Delta-hedging Correlation Risk?

Areski Cousin[1*] (areski.cousin@univ-lyon1.fr)
Stéphane Crépey[2†] (stephane.crepey@univ-evry.fr)
Yu Hang Kan[3,†] (gabriel.kan@gmail.com)

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

[2] Laboratoire Analyse et Probabilités,
Université d'Évry Val d'Essonne, 91025 Évry Cedex, France

[3] IEOR Department
Columbia University, New York

June 6, 2011

Abstract

While the Gaussian copula model is commonly used as a static quotation device for CDO tranches, its use for hedging is questionable. In particular, the spread delta computed from the Gaussian copula model assumes constant base correlations, whereas we show that the correlations are dynamic and correlated to the index spread. It might therefore be expected that a dynamic model of credit risk, which is able to capture the dependence between the base correlations and the index spread, will have better hedging performances. In this paper, we compare delta hedging of spread risk based on the Gaussian copula model, to the implementation of jump-to-default ratio computed from the dynamic local intensity model. Theoretical and empirical analysis are illustrated by using the market data in both before and after the subprime crisis. We observe that delta hedging of spread risk outperforms the implementation of jump-to-default ratio in the pre-crisis period associated with CDX.NA.IG series 5, and the two strategies have comparable performance for crisis period associated with CDX.NA.IG series 9 and 10. This shows that, although the local intensity model is a dynamic model, it is not sufficient to explain the joint dynamic of the index spread and the base correlations, and a richer dynamic model is required to obtain better hedging results. Moreover, although different specifications of the local intensity can be fitted to the market data equally well, their hedging results can be significant different. This reveals substantial model risk when hedging CDO tranches.

* 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', FBF.

The authors warmly thank Rama Cont and Jean-Paul Laurent for stimulating discussions throughout the preparation of this work.

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 model parameters. With portfolio credit derivatives these difficulties are exacerbated. Regarding the first point, it is enough to think of the complexity of the 'universe' underlying a collateralized debt obligation (CDO), not to mention a CDO of ABSs or CDO square. As for the second point, one must mention the rarity of default events and also the small number of liquid instruments, e.g., 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. We refer readers to [5, 9, 21, 22, 26, 28] for a review of market practices regarding risk management of index CDO tranches.

In practice, the most commonly used hedging strategy for CDO tranches is delta hedging small movements in the underlying credit index or credit default swap (CDS) spread based on the Gaussian copula model, which is also known as hedging of spread risk. Its corresponding hedge ratio, the spread delta, is defined as the ratio of the change in tranche value over the change in the underlying index or CDS value with respect to small changes in the underlying index or CDS spread, while assuming constant correlations. However, the Gaussian copula model is essentially a static quotation device and its use for hedging is questionable. As we will show in Section 3, the correlations are dynamic and correlated to the index spread. It might therefore be expected that a dynamic model of credit risk, which is able to capture the dependence between the base correlations and the index spread, will have better hedging performances.

In Cont and Kan [5], a wide variety of dynamic models were considered, and one of the conclusions is that essentially two concepts of hedging strategy emerge: hedging of spread risk and default risk. While delta hedging under the Gaussian copula model as described above is the common strategy for hedging of spread risk, the natural hedge ratio for hedging of default risk is the jump-to-default ratio, which is defined as the change in tranche value over the change in the underlying index or CDS value with respect to one additional default. Therefore, also inspired by the analogous studies of equity derivatives by Derman [13] and Crépey [12] who study hedging under the Black-Scholes model and the local volatility model, our goal is to compare delta hedging of spread risk based on the Gaussian copula model (as an analog to Black-Scholes model in [12, 13]) to the hedging of default risk based on the dynamic local intensity model (as an analog to local volatility model in [12, 13]).

The hedging of CDO tranches in local intensity models has been, among others, studied by Frey and Backhaus [17, 18], Laurent, Cousin and Fermanian [23], Cousin, Jeanblanc and Laurent [11], Cont and Kan [5], and Cont, Deguest and Kan [4]. As far as index CDO tranches are concerned, only few empirical papers analyze the performance of alternative quantitative methods for hedging. Cont and Kan [5] perform a comprehensive backtest of hedging performances using different frameworks including the Gaussian copula model and the local intensity model. 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. In this empirical study, the authors compare the compound and the base correlation methods to hedge an iTraxx tranche with other tranches. They find that hedging based on base correlation method outperforms compound correlation method and that adjacent tranches give better hedge than other tranches. Cousin and Laurent [10] review the main theoretical and operational issues associated with hedging in the Gaussian copula and the local intensity approaches.

With respect to the above references, our contributions include:

- We perform an empirical analysis on the hedging performance by using dataset before and after the subprime crisis, whereas [5] only used data during the crisis. Moreover, we characterize two market regimes in our dataset, normal/steady as opposed to crisis/systemic.
- We provide a theoretical explanation of the relative position of the spread delta with respect to the jump-to-default ratio depending on the market regime.
- We illustrate the results on both one-day and five-day time intervals. While the time scale of trading and risk managing credit derivatives is typically of the order of 5 business days, the results obtained from one-day time interval as in [5] may be dominated by 'noise' and short-term volatility. Therefore we systematically present all the numerical results for two time horizons.
- While Cont, Deguest and Kan [4] study the differences of the hedge ratios implied by the local intensity among different calibration schemes, we also backtest their hedging performance which reveals significant model risk when hedging CDO tranches.

The paper is organized as follows. Section 2 presents a brief introduction to CDS index and CDO tranches, and also the dataset that will be used for analysis. Sections 3 and 4 respectively study the Gaussian copula model and the local intensity model, as well as the related hedging strategies. In Sections 5 and 6 we compare the two models from a theoretical point of view. In Sections 5 we identify different market conditions in which the spread delta and the local deltas can be ordered. Building on the dynamic feature of the local intensity model, this ordering of the deltas is then used in Section 6 for comparing the resulting P&Ls, in case either delta is used for hedging a CDO tranche. Backtesting experiments on the real dataset are conducted and discussed in Section 7. We conclude our results in Section 8.

2 CDS index and CDO tranches

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 or buy 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 North-American corporations. Market makers of this index have also agreed to quote prices for the standard tranches on these portfolios. 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.

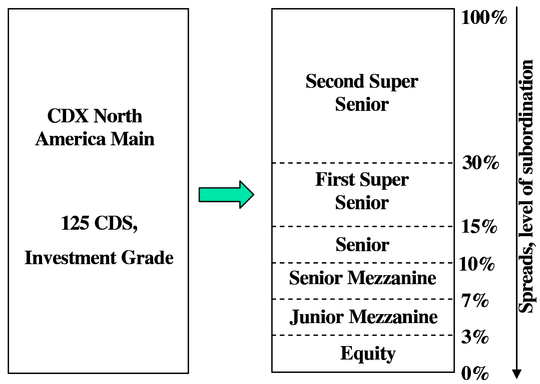

Figure 1: Standardized CDO tranches on CDX.NA.IG.

Considering a pool (portfolio) of n credit names, we denote by τ_i the default time corresponding to the i-th name, and by R an homogeneous and constant recovery rate at default, taken as 40% in all our numerics. We define the cumulative default process N and the cumulative loss process L by $N_t = \sum_{i=1}^{n} \mathbb{1}_{\tau_i \le t}$ 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.

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

A CDO tranche is a swap with two legs, a default protection leg and a fee leg, and a notion of fair spread defined much as in the case of interest rate swaps, so that the two legs of the contract would have equal values. Note that the contractual spread of the CDO tranche is fixed once a particular contract is traded, and the changes in value of the contract can be presented by the market par spread. We refer the reader to, for instance, Cousin and Laurent [8] or Cont and Kan [5] for more details on the cash-flows of synthetic CDO tranches and related products.

For the theoretical aspects of the paper we assume nil interest rates to simplify the notation. A constant interest rate $r = 3\%$ is throughout used in the numerics, the extension of all results to constant or time-deterministic interest rates being straightforward (but more cumbersome notationally, especially regarding hedging).

2.2 Data

For numerical illustration throughout this paper, we consider three 5-year CDX.NA.IG indexes and the corresponding tranches data in the period

- 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 'steady' sample, as opposed to the 'systemic' or 'crisis' period of series 9 and 10. From Figure 2, we can see that spreads are low and little volatile in the case of the pre-crisis series 5, increasing and volatile in the series 9, and high and volatile in the series 10.

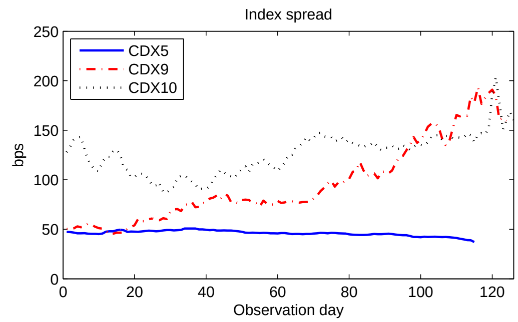

Figure 2: Index spread of 5-year CDX.NA.IG series 5 from 20 September 2005 to 20 March 2006, series 9 from 20 September 2007 to 20 March 2008 and series 10 from 21 March 2008 to 20 September 2008.

