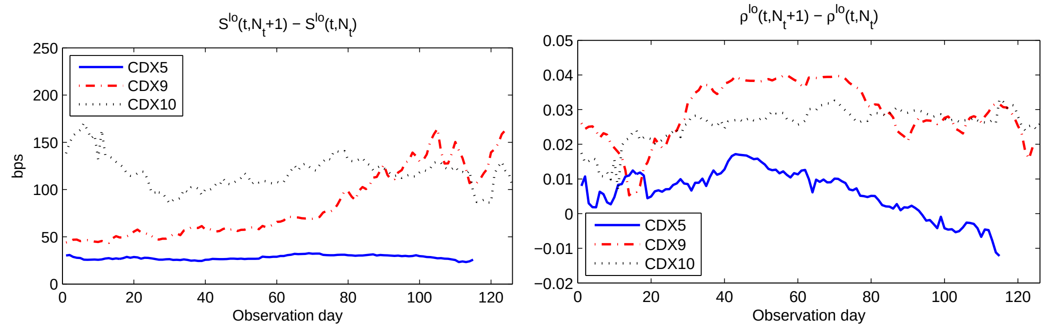

Figure 8: Left: Differences between the index spread $S^{lo}(t,N_t)$ and the values implied by the local intensity model with one additional default $S^{lo}(t,N_t+1)$. Right: Differences between market equity tranche base correlation $\rho^{lo}(t,N_t)$ and the base correlation implied from equity tranche prices computed in the local intensity model with one additional default. Model: Time-homogeneous local intensity model where the dependence of defaults are piecewise linear with grid points close to the attachment points.

is significantly smaller than the parametric deltas throughout the whole time series, though they are both 'local deltas' in local intensity models calibrated to the same data sets. For the sake of comparison, we also provide time series of the industry standard 'Gaussian copula' deltas (defined by equation (13)).

The parametric deltas lies somewhere between the entropic deltas and the Gaussian copula deltas. More specifically the second panel of Table 3 shows the deltas in the Gaussian copula model and in both specifications of the local intensity model, calibrated on the CDX series 9 spreads of 20 September 2007. Note the gap between the entropic delta and the parametric delta, whereas the related calibrated spreads are nearly identical (cf. the first panel).

This is a striking example of model risk. Note that the local intensity model is, in a sense, the simplest dynamic model of credit risk. In case more complex models would be considered, one should be extremely careful about the issue of model risk.

As can be seen from Figure 9, the entropic deltas are substantially more stable than the parametric deltas for CDX series 9 and 10. The stability of the entropic deltas is consistent with the observation by Cont et al. [3] that the local intensity function calibrated by entropy minimization is significantly more stable to the changes in the market data than the parametric local intensity function. Since the computation of the local intensity delta requires the full local intensity function, it is not surprising that the entropic deltas are less sensitive to the changes in the market data than parametric deltas. This difference is more significant in the volatile periods for CDX series 9 and 10.

However, the stability of entropic deltas does not necessarily imply a better hedging strategy. Indeed, the entropic deltas may fail to reflect the market information while they are very much indifferent to the market spreads changes. We will see in Section 5 that the entropic deltas in fact lead to poor hedging performance.

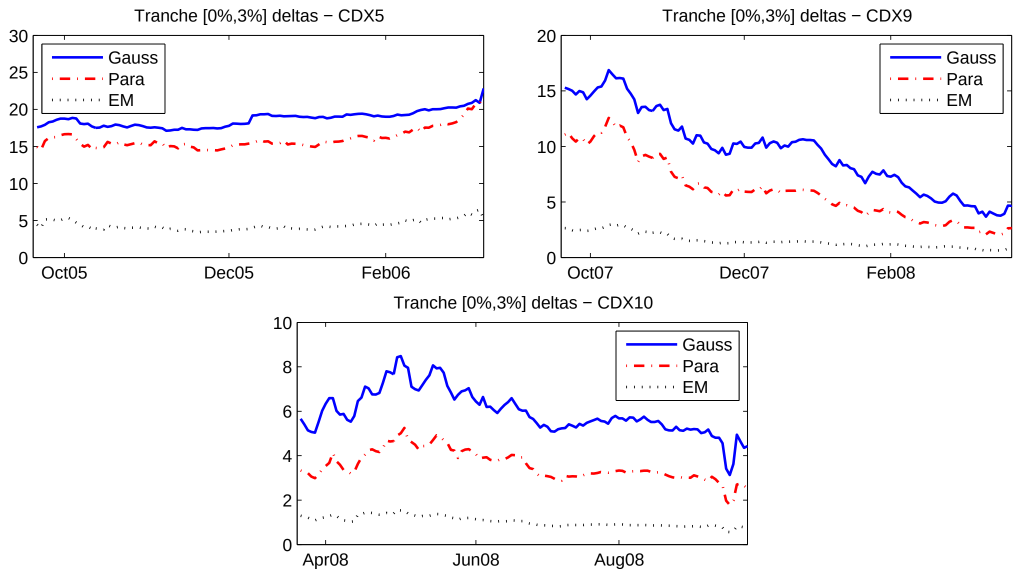

Figure 9: Equity tranche deltas for CDX series 5, 9 and 10.

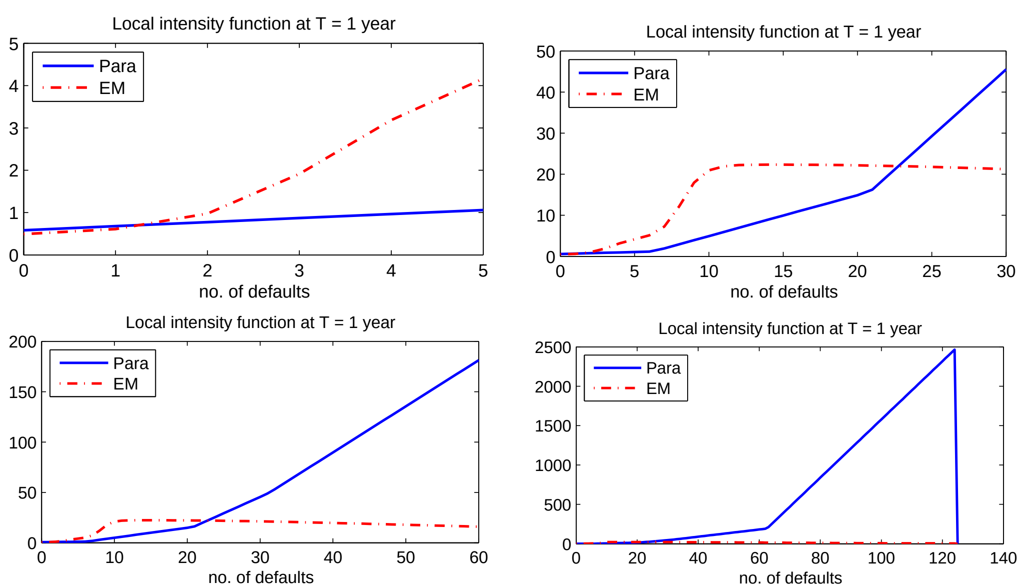

Figure 10: Local intensity function at $T = 1$ year. Data: CDX series 9 on 20 September 2007.

