Dynamic Hedging of Portfolio Credit Derivatives[*]

Rama Cont[†‡] and Yu Hang Kan[‡]

Abstract. We compare the performance of various hedging strategies for index collateralized debt obligation (CDO) tranches across a variety of models and hedging methods during the recent credit crisis. Our empirical analysis shows evidence for market incompleteness: a large proportion of risk in the CDO tranches appears to be unhedgeable. We also show that, unlike what is commonly assumed, dynamic models do not necessarily perform better than static models, nor do high-dimensional bottom-up models perform better than simpler top-down models. When it comes to hedging, top-down and regression-based hedging with the index provide significantly better results during the credit crisis than bottom-up hedging with single-name credit default swap (CDS) contracts. Our empirical study also reveals that while significantly large moves—"jumps"—do occur in CDS, index, and tranche spreads, these jumps do not necessarily occur on the default dates of index constituents, an observation which shows the insufficiency of some recently proposed portfolio credit risk models.

1. Introduction. Static factor models, in particular the Gaussian copula model [18], have been widely used for hedging portfolio credit derivatives such as collateralized debt obligations (CDOs). In such models, the risk of a CDO tranche is characterized in terms of sensitivities to shifts in risk factors [10, 24]. Accordingly, hedging practices have typically been based on such measures of sensitivity. The most common hedging approach has been to "delta hedge" spread fluctuations using credit default swaps (CDSs).

However, the recent turmoil in credit derivatives markets shows that these commonly used hedging approaches are inefficient. One of the main criticisms has been the lack of well-defined dynamics for the risk factors in such static models, which prevents any model-based assessment of hedging strategies. In particular, delta hedging of spread risk is loosely justified using a Black–Scholes analogy which does not necessarily hold, and the corresponding hedge ratios, the spread-deltas, are in fact computed from a static model without spread risk. Indeed, delta hedging of spread risk is not deduced from any theory of derivative replication. Furthermore, delta hedging of spread risk ignores default risk and jumps in the spreads, which appeared to be critical for risk management during the difficult market environment in 2008. Although gamma hedging can improve performance slightly, it is not sufficient to solve these issues.

[†]Laboratoire de Probabilités et Modèles Aléatoires, CNRS-UPMC, 4 Place Jussieu, Paris 75252, France (Rama.Cont@upmc.fr).

[‡]IEOR Department, Columbia University, New York, NY (yk2246@columbia.edu).

Finally, the common approach to price portfolio credit derivatives using copula-based models does not guarantee the absence of arbitrage. Cont, Deguest, and Kan [6] show that pricing CDO tranches based on linear interpolation of the base correlations in a one-factor Gaussian copula model can lead to static arbitrage.

Given the deficiencies of copula-based hedging methods, alternatives have been proposed to tackle the problem of hedging portfolio credit derivatives. Durand and Jouanin [10] describe common hedging practices for credit derivatives and correctly point out the inconsistency between most of the pricing models, where the sole risk is in the occurrence of defaults, and delta hedging strategies, where the trader seeks to protect his/her portfolio against small movements in CDS spreads. Bielecki, Jeanblanc, and Rutkowski [3] show that, in a bottom-up hazard process framework driven by a Brownian motion, perfect replication is possible by continuously trading a sufficient number of liquid CDS contracts. Bielecki, Crépey, and Jeanblanc [2] discuss hedging performance in bottom-up and top-down models using simulation but do not comment on the performance of such strategies in a real market setting.

Laurent, Cousin, and Fermanian [17] study hedging of synthetic CDO tranches in a local intensity framework without spread risk, and show that CDO tranches can then be replicated by a self-financing portfolio consisting of the index default swap and a risk-free bond. However, as we will show in section 3, spread fluctuation is a major source of risk even in the absence of defaults, so failure to incorporate spread risk can lead to unrealistic conclusions.

Using a more realistic approach which acknowledges market incompleteness and incorporates both spread risk and default contagion, Frey and Backhaus [14] observe significant differences between the sensitivity-based hedging strategies computed in the Gaussian copula framework and the dynamic hedging strategies derived in their setup. They also show that variance-minimization hedging provides a model-based endogenous interpolation between the hedging against spread risk and default risk.

Giesecke, Goldberg, and Ding [15] discuss an alternative hedging approach based on a self-exciting process for portfolio defaults and compare the hedging performance for equity CDO tranches in September 2008 with a Gaussian copula model.

The hedging methods in these studies approach the problem from different, often incompatible, standpoints, and a systematic comparison of the resulting hedge ratios and the subsequent hedging performance has not been done in a realistic setting with market data. Needless to say, in order for such a comparison to be meaningful, the models need to be calibrated to the same data set. The very feasibility of this calibration is a serious (computational) constraint which excludes many models discussed in the literature, leading us to focus on the class of tractable models.

Motivated by previous studies indicating the impact of model uncertainty on the pricing and hedging derivative instruments [5], our objective is to assess the performance of hedging strategies of index CDO tranches derived under various model assumptions. We compare the performance of different dynamic hedging strategies across a range of models including the Gaussian copula model, a multiname reduced-form model introduced by Duffie and Gârleanu [9], a Markovian portfolio default model [16], and a two-factor model with spread and default risk [1]. Strategies considered include delta hedging of spread risk, hedging of default risk, variance minimization (quadratic hedging), and regression-based hedging.

In particular we shall attempt to address some important questions which have been left

unanswered by previous studies:

- How did various hedging strategies for CDOs perform during major credit events in 2008?
- Do complete market models provide the right insight for hedging credit derivatives?
- How good are delta hedging strategies for CDO tranches?
- Does gamma hedging improve hedging performance?
- Do hedge ratios based on jump-to-default fare better than sensitivity-based hedge ratios?
- Are dynamic models better for hedging than static models?
- Do hedging strategies using single-name CDSs perform better than hedges using the index?
- Are bottom-up models more suitable for hedging than top-down models?

This article is structured as follows. Section 2 describes the cash flow structure of credit default swaps, index default swaps, and index CDO tranches. Section 3 presents the dataset used for the empirical analysis and describes some important statistical features of the CDO and CDS markets. Section 4 introduces the models under consideration and discusses procedures used for parameter calibration. Section 5 discusses the hedging strategies under consideration. Section 6 compares the performance of different strategies for the hedging of index tranches in 2008. Section 7 summarizes our main findings and discusses some implications.

2. Credit derivatives. A credit derivative is a financial instrument whose payoff depends on the losses due to defaults of the reference obligors (debt instruments). A portfolio credit derivative is a credit derivative whose payoffs depend on default losses in a reference portfolio of obligors. We will consider here index credit derivatives, for which the underlying portfolio is an equally weighted portfolio, such as the CDX or iTraxx indices. Typically, the payoffs depend only on the aggregate loss due to defaults in the index, not on the identity of the defaulting firm.

Consider an equally weighted portfolio consisting of n obligors, and assume for simplicity a constant recovery rate R (typically assumed to be 40%) and deterministic interest rates. Let τ_i be the default time of obligor i. The portfolio loss (in fraction of total notional value) at time t is equal to

$$L_t = \frac{1-R}{n} \sum_{i=1}^{n} 1_{\tau_i \leq t} = \frac{1-R}{n} N_t,$$

where N_t is the number of defaults by time t. The portfolio loss (L_t) is modeled as a stochastic process on a (filtered) probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbb{Q})$, where Ω is the set of market scenarios, $(\mathcal{F}_t)$ represents the flow of information, and $\mathbb{Q}$ is a risk-neutral probability measure representing the market pricing rule.

We will consider the three most commonly traded credit derivatives: credit default swaps (CDSs), index default swaps (index), and collateralized debt obligations (CDOs). All three derivatives are swap contracts between two parties, a protection buyer and a protection seller, whereby the protection buyer is compensated for the loss generated by the default of a reference obligor (CDS) or defaults from a pool of obligors (index and CDO). In return, the protection buyer pays a premium to the protection seller. A more detailed description of these products can be found in [7, 14, 20].

2.1. Credit default swaps. Consider a reference obligor i and its corresponding CDS contract initiated at time 0 with unit notional and payment dates $T_1 < T_2 < \cdots < T_M$, where $T = T_M$ is the maturity date. Assume that the default payments are made on the next payment date;[1] then if obligor i defaults between time T_{m-1} and T_m, the default payment at time T_m is equal to $1 - R$. On the other hand, if obligor i has not defaulted yet at time T_m, the protection seller will receive a premium payment $s_0^i(T_m - T_{m-1})$, where s_0^i is the CDS spread that has been determined at the inception.