3 Gaussian Copula Model

The one factor Gaussian copula model introduced in Li [24] is the market quotation standard for multi-name credit derivatives. Under this model, the prices of CDO tranches depend on the current time t, a correlation parameter $\rho_t \in [0, 1]$, and a family $F_t = (F_t^i)_{1 \leq i \leq n}$ 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 marginal time-to-default distribution functions are equal, i.e., $F_t^i = F_t$ for all i. Moreover, since we only consider the 5-year maturity, we assume that the CDS term structures are constant. Under this assumption, the marginal time-to-default function can be uniquely represented by the index spread S_t. The homogeneity assumption is motivated by two reasons. First, we consider the hedging of CDO tranches by trading the underlying CDS index, so marginal effects are not of primary importance (as opposed to 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, where the latter is a top-down model for which the dispersion of individual defaults is embedded in the dynamics of the aggregate loss process. The homogeneity assumption on the Gaussian copula model allows us to illustrate more comparable results to the top-down local intensity model.

3.1 Market Implied Correlation Parameters

As the Black-Scholes formula for the volatility markets, the Gaussian copula model is usually used in reverse engineering for quoting CDO tranches in terms of their implied correlations. More precisely, let $\Sigma_t(a, b)$ be the market spread of a CDO tranche $[a, b]$ at time t, and $\Sigma^{gc}(a, b; t, S_t, \rho_t)$ be the model spread of the CDO tranche $[a, b]$ computed from the Gaussian copula model:

- The *compound correlation* of the tranche $[a, b]$ is the value of the correlation $\widetilde{\rho}_t^{a,b}$ in the

Gaussian copula model such that

$$\Sigma^{gc}(a, b; t, S_t, \widetilde{\rho}_t^{a,b}) = \Sigma_t(a, b)\,; \tag{1}$$

- the *base correlation* of level b is the value of the correlation ρ_t^b such that

$$\Sigma^{gc}(0, b; t, S_t, \rho_t^b) = \Sigma_t(0, b). \tag{2}$$

Since CDO tranches are usually quoted in non-overlapping tranches (see Section 2), the equity spreads at different attachment levels in (2) have to be bootstrapped from various mezzanine and senior tranches. We refer readers to [27] for details about the base correlation calibration method.

In this paper, we calibrate the Gaussian copula model based on the base correlation method because of its stability in calibration and popularity among market participants[1]. In terms of stability, spreads of equity tranches can be expressed as a decreasing function of the correlation parameter in the one-factor Gaussian copula model (see [7] for a formal proof). As a result, there is a unique base correlation for each attachment level, given that it exists. On the other hand, this uniqueness property does not necessarily hold for the mezzanine tranches if we use the compound correlation calibration.

Since each base correlation corresponds only to an equity tranche value, values of mezzanine tranches have to be represented by two base correlations at their attachment and detachment levels. In terms of the base correlation, the market price of a CDO tranche is thus eventually represented as $u^{gc}(a, b; t, S_t, \rho_t^a, \rho_t^b)$. In case of an equity (resp. senior) tranche with $a = 0$ (resp. $b = 1$), this reduces to $u^{gc}(0, b; t, S_t, 0, \rho_t^b)$ (resp. $u^{gc}(a, 1; t, S_t, \rho_t^a, 0)$), which one further simplifies to $u^{gc}(t, S_t, \rho_t)$ when the context clearly specifies the equity or senior tranche under consideration. As the index is not affected by the correlation level, the index price is denoted by $v^{gc}(t, S_t)$.

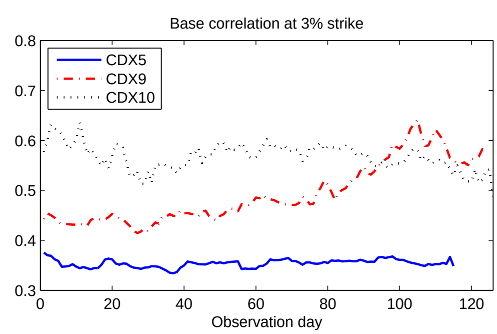

Figure 3: Base correlation at 3% strike of 5-year CDX.NA.IG series 5 from 20 September 2005 to 20 March 2006, series 9 from 20 September 2007 to 20 March 2008 and series 10 from 21 March 2008 to 20 September 2008.

Figure 3 shows the base correlation at 3% strike level during the three sample periods. In the pre-crisis period associated with CDX series 5, the base correlation appears to be rather low and globally stable whereas it clearly increases during crisis period of series 9 and

[1]The calibrated base correlations under the homogeneity assumption may differ from the market-implied ones as the calibration of the Gaussian copula model is typically not made under the pool homogeneity assumption. Therefore, the indicated correlation levels also embed the impact of the individual spread changes or dispersion.

remains at high and volatile levels during the more recent period of series 10. When time series of index spreads (Figure 2) are compared with time series of base correlations (Figure 3), it is interesting to remark that, in each period, the trend is similar. On these data sets, index spreads and base correlations seem to evolve in the same direction.

3.2 Delta Hedging of Spread Risk

Index delta hedging of spread risk for a tranche consists in entering a specific position in the CDS index, in such a way that changes in market price of a tranche $[a, b]$ due to moves in the index spread, are compensated by changes in market price of the index. The (index) *spread delta* for the tranche $[a, b]$, which represents the hedging position in the CDS index, is determined by computing index spreads sensitivities of both the tranche and the CDS index values using the one-factor Gaussian copula model. The spread delta is thus defined as

$$\Delta_t^{gc} = \Delta^{gc}(a, b; t, S_t, \rho_t^a, \rho_t^b) = \frac{u^{gc}(a, b; t, S_t + \delta S, \rho_t^a, \rho_t^b) - u^{gc}(a, b; t, S_t, \rho_t^a, \rho_t^b)}{v^{gc}(t, S_t + \delta S) - v^{gc}(t, S_t)} \tag{3}$$

where δS is typically equal to 1bp. Note that the market convention as in (3) is to keep constant the base correlations when recalculating the tranche prices. This corresponds to the so-called 'sticky strike' rule. The rationale behind this rule is related to the static nature of the Gaussian copula model: the sensitivity of base correlation with respect to a change in index spread cannot be predicted in this model, because this model does not capture any dynamic relationship between index spreads and the base correlations.

	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 returns of the base correlation at 3% strike level.

However, Table 1 shows that the level of correlation between changes in index spread and changes in base correlation is significant during the sample periods. 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.

Consequently, this empirical study suggests that, at least for series 9 and 10, delta hedging of spread risk is not able to account for the change in correlation when only the index spreads are bumped for the hedge ratio computation.

3.3 Greeks in the Gaussian Copula Model

Recall that the value of a CDO tranche $[a, b]$ computed in the Gaussian copula model can be expressed as a function of the current time t, the index spread S_t and the base correlations ρ_t^a and ρ_t^b. Thanks to the Itô formula, tranche price increment can be decomposed as a sum of terms that involve sensitivities (Greeks) with respect to the variables. In this section, we empirically study the relative contribution of these 'Greeks' to tranche price increments.

Consider a buy-protection position on an equity tranche. Let θ^{gc} be the first order partial derivative of the Gaussian copula pricing function u^{gc} with respect to time t and let $\delta_x^{gc} = \partial_x u^{gc}$, $\gamma_{x,y}^{gc} = \partial_{xy}^2 u^{gc}$ be the first and second order partial derivatives of u^{gc} with respect to the variables x and y. By applying the Itô formula, neglecting here jumps in the variables for simplicity:

$$\begin{aligned} du^{gc}(t, S_t, \rho_t) &= \theta^{gc}(t, S_t, \rho_t)dt + \delta_S^{gc}(t, S_t, \rho_t)dS_t \\ &\quad + \delta_\rho^{gc}(t, S_t, \rho_t)d\rho_t \\ &\quad + \frac{1}{2}\gamma_S^{gc}(t, S_t, \rho_t)d\langle S\rangle_t + \frac{1}{2}\gamma_\rho^{gc}(t, S_t, \rho_t)d\langle\rho\rangle_t \\ &\quad + \gamma_{S,\rho}^{gc}(t, S_t, \rho_t)d\langle S, \rho\rangle_t. \end{aligned} \tag{4}$$

By using this expression, we investigate the decomposition of value for the equity tranche $[0\%, 3\%]$ over a 6-month period corresponding to CDX series 9. In particular, we check the changes in tranche prices by using a discrete-time version of (4) where the Greeks are estimated by finite differencing with changes in time, index spread and base correlation are taken to be 1 day, 1bp and 0.1% respectively, and the variations of the spread and correlations are approximated by taking the differences in one and five days.

Figure 4 shows the decomposition of the changes in the [0%,3%] equity tranche value of CDX series 9 due to first and second order changes in index spread and base correlation. As visible in Figure 5, actual changes in the equity tranche and changes by summation of all terms of the Itô formula are very similar. According to Figure 4, the most influential terms are the first order sensitivities with respect to changes in the index spread and the base correlation. Note that these two sensitivities are about of the same order of magnitude, so that delta hedging of spread risk leaves a significant exposure to correlation risk. On the other hand, if index spread and base correlation would correspond to market prices of some tradable assets, it would be then possible to design an effective hedge using the one-factor Gaussian copula model.

Note that the second order changes of equity tranche value with respect to changes in the index 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[2].