Calibrated spreads

Tranche	Market	Gauss	Para	EM
Index	50.38	50.36	47.58	47.58
0%-3%	35.55	35.55	36.35	36.35
3%-7%	131.44	131.44	132.04	132.07
7%-10%	45.51	45.51	45.54	45.56
10%-15%	25.28	25.28	25.30	25.31
15%-30%	15.24	15.24	15.36	15.36

Hedge ratios

Tranche	Gauss	Para	EM
0%-3%	15.29	11.05	2.64
3%-7%	5.03	4.59	2.70
7%-10%	1.94	2.26	2.29
10%-15%	1.10	1.47	1.99
15%-30%	0.60	1.01	1.74

Table 3: Calibrated spreads and hedge ratios. Data: 5-year CDX series 9 on 20 September 2007.

4 Analysis of the Hedging Problem within a Local Intensity Model

In this section, we consider a theoretical and dynamic market driven by a local intensity model calibrated on market quotes. We identify different market conditions in which deltas computed in the Gaussian copula and in the local intensity frameworks can be ordered. We then propose to compare the P&L increments associated with these two approaches. For the sake of simplicity, we assume in this section that the short rate interest rate is equal to zero.

4.1 Delta-Hedging in Discrete Time

Delta-hedging in discrete time the tranche (long position, that is protection bought, as typical for banks) with the index and the riskless asset over the time interval $[0, T]$, consists in rebalancing in a self-financed way, at every point in time of a subdivision $0 = t_0 \leq t_1 \leq \ldots \leq t_p = T$ of $[0, T]$, a complementary position Δ in the index, in order to minimize the overall exposure to 'small' moves in the index.

The *tracking error*, or *profit-and-loss* process $e = (e_{t_k})_{0 \leq k \leq p}$, is obtained by adding up the following profit-and-loss increments, starting with $e_0 = 0$, from $k = 0$ to $p - 1$:

$$\delta_k e = \delta_k \Pi - \Delta_{t_k} \, \delta_k P \, , \tag{12}$$

where:
• $\delta_k \Pi$ and $\delta_k P$ are the increments of the tranche and index values between times t_k and t_{k+1},
• Δ_{t_k} is the index delta (number of units of index contract in the hedging portfolio over the time interval $(t_k, t_{k+1}]$).

Our aim in this paper is to compare the profit-and-loss processes (e_t) obtained using two strategies, with Δ_t given by:

• The *Gaussian copula implied delta* of the option, that is

$$\Delta_t^{gc} = \frac{u^{gc}(t, S_t + \delta S, \rho_t) - u^{gc}(t, S_t, \rho_t)}{v^{gc}(t, S_t + \delta S) - v^{gc}(t, S_t, \rho_t)}, \tag{13}$$

for $\delta S = 1\text{bp}$ (typically). Here the functions $u^{gc}(t, S, \rho)$ and $v^{gc}(t, S)$ stand for the Gaussian copula pricing function of the tranche and of the index, and the number ρ_t stands for the Gaussian copula implied correlation of the tranche at time t as of (5);

• Or, alternatively, the *local intensity delta* of the tranche, that is (cf. (10), (11))

$$\Delta_t^{lo} = \Delta_i^{lo}(t) = \frac{u^{lo}(t, i+1) - u^{lo}(t, i)}{v^{lo}(t, i+1) - v^{lo}(t, i)} \tag{14}$$

evaluated at $i = N_{t-}$, where $u^{lo}(t, i)$ and $v^{lo}(t, i)$ with $t \in [0, T]$ and $i \in \{0, \ldots, n\}$ refer to the tranche and index pricing functions in the sense of Section 3, in a local intensity model calibrated to the full set of standard tranche quotes at time t.

To conduct our analysis we shall operate in this section in the set-up of a theoretical market given as a local intensity model. Recall from equation (9) in Section 3 that in this set-up, the strategy Δ^{lo}, if applied in continuous time, would provide a perfect replication of the tranche by the index (profit-and-loss process e identically equal to zero), so, for $t \in [0, T]$,

$$d\Pi_t = \Delta_t^{lo} dP_t \ . \tag{15}$$

But in practice we only hedge in discrete time.

As can be seen on Figure 8, all local intensity models calibrated to market quotes exhibit a surge of index spreads at the arrival of a default, the jump in credit spreads being smaller for Series 5 compared with Series 9 and 10. However, a new default tends to increase (equity tranche) base correlation across all Series 9 and 10 whereas it has a smaller (or even negative) impact on correlation for pre-crisis Series 5. Therefore, we distinguish two stylized market regimes: *pre-crisis or steady* market represented by data sample of Series 5 and *crisis or systemic* market represented by data sample of Series 9 and 10. Within each regime and for a given hedging period, we also consider two stylized market scenarios: *widening* and *tightening*, corresponding to values of the index spread increasing and decreasing, respectively.

Let us first consider the case of an equity tranche (protection bought) in a 'systemic' local intensity market, as of those calibrated on Series 9 and 10 (see Section 3.2 above).

4.2 Equity Tranche Convexity

One has that

$$\delta e^{lo} \text{ is negative in tightening scenarios and positive in widening scenarios.} \tag{16}$$

Indeed, (15) yields,

$$\delta_k e^{lo} = \delta_k \Pi - \Delta_{t_k}^{lo} \, \delta_k P = \int_{t_k}^{t_{k+1}} (\Delta_t^{lo} - \Delta_{t_k}^{lo}) \, dP_t \ . \tag{17}$$

Now, it is well known that for an equity tranche, $\Delta^{lo}(t,i)$ is an increasing function of time t (see Figure 11 for a typical example). As we get nearer to maturity, the time-value of both the tranche and the index vanishes. Therefore, the change in value at the arrival of a new default is only the consequence of a protection payment that is essentially the same for the tranche and the index, provided that the tranche is fully in the money. This explains why the deltas tend to $1/0.03 \simeq 33.33$ as time goes to maturity (recall that deltas are computed by unit of nominal exposure).