The par CDS spread s_t^i quoted in the market at date t is defined as the value of the spread which sets the present values of the default leg and the premium leg equal. The mark-to-market value of a protection seller's position at time t is equal to the difference between the net present values of the two legs:

$$(1) \qquad V_t^i = (s_0^i - s_t^i) \sum_{T_m > t} B(t, T_m)(T_m - T_{m-1})\mathbb{Q}(\tau_i > T_m | \mathcal{F}_t),$$

where $B(t, T_m)$ is the discount factor from time t to T_m. In what follows we will refer to the value of the protection seller's position as the mark-to-market value, and we will use $\mathbf{s}_t^{cds} = (s_t^1, \ldots, s_t^n)$ to denote the vector of constituent CDS spreads and $\mathbf{D}_t = (1_{\tau_1 \leq t}, \ldots, 1_{\tau_n \leq t})$ to denote the vector of default indicators at time t.

2.2. Index default swap. Index default swaps are now commonly traded on various credit indices such as iTraxx and CDX series which are equally weighted indices of CDSs. In an index default swap transaction initiated at time 0, a protection seller agrees to pay all default losses in the index in return for a fixed periodic spread s_0^{idx} paid on the total notional of obligors remaining in the index.

The index default swap par spread s_t^{idx} quoted in the market at time t is defined as the value of the spread which balances the present values of the default leg and the premium leg. The mark-to-market value of a protection seller's position at time t is equal to the difference between the two legs, which can be expressed as

$$(2) \qquad V_t^{idx} = \left(s_0^{idx} - s_t^{idx}\right) \sum_{T_m > t} B(t, T_m)(T_m - T_{m-1}) E^{\mathbb{Q}}\left[1 - \frac{N_{T_m}}{n} \Big| \mathcal{F}_t\right].$$

Here, we assume that the outstanding notional value is calculated at payment dates.[2]

2.3. Collateralized debt obligations. Consider a tranche defined by an interval $[a, b]$, $0 \leq a < b \leq 1$, for the portfolio loss normalized by the total notional value of the underlying portfolio. We call a (resp., b) the attachment (resp., detachment) point of the tranche. A synthetic CDO tranche swap is a bilateral contract in which the protection seller agrees to pay all portfolio loss within the interval $[a, b]$ in return for a periodic spread $s_0^{[a,b]}$, which is determined at inception $t = 0$, on the remaining tranche notional value.

---

[1]The default payment is sometimes assumed to be made immediately after the default. Nevertheless, the choice of payment schedule has negligible effects on our analysis in this paper.

[2]A more precise valuation would consider the average outstanding notional value over the time period between the payment dates [4], but the approximation above has negligible effects on our analysis.

The par tranche spread $s_t^{[a,b]}$ quoted in the market at time t is defined as the spread which sets the present values of the default leg and the premium leg to be equal. The mark-to-market value of a protection seller's position (normalized by the total tranche notional value) at time t is equal to the difference between the two legs, which can be expressed as

$$(3) \quad V_t^{[a,b]} = \left(s_0^{[a,b]} - s_t^{[a,b]}\right) \sum_{T_m > t} \frac{B(t, T_m)}{b - a}(T_m - T_{m-1})E^{\mathbb{Q}}\left[(b - L_{T_m})^+ - (a - L_{T_m})^+\Big|\mathcal{F}_t\right].$$

3. Data analysis. Our dataset contains the 5-year CDX North America Investment Grade Series 10 (CDX) index spreads; the standard tranche spreads with attachment/detachment points $0\%, 3\%, 7\%, 10\%, 15\%, 30\%, 100\%$; and the constituent 5-year CDS spreads, all obtained from Bloomberg. The time series runs from 25 March, 2008 until 25 September, 2008. Figure 1 illustrates the time series of the index, tranche $[10\%, 15\%]$, CDS of IBM and Disney Corp.
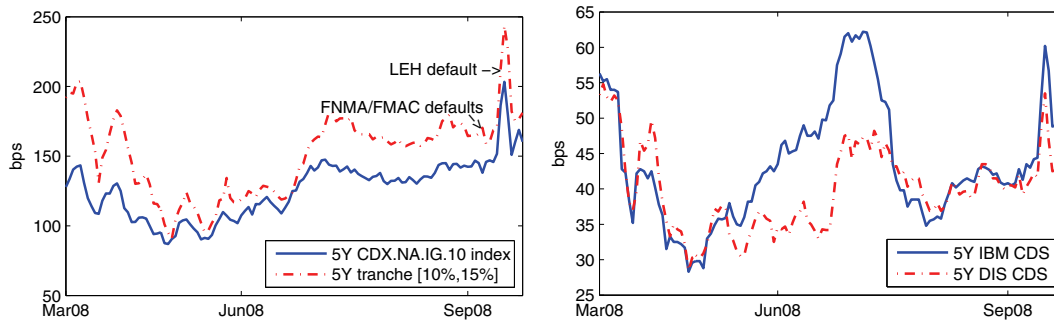

Figure 1. *Left: 5-year CDX.NA.IG.10 index and tranche* $[10\%, 15\%]$ *spreads. Right: 5-year CDS spreads of IBM and Disney Corp. (DIS).*

3.1. Comovements in CDSs and CDO tranches. A CDO hedging strategy should be based on a good understanding of the relation between the profit and loss (P&L) of the hedging instruments, namely the CDSs and the index, and that of the target instruments, the CDO tranches. Given that the P&L is driven mainly by the changes in spreads, this requires a correct representation of comovements in the credit spreads. Figure 2 shows the tranche $[10\%, 15\%]$ daily spread returns against the index and IBM CDS daily spread returns. The crosses and circles represent the data points where spread returns of the two credit derivatives move in the same and opposite directions, respectively. Here we use IBM CDS and the $[10\%, 15\%]$ tranche data for illustration, but similar results are obtained by looking at other constituent CDSs and tranches. From the figure, we can immediately observe two important properties:

1. CDS/index spreads tend to move *together* with the tranche spreads when the movements are *large*.
2. In many cases, CDS/index spreads and tranche spreads move in *opposite* directions, especially when the movements are *small*.

Large comovements of the spreads, or common jumps, can be explained by the exposures of the credit derivatives to common risk factors which undergo large movements. This phe-
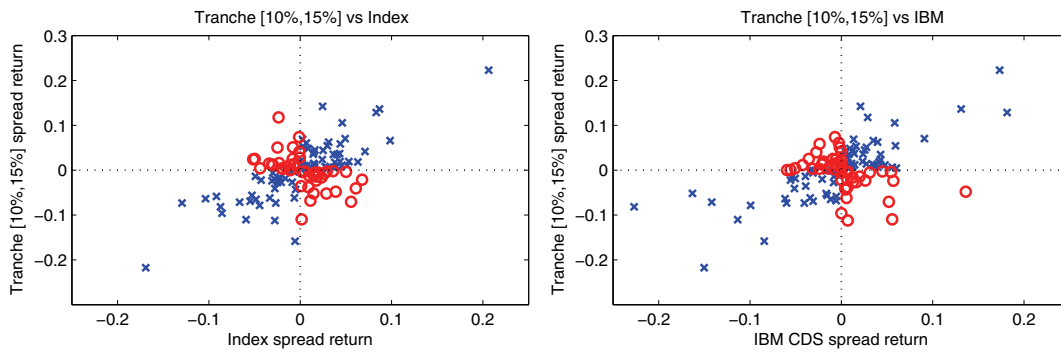
Figure 2. *Tranche* [10%, 15%] *daily spread returns versus index daily spread returns and IBM CDS daily spread returns. Crosses represent data points where spread returns have the same signs (movements in the same direction), and circles represent data points where spread returns have opposite signs (movements in opposite directions).*

Table 1

Conditional correlations between daily spread returns of the index and tranche [10%, 15%] *show evidence of a common heavy-tailed factor.*

Index spread return	Unconditional	> 8%	> 5%	> 1%	< −1%	< −5%	< −8%
Correlation	0.63	0.82	0.82	0.60	0.68	0.71	0.81
Observations	126	4	12	52	43	11	6

nomenon can be seen more clearly in Table 1, which shows that the correlation between the index and tranche spread returns increases substantially if we condition on larger observations.