While we observe significant contribution from the correlations change to the changes in the CDO tranche value, delta hedging of spread risk is not sufficient to provide an effective hedge. In order to overcome this problem, the question becomes whether we can benefit from the use of a dynamic model that may capture the observed dynamic feature of correlation risk, or at least, the component of this risk which is correlated to the index. Our analysis will focus on the simplest Markovian model of portfolio credit risk, namely the dynamic local intensity framework, which is presented in the next section.

[2]The second order changes with respect to spread appear to be consistently negative for a buy protection position. This feature has been proved formally by [25] for an equity tranche.

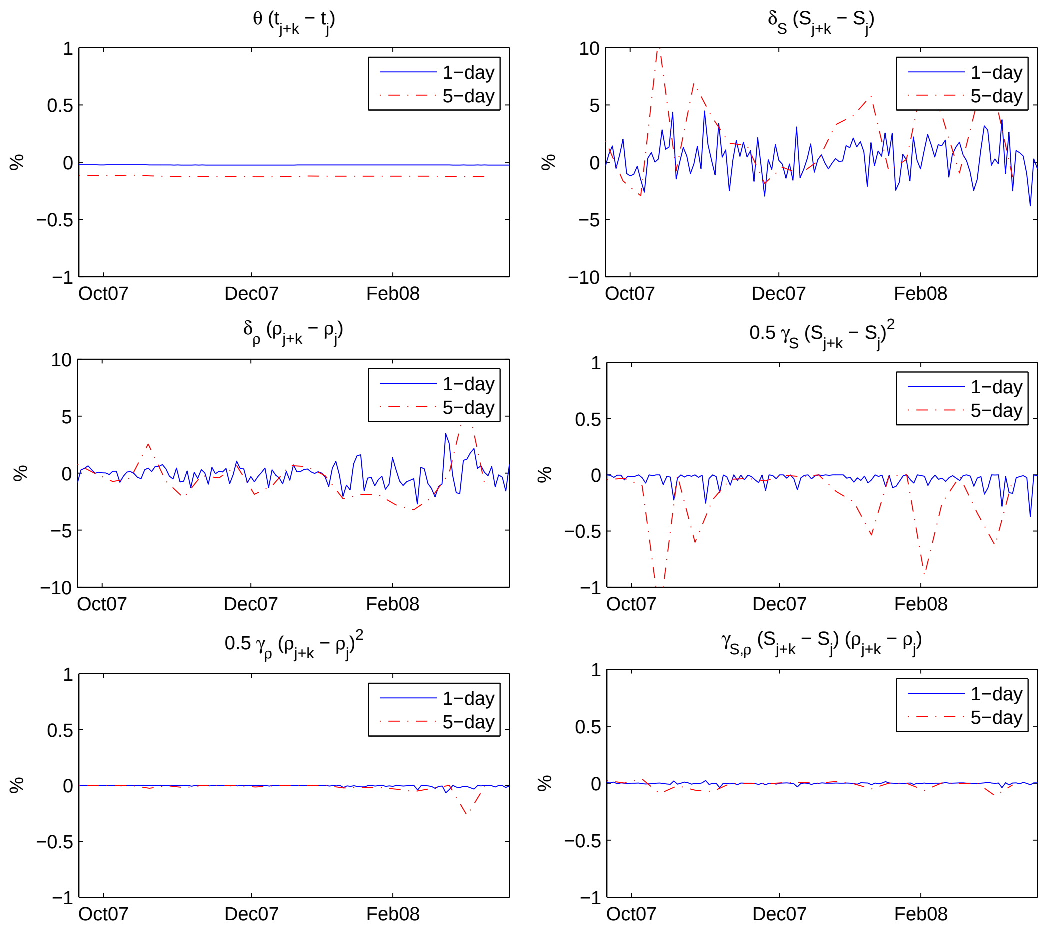

Figure 4: Changes of equity tranche [0%, 3%] value (100% notional) of CDX series 9 with respect to first and second order changes in time, index spread and base correlation.

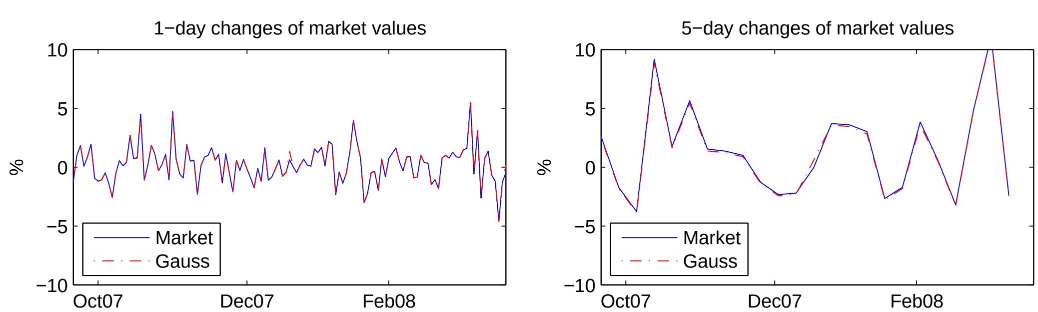

Figure 5: Actual changes in equity tranche [0%, 3%] value (100% notional) of CDX series 9 and the estimation of changes in equity tranche values based on discrete-time approximation of Itô formula (4) under the Gaussian copula model.

4 Local Intensity Model

In the local intensity model, the cumulative number of defaults $N = \{N_s, s \geq t\}$ of a credit portfolio of n names is a Markov point process (see, e.g., Laurent, Cousin and Fermanian [23], Cont and Minca [6] or Cont, Deguest and Kan [4]). More specifically, we assume that $(N_s)_{s\geq t}$, which represents the number of defaults in the portfolio beyond the current time t, is a pure birth process with an intensity $\{\lambda_t(s, N_s), s \geq t\}$ given by a deterministic function $\{\lambda_t(s,i)\}_{s\geq t,\, i=N_t,\ldots,n}$ satisfying $\lambda_t(s,i) = 0$ for $i \geq n$. This last condition guarantees that the process N is stopped at the level n, as there are n names in the pool. 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.

Conditionally on the information $\mathcal{F}_s = \mathcal{F}_s^N$ available at time s, the probability of a jump in the infinitesimal time interval $(s, s+ds)$ is given by $\lambda_t(s, N_s)ds$. The infinitesimal generator, which is also known as the *local intensity*, is thus given by the $(n+1) \times (n+1)$ matrix

$$\Lambda_t(s) = \begin{pmatrix} -\lambda_t(s,N_t) & \lambda_t(s,N_t) & 0 & 0 & 0 \\ 0 & -\lambda_t(s,N_t+1) & \lambda_t(s,N_t+1) & 0 & 0 \\ & & \cdots & & \\ 0 & 0 & 0 & -\lambda_t(s,n-1) & \lambda_t(s,n-1) \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

For notational simplicity, let us consider a stylized 'European-type' CDO tranche $[a, b]$ that provides the payoff $\phi(N_T)$ at maturity T but with no premium or default payment before maturity. Let $u^{lo}(a, b; s, N_s, \Lambda_t)$, or simply $u^{lo}(s, N_s, \Lambda_t)$ if the context is clear, be the value of the CDO tranche $[a, b]$ computed from the local intensity function $\Lambda_t = \{\Lambda_t(s), t \leq s \leq T\}$, given that there are N_s number of defaults by time s. The pricing function u^{lo} is characterized as the solution to the following system of backward Kolmogorov differential equations: $u^{lo}(T, i, \Lambda_t) = \phi(i)$ for $i = N_t, ..., n$, and for $t \leq s \leq T$,

$$\begin{cases} \partial_s u^{lo}(s,i,\Lambda_t) - \lambda_t(s,i)u^{lo}(s,i,\Lambda_t) + \lambda_t(s,i)u^{lo}(s,i+1,\Lambda_t) = 0\,, & i = N_t, \ldots, n-1\,, \\ \partial_s u^{lo}(s,n,\Lambda_t) = 0\,, & i = n\,. \end{cases} \tag{5}$$

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

$$u^{lo}(s,N_s,\Lambda_t) = u^{lo}(t,N_t,\Lambda_t) + \int_t^s \Big(u^{lo}(\zeta,N_{\zeta-}+1,\Lambda_t) - u^{lo}(\zeta,N_{\zeta-},\Lambda_t)\Big)\, dM_\zeta$$

where M is the compensated jump martingale of N, i.e., $dM_s = dN_s - \lambda_t(s, N_{s-})ds$. Using the analogous martingale representation for the price process $v^{lo}(s, N_s, \Lambda_t)$ of the credit index, it follows that in the local intensity model Λ_t, one can dynamically replicate the tranche by the CDS index (and the riskless asset) over the time interval $[t, T]$, by trading the index according to the *jump-to-default ratio* $\Delta^{lo}(a, b; s, N_{s-}, \Lambda_t)$, for $s \in [t, T]$, where for $i = N_{t-}, \ldots, n$,

$$\Delta^{lo}\,(a,b;s,i,\Lambda_t) = \frac{u^{lo}(a,b;s,i+1,\Lambda_t) - u^{lo}(a,b;s,i,\Lambda_t)}{v^{lo}(s,i+1,\Lambda_t) - v^{lo}(s,i,\Lambda_t)}\,. \tag{6}$$

As an analog to the hedging position in stocks under the local volatility model, we also refer this hedge ratio as the *local delta* and we call the implementation of this hedge ratio *delta hedging of default risk*.