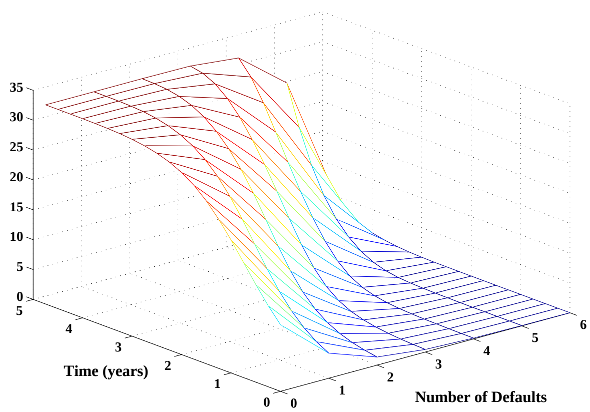

Figure 11: Equity tranche deltas $\Delta^{lo}(t,i)$ computed in a local intensity model as a function of time and number of defaults. Parametric loss intensities calibrated on market spreads of 5-year CDX series 9 on 20 September 2007.

Therefore, on a (small) time interval $[t_k, t_{k+1}]$,
• if *no default* occurs, the change in value of the index is only due to a decrease in time-value, then $\delta_k P \leq 0$. This corresponds to a *tightening scenario*. Then, since from (17) $\delta_k e^{lo} \simeq (\Delta^{lo}_{t_{k+1}} - \Delta^{lo}_{t_k})\,\delta_k P$, the P&L increment $\delta_k e^{lo}$ is negative in this period.
• if *one default* occurs, the decrease in time-value is dominated by a surge in index spreads due to contagion effects, then $\delta_k P \geq 0$. This corresponds to a *widening scenario*. Note that this feature has been checked empirically for all sample periods (see left panel of Figure 8). Then, thanks to representation (17), the P&L increment $\delta_k e^{lo}$ is positive in this period.

4.3 Ordering Between the Equity Tranche Deltas

We have seen in Figure 9 that the local intensity deltas are consistently smaller than the Gaussian copula deltas across the three sample periods, i.e.,

$$\Delta^{lo}_i(t) \leq \Delta^{gc}_t. \tag{18}$$

We attempt here to give some theoretical arguments in favor of this empirical observation. First note that in the case of a systemic local intensity model (similar to those calibrated on Series 9 and 10, see right panel of Figure 8), one has

$$\rho^{lo}(t,i) \leq \rho^{lo}(t,i+1)\;. \tag{19}$$

Moreover it is well known (see, e.g., [6] for a formal proof) that for an equity tranche, one has,

$$\partial_\rho u^{gc}(t, S, \rho) \le 0 \ .$$

So

$$u^{gc}(t, S^{lo}(t, i+1), \rho^{lo}(t, i+1)) \le u^{gc}(t, S^{lo}(t, i+1), \rho^{lo}(t, i)) \ .$$

Now, one has by definition of the Gaussian copula base implied correlation, for every $t \in [0, T]$ and $i \in \{0, \dots, n-1\}$:

$$\begin{aligned} & u^{lo}(t, i+1) - u^{lo}(t, i) = u^{gc}(t, S^{lo}(t, i+1), \rho^{lo}(t, i+1)) - u^{gc}(t, S^{lo}(t, i), \rho^{lo}(t, i)) \\ & = \big(u^{gc}(t, S^{lo}(t, i+1), \rho^{lo}(t, i+1)) - u^{gc}(t, S^{lo}(t, i+1), \rho^{lo}(t, i))\big) \\ & + \big(u^{gc}(t, S^{lo}(t, i+1), \rho^{lo}(t, i)) - u^{gc}(t, S^{lo}(t, i), \rho^{lo}(t, i))\big) \ . \end{aligned} \tag{20}$$

Therefore, by (20) used at $i = N_{t-}$:

$$\Delta_t^{lo} = \Delta_i^{lo}(t) = \frac{u^{lo}(t, i+1) - u^{lo}(t, i)}{v^{lo}(t, i+1) - v^{lo}(t, i)} \tag{21}$$

$$\le \frac{u^{gc}(t, S^{lo}(t, i+1), \rho^{lo}(t, i)) - u^{gc}(t, S^{lo}(t, i), \rho^{lo}(t, i))}{v^{lo}(t, i+1) - v^{lo}(t, i)} \approx \Delta_t^{gc} \ , \tag{22}$$

consistently with (18).

4.4 Equity Tranche P&L Analysis

In view of (12), it follows that

$$\delta e^{gc} = \delta e^{lo} - \left(\Delta^{gc} - \Delta^{lo}\right) \delta P \tag{23}$$

and then

$$\begin{cases} \delta e^{lo} \le \delta e^{gc} & \text{for } \delta P \le 0, \\ \delta e^{lo} \ge \delta e^{gc} & \text{for } \delta P \ge 0. \end{cases} \tag{24}$$

Combining (16) and (24), we get the picture depicted in Table 4. It might thus be so that

Tightening	**Widening**
$\delta e^{lo} \le -(\delta e^{gc})^-$	$(\delta e^{gc})^+ \le \delta e^{lo}$

Table 4: Equity tranche in a systemic local intensity model.

in some scenarios the Gaussian copula delta provides a better hedge than the local intensity delta. Indeed, using simulations of default times in a local intensity model calibrated to CDX Series 9 spreads, we observe in Figure 12 that an increase of the hedging horizon may effectively worsen hedging performance compared to the Gaussian delta.

However, recall that we are in a local intensity model, in which the strategy Δ^{lo}, if applied in continuous time, would provide a perfect replication of the tranche by the index. This means that for hedge rebalancing frequencies large enough (like one week or less) δe^{lo} is very close to 0, and Table 4 reduces to Table 5:

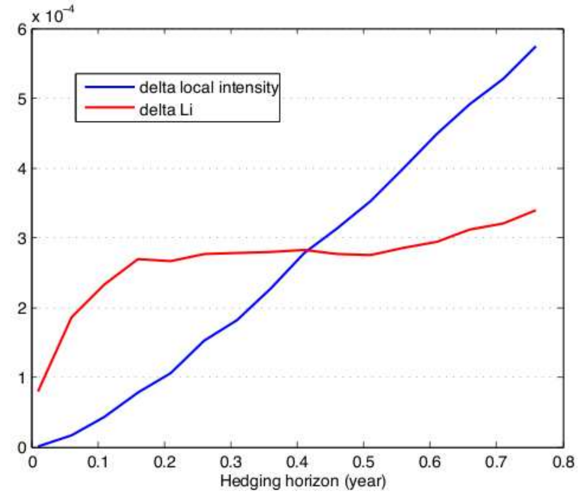

Figure 12: Standard deviation of equity tranche P&L increments δe^{lo} and δe^{gc} as a function of the hedging horizon. Default times are simulated in a local intensity market calibrated on market spreads typical of CDX Series 9.