From a hedging perspective, frequent opposite movements between the CDS/index and the tranche spreads can lead to serious problems, because most hedging strategies imply positive hedge ratios with respect to the CDSs or the index. When the values of the hedging instruments and the tranche move in opposite directions, those strategies may fail to reduce the exposure of the tranche positions or, more seriously, can substantially amplify the overall exposure. As we will see in our empirical study in section 6, this problem frequently arises in common hedging strategies.

3.2. Impact of defaults on credit spreads. Our sample period covers several important credit events: the takeover of Fannie Mae and Freddie Mac, which led to losses in the CDX, and the bankruptcy of Lehman Brothers, which led to a significant shock to the market.

During the sample period, Fannie Mae and Freddie Mac were taken over by the U.S. government on 7 September, 2008, which generated a credit event in the CDX reference portfolio. According to Bloomberg, the recovery rates of the 5-year senior CDS contracts of Fannie Mae and Freddie Mac were 92% and 94%, respectively, which will be used to determine the losses in our empirical study. On the other hand, although Lehman Brothers is *not* a reference obligor in the CDX, Figure 1 shows that there is considerable upward movement of the spreads on the next business day after Lehman Brothers announced bankruptcy.

Table 2 shows the daily spread returns on the next business day after Fannie Mae/Freddie Mac and Lehman Brothers credit events in units of sample standard deviation. Interestingly, we observe that the IBM, the index, and the super senior tranche [30%, 100%] spreads *decrease*

	IBM	Index	0%–3%	3%–7%	7%–10%	10%–15%	15%–30%	30%–100%
8-Sep-08	−0.66	−1.10	0.16	0.05	0.04	0.44	0.02	−0.06
16-Sep-08	3.23	4.55	3.51	4.92	4.45	3.91	4.14	4.05

on the next business day after Fannie Mae/Freddie Mac were taken over. Although the spreads of other tranches do increase, the magnitudes of these changes are rather small, less than 0.5 standard deviations.

On the other hand, we *do* observe jumps in CDX spreads, but not necessarily on the dates corresponding to constituent defaults. The typical example is on 16 September, 2008 when Lehman Brothers filed for bankruptcy. Although Lehman Brothers is not a constituent of the CDX, the IBM, index, and tranche spreads increase by as much as 4.9 standard deviations, which are substantial upward moves and can be attributed to jumps.

These observations have two important implications. First, they show that jumps in the spreads are not necessarily tied to defaults in the underlying portfolio, as is the case in Markovian contagion models [1, 16] and self-exciting models [12, 15], where jumps occur only on portfolio default dates. Jumps may be caused by information external to the portfolio, such as macroeconomic events, of which the Lehman credit event is an example. Second, jump sizes at default dates appear to depend on the severity of the events, with lower recovery rate implying fewer or no upward jumps in the spreads. This suggests that models with constant jumps in the default intensity at each default are insufficient for capturing the impact of defaults on the spread movements: this impact should depend on the severity of loss in the given default, as suggested in [15]. We note that this may be difficult in practice, since recovery rates are usually not observable immediately after default and can be determined only after liquidation.

4. Models for portfolio credit derivatives. We will consider four different modeling approaches in our analysis: the one-factor Gaussian copula model [18], a bottom-up affine jump-diffusion model [9], a local intensity model [6, 7, 16, 21, 25], and a top-down bivariate spread-loss model [1].

4.1. Gaussian copula model. The one-factor Gaussian copula model [18] is a standard market reference for pricing CDO tranches, in which the default times are constructed as

$$\tau_i = F_i^{-1}(\Phi(\rho M_0 + \sqrt{1 - \rho^2} M_i)),$$

where M_0, M_i are independent standard normal random variables, $\Phi(.)$ is the standard normal distribution function, F_i is the marginal distribution of τ_i, and ρ is a correlation parameter. The distribution function $F_i(.)$ is calibrated to the single-name CDS spreads by assuming a constant hazard rate.[3] Then, we fit one correlation to each tranche, which is a situation known as *compound correlations*. If multiple correlations give the same tranche spread, we will choose the smallest one.

---

[3]The hazard rate term structure is usually assumed to be piecewise constant. Since we consider only CDS with one maturity for each obligor, it reduces to a constant hazard rate.

There are two reasons why we consider compound correlations instead of the base correlations [19] for calibration. As noted by Morgan and Mortensen [22], we found that computing the spread-deltas while keeping base correlations fixed can lead to a negative sensitivity of a tranche with respect to a change in the CDS spreads. Therefore, even if the CDS and the tranche spreads move in the same direction, especially when the movement is large, the negative spread-deltas will have the wrong sign and give poor hedging results. Second, unlike the Black–Scholes implied volatility, which is in a one-to-one correspondence to the vanilla options prices, base correlations are not guaranteed to exist. For instance, we were not able to calibrate the base correlation for 15% strike on many of the dates in our sample.

In order to express the hedging positions in later sections, it is convenient to write the mark-to-market values of the credit derivatives as functions of the modeling variables. Given the CDS spreads $\mathbf{s}_t^{cds} = \mathbf{s}^{cds} = (s^1, \ldots, s^n)$, the default indicators $\mathbf{D}_t = \mathbf{D}$, and the set of compound correlations $\boldsymbol{\rho}_t = \boldsymbol{\rho}$, we write the mark-to-market values of CDS i, the index, and a tranche $[a, b]$ at time t computed under the Gaussian copula model as $V_{gc}^i(t, s^i)$, $V_{gc}^{idx}(t, \mathbf{s}^{cds}, \boldsymbol{\rho}, \mathbf{D})$, and $V_{gc}^{[a,b]}(t, \mathbf{s}^{cds}, \boldsymbol{\rho}, \mathbf{D})$, respectively.

4.2. Affine jump-diffusion model. Various dynamic reduced-form models have been proposed to overcome some of the shortcomings of static copula-based models. An example of such a model used in industry is the affine jump-diffusion model introduced by Duffie and Gârleanu [9]. In this model the default time τ_i of an obligor i is modeled as a random time with a stochastic intensity (λ_t^i) given by

$$(4) \qquad \lambda_t^i = X_t^i + a^i X_t^0,$$

where the idiosyncratic risk factors $(X_t^1, \ldots, X_t^n)$ and the common (macro) risk factor (X_t^0) are independent affine jump-diffusion processes

$$dX_t^i = \kappa(\theta - X_t^i)dt + \sigma\sqrt{X_t^i}dW_t^i + dJ_t^i,$$

where (W_t^i) are standard Brownian motions and (J_t^i) are compound Poisson process with exponentially distributed jump sizes. The conditional survival probability is then given by

$$\mathbb{Q}(\tau_i > T|\mathcal{F}_t) = E^{\mathbb{Q}}\left[\exp\left(-\int_t^T \lambda_u^i du\right)\Big|\mathcal{F}_t\right].$$

We will denote by $\mathbf{X}_t = (X_t^0, X_t^1, \ldots, X_t^n)$ the risk factor values at time t. In order to calibrate the model, we follow the algorithm proposed by Eckner [11]. The tractability of this model relies on the conditional independence assumption of the default processes and also an analytical formula for the characteristic function of the affine jump-diffusion process. We refer readers to [9, 11] for the details of the calibration procedure.

Since $(\mathbf{X}_t, \mathbf{D}_t)$ is a Markov process, given the values $\mathbf{X}_t = \mathbf{X} = (X^0, X^1, \ldots, X^n)$ and the default indicators $\mathbf{D}_t = \mathbf{D}$, we can write the mark-to-market values of CDS i, the index, and a tranche $[a, b]$ at time t computed under the affine jump-diffusion model as $V_{af}^i(t, X^0, X^i)$, $V_{af}^{idx}(t, \mathbf{X}, \mathbf{D})$, and $V_{af}^{[a,b]}(t, \mathbf{X}, \mathbf{D})$, respectively. The mark-to-market values are computed as given in (1), (2), and (3).