In practice, the local intensity function Λ_t is calibrated to the available tranches and index spreads. Unlike the base correlation calibration for the Gaussian copula model where one correlation parameter is calibrated to each attachment level, the local intensity function provides a global fit to all tranches. If the local intensity function is re-calibrated at each time, the time-t local delta which is used for hedging is given by

$$\Delta_t^{lo} = \Delta^{lo}(a, b; t, N_t, \Lambda_t) = \frac{u^{lo}(a, b; t, N_t + 1, \Lambda_t) - u^{lo}(a, b; t, N_t, \Lambda_t)}{v^{lo}(t, N_t + 1, \Lambda_t) - v^{lo}(t, N_t, \Lambda_t)}. \tag{7}$$

In all of our later numerical analysis, prices computed in the local intensity model take into account the real cash-flows as opposed to the above stylized presentation. Moreover, in order to be consistent with market conventions, all prices, and also later the hedge ratios, are computed from 100% notional value. Therefore, both the tranche values and deltas are 'scaled' by the tranche width. The price for any tranche or for the CDS index is thus 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 $[0, 33.33]$.

4.1 Model Calibration

Various methods have been proposed to recover the local intensity function (see for instance Chapter 2 of [11] for a review of such approaches). However, Cont, Deguest and Kan [4] show that even if the local intensity function is calibrated to the same set of market data, model dependent quantities such as the local 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 the local intensity;
- Non-parametric: Entropy minimization calibration introduced by Cont and Minca [6].

Regarding the parametric specification, note that in Herbertsson [20] a piecewise constant parametrization of $i \mapsto \lambda(s, i)$ is used, whereas we use for our parametrization in this work a piecewise linear form for $i \mapsto \lambda(s, i)$, which we found to be more robust and to provide a better fit at the time of daily recalibration of the model on our time series of data[3]. More specifically, the time-homogeneous local intensity function is assumed to be linear in the number of default variable where the grid points are equal to the attachment levels divided by the loss given default, rounded to the closest integers.

One advantage of the time-homogeneous parametric model is that the shape of the calibrated local intensity function can be easily interpreted in terms of default dependence. On the other hand, the non-parametric approach, as shown by Cont et al. [4], usually produces an irregular local intensity function which is difficult to interpret. However, Cont et al. [4] shows that the non-parametric approach is more stable with respect to shifting in the market spreads.

Table 2 shows the relative root mean squared errors of the calibrated spreads across quotation dates of each series. The per tranche calibration of the Gaussian copula model is of course nearly perfect by the base correlation approach. For the local intensity models (global fit), the errors are about 2% for the tranches and about 5% for the index.

[3] We have also tried piecewise constant formulation and the hedge ratios do not appear to be significantly different from the ones obtained with the piecewise linear formulation (at least not as much as from the entropy minimization calibration).

	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: Relative 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.

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. Similar to the findings by Cont et al [4], the calibrated spreads are nearly identical for both specifications of the local intensity.

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

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

Figure 6 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.

4.2 Impact of a Default on Index Spreads and Base Correlations

From the empirical study of Section 3.2, we know that, especially for CDX series 9 and 10, index spreads and base correlations usually move in the same direction. We have insisted on the fact that the Gaussian copula delta computed under the 'sticky-strike' rule does not capture the dynamic feature of base correlations although, according to Section 3.3, the correlation risk may strongly contribute to change in tranche market prices. Given these observations, our aim is to explore whether the local intensity model is able to account for a dynamic evolution of the correlation among defaults in the portfolio.

Recall that the dynamic aspects of the local intensity model are entirely captured by the default filtration, that is, the information associated with timing of defaults in the pool. Moreover, the local delta defined by expression (7) corresponds to sensitivity with respect to default risk. It is the ratio of the tranche price sensitivity over the index price sensitivity

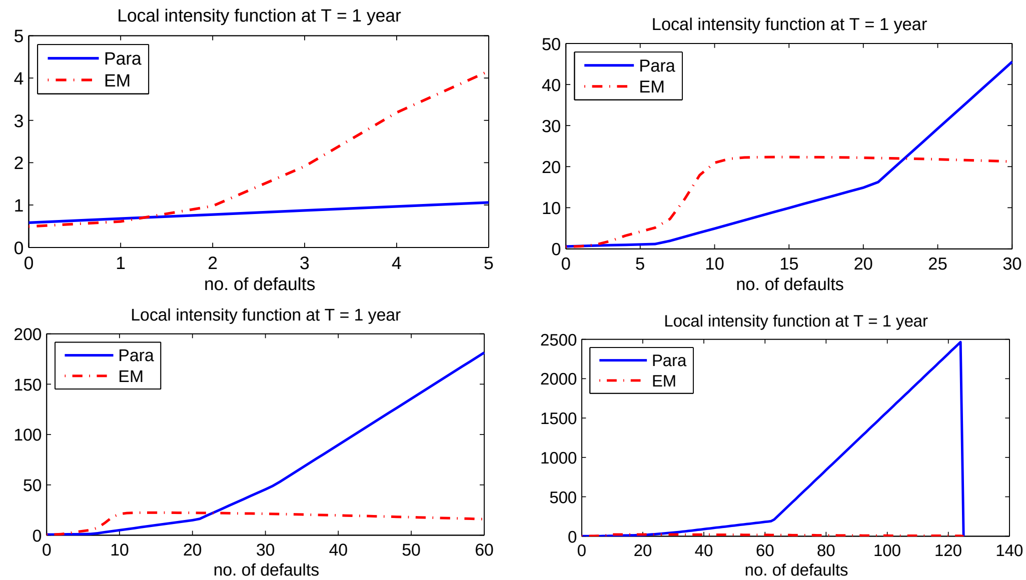

Figure 6: Local intensity function at $T = 1$ year. Same data as Table 3.

with respect to the occurrence of a new default. So, in the same vein as Section 3.2 for the Gaussian copula model, we aim here at comparing the correlation between variations in index spreads and base correlation levels as a result of a new default in the local intensity model, with the actual or realized correlation between variations of these two quantities.

In order to do so, we need to introduce a concept of correlation as implied by prices computed in the local intensity model. Let us denote by $S^{lo}(t, N_t, \Lambda_t)$ the current CDS index spread computed in the local intensity model calibrated at time t. We also denote by $\rho^{lo}(b; t, N_t, \Lambda_t)$, or simply $\rho^{lo}(t, N_t, \Lambda_t)$ if the context is clear, the base correlation implied from the equity tranche price with detachment level b computed in a local intensity model pre-calibrated at time t on market spreads.

Figure 7 shows the time series of the differences between the index spread $S^{lo}(t, N_t, \Lambda_t)$ (resp. market base correlation $\rho^{lo}(b; t, N_t, \Lambda_t)$ for $b = 3\%$) and the values implied by the local intensity model with one additional default $S^{lo}(t, N_t + 1, \Lambda_t)$ (resp. $\rho^{lo}(b; t, N_t + 1, \Lambda_t)$ for $b = 3\%$). Observe that index spreads implied by the local intensity model with one more default are always higher than the market spreads,[4] so the index spread is increasing in the number of defaults (at least, for the first default) across all the data sets. Now, rather consistently with Table 1 in Section 3.2, 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 or lower than 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 is always greater than the market value.

In the case of series 9 and 10, index spreads and base correlations thus move in the same direction when a default occurs in the local intensity model, whereas they may move in opposite direction for CDX series 5. From this point of view, the dynamic nature of base correlations and index spreads as generated in the local intensity model seems to be rather consistent with what we observed in Table 1.

However, Table 4, which presents the correlations between index spread returns implied

[4]Recall again that $S^{lo}(t, N_t, \Lambda_t) \simeq S_t$ as a result of the calibration procedure.

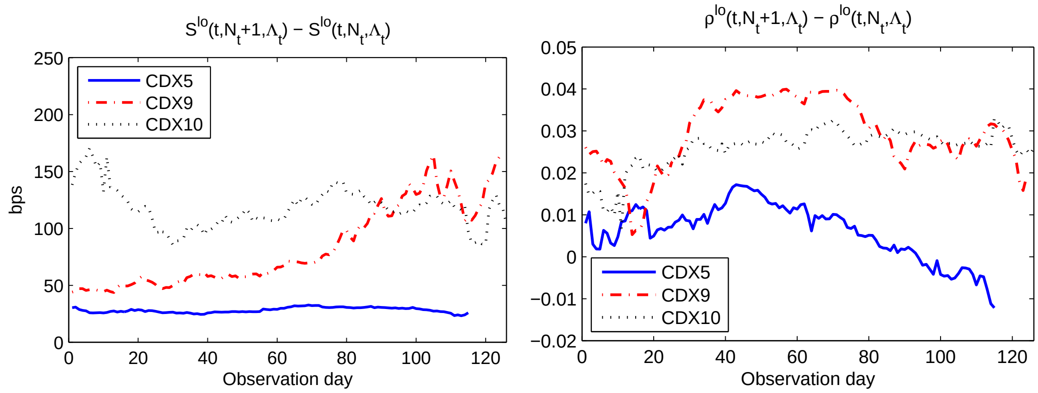

Figure 7: Left: Differences between the index spread $S^{lo}(t, N_t, \Lambda_t)$ and the values implied by the local intensity model with one additional default $S^{lo}(t, N_{t+1}, \Lambda_t)$. Right: Differences between market equity tranche base correlation $\rho^{lo}(t, N_t, \Lambda_t)$ and the base correlation implied from equity tranche $[0\%, 3\%]$ 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.

by the calibrated local model with one additional default, and base correlation returns implied by the calibrated local model with one additional default, illustrates the results as opposed to what we observe empirically in Table 1. The correlations of the index spread return and base correlation returns implied by the local intensity model with one additional default are all negative, including the crisis data of series 9 and 10. It might thus eventually be so that, even if the local intensity model builds in some dynamics of the joint evolution of base correlations and index spreads, these dynamics are in fact misspecified. The local delta would then incorporate a 'wrong correction' with respect to the spread delta. This kind of phenomenon is reminiscent of similar difficulties met with local volatility models on certain derivatives markets, which led to the introduction of models with richer dynamics, like the SABR model of Hagan et al. [19].