Tightening	**Widening**
$\delta e^{lo} \simeq 0 \leq \delta e^{gc}$	$\delta e^{gc} \leq 0 \simeq \delta e^{lo}$

Table 5: Case of a moderate to high rebalancing frequency in Table 4.

Remark 4.1 In Section 5, we confirm that the latter ordering of P&L really holds for CDX series 9 (with mainly spread widening periods) when realized (cumulative) P&L are backtested using hedging experiments (see Figure 16).

4.5 Senior Tranche P&L Analysis

The previous analysis focused on an equity tranche in a systemic local intensity model. However, the same analysis can be made for a super senior tranche that protects against last losses, i.e., a CDO tranche $[a, 100\%]$, with $0 < a < 100\%$. We have checked that contrary to equity tranche deltas, senior tranche deltas $\Delta^{lo}(t, i)$ computed in the local intensity model calibrated on Series 9 and 10 are typically decreasing functions of time as far as the tranche is deeply out-of-the-money, as this is typically the case for a $[15\%, 30\%]$ CDO tranche in normal circumstances (see Figure 13).

Then, contrary to equity tranche, the P&L increment δe^{lo} of a super-senior tranche is expected to be positive in tightening scenarios and negative in widening scenarios.

Moreover, the Gaussian copula pricing function associated with this tranche is *increasing* with respect to correlation, i.e., $\partial_\rho u^{gc}(t, S, \rho) \geq 0$ (see, e.g., [6] for a formal proof) and the opposite ordering holds between Gaussian copula deltas and local intensity deltas, i.e., $\Delta_t^{lo} \geq \Delta_t^{gc}$. It turns out that cells in each of the Tables 4 and 5, can be simply exchanged for the senior tranche, leading to Tables 6 and 7.

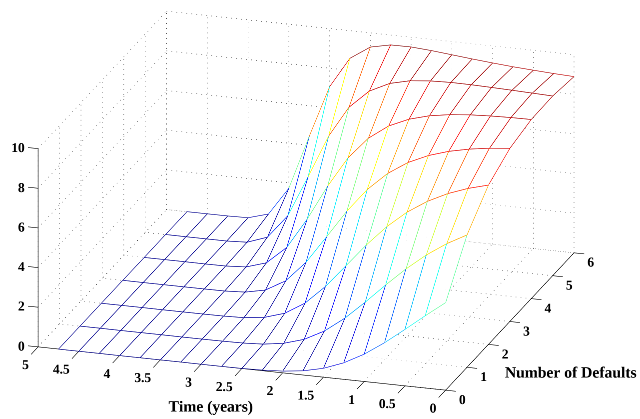

Figure 13: $[15\%, 30\%]$ senior tranche deltas $\Delta^{lo}(t, i)$ computed in a local intensity model as a function of time and number of defaults. Parametric loss intensities calibrated on market spreads of 5-year CDX series 9 on 20 September 2007.

Tightening	**Widening**
$(\delta e^{gc})^+ \leq \delta e^{lo}$	$\delta e^{lo} \leq -(\delta e^{gc})^-$

Table 6: Senior tranche in a systemic local intensity model.

Tightening	**Widening**
$\delta e^{gc} \leq 0 \simeq \delta e^{lo}$	$\delta e^{lo} \simeq 0 \leq \delta e^{gc}$

Table 7: Case of a moderate to high rebalancing frequency in Table 6.

4.6 Analysis in a Steady Market

In the situation of a steady (or pre-crisis) market driven by a local intensity model calibrated on CDX series 5 quotes, the impact of a default may have a negative (or slightly positive) effect on base correlation (see right panel Figure 8) contrary to the crisis period associated with Series 9 and 10. In view of Section 4.3, one may expect that, at least for some quotation dates, the ordering between the two deltas changes during this pre-crisis period. However, as can be seen on Figure 9, the Gaussian copula deltas are consistently greater than the local intensity deltas for the whole Series 5. Therefore, the reverse implication associated with (19) and (22), i.e., the fact that $\rho^{lo}(t,i) \geq \rho^{lo}(t,i+1) \Rightarrow \Delta_i^{lo}(t) \geq \Delta_t^{gc}$, is not satisfied by empirical observation (see also Table 8 for a case study). Nevertheless, the two deltas are quite close in this pre-crisis period with a discrepancy which tends to vanish at the end of Series 5. Then, from relation 23, one may expect that $\delta e^{gc} \simeq \delta e^{lo}$. In Section 5, we confirm that realized P&L of equity and senior tranches are comparable for at least the end of CDX series 5 (see Figure 16).

Example 4.2 Table 8 shows the changes in base correlation of the equity tranche and compares the deltas under the Gaussian copula and the local intensity model.
• On 17 March 2006 (end of Series 5, closest example of a steady market in our data), the base correlation predicted by the local intensity model with one additional default decreases. In this case, the Gaussian copula delta is almost the same as the local intensity delta;
• On 16 September 2008 (next business day after Lehman Brothers defaulted, representative example of a crisis market), the base correlation predicted by the local intensity model increases. In this case, the Gaussian copula delta is significantly larger than the local intensity delta, as suggested by (18).

Date	$\rho^{lo}(t,N_t)$	$\rho^{lo}(t,N_t+1)$	Δ^{gc}	Δ^{lo}
17-Mar-2006	28.55%	27.44%	20.89	20.86
16-Sep-2008	48.27%	51.08%	3.41	1.97

Table 8: Base correlation implied by market data $\rho^{lo}(t,N_t)$, base correlation implied by one additional default in the local intensity model $\rho^{lo}(t,N_t+1)$, Gaussian copula delta Δ^{gc} and the local intensity delta Δ^{lo} for the equity tranche $[0\%,3\%]$ of 5-year CDX series 5 on 17 March 2006 and Series 10 on 16 September 2008.

Remark 4.3 One can image a market regime where the ordering between the two deltas would be inverted for both equity and senior tranches. This might be the case for instance for CDX series 6 but this point has not been checked. Then, the results analogous to those of Tables 4 and 6 are displayed in Table 9. Note that the conclusion is in fact even clearer in that case since the ordering of $P\&L$ holds irrespectively of the frequency of the hedge rebalancing.