4.3. Local intensity models. Local intensity models [6, 7, 16, 21, 25] are top-down models in which the number of defaults N_t in a reference portfolio is modeled as a Markov point process with an intensity $\lambda_t = f(t, N_{t-})$: the *portfolio default intensity* is a (positive) function of time and the number of defaulted obligors.

We consider the following parametrization of the local intensity function, introduced by Herbertsson [16]:

$$(5) \qquad \lambda_t = (n - N_{t-}) \sum_{k=0}^{N_{t-}} b_k,$$

where $\{b_k\}$ are the parameters. The interpretation of (5) is that the portfolio default intensity jumps by an amount b_k when the kth default happens. There is no sign restriction on $\{b_k\}$ as long as the portfolio default intensity remains positive. As in [16], we parameterize $\{b_k\}$ as

$$(6) \qquad b_k = \begin{cases} b^{(1)}, & 1 \leq k < \mu_1, \\ b^{(2)}, & \mu_1 \leq k < \mu_2, \\ \vdots \\ b^{(I)}, & \mu_{I-1} \leq k < \mu_I = n, \end{cases}$$

where $1, \mu_1, \ldots, \mu_I$ is a partition of $\{1, \ldots, n\}$ which includes the attachment points of the tranches. The local intensity model is a Markovian top-down model in which the only risk factor is the loss process. Therefore we can express the mark-to-market values of the index and a tranche $[a, b]$ at date t computed as functions $V_{lo}^{idx}(t, N)$ and $V_{lo}^{[a,b]}(t, N)$ of the number of defaults and time.

4.4. Bivariate spread-loss model. One major shortcoming of the local intensity model is that spreads have piecewise-deterministic dynamics—i.e., no "volatility"—between defaults. As we have seen in Figure 1, credit derivative positions fluctuate substantially in value even in the absence of defaults in the underlying credit portfolio, so a hedging strategy based on the jump-to-default ratio may lead to poor performance. A more realistic picture is given by a two-factor top-down model [1, 21] which accounts for both default risk and spread volatility by allowing the portfolio default intensity to depend on the number of defaults and a factor driving spread volatility:

$$\lambda_t = F(t, N_{t-}, Y_t).$$

Arnsdorff and Halperin [1] model the number of defaults N_t in a reference portfolio as a point process which has a (portfolio default) intensity (λ_t) that follows

$$(7) \qquad \lambda_t = e^{Y_t}(n - N_{t-}) \sum_{k=0}^{N_{t-}} b_k,$$

where $\{b_k\}$ are parameters and the second factor (Y_t) generates spread volatility between default dates which follows an Ornstein–Uhlenbeck process,

$$dY_t = -\kappa Y_t dt + \sigma dW_t,$$

Model	CDS	Index	0%–3%	3%–7%	7%–10%	10%–15%	15%–30%
Gaussian copula	6.04	0.00	0.00	0.18	0.05	0.00	0.00
Affine jump-diffusion	14.53	21.67	14.56	5.05	10.40	13.91	6.17
Local intensity	-	6.29	1.60	0.95	0.46	0.34	1.71
Bivariate spread-loss	-	6.67	1.67	1.02	0.42	0.39	1.79

where (W_t) is a standard Brownian motion. Notice that the parameters $\{b_k\}$ will provide enough degrees of freedom to fit the CDO tranche spreads on a given date, so the remaining parameters (κ, σ) are estimated from time series of tranche spreads as follows:

- On the first sample day: Set $Y_0 = 0$, $\kappa = 0.3$, and $\sigma = 0.7$ and calibrate $\{b_k\}$ to index and tranche spreads on day 1.
- On the jth sample day:
    1. Fix $\{b_k\}$ as those calibrated on day $j - 1$.
    2. Calibrate Y_0, κ, and σ by minimizing the mean square pricing error of day $j - 10$, $j - 9, \ldots, j$.
    3. Calibrate $\{b_k\}$ to the index and the tranche spreads on day j.

Since (N_t, Y_t) is a Markov process, given the values $N_t = N$ and $Y_t = Y$, we can write the mark-to-market values of the index and a tranche $[a, b]$ at time t computed under the bivariate spread-loss model as functions $V_{bi}^{idx}(t, Y, N)$ and $V_{bi}^{[a,b]}(t, Y, N)$ of the state variables.

4.5. Calibration results. All models are calibrated to the same market data using a 40% recovery rate. Table 3 shows the root mean square calibration error (RMSE). Since we calibrate the Gaussian copula model to each tranche with a different correlation (compound correlations), it by design gives good fits to the tranche spreads. The discrepancy of CDS spreads is due to the adjustment to match the index spreads. Top-down models are amendable to calibration to the market data as well. The RMSE are well within 2% for all tranches and around 5% for the CDS and the index. On the other hand, the Duffie–Gârleanu affine jump-diffusion does not calibrate market data as well as the top-down models. The CDS and the tranche spreads have RMSE at about 10%, which is still reasonable, but the fit to the index spread has RMSE larger than 20%, which is a poor fit. This is due to the fact that its calibration involves a high dimensional nonlinear optimization problem which is not guaranteed to converge. Therefore, we will consider only a hedging strategy using the single-name CDS as the hedging instruments in the affine jump-diffusion framework, so that the poor calibration to the index will not affect our analysis significantly. Note that, for all models under consideration, we have experienced a poor fit to the super senior tranche $[30\%, 100\%]$, even with the Gaussian copula model. Since poor calibration leads to inaccurate computation of the mark-to-market values and the hedge ratios, we will omit the $[30\%, 100\%]$ tranche in what follows.

5. Hedging strategies. Our objective is to hedge a position in a tranche $[a, b]$ using the constituent CDSs and the index, and sometimes with an additional tranche $[l, u]$. We will now introduce different dynamic hedging strategies that aim to achieve this task.

We assume a continuously rebalancing framework and let (ϕ_t^i), (ϕ_t^{idx}), and $(\phi_t^{[l,u]})$ be predictable processes which denote the hedging positions in CDS i, the index, and a tranche

$[l, u]$, respectively. In addition, we will use the same notation as in section 4 to represent the mark-to-market values of the credit derivatives computed under different models. All hedging strategies are implemented using daily rebalancing.

5.1. Delta hedging of single-name spread movements. The most common approach for hedging CDO tranches is to hedge against small changes in the single-name CDS spreads [10, 24]. In practice, traders usually consider delta hedging under the Gaussian copula model where the corresponding hedging position in CDS i is known as the *spread-delta*:

$$(8) \qquad \phi_t^i = \frac{\delta_{s^i} V_{gc}^{[a,b]}(t, \mathbf{s}_{t-}^{cds}, \boldsymbol{\rho}_{t-}, \mathbf{D}_{t-})}{\delta_{s^i} V_{gc}^i(t, \mathbf{s}_{t-}^i)},$$

where

$$\delta_{s^i} V_{gc}^{[a,b]}(t, \mathbf{s}^{cds}, \boldsymbol{\rho}, \mathbf{D}) = V_{gc}^{[a,b]}(t, \mathbf{s}^{cds} + \mathbf{e}_i, \boldsymbol{\rho}, \mathbf{D}) - V_{gc}^{[a,b]}(t, \mathbf{s}^{cds}, \boldsymbol{\rho}, \mathbf{D}),$$

$$\delta_{s^i} V_{gc}^i(t, s^i) = V_{gc}^i(t, s^i + 1bp) - V_{gc}^i(t, s^i)$$

are the changes in value of a tranche $[a, b]$ and CDS i with respect to an increase in the CDS spread of obligor i by 1 basis point, while the correlations and other CDS spreads remain unchanged. $\mathbf{e}_i \in \mathbb{R}^n$ is a vector with all entries equal to 0 except for the ith entry equal to 1 basis point.

In order to compute the spread-deltas, we first calibrate the one-factor Gaussian copula model [18] to the market CDS and the tranche spreads as described in section 4.1. After that, we perturb the CDS spread of, say, obligor i by 1 basis point while keeping all other CDS spreads and the correlations unchanged. Then, we *recalibrate* the hazard rate function of obligor i and compute the new values for CDS i and the tranche. The spread-delta defined by (8) is the ratio of the change in the tranche value to the change in the CDS value.