Strike	CDX5	CDX9	CDX10
3%	-0.57	-0.72	-0.76
7%	-0.64	-0.47	-0.75
10%	-0.62	-0.31	-0.62
15%	-0.55	-0.12	-0.36
30%	-0.37	0.06	0.07

Table 4: Correlations between index spread returns implied by the calibrated local model with one additional default, and base correlation returns implied by the calibrated local model with one additional default

4.3 Stability of Hedge Ratios

Figure 8 shows time series of local deltas for the equity tranche $[0\%, 3\%]$ in each sample period. Observe that the local deltas computed from entropy minimization, or 'entropic local deltas' for short, is significantly smaller than the local deltas computed from the parametric model, or 'parametric local deltas' in short-hand, throughout the whole time series, even though they are both local deltas in local intensity models calibrated to the same data sets. When comparing to the spread deltas computed from the Gaussian copula model, the parametric local deltas lies somewhere between the entropic local deltas and the spread deltas.

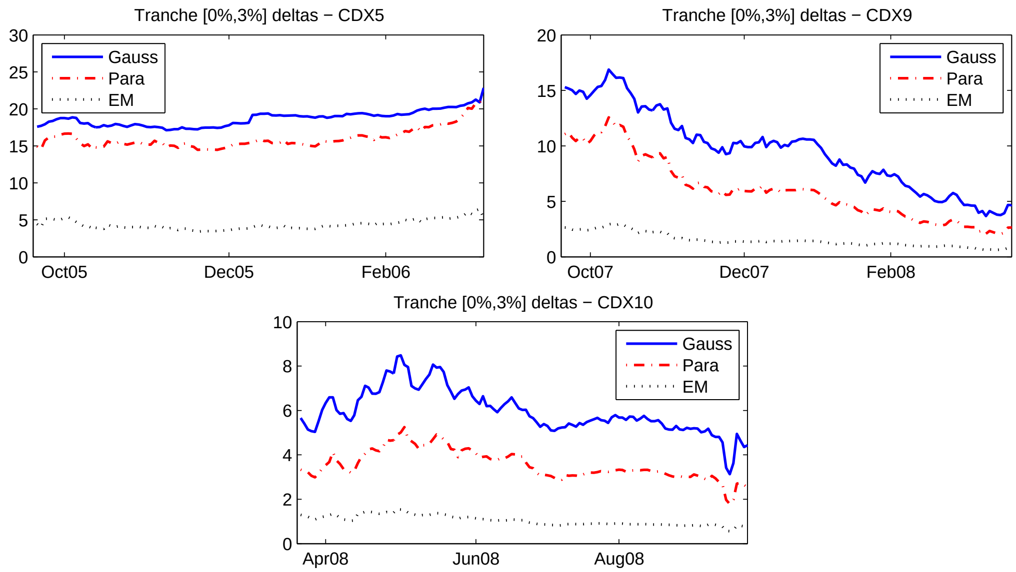

Figure 8: Equity tranche deltas for CDX series 5, 9 and 10.

Table 5 shows the spread deltas and the local deltas, computed by models which are calibrated on the same data as those used in Table 3 (CDX series 9 spreads of 20 September 2007). Note the gap between the parametric local delta and the entropic local delta, whereas the related calibrated spreads are nearly identical. This is a striking example of model risk.

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 5: Hedge ratios. Same data as Table 3

As can be seen from Figure 8, the entropic local deltas are substantially more stable than the parametric local deltas for CDX series 9 and 10. The stability of the entropic local deltas is consistent with the observation by Cont et al. [4] 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 delta requires the full local intensity function, it is not surprising that the entropic local deltas are less sensitive to the changes in the market data than parametric local deltas. This difference is more significant in the volatile periods for CDX series 9 and 10.

However, the stability of entropic local deltas does not necessarily imply a better hedging strategy when implementing the local delta. Indeed, the entropic local deltas may fail to reflect the market information while they are very much indifferent to the market spreads changes. We will see in Section 7 that the entropic local deltas in fact lead to poor hedging performance.

5 Ordering Between the Deltas

In this section we study the ordering between the spread delta Δ_t^{gc} in (3) and the local delta Δ_t^{lo} in (7), depending on the seniority of the tranche and the market regime.

5.1 Market Regimes

As we have seen in Figure 7, all local intensity models calibrated to market quotes exhibit a surge of index spreads at the arrival of a default, the jump in index 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. This should also be connected[5] with our observation of Table 1 in Section 3.2 that 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.

Accordingly, for the purpose of the theoretical study of Sections 5 and 6, we distinguish two stylized market regimes: a *systemic* regime, in which, by definition, a new default in a local intensity model calibrated to the market, tends to increase (equity tranche) base correlation, and a *steady* regime, in which a new default in a local intensity model calibrated to the market, has a smaller (or even negative) impact on correlation.

5.2 Equity Tranche in a Systemic Market

Let us first consider the case of a buy-protection equity tranche ($a = 0$) in a systemic market, as of those calibrated on Series 9 and 10 (see Section 4.2 above).

We have seen in Figure 8 that the local intensity deltas are consistently smaller than the Gaussian copula deltas across the two crisis Series 9 and 10, i.e.,

$$\Delta_t^{lo} \leq \Delta_t^{gc}. \tag{8}$$

We attempt here to give some theoretical arguments in favor of this empirical observation. First note that from the definition of a systemic regime defined in Section 5.1 (similar to those calibrated on Series 9 and 10, see right panel of Figure 7), one has

$$\rho^{lo}(t, N_t, \Lambda_t) \leq \rho^{lo}(t, N_t + 1, \Lambda_t). \tag{9}$$

Moreover it is well known (see, e.g., [7] for a formal proof) that the price of an equity tranche computed in the one-factor Gaussian copula model is a decreasing function of the correlation parameter ρ:

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

So

$$u^{gc}\left(t, S^{lo}(t, N_t + 1, \Lambda_t), \rho^{lo}(t, N_t + 1, \Lambda_t)\right) \leq u^{gc}\left(t, S^{lo}(t, N_t + 1, \Lambda_t), \rho^{lo}(t, N_t, \Lambda_t)\right) \ .$$

[5] With the reservation made in the comments to Table 4 however.

Now, one has by definition of the Gaussian copula implied correlation, for every $t \in [0, T]$:

$$\begin{aligned} &u^{lo}(t, N_t+1, \Lambda_t) - u^{lo}(t, N_t, \Lambda_t) \\ &= u^{gc}\left(t, S^{lo}(t, N_t+1, \Lambda_t), \rho^{lo}(t, N_t+1, \Lambda_t)\right) - u^{gc}\left(t, S^{lo}(t, N_t, \Lambda_t), \rho^{lo}(t, N_t, \Lambda_t)\right) \\ &= u^{gc}\left(t, S^{lo}(t, N_t+1, \Lambda_t), \rho^{lo}(t, N_t+1, \Lambda_t)\right) - u^{gc}\left(t, S^{lo}(t, N_t+1, \Lambda_t), \rho^{lo}(t, N_t, \Lambda_t)\right) \\ &\quad + \left(u^{gc}\left(t, S^{lo}(t, N_t+1, \Lambda_t), \rho^{lo}(t, N_t, \Lambda_t)\right) - u^{gc}\left(t, S^{lo}(t, N_t, \Lambda_t), \rho^{lo}(t, N_t, \Lambda_t)\right)\right) . \end{aligned} \tag{10}$$

Therefore, by (10):

$$\begin{aligned} \Delta_t^{lo} &= \frac{u^{lo}(t, N_t+1, \Lambda_t) - u^{lo}(t, N_t, \Lambda_t)}{v^{lo}(t, N_t+1, \Lambda_t) - v^{lo}(t, N_t, \Lambda_t)} \\ &\leq \frac{u^{gc}(t, S^{lo}(t, N_t+1, \Lambda_t), \rho^{lo}(t, N_t, \Lambda_t)) - u^{gc}(t, S^{lo}(t, N_t, \Lambda_t), \rho^{lo}(t, N_t, \Lambda_t))}{v^{lo}(t, N_t+1, \Lambda_t) - v^{lo}(t, N_t, \Lambda_t)} \\ &= \frac{u^{gc}(t, S^{lo}(t, N_t+1, \Lambda_t), \rho^{lo}(t, N_t, \Lambda_t)) - u^{gc}(t, S^{lo}(t, N_t, \Lambda_t), \rho^{lo}(t, N_t, \Lambda_t))}{v^{gc}(t, S^{lo}(t, N_t+1, \Lambda_t)) - v^{gc}(t, S^{lo}(t, N_t, \Lambda_t))} \\ &= \frac{u^{gc}(t, S^{lo}(t, N_t, \Lambda_t) + \varepsilon, \rho^{lo}(t, N_t, \Lambda_t)) - u^{gc}(t, S^{lo}(t, N_t, \Lambda_t), \rho^{lo}(t, N_t, \Lambda_t))}{v^{gc}(t, S^{lo}(t, N_t, \Lambda_t) + \varepsilon) - v^{gc}(t, S^{lo}(t, N_t, \Lambda_t))} \\ &\approx \Delta_t^{gc} , \end{aligned} \tag{11}$$

where $\varepsilon = S^{lo}(t, N_t+1, \Lambda_t) - S^{lo}(t, N_t, \Lambda_t)$, which yields (8). Note that denominator of (11) does not depend on the implied correlation.