Market regime	Tightening	Widening
Equity tranche	$\delta e^{gc} \leq \delta e^{lo} \leq 0$	$0 \leq \delta e^{lo} \leq \delta e^{gc}$
Senior tranche	$0 \leq \delta e^{lo} \leq \delta e^{gc}$	$\delta e^{gc} \leq \delta e^{lo} \leq 0$

Table 9: *Ordering of equity and senior tranches P&L in a steady local intensity model.*

5 Backtesting Hedging Experiments

Let us now check the actual performance of the two deltas by backtesting with historical data. We use the following two metrics to compare the hedging strategies:

$$\begin{aligned} \text{Relative hedging error} &= \left|\frac{\text{Average P\&L increment of the hedged position}}{\text{Average P\&L increment of the unhedged position}}\right|, \\ \text{Residual volatility} &= \frac{\text{P\&L increment volatility of the hedged position}}{\text{P\&L increment volatility of the unhedged position}}. \end{aligned}$$

We consider two cases where the hedging portfolio is rebalanced every day and every 5 days. The profit-and-loss is evaluated in the same frequency as rebalancing. Tables 10 to 13 and Figures 14 and 15 illustrate the hedging performance for 1-day and 5-day rebalancing. Note that in most cases the change in rebalancing frequency does not significantly perturb the comparison with respect to hedging performance.

Unsurprisingly, hedging based on the local intensity model with entropy minimization calibration, which gives admittedly stable but also very low equity tranche hedge ratios in 2007-08, performs worse than other two approaches. Regarding CDX series 5, the Gaussian copula delta outperforms the parametric local intensity delta for nearly all tranches and for both 1-day and 5-day rebalancing. For CDX series 9 and 10, there is no clear evidence to distinguish the performance based on the Gaussian copula model and the parametric local intensity model. For most tranches, the Gaussian copula delta provides a reduction in volatility for about 50% which appear to be larger than those observed by Cont and Kan [4]. This may cause by the differences in the Gaussian copula model implementation in which we assume a homogeneous model with base correlation calibration while Cont and Kan [4] assume a inhomogeneous model with compound correlation calibration. This observation is consistent with the result obtained by Ammann and Brommundt [1] who also find that the base correlation method is better for hedging than the compound correlation method.

On the other hand, the comparison results obtained in Section 4 within the local intensity model for equity and senior tranches are somehow consistent with empirical observations. In Figure 16, we plot the path of (cumulative) P&L associated with hedged and unhedged positions in tranche [0%,3%] and [15%,30%] based on daily rebalancing. One can see that, for the spread widening period of CDX series 9, the P&L orderings predicted in Table 5 (equity tranche) and in Table 7 (senior tranche) behave according to the ordering of (cumulative) P&L trajectories. As for CDX series 5, representative of a steady market period, we also confirm that realized P&L of equity and senior tranches are comparable due to very close deltas.

As a final remark, Table 14 illustrates that defaults among names of the index do not necessarily positively impact intensities, since they may have been anticipated by the market. This observation had already been made in Cont and Kan [4], who also noted that spreads may jump on the other hand at times others then index' names defaults. This was for

Relative hedging error

	CDX5			CDX9			CDX10		
Tranche	Li	Para	EM	Li	Para	EM	Li	Para	EM
0%-3%	4	5	73	80	10	72	33	55	90
3%-7%	1	3	35	0.4	19	59	48	49	75
7%-10%	10	10	43	15	13	37	49	25	44
10%-15%	7	27	131	27	18	14	139	181	208
15%-30%	0.54	61	324	3	32	89	172	269	396

Residual volatility

	CDX5			CDX9			CDX10		
Tranche	Gauss	Para	EM	Gauss	Para	EM	Gauss	Para	EM
0%-3%	45	47	79	59	59	87	105	91	93
3%-7%	70	72	68	58	47	64	85	74	78
7%-10%	90	101	120	53	50	46	83	79	70
10%-15%	90	107	188	61	63	60	91	93	86
15%-30%	93	110	256	37	49	77	84	99	127

Table 10: *Relative hedging error and residual volatility (both in percentage) for 1-day rebalancing. Gauss: Gaussian copula model; Para: Parametric local intensity model; EM: Local intensity model with entropy minimization calibration.*

Relative hedging error

	CDX5			CDX9			CDX10		
Tranche	Gauss	Para	EM	Gauss	Para	EM	Gauss	Para	EM
0%-3%	6	10	77	59	2	73	24	48	88
3%-7%	16	16	51	2	18	58	48	43	72
7%-10%	19	1	15	11	12	36	50	15	41
10%-15%	22	8	75	13	5	5	141	198	209
15%-30%	21	30	207	1	35	86	127	227	382

Residual volatility

	CDX5			CDX9			CDX10		
Tranche	Gauss	Para	EM	Gauss	Para	EM	Gauss	Para	EM
0%-3%	42	46	83	50	56	86	71	72	89
3%-7%	75	75	66	73	65	71	43	40	64
7%-10%	99	118	135	57	56	54	40	38	44
10%-15%	82	110	202	94	98	95	42	44	40
15%-30%	77	108	298	46	69	108	31	33	54

Table 11: *Relative hedging error and residual volatility (both in percentage) for 5-day rebalancing. Gauss: Gaussian copula model; Para: Parametric local intensity model; EM: Local intensity model with entropy minimization calibration.*