The main drawback of implementing the spread-deltas (8) is the absence of well-defined dynamics for the single-name CDS spreads in the Gaussian copula framework. On the other hand, we can consider delta hedging under the dynamic affine jump-diffusion model [9]. Hedging moves in the single-name CDS spreads is then equivalent to hedging against changes in the idiosyncratic risk factor. The corresponding position in CDS i is equal to

$$(9) \qquad \phi_t^i = \frac{\partial_{X^i} V_{af}^{[a,b]}(t, \mathbf{X}_{t-}, \mathbf{D}_{t-})}{\partial_{X^i} V_{af}^i(t, X_{t-}^0, X_{t-}^i)},$$

where $\partial_{X^i} V_{af}^{[a,b]}(t, \mathbf{X}, \mathbf{D})$ and $\partial_{X^i} V_{af}^i(t, X^0, X^i)$ are the partial derivatives with respect to X^i which can be approximated by finite differences. Note that one of the main differences between the hedge ratio (9) and the spread-delta (8) is that there is no recalibration involved when computing (9) under the affine jump-diffusion model.

5.2. Delta hedging of index spread movements. In section 3.1, we observe that the CDS and the tranche spreads appear to be driven by some common risk factors. Therefore, we may argue that it is also important to hedge against *global* movements in the CDS spreads. We use the Gaussian copula model and enter positions in the index to neutralize the *index spread-delta*:

$$(10) \qquad \phi_t^{idx} = \frac{\Delta_{gc}^{[a,b]}(t)}{\Delta_{gc}^{idx}(t)},$$
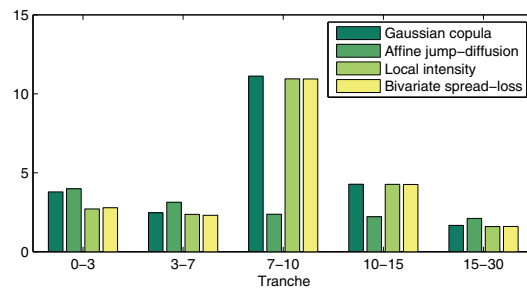
Figure 3. *Index spread-deltas. Data: CDX.NA.IG.S10 on 25 March, 2008.*

where

$$\Delta_{gc}^{[a,b]}(t) = V_{gc}^{[a,b]}(t, \mathbf{s}_{t-}^{cds} + \mathbf{e}, \boldsymbol{\rho}_{t-}, \mathbf{D}_{t-}) - V_{gc}^{[a,b]}(t, \mathbf{s}_{t-}^{cds}, \boldsymbol{\rho}_{t-}, \mathbf{D}_{t-}),$$
$$\Delta_{gc}^{idx}(t) = V_{gc}^{idx}(t, \mathbf{s}_{t-}^{cds} + \mathbf{e}, \mathbf{D}_{t-}) - V_{gc}^{idx}(t, \mathbf{s}_{t-}^{cds}, \mathbf{D}_{t-}),$$

and $\mathbf{e} \in \mathbb{R}^n$ is a vector with all entries equal to 1 basis point. Notice that $\Delta_{gc}^{[a,b]}(t)$ and $\Delta_{gc}^{idx}(t)$ are the changes of a tranche $[a, b]$ and the index values with respect to a parallel shift in all CDS spreads by 1 basis point while keeping the correlations unchanged. Computation of the index spread-delta is the same as for the spread-deltas, except that we need to shift all CDS spreads by 1 basis point.

This strategy also has the advantage of being cost-effective. Unlike delta hedging individual CDS fluctuations, which requires rebalancing multiple hedging positions, this strategy only requires adjusting the position in the index.

If we consider the CDS/index spreads and the correlations as the market inputs, which is analogous to the stock price and implied volatility moves in equity derivatives markets, the index spread-delta (10) can also be computed by models other than the Gaussian copula model. The procedure is similar to the case described above in which we first calibrate the models to the CDS/index spreads and the correlations. Then, we recalibrate the models to the perturbed CDS/index spreads while keeping the correlations unchanged. The index spread-delta is the ratio of the change in tranche value over the change in the index value.

Figure 3 shows the index spread-deltas (10) computed under different models. Interestingly, we observe that the index spread-deltas are very similar across the models, except those for tranches $[7\%, 10\%]$ and $[10\%, 15\%]$ computed from the affine jump-diffusion model. In fact, this discrepancy is due only to the fact that the affine jump-diffusion model does not calibrate well to the market data.

The similarity of the index spread-deltas across the models implies that there is no point in using a more sophisticated model if its only use is to delta hedge spread risk. The standard one-factor Gaussian copula model would be sufficient to carry out this strategy. Indeed, a more meaningful hedging strategy for the dynamic models is to hedge against the underlying risk factors specified in the modeling framework, taking into account the dynamics of these factors and their correlations.

5.3. Delta and gamma hedging of index spread movements. By analogy with gamma hedging of equity derivatives, one may consider hedging the second-order changes in the tranche values due to fluctuations in the CDS spreads. We consider positions in the index

and a tranche $[l, u]$ such that

$$\Delta_{gc}^{[a,b]}(t) = \phi_t^{[l.u]} \, \Delta_{gc}^{[l,u]}(t) + \phi_t^{idx} \, \Delta_{gc}^{idx}(t),$$
$$\Gamma_{gc}^{[a,b]}(t) = \phi_t^{[l.u]} \, \Gamma_{gc}^{[l,u]}(t) + \phi_t^{idx} \, \Gamma_{gc}^{idx}(t),$$

where

(11)
$$\Gamma_{gc}^{[a,b]}(t) = V_{gc}^{[a,b]}(t, \mathbf{s}_{t-}^{cds} + \mathbf{e}, \boldsymbol{\rho}_{t-}, \mathbf{D}_{t-}) - 2V_{gc}^{[a,b]}(t, \mathbf{s}_{t-}^{cds}, \boldsymbol{\rho}_t, \mathbf{D}_{t-}) + V_{gc}^{[a,b]}(t, \mathbf{s}_{t-}^{cds} - \mathbf{e}, \boldsymbol{\rho}_{t-}, \mathbf{D}_{t-}),$$

(12)
$$\Gamma_{gc}^{idx}(t) = V_{gc}^{idx}(t, \mathbf{s}_{t-}^{cds} + \mathbf{e}, \mathbf{D}_{t-}) - 2V_{gc}^{idx}(t, \mathbf{s}_{t-}^{cds}, \mathbf{D}_{t-}) + V_{gc}^{idx}(t, \mathbf{s}_{t-}^{cds} - \mathbf{e}, \mathbf{D}_{t-})$$

are the *gammas* of a tranche $[a, b]$ and the index, or equivalently the second-order finite differences in the values with respect to 1 basis point shifting of all CDS spreads. Solving for the hedge ratios, we have

(13)
$$\phi_t^{idx} = \frac{\Delta_{gc}^{[a,b]}(t)\Gamma_{gc}^{[l,u]}(t) - \Delta_{gc}^{[l,u]}(t)\Gamma_{gc}^{[a,b]}(t)}{\Delta_{gc}^{idx}(t)\Gamma_{gc}^{[l,u]}(t) - \Delta_{gc}^{[l,u]}(t)\Gamma_{gc}^{idx}(t)},$$

(14)
$$\phi_t^{[l,u]} = \frac{\Delta_{gc}^{[a,b]}(t)\Gamma_{gc}^{idx}(t) - \Delta_{gc}^{idx}(t)\Gamma_{gc}^{[a,b]}(t)}{\Delta_{gc}^{[l,u]}(t)\Gamma_{gc}^{idx}(t) - \Delta_{gc}^{idx}(t)\Gamma_{gc}^{[l,u]}(t)}.$$

Note that we can also delta hedge against movements in the single-name CDS spreads in the case of gamma hedging. However, we need to solve an ill-conditioned linear system which may lead to unstable hedge ratios. Moreover, as we will see in section 6, the main component of changes in the single-name CDS spreads is a parallel move which is already reflected in the index spread. Thus, the inclusion of single-name CDS corresponds to hedging higher-order principal components, which have a smaller impact on the variance of the portfolio. Therefore, we do not include single-name CDS hedges in our gamma hedging analysis.