5.3 Senior Tranche in a Systemic Market

The previous analysis focused on an equity tranche in a systemic market. However, the same analysis can be made for a senior tranche that protects against last losses, i.e., a CDO tranche $[a, b]$, with $0 < a < 100\% = b$. The Gaussian copula pricing function associated with a senior tranche is *increasing* with respect to correlation, i.e., one has in this case that $\partial_\rho u^{gc}(t, S, \rho) \geq 0$ (see, e.g., [7] for a proof). As a result, in case of a senior tranche, the above arguments yield the opposite ordering between Gaussian copula deltas and local deltas, i.e., $\Delta_t^{lo} \geq \Delta_t^{gc}$.

5.4 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 7), contrary to the crisis period associated with series 9 and 10. In view of Section 5.2, 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 8, the Gaussian copula equity deltas are consistently greater than the equity local deltas for the whole series 5. Therefore, the reverse implication associated with (9) and (11), i.e., the fact that $\rho^{lo}(t, N_t, \Lambda_t) \geq \rho^{lo}(t, N_t+1, \Lambda_t)$ would imply $\Delta_t^{lo} \geq \Delta_t^{gc}$, is not satisfied by empirical observation (see also Table 6 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.

Example 5.1 Table 6 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 spread delta is almost the same as the local delta;

- On 16 September 2008 (next business day after Lehman Brothers defaulted, representative example of a crisis market; note however that Lehman Brothers was not part of the pool underlying series 10), the base correlation predicted by the local intensity model increases. In this case, the spread delta is significantly larger than the local delta, as suggested by (8).

Date	$\rho^{lo}(t, N_t, \Lambda_t)$	$\rho^{lo}(t, N_t + 1, \Lambda_t)$	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 6: Base correlation implied by market data $\rho^{lo}(t, N_t, \Lambda_t)$, base correlation implied by one additional default in the local intensity model $\rho^{lo}(t, N_t + 1, \Lambda_t)$, spread delta Δ^{gc} and local delta Δ^{lo}. Equity tranche $[0\%, 3\%]$ of 5-year CDX series 5 on 17 March 2006 and series 10 on 16 September 2008.

6 Ordering Between the P&Ls

Delta hedging in discrete time the tranche 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 (buy-protection) tranche and index market values between times t_k and t_{k+1}, and 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 main aim in this paper is to compare the profit-and-loss processes obtained by using two strategies, with hedge ratios given by (i) the spread delta in (3) and (ii) the local delta in (6).

Let δe^{lo} (resp. δe^{gc}) represent the P&L increment (12) while we implement the local delta (resp. the spread delta). After having studied the ordering of the deltas in the previous section, we study that of the resulting δe's in the current one. Toward this end, within each market regime: systemic or steady, for each hedging rebalancing period (one day or five days), we shall also consider two stylized market scenarios: *widening and tightening*, corresponding to values of the index spread increasing and decreasing, respectively.

In most of this section, we consider a theoretical market given in the form of a local intensity model. Equivalently, we assume $\Lambda_t = \Lambda_0$ for every $t \in [0, T]$. In a local intensity model, we know from Section 4 that the strategy implementing the local delta (7) in continuous time, provides a perfect replication of the tranche by the index. One thus has for

$t \in [0, T]$, Π and P denoting the tranche and index price process in the local intensity model with intensity Λ_0,

$$d\Pi_t = \Delta_t^{lo} dP_t \; . \tag{13}$$

However, such a perfect hedge cannot be achieved in discrete time. We thus propose in Subsect. 6.1 to 6.3 to compare the P&L increments associated with delta hedging in discrete time of spread risk and default risk in this setup.

Extension of the analysis to a 'real market' is then discussed in Subsect. 6.4.

6.1 Equity Tranche in a Systemic Local Intensity Model

We first consider the case of an equity tranche ($a = 0$) in a systemic local intensity model. One then has that

$$\delta e^{lo} \text{ is negative in tightening scenarios and positive in widening scenarios.} \tag{14}$$

Indeed, (13) yields,

$$\delta_k e^{lo} = \delta_k \Pi - \Delta_{t_k}^{lo}\, \delta_k P = \int_{t_k}^{t_{k+1}} (\Delta_t^{lo} - \Delta_{t_k}^{lo})\, dP_t \; . \tag{15}$$

Now, it is quite intuitive that the equity tranche delta $\Delta^{lo}(0, b; t, i, \Lambda_0)$ is an increasing function of time t (see Figure 9 for a typical example). As we get closer to maturity, the time-value of both the tranche and the index vanishes. Therefore, the change in value at the arrival of a default[6] is only the consequence of a protection payment. This protection payment is $1/b$ times larger for the tranche than for the index[7]. 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).

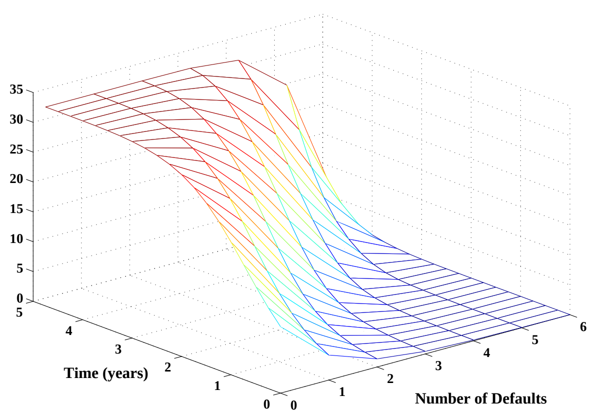

Figure 9: Equity tranche deltas $\Delta^{lo}(0, 3\%; t, i, \Lambda_0)$ computed in a local intensity model as a function of time t and number of defaults i. The parametric local intensity function Λ_0 is calibrated on market spreads of 5-year CDX series 9 on 20 September 2007.

Therefore, on a small time interval $(t_k, t_{k+1}]$,

[6] Assuming that the default fully affects the tranche.

[7] If the nominal is the same for the tranche and the index.

- 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 (15) $\delta e^{lo} \simeq (\Delta^{lo}_{t_{k+1}} - \Delta^{lo}_{t_k})\, \delta_k P$, the P&L increment δe^{lo} would be 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 7). Then, thanks to representation (15), the P&L increment δe^{lo} would be positive in this period.

By definition (12), one has

$$\delta e^{gc} = \delta e^{lo} - \left(\Delta^{gc} - \Delta^{lo}\right) \delta P \tag{16}$$

Thus, in view of the ordering (8) of the deltas, the following comparisons holds

$$\begin{cases} \delta e^{lo} \leq \delta e^{gc} & \text{for } \delta P \leq 0, \\ \delta e^{lo} \geq \delta e^{gc} & \text{for } \delta P \geq 0. \end{cases} \tag{17}$$

Combining (14) and (17), we get the picture depicted in Table 7. It might thus be so that, in

Tightening	**Widening**
$\delta e^{lo} \leq \min(\delta e^{gc}, 0)$	$\max(\delta e^{gc}, 0) \leq \delta e^{lo}$

Table 7: Equity tranche in a systemic local intensity model.

some scenarios, the spread delta provides a better hedge than the local delta. Indeed, using simulations of default times in a local intensity model calibrated to CDX series 9 spreads, we observe in Figure 10 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 every week or more) δe^{lo} is very close to 0, and Table 7 reduces to Table 8:

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

Table 8: Case of a moderate to high rebalancing frequency in Table 7.

In Section 7, we shall confirm that the latter ordering of P&L really holds for CDX series 9 (with mainly spread widening periods) when realized cumulative P&Ls are backtested using hedging experiments (see Figure 14).