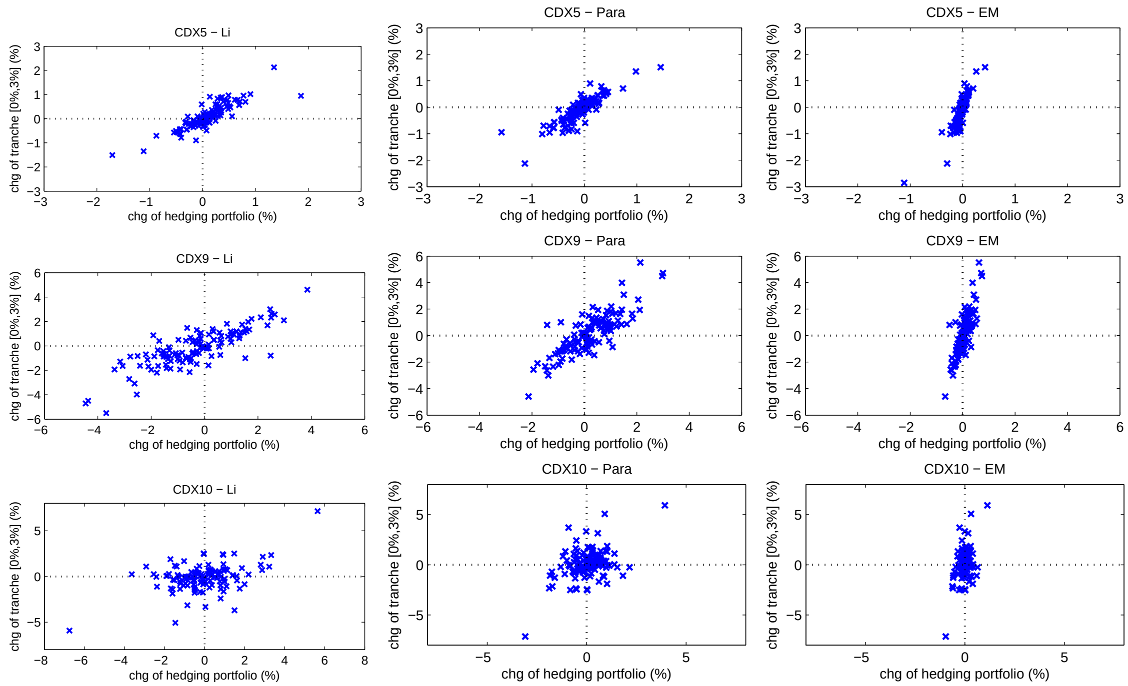

Figure 14: *Daily changes of tranche [0%,3%] value against daily changes of hedging portfolio value (100% notional values).*

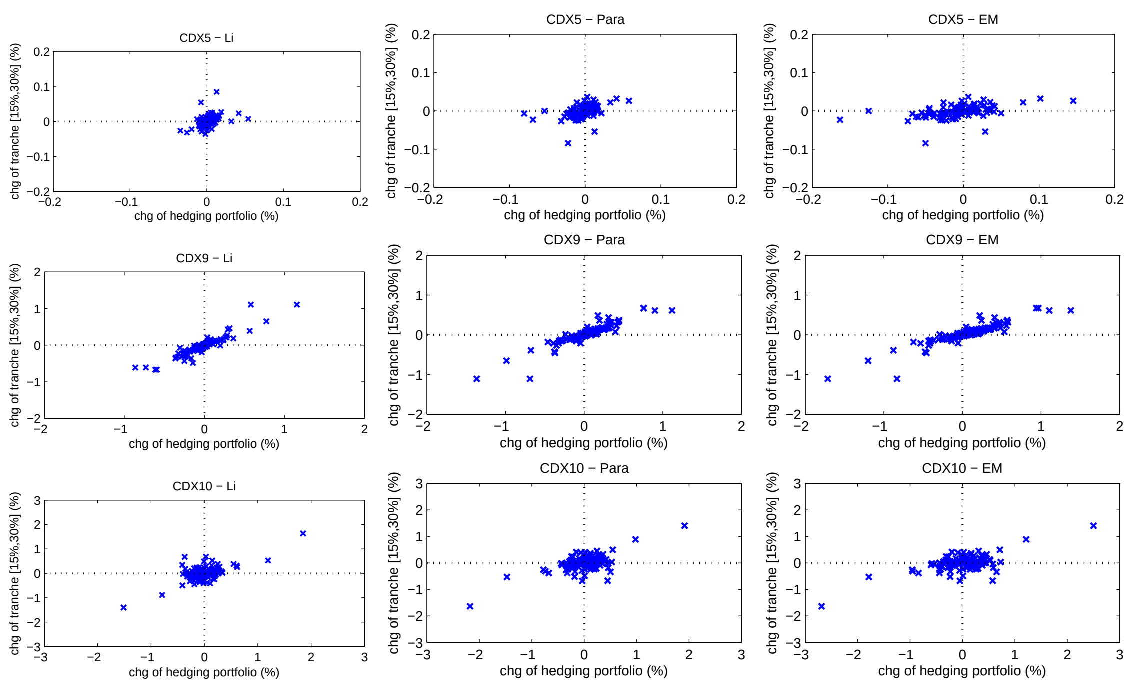

Figure 15: *Daily changes of tranche [15%,30%] value against daily changes of hedging portfolio value (100% notional values).*

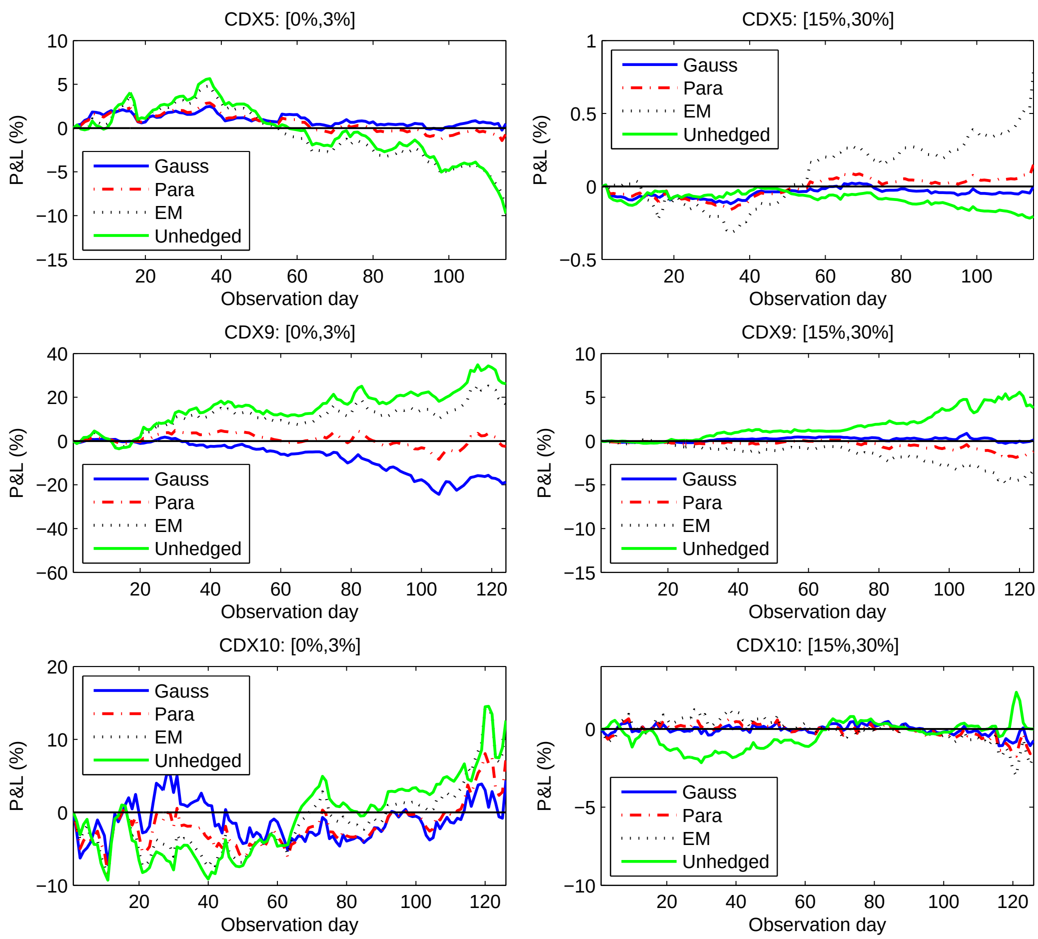

Figure 16: *Path of (Cumulative) P&L of hedged and unhedged positions in tranche [0%,3%] and [15%,30%] based on daily rebalancing (100% notional values).*