Unlike the situation in a Black–Scholes model, where the gamma of a call or put option is always positive and the gamma of a long position in a call or put option can be neutralized by shorting another call or put option, such simple relations fail to hold in the Gaussian copula framework for CDO tranches. Figure 4 shows the gammas of tranches $[0\%, 3\%]$ and $[3\%, 7\%]$ computed for various days in the sample. Observe that the gammas can be positive or negative, even for the equity tranche. Moreover, the gammas of the two tranches do not have any clear relationship, in the sense that they do not always have the same sign. These results also suggest that the empirical performance of gamma hedging may depend on the choice of the hedging tranches, so we will consider below different choices of tranches as hedging instruments.

5.4. Hedging parallel shifts in correlations. In addition to hedging against changes in spreads, by analogy with "vega" hedging in the Black–Scholes model, one can argue that it is also important to manage the risk of the fluctuation in another parameter of the Gaussian copula model, the implied correlation. We will consider scenarios where all (compound) correlations shift by the same magnitude.
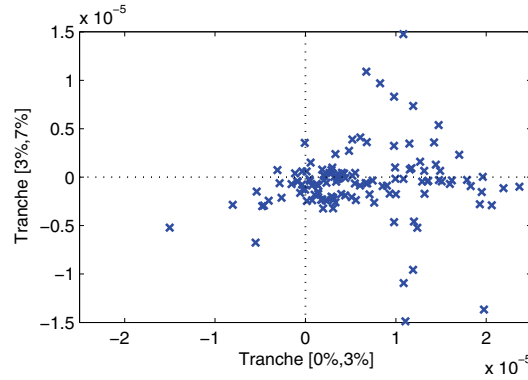
Figure 4. *Values of gammas* $\Gamma_{gc}^{[a,b]}(t)$ *for* $[0\%, 3\%]$ *and* $[3\%, 7\%]$ *CDX tranches. Each point represents one day in the sample.*

A joint hedge with respect to small changes in spreads and correlation then requires us to enter positions in the CDSs and a tranche $[l, u]$ such that

$$\delta_{s^i} V_{gc}^{[a,b]}(t, \mathbf{s}_{t-}^{cds}, \boldsymbol{\rho}_{t-}, \mathbf{D}_{t-}) = \phi_t^{[l,u]} \delta_{s^i} V_{gc}^{[l,u]}(t, \mathbf{s}_{t-}^{cds}, \boldsymbol{\rho}_{t-}, \mathbf{D}_{t-}) + \phi_t^i \delta_{s^i} V_{gc}^i(t, s_{t-}^i), \quad i = 1, \dots, n,$$

$$\delta_{\rho} V_{gc}^{[a,b]}(t, \mathbf{s}_{t-}^{cds}, \boldsymbol{\rho}_{t-}, \mathbf{D}_{t-}) = \phi_t^{[l,u]} \delta_{\rho} V_{gc}^{[l,u]}(t, \mathbf{s}_{t-}^{cds}, \boldsymbol{\rho}_{t-}, \mathbf{D}_{t-}),$$

where

$$\delta_{\rho} V_{gc}^{[a,b]}(t, \mathbf{s}^{cds}, \boldsymbol{\rho}, \mathbf{D}) = V_{gc}^{[a,b]}(t, \mathbf{s}^{cds}, \boldsymbol{\rho} + 0.1\%, \mathbf{D}) - V_{gc}^{[a,b]}(t, \mathbf{s}^{cds}, \boldsymbol{\rho}, \mathbf{D})$$

is the change in the tranche value with respect to an increase in all compound correlations by 0.1%. Since the index is insensitive to the correlations, we must consider another tranche as a hedging instrument. Solving for the hedge ratios, we have

$$\phi_t^{[l,u]} = \frac{\delta_{\rho} V_{gc}^{[a,b]}(t, \mathbf{s}_{t-}^{cds}, \boldsymbol{\rho}_{t-}, \mathbf{D}_{t-})}{\delta_{\rho} V_{gc}^{[l,u]}(t, \mathbf{s}_{t-}^{cds}, \boldsymbol{\rho}_{t-}, \mathbf{D}_{t-})},$$

$$\phi_t^i = \frac{\delta_{s^i} V_{gc}^{[a,b]}(t, \mathbf{s}_{t-}^{cds}, \boldsymbol{\rho}_{t-}, \mathbf{D}_{t-}) \delta_{\rho} V_{gc}^{[l,u]}(t, \mathbf{s}_{t-}^{cds}, \boldsymbol{\rho}_{t-}, \mathbf{D}_{t-}) - \delta_{\rho} V_{gc}^{[a,b]}(t, \mathbf{s}_{t-}^{cds}, \boldsymbol{\rho}_{t-}, \mathbf{D}_{t-}) \delta_{s^i} V_{gc}^{[l,u]}(t, \mathbf{s}_{t-}^{cds}, \boldsymbol{\rho}_{t-}, \mathbf{D}_{t-})}{\delta_{s^i} V_{gc}^i(t, s_{t-}^i) \, \delta_{\rho} V_{gc}^{[l,u]}(t, \mathbf{s}_{t-}^{cds}, \boldsymbol{\rho}_{t-}, \mathbf{D}_{t-})}.$$

Note that this strategy does not take into account the *comovements* in the index and in correlations. This is simply a first-order sensitivity-based hedge against moves in the CDS spreads and the correlation movements.

5.5. Hedging default risk. Although the common hedging practice is to protect against CDS spread fluctuations, the occurrence of defaults is also a major source of risk of a CDO tranche. The natural strategy to hedge against constituent defaults is to enter positions in the CDS or the index according to the *jump-to-default ratio*, which is defined as the ratio of the change in the tranche value over the change in the index (or CDS) value with respect to one additional default.

Since the top-down models, such as the local intensity model, focus on the next-to-default rather than the default of a specific obligor, there is only one jump-to-default ratio to consider,
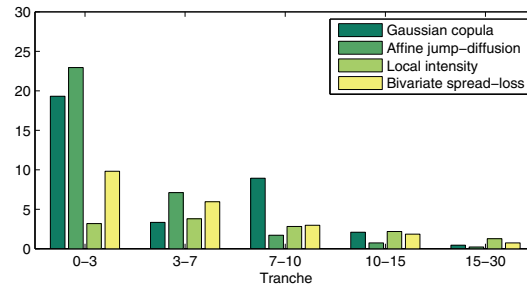
Figure 5. *Jump-to-default ratios. Data: CDX.NA.IG.S10 on 25 March, 2008. Ratios of Gaussian copula model and affine jump-diffusion model are computed by assuming that IBM defaults.*

which is equal to

$$(15) \qquad \phi_t^{idx} = \frac{V_{lo}^{[a,b]}(t, N_{t-} + 1) - V_{lo}^{[a,b]}(t, N_{t-})}{V_{lo}^{idx}(t, N_{t-} + 1) - V_{lo}^{idx}(t, N_{t-})}.$$

The jump-to-default ratio under the bivariate spread-loss model is defined in the same manner.

On the other hand, since the bottom-up models specify the default probabilities of each obligor, there are n possible jump-to-default ratios to be considered. For the Gaussian copula model, the jump-to-default ratio corresponding to obligor i, using the index as the hedging instrument, is equal to

$$(16) \qquad \phi_t^{idx} = \frac{V_{gc}^{[a,b]}(t, \mathbf{s}_{t-}^{cds}, \boldsymbol{\rho}_{t-}, \mathbf{D}_{t-} + \mathbf{u}_i) - V_{gc}^{[a,b]}(t, \mathbf{s}_{t-}^{cds}, \boldsymbol{\rho}_{t-}, \mathbf{D}_{t-})}{V_{gc}^{idx}(t, \mathbf{s}_{t-}^{cds}, \mathbf{D}_{t-} + \mathbf{u}_i) - V_{gc}^{idx}(t, \mathbf{s}_{t-}^{cds}, \mathbf{D}_{t-})},$$

where $\mathbf{u}_i \in \mathbb{R}^n$ is a vector with all entries equal to 0 except for the ith entry equal to 1. We can define the jump-to-default ratio for the affine jump-diffusion model in the same manner.

Figure 5 shows the jump-to-default ratios computed from different models using the index as the hedging instrument as in (15) and (16). For the bottom-up models, the ratios are computed in the scenario where IBM defaults. Unlike the index spread-deltas, the jump-to-default ratios are substantially different across models. This implies that hedging against occurrence of default is also exposed to substantial model risk, because the jump-to-default ratio, which is computed by assuming the occurrence of one additional default, depends on the credit portfolio loss dynamic in the modeling framework. Although each model is calibrated to the same CDS, index, and CDO market data, those credit derivatives provide information only on the *marginal distribution* of the loss process at some fixed times. Therefore, the dynamic of the loss process cannot be uniquely determined by the market data, and thus the jump-to-default ratios are substantially different across models.