6.2 Senior Tranche in a Systemic Local Intensity Model

It can be checked numerically and analyzed as in the case of an equity tranche (see Figure 9) that contrary to equity tranche deltas, senior tranche deltas $\Delta^{lo}(a, 1; t, i, \Lambda_0)$ computed in the local intensity model calibrated on series 9 and 10 are typically decreasing functions

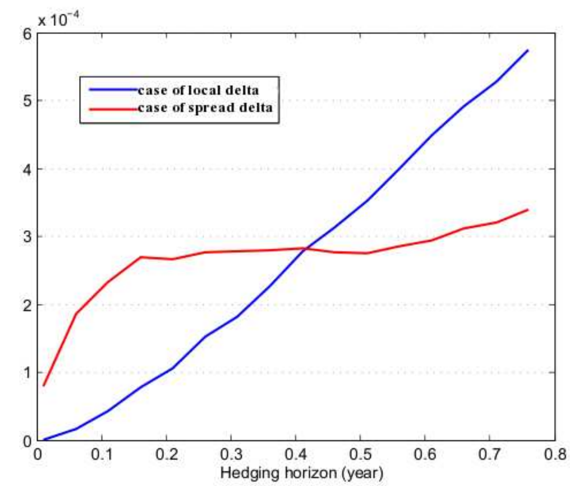

Figure 10: Standard deviation of equity tranche P&L increments δe^{lo} and δe^{gc} as a function of the hedging horizon. Default times are simulated in a local intensity market calibrated on market spreads typical of CDX series 9.

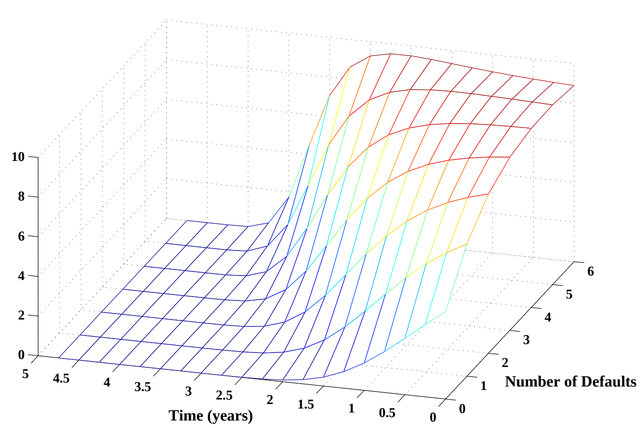

Figure 11: Senior tranche deltas $\Delta^{lo}(15\%, 30\%; t, i, \Lambda_0)$ computed as a function of time t and number of defaults i. The parametric local intensity function Λ_0 is calibrated on market spreads of 5-year CDX series 9 on 20 September 2007.

of time. As can be seen on Figure 11, this is also typically the case for a [15%, 30%] CDO tranche.

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, in the context of a systemic local intensity model.

Moreover, one has from Sect. 5.3 that $\Delta_t^{lo} \geq \Delta_t^{gc}$ for a senior tranche in a systemic market. It thus turns out that cells in each of the Tables 7 and 8, can be simply exchanged for the senior tranche, leading to Tables 9 and 10.

Tightening	**Widening**
$\max(\delta e^{gc}, 0) \leq \delta e^{lo}$	$\delta e^{lo} \leq \min(\delta e^{gc}, 0)$

Table 9: 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 10: Case of a moderate to high rebalancing frequency in Table 9.

6.3 Analysis in a Steady Local Intensity Model

One can imagine a local intensity model where the ordering between the two deltas would be inverted for both equity and senior tranches. Then, the results analogous to those of Tables 7 and 9 are displayed in Table 11. Note that the conclusion is in fact even clearer in that case since the ordering of δe holds irrespectively of the frequency of the hedge rebalancing.

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

6.4 Analysis in a Real Market

Mimicking the volatility approach of Crépey [12], it can be checked that the conclusions of Tables 8, 10 and 11 would still hold true in the real market and not only in a fixed local intensity model, under the assumption of Table 12 regarding the dynamics of implied correlations in the real market: Gaussian copula implied base correlations moving upwards (resp. downwards) compared to those predicted by the local intensity model calibrated at date t_k, in widening scenarios in systemic markets or in tightening scenarios in steady markets (resp. in tightening scenarios in systemic markets or in widening scenarios in steady markets).

In a market satisfying this theoretical assumption about the dynamics of implied correlations of CDO tranches, the local delta would then provide a better hedge than the spread delta. But this theoretical assumption might of course not be satisfied in practice. As we explained in the last paragraph of Section 4.2, even though the local intensity model is a

Regime / Scenario	Tightening	Widening
Systemic	$\rho_{t_{k+1}} \leq \rho^{lo}(t_{k+1}, N_{t_{k+1}}, \Lambda_{t_k})$	$\rho_{t_{k+1}} \geq \rho^{lo}(t_{k+1}, N_{t_{k+1}}, \Lambda_{t_k})$
Steady	$\rho_{t_{k+1}} \geq \rho^{lo}(t_{k+1}, N_{t_{k+1}}, \Lambda_{t_k})$	$\rho_{t_{k+1}} \leq \rho^{lo}(t_{k+1}, N_{t_{k+1}}, \Lambda_{t_k})$

Table 12: Theoretical assumption about the dynamics of implied correlations in a real market, under which the local delta would yield a better hedge than the spread delta.

consistent dynamic model of portfolio credit risk providing a perfect fit to the market in static terms (at any fixed time t_k), it might well be misspecified in dynamical terms. As a matter of fact, the market data analysis of the following section suggests that this is indeed the case on our datasets.

7 Backtesting Experiments

Let us now examine the actual performance of implementing the spread delta and the local 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 five days. The profit-and-loss is evaluated in the same frequency as rebalancing. Tables 13-16 and Figures 12-13 illustrate the hedging performance for 1-day and 5-day rebalancing. We find that in most cases the change in rebalancing frequency does not significantly perturb the comparison with respect to hedging performance.

Interestingly, hedging based on the entropic local deltas, which gives admittedly stable but also very low equity tranche hedge ratios in 2007-08, performs worse than the implementation of spread delta and parametric local delta. Regarding CDX series 5, the Gaussian copula delta outperforms the local deltas 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 [5]. This may be caused by the differences in the Gaussian copula model implementation in which we assume a homogeneous model with base correlation calibration while Cont and Kan [5] 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.

In Figure 14, 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 right panel of Table 7 (equity tranche) and Table 9 (senior tranche) behave according to the ordering of cumulative P&L trajectories.

As a final remark, Table 17 illustrates that defaults among names of the index do not necessarily positively impact intensities, since they may have been anticipated by the

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 13: 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 14: 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.

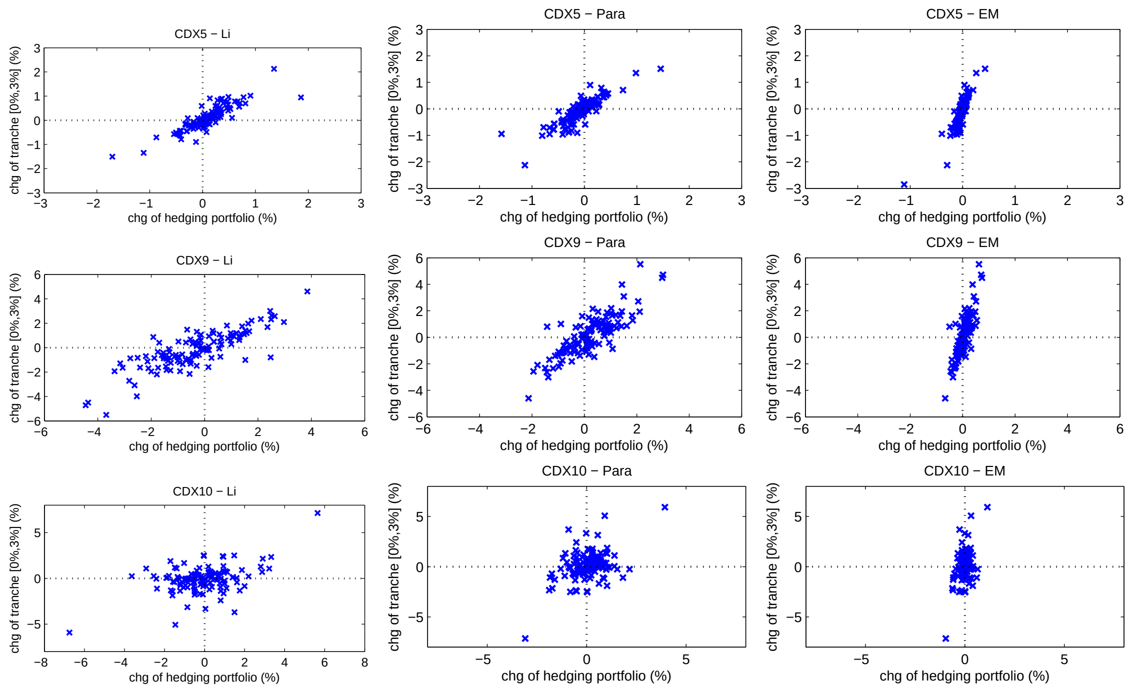

Figure 12: Daily changes of tranche [0%,3%] value against daily changes of hedging portfolio value (100% notional values).