Slope estimate: 1-Day

	CDX5			CDX9			CDX10		
Tranche	Gauss	Para	EM	Gauss	Para	EM	Gauss	Para	EM
0%-3%	0.94	1.02	3.44	0.81	1.34	5.01	0.45	0.75	2.44
3%-7%	0.76	0.73	1.15	0.72	0.95	1.88	0.60	0.81	1.53
7%-10%	0.61	0.49	0.39	0.74	0.78	1.14	0.60	0.64	0.99
10%-15%	0.61	0.45	0.25	0.76	0.72	0.78	0.54	0.53	0.58
15%-30%	0.63	0.40	0.15	0.95	0.75	0.58	0.63	0.50	0.39

Slope estimate: 5-Day

	CDX5			CDX9			CDX10		
Tranche	Gauss	Para	EM	Gauss	Para	EM	Gauss	Para	EM
0%-3%	1.10	1.29	4.63	0.93	1.41	5.56	0.82	1.34	4.54
3%-7%	0.69	0.69	1.16	0.69	0.87	1.74	0.86	1.13	2.27
7%-10%	0.50	0.39	0.34	0.79	0.81	1.14	0.84	0.87	1.38
10%-15%	0.63	0.46	0.27	0.55	0.51	0.54	0.83	0.79	0.87
15%-30%	0.76	0.46	0.19	0.79	0.62	0.48	1.10	0.86	0.68

Table 12: *Slope estimates of the OLS regression* $y_i = \alpha + \beta x_i + \varepsilon_i$ *where* y_i *are the 1-day/5-day changes in tranche values and* x_i *is the 1-day/5-day changes in hedging portfolio value. Estimates in italic font (*'good hedges'*) represent the failure to reject the hypothesis* $H_0 : \beta = 1$ *at a 95% confidence level.*

instance the case the day following Lehman's default on September 15 2008, whereas Lehman was not part of CDX10.

6 Conclusions

It appears in Section 5 that the dynamic local intensity model of credit risk does not succeed in outperforming the static Gaussian copula model in terms of hedging CDO tranches by the credit index. Regarding pre-crisis CDX series 5, the Gaussian copula delta outperforms the parametric local intensity delta for nearly all tranches and for both 1-day and 5-day rebalancing. However, for CDX series 9 and 10, there is no clear evidence to distinguish the performance based on the Gaussian copula model and the parametric local intensity model. This can be considered surprising inasmuch as the local intensity model fits the market over the full set of CDO tranches at every point in time, whereas the Gaussian copula model only provides a per tranche fit.

Apparently, the local intensity model, even if calibrated to the market, is definitely 'too far from the market' in dynamic terms, so that the re-calibration shift of the local intensity model badly affects its hedging performances with respect to what one could expect judging by the analysis of section 4 within the set-up of a local intensity model. The fact that a local intensity model is far from the market in dynamic terms is in a sense obvious since it contains no spread risk, which is one of the main sources of risk of credit derivatives (along with default risk, and also recovery risk, the latter being also absent from the local intensity model). But we believe that beyond this obvious fact, the failure of a dynamic model in the context of portfolio credit risk can also be connected with the intrinsic incompleteness

R^2: 1-Day

	CDX5			CDX9			CDX10		
Tranche	Gauss	Para	EM	Gauss	Para	EM	Gauss	Para	EM
0%-3%	0.79	0.77	0.74	0.68	0.69	0.66	0.20	0.19	0.19
3%-7%	0.56	0.56	0.54	0.78	0.78	0.75	0.50	0.47	0.43
7%-10%	0.30	0.30	0.31	0.81	0.81	0.79	0.55	0.53	0.51
10%-15%	0.31	0.32	0.31	0.70	0.70	0.69	0.53	0.52	0.52
15%-30%	0.18	0.19	0.17	0.86	0.85	0.84	0.45	0.44	0.44

R^2: 5-Day

	CDX5			CDX9			CDX10		
Tranche	Gauss	Para	EM	Gauss	Para	EM	Gauss	Para	EM
0%-3%	0.83	0.83	0.80	0.75	0.74	0.77	0.52	0.51	0.52
3%-7%	0.54	0.55	0.57	0.58	0.58	0.60	0.83	0.85	0.84
7%-10%	0.29	0.29	0.30	0.72	0.72	0.71	0.86	0.87	0.87
10%-15%	0.49	0.51	0.51	0.32	0.32	0.33	0.86	0.86	0.86
15%-30%	0.45	0.45	0.45	0.84	0.85	0.87	0.91	0.91	0.91

Table 13: *R^2 of the OLS regression $y_i = \alpha + \beta x_i + \varepsilon_i$ where y_i are the 1-day/5-day changes in tranche values and x_i is the 1-day/5-day changes in hedging portfolio value.*

Index	0%-3%	3%-7%	7%-10%	10%-15%	15%-30%	30%-100%
-1.10	0.16	0.05	0.04	0.44	0.02	-0.06

Table 14: *Daily spread returns on the next business day after Fannie Mae/Freddie Mac defaults on 8 September 2008, normalized by unconditional sample standard deviation of CDX.IG.NA Series 10.*

of credit markets. This incompleteness is in a sense a good reason for traders to favor static models, rather than dynamic models whose interest and justification relies to a large extent on the theory of pricing by replication.
Also, the contagious nature of the default dependence in the local intensity model, does not seem to be that well reflected by the data (see end of section 5).
Maybe another reason of the above results is a 'self-fulfilling' property of the Gaussian copula model, namely the fact that this model, which emerged along the years and still is the market model of portfolio credit risk, more or less 'imposes' its deltas and therefore a consistently related price dynamics.