5.6. Variance minimization. When spread risk and default risk are considered simultaneously, we are in an incomplete market setting, and hedging strategies in this setting need to be determined by an optimality criterion which takes both spread risk and default risk into account. A well-known approach to hedging in incomplete markets is the *variance-minimizing*

strategy, introduced by Föllmer and Sondermann [13]. Unlike many other approaches to hedging in incomplete markets, it has been shown that this approach actually leads to analytically tractable hedging strategies [8, 14].

Definition 5.1. *Let H be a square-integrable contingent claim at maturity T, and (X_t) be the discounted price process of the hedging instrument which is a square-integrable martingale under $\mathbb{Q}$. Let $\mathcal{S}$ be the set of admissible self-financing strategies with $E^{\mathbb{Q}}[(\int_0^T \phi_t dX_t)^2] < \infty$. A variance-minimizing strategy is a choice of initial capital c and a self-financing trading strategy $(\phi_t) \in \mathcal{S}$ which minimizes the quadratic hedging error:*

$$(17) \qquad \inf_{c \in \mathbb{R}, (\phi_t) \in \mathcal{S}} E^{\mathbb{Q}} \left[\left(c + \int_0^T \phi_t dX_t - H \right)^2 \right].$$

The variance (17) is computed under the risk-neutral probability measure $\mathbb{Q}$, because the models are calibrated to the observed credit spreads which only provide information on the risk factor dynamics under the risk-neutral measure. If we want to minimize the hedging error under the real-world measure, we will need a statistical model. One example is the regression-based hedging strategy in section 5.7.

Föllmer and Sondermann [13] characterize the variance-minimizing strategy in terms of a Galtchouk–Kunita–Watanabe projection of the claim on the set of replicable payoffs (see Appendix A). One nice property of the variance-minimizing strategy is that it coincides with the self-financing hedging strategy, which replicates the contingent claim in a complete market. Furthermore, in many Markovian models with jumps, the variance-minimizing hedge ratios can be explicitly computed [8, 14, 23]. The justification for this approach, which is not specific to credit derivatives, is discussed in [8, 14] from a methodological standpoint.

We will show that these variance-minimizing hedge ratios can be explicitly computed for the local intensity model and the bivariate spread-loss model. Our analysis will omit the Gaussian copula model and the affine jump-diffusion model. The reason is that since the Gaussian copula model defines the marginal distribution of the portfolio losses only at fixed times, there is no intrinsic dynamic for the loss process. For the affine jump-diffusion model, the computation requires inverting a high dimensional matrix (125×125) which is numerically unstable.

5.6.1. Local intensity model. Laurent, Cousin, and Fermanian [17] show that the local intensity framework generates a complete market in the sense that we can replicate the payoff of a tranche $[a, b]$ by means of a self-financing portfolio with positions in the index default swap and a numeraire. The hedge ratio in the replication strategy then coincides with the variance-minimizing hedge ratio, which is equal to the jump-to-default ratio.

Proposition 5.2. *Consider a local intensity model in which the portfolio default intensity (λ_t) under the risk-neutral measure $\mathbb{Q}$ has the form $\lambda_t = f(t, N_{t-})$ for some positive function $f(.,.) > 0$. Let $V_{lo}^{idx}(t, N)$ and $V_{lo}^{[a,b]}(t, N)$ be the mark-to-market values of the index and a tranche $[a, b]$ at time t given N number of defaults which satisfy the following:*

- *For $N = 0, \ldots, n$, the functions $t \mapsto V_{lo}^{idx}(t, N)$ and $t \mapsto V_{lo}^{[a,b]}(t, N)$ belong to $C^1([0, T])$.*
- *$|V_{lo}^{idx}(t, N+1) - V_{lo}^{idx}(t, N)| > 0$ for all $t \in [0, T]$, $N = 0, \ldots, n-1$.*

Then, the variance-minimizing hedge defined in Definition 5.1 *for a tranche* $[a,b]$ *using the index as the only hedging instrument is given by*

(18)
$$\phi_t = \frac{V_{lo}^{[a,b]}(t, N_{t-} + 1) - V_{lo}^{[a,b]}(t, N_{t-})}{V_{lo}^{idx}(t, N_{t-} + 1) - V_{lo}^{idx}(t, N_{t-})}.$$

The second condition in Proposition 5.2 implies that the index default swap value is always sensitive to defaults in the underlying portfolio. The proof of this proposition is shown in Appendix B.1.

5.6.2. Bivariate spread-loss model. In the bivariate spread-loss model, the portfolio default intensity is driven by the loss process (L_t) and a diffusion process (Y_t). Since we consider the index as the only hedging instrument, the market is incomplete in this two-factor framework. We use variance minimization to compute a trade-off between default risk and spread risk, as follows.

Proposition 5.3. *Consider the bivariate spread-loss model* [1] *in which the portfolio default intensity* (λ_t) *under the risk-neutral measure* $\mathbb{Q}$ *follows* (7). *Let* $V_{bi}^{idx}(t, Y, N)$ *and* $V_{bi}^{[a,b]}(t, Y, N)$ *be the values of the index and a tranche* $[a,b]$ *at time* t *given* N *number of defaults and risk factor value* Y *which satisfy the following:*
- *For* $N = 0, \ldots, n$, *the functions* $(t, Y) \mapsto V_{bi}^{idx}(t, Y, N)$ *and* $(t, Y) \mapsto V_{bi}^{[a,b]}(t, Y, N)$ *belong to* $C^{1,2}([0, T) \times \mathbb{R})$;
- $[\partial_Y V_{bi}^{idx}(t, Y, N)]^2 + [\delta_N V_{bi}^{idx}(t, Y, N)]^2 \lambda > 0$ *for all* $t \in [0, T]$, $Y \in \mathbb{R}$, $N = 0, \ldots, n-1$;
- $E^{\mathbb{Q}} \left[\int_0^T \left(\partial_Y V_{bi}^{[a,b]}(t, Y_{t-}, N_{t-}) \right)^2 dt \right] < \infty$,

where $\lambda = (n - N) \sum_{k=0}^N b_k$, $\partial_Y V_{bi}^{\cdot}(t, Y, N)$ *is the partial derivative with respect to* Y, *and* $\delta_N V_{bi}^{\cdot}(t, Y, N) = V_{bi}^{\cdot}(t, Y, N + 1) - V_{bi}^{\cdot}(t, Y, N)$ *is the change of value with respect to one additional default. Then, the variance-minimizing hedge defined in Definition* 5.1 *for a tranche* $[a, b]$ *using the index as the only hedging instrument is given by*

(19)
$$\phi_t = \frac{\partial_Y V_{bi}^{[a,b]}(t, Y_{t-}, N_{t-}) \, \partial_Y V_{bi}^{idx}(t, Y_{t-}, N_{t-}) \, \sigma^2 + \delta_N V_{bi}^{[a,b]}(t, Y_{t-}, N_{t-}) \, \delta_N V_{bi}^{idx}(t, Y_{t-}, N_{t-}) \, \lambda_t}{\left[\partial_Y V_{bi}^{idx}(t, Y_{t-}, N_{t-}) \, \sigma \right]^2 + \left[\delta_N V_{bi}^{idx}(t, Y_{t-}, N_{t-}) \right]^2 \lambda_t}.$$

The second condition in Proposition 5.3 implies that the index default swap is always sensitive to either the change of the risk factor (Y_t) or defaults in the underlying portfolio. The third condition is an integrability condition to ensure that the optimal hedging portfolio has finite variance. The proof of this proposition is shown in Appendix B.2.

Unlike the case of the local intensity model, the variance-minimizing hedge ratio for the bivariate spread-loss model involves not only the jump-to-default values but also the partial derivatives of the values with respect to the additional risk factor (Y_t). This reflects the fact that variance-minimization hedging is a strategy that takes both default risk and spread risk into account.