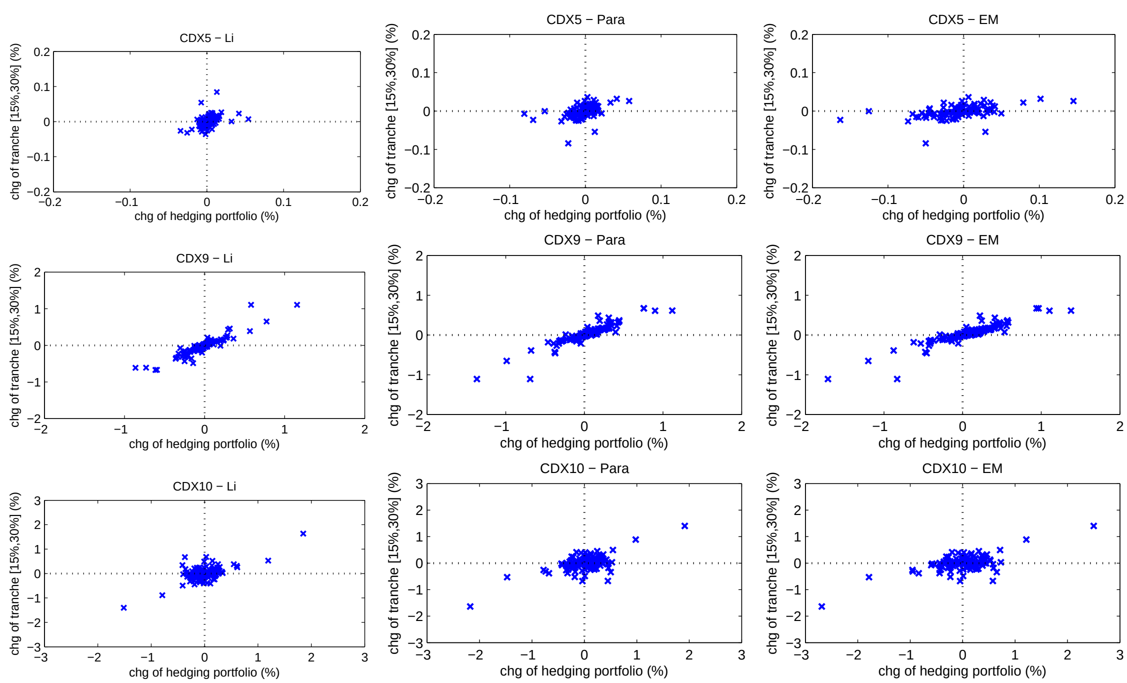

Figure 13: Daily changes of tranche [15%,30%] value against daily changes of hedging portfolio value (100% notional values).

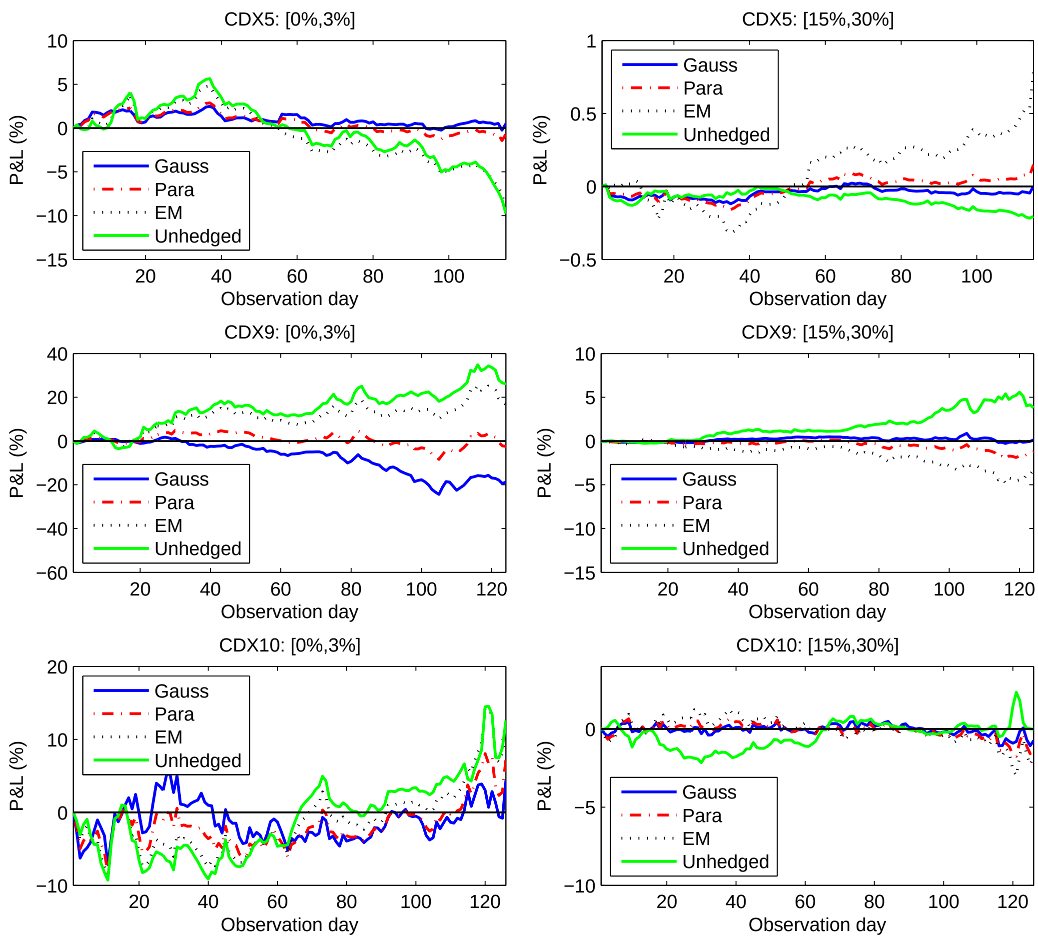

Figure 14: Path of cumulative P&L of hedged and unhedged positions in tranche [0%,3%] and [15%,30%] based on daily rebalancing (100% notional values).

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 15: 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.

market. This observation had already been made in Cont and Kan [5], who also noted that spreads may jump at times others than constituent defaults. This is for instance the case the day following Lehman's default on September 15 2008, whereas Lehman was not part of the CDX series.

8 Conclusions

Since we observe significant correlations between the index spreads and the base correlations, the commonly used delta hedging of spread risk by using the Gaussian copula model is insufficient to provide an effective hedge. Therefore, we attempt to use the dynamic local intensity model to capture the observed joint dynamic of the index spread and base correlations. However, our empirical study shows that implementation of the local delta under the dynamic local intensity model does not necessarily outperform delta hedging of spread risk by using the static Gaussian copula model. This observation shows the insufficiency of the local intensity model which can be explained as follows.

Given the calibrated local intensity function, the implementation of the local delta assumes that the only source of risk is the occurrence of underlying defaults. In this case, the model implies a deterministic movements in the index and CDO tranche values between defaults, which is unrealistic. Moreover, 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 7).

By using the local intensity model as a starting point, we propose two future research directions which can potentially solve the above problems. First, the local intensity model in fact accommodates the change in the index and CDO tranche spreads not only by the

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 16: 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 17: 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.

number of defaults, but also by re-calibrating the local intensity function. Therefore, the arguably more appropriate strategy for hedging CDO tranches is to consider the sensitivity with respect to the changes in the local intensity function instead of the occurrence of additional defaults. However, this approach may face substantial model risk as we observe that different parametrization scheme of the local intensity function can lead to significantly different hedging results. Second, we can introduce additional risk factors in the credit risk model to explain the spread volatility between successive defaults. This approach requires a more extensive study of the joint dynamics of the index spread and of CDO tranche base correlations so that the additional risk factors can capture the appropriate dependence between the two.

Incidentally the present work also puts into evidence an important difficulty with dynamic credit risk modeling, namely model risk. That critical issue is somewhat expected, in regard to the scarceness of market data that can be used for the model calibration. Yet we saw in Section 4.3 two specifications of local intensity models giving exactly the same calibrated spreads, but significantly different local deltas. 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 recent attempt to explain a certain robustness of the Gaussian copula model. Following a 'reverse-engineering' approach, Fermanian and Vigneron thus consider in [16] the question of determining dynamic models of credit risk in which the Gaussian copula spread delta would happen to be the 'right' hedge ratio when no defaults occur. Their answer is that such models can be found in the form a certain class of conditional density (as opposed to intensity) models as developed independently at an abstract level by El Karoui et al. [14]. However, the recent crisis in 2008 shows that defaults can happen frequently in a short period of time. Therefore, adequacy of those hedging strategies, which are proven to be valid only in the absence of defaults, has to be further examined by testing on empirical data with credit events.

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] Bielecki, T.R., Crépey, S. and Jeanblanc, M.: Up and Down Credit Risk. *Quantitative Finance*, p. 1-15, April 2010.

[4] 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, p. 555-585, 2010.

[5] Cont, R. and Kan, Y.H.: Dynamic Hedging of Portfolio Credit Derivatives, *SIAM Journal on Financial Mathematics*, 2, p. 112-140, 2011.

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

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

[8] 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.

[9] 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.

[10] 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.

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

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

[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] Hagan, P. , Kumar, D., Lesniewski, A. and Woodward, D.: Managing Smile Risk. *Wilmott Magazine*, 2002.

[20] 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.

[21] 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, 295-338, 2007.

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

[23] Laurent, J.-P., Cousin, A. and Fermanian, J-D.: Hedging default risks of CDOs in Markovian contagion models. *Quantitative Finance*, 1469-7696, April 2010.

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

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

[26] 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.

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

[28] 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, 339-371, 2007