Incidentally the present work also puts into evidence an important difficulty with dynamic credit risk modeling, namely model risk. That critical issue is of course expected, in regard to the scarceness of market data that can be used for the credit models' calibration. Yet we saw in section 3.3 two specifications of local intensity models giving exactly the same calibrated spreads, but completely different deltas, an unheard-of situation in the context, say, of equity or fixed-income derivatives. If this can happen with local intensity models which are, in a sense, the simplest dynamic models of credit risk, a lesson of the present paper is that in case more complex dynamic models would be considered, one should always be extremely careful about the issue of model risk.

In final word of this paper, we would like to mention a maybe more 'positive' perspective on the above results consisting in a recent attempt to explain a certain robustness of the Gaussian copula model. Here it is worth stressing that this model which, according to some people, 'killed Wall Street', is still the standard for the risk management of credit derivatives (in the form of the basic model or of variants entailing in particular random recoveries). In an interesting 'reverse-engineering' approach, Fermanian and Vigneron thus consider in [16] the question of determining dynamic models of credit risk in which the Gaussian copula delta would happen to be the 'right' delta. Their answer is that such models do exist, and for the purpose of their analysis, they resort to a conditional density (as opposed to intensity) modeling approach developed independently at an abstract level by El Karoui et al. [14]. Interestingly enough, Fermanian and Vigneron's model can also be given a structural interpretation, as opposed to the intensity set-up of the local intensity model, possibly implying that after all, credit markets are maybe more 'structural' (like the Gaussian copula) than driven by 'intensities'.

References

[1] Ammann, M. and Brommundt, B.: Hedging Collateralized Debt Obligations. *Working paper*, University of St Gallen, 2009

[2] Bielecki, T.R. , Cousin, A., Crépey, S. and Herbertsson, A.: Pricing and Hedging Portfolio Credit Derivatives in a Bottom-up Model with Simultaneous Defaults. *In Preparation.*

[3] Cont, R., Deguest, R. and Kan, Y. H.: Recovering Default Intensity from CDO Spreads: Inversion Formula and Model Calibration, *SIAM Journal on Financial Mathematics*, 1 (2010), pp. 555-585

[4] Cont, R. and Kan, Y.H.: Dynamic Hedging of Portfolio Credit Derivatives, *Financial Engineering Report 2008-08*, Columbia University, 2008.

[5] Cont, R. and Minca, A.: Recovering Portfolio Default Intensities Implied by CDO Quotes, *Working Paper*, 2008.

[6] Cousin, A. and Laurent, J.-P.: Comparison Results for Exchangeable Credit Risk Portfolios, *Insurance: Mathematics and Economics*, 42(3), 1118-1127, 2008.

[7] Cousin, A. and Laurent, J.-P.: An Overview of Factor Modeling for Pricing CDO Tranches, *Frontiers In Quantitative Finance, Ed. R. Cont, Wiley Finance*, 2008.

[8] Cousin, A. and Laurent, J.-P.: Hedging Issues for CDOs, in *The Definitive Guide to CDOs*, chapter 17, 461–480, G. Meissner (ed.), Risk Books, 2008.

[9] Cousin, A. and Laurent, J.-P.: Dynamic Hedging of Synthetic CDO Tranches: Bridging the Gap between Theory and Practice, to appear as a book chapter in the volume *Credit Risk Frontiers: Subprime crisis, Pricing and Hedging, CVA, MBS, Ratings and Liquidity, Wiley*, 2010.

[10] Cousin, A., Jeanblanc, M. and Laurent, J.-P.: Hedging CDO tranches in a Markovian environment, Forthcoming in *Paris-Princeton Lectures in Mathematical Finance*, Lecture Notes in Mathematics, Springer, 2010.

[11] S. Crépey:. *Delta-Hedging Vega Risk? Quantitative Finance*, 4 (October 2004), pp. 559–579.

[12] B. Dupire:. Pricing with a Smile, *Risk*, 7 (1994), pp. 18–20.

[13] Derman, E.: Regimes of Volatility. *Risk magazine*, April 1999.

[14] El Karoui, N., Jeanblanc, M., Jiao, Y.: What Happens After a Default: the Conditional Density Approach. Forthcoming in *Stochastic Processes and Applications.*

[15] El Karoui, N., Jeanblanc, M. and Shreve, S.: Robustness of the Black and Scholes formula. *Mathematical Finance* 7 (1997), 1–71.

[16] Fermanian, J.-D., Vigneron, O.: On Break-Even Correlation: The Way to Price Structured Credit Derivatives by Replication. *Risk Magazine*, 2010.

[17] Frey, R., Backhaus, J.: Dynamic hedging of synthetic CDO-tranches with spread- and contagion risk, Forthcoming in *Journal of Economic Dynamics and Control.*

[18] Frey, R., Backhaus, J.: Pricing and hedging of portfolio credit derivatives with interacting default intensities. *International Journal of Theoretical and Applied Finance*, 11 (6), 611-634, 2008.

[19] Herbertsson, A.: Pricing Synthetic CDO Tranches in a Model with Default Contagion using the Matrix-Analytic Approach. *Journal of Credit Risk*, 4(4), 3–35, 2008.

[20] Jobst, N.:, An Introduction to the Risk Management of Collateral Debt Obligations, in *The Handbook of Structured Finance, A. de Servigny and N. Jobst (eds)*, McGraw Hill, 295Ű 338, 2007.

[21] Kakodkar, A., Galiani, S., Jonsson, J. and Gallo, A.: Credit Derivatives Handbook - Vol. 2, Merrill Lynch, 2006

[22] Laurent, J.-P., Cousin, A. and Fermanian, J-D.: Hedging default risks of CDOs in Markovian contagion models. Forthcoming in *Quantitative Finance.*

[23] Li, D.-X.: On default correlation: a copula approach. *Journal on Fixed Income* 9, March, 43–54, 2000.

[24] Mang, C. and Sengupta, A. N.: CDO tranche sensitivities in the Gaussian copula model. *Working Paper*, 2008.

[25] Meissner, G., Hector, R., and T. Rasmussen, Hedging CDOs within the Gaussian Copula Framework, in *The Definitive Guide to CDOs - Market, Application, Valuation, and Hedging*, RISK books, 2008.

[26] O'Kane, D. and Livesey, M.: Base Correlation Explained, *Lehman Brothers Fixed Income Quantitative Credit Research*, Nov 2004.

[27] Petrelli, A., Zhang, J., Jobst, N., and V. Kapoor:,, A Practical Guide to CDO Trading Risk Management, in *The Handbook of Structured Finance*, A. de Servigny and N. Jobst (eds), McGraw Hill, 339Ű371, 2007