Figure 6 shows the variance-minimizing hedge ratios of the local intensity model and the bivariate spread-loss model. Similar to the case of comparing the jump-to-default ratios, the variance-minimizing hedge ratios are substantially different across the models, especially for the junior tranches. The reason for the differences is the same as the case for the jump-to-
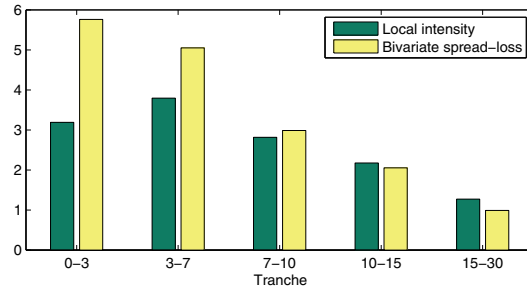
Figure 6. *Variance-minimization hedge ratios. Data: CDX.NA.IG.S10 on 25 March, 2008.*

default ratios: from (18) and (19), we can see that the variance-minimizing hedge ratio is a model-dependent quantity.

5.7. Regression-based hedging. One drawback of hedging strategies based on pricing models is that it is not clear how well the models can capture the dynamics of the credit spreads under the real-world measure, which is an important issue for hedging in practice. We now discuss a model-free, regression-based hedging strategy based on the observed dynamics of the credit spreads.

Consider a simple regression model relating the daily changes in the index and tranche values:

$$\delta V_{t_j}^{[a,b]} = \alpha^{[a,b]} + \beta^{[a,b]} \, \delta V_{t_j}^{idx} + \epsilon_j,$$

where $\delta V_{t_j}^{\cdot} = V_{t_j}^{\cdot} - V_{t_{j-1}}^{\cdot}$ is the daily change of value from time t_{j-1} to t_j. $\alpha^{[a,b]}$ and $\beta^{[a,b]}$ can be estimated by the ordinary least squares regression over a rolling window. Choosing the hedging position in the index default swap as

$$(20) \qquad \hat{\beta}_t^{[a,b]} := \frac{\sum_{t_j \leq t}(\delta V_{t_j}^{idx} - \overline{\delta V_t^{idx}})(\delta V_{t_j}^{[a,b]} - \overline{\delta V_t^{[a,b]}})}{\sum_{t_j \leq t}(\delta V_{t_j}^{idx} - \overline{\delta V_t^{idx}})^2}$$

yields a model-free hedging strategy which we call the *regression-based hedge*. Here $\overline{\delta V_t^{\cdot}}$ is the average of daily P&L on $[0,t]$, and $\hat{\beta}_t^{[a,b]}$ is the estimate of $\beta^{[a,b]}$ using observations over the period $[0,t]$.

The main advantages of this strategy are its ability to directly capture the actual dynamics of comovements in credit spreads and its simplicity.

6. Empirical performance of hedging strategies. We now present an empirical assessment of the performance of the eight hedging strategies described in section 5 using the dataset in section 3. Table 4 summarizes all the strategies that will be considered. Note that we will consider two choices of hedging tranches for gamma hedging where the details will be presented in the following subsections.

We consider the hedging of the protection seller's position on a CDX tranche, initiated on the first day of the sample period. On each day, we calibrate the models to the market data, as stated in section 5, and compute the corresponding hedging positions. A successful

Table 4
Overview of hedging strategies.

Strategy	Underlying risk	Model	Type	Nature
1	Single-name CDS fluctuation	Gaussian copula	Bottom-up	Static
2	Global CDS/index fluctuation	Gaussian copula	Bottom-up	Static
3	1st- and 2nd-order global CDS fluctuation	Gaussian copula	Bottom-up	Static
4	Single-name CDS fluctuation + correlations shifts	Gaussian copula	Bottom-up	Static
5	Single-name CDS fluctuation	Affine jump-diffusion	Bottom-up	Dynamic
6	Default risk	Local intensity	Top-down	Dynamic
7	Variance minimization (risk-neutral measure)	Bivariate spread-loss	Top-down	Dynamic
8	Variance minimization (statistical measure)	Ordinary least squares regression	Statistical	-
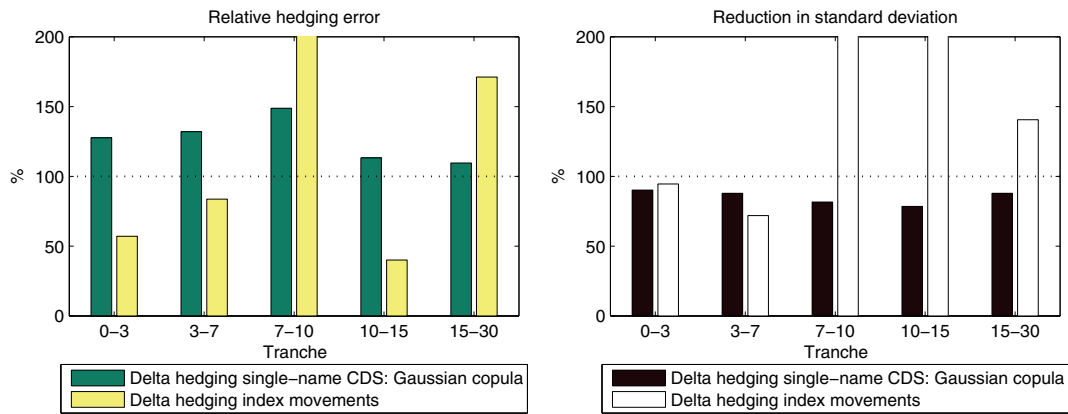
Figure 7. *Comparison among delta hedging strategies based on the Gaussian copula model.*

strategy should (substantially) reduce dispersion of the P&L distribution with respect to an unhedged position. Here, we use the following metrics to assess the reduction in magnitude and volatility of the daily P&L:

$$\text{Relative hedging error} = \left| \frac{\text{Average daily P\&L of hedged position}}{\text{Average daily P\&L of unhedged tranche position}} \right|,$$

$$\text{Residual volatility} = \frac{\text{Daily P\&L volatility of hedged position}}{\text{Daily P\&L volatility of unhedged tranche position}}.$$

Note that the two ratios should be close to 0 for a good hedging strategy.

6.1. Does delta hedging work? Our first analysis is to check whether the commonly criticized delta hedging strategies under the Gaussian copula model work. In Figure 7, we see that delta hedging does not work well. Indeed, the only effective strategy is a delta hedge against index movements, which reduces the absolute exposures of tranches $[0\%, 3\%]$, $[3\%, 7\%]$, and $[10\%, 15\%]$. On the other hand, delta hedging against single-name CDS movements fails to reduce absolute exposures. In terms of reduction in P&L volatility, delta hedging of single-

Table 5
Hedging tranches for gamma hedging strategy.

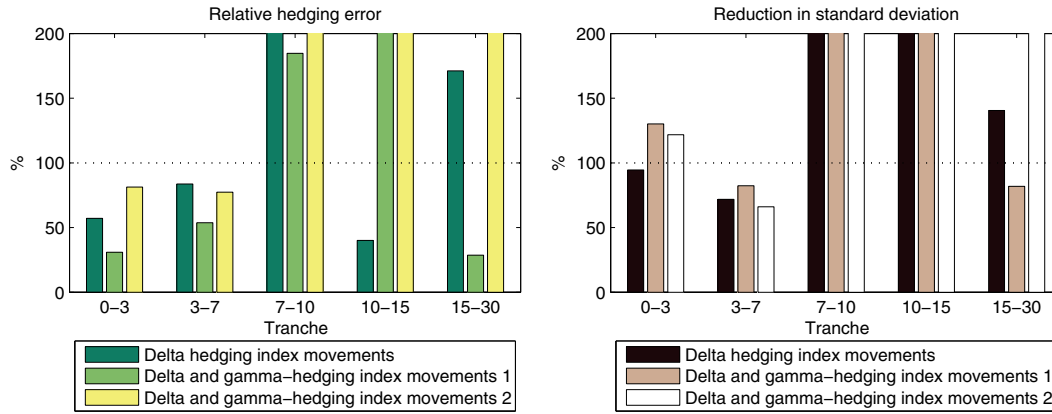| Tranche being hedged	0%–3%	3%–7%	7%–10%	10%–15%	15%–30%
Hedging tranche: Case 1	7%–10%	7%–10%	10%–15%	15%–30%	10%–15%
Hedging tranche: Case 2	15%–30%	15%–30%	15%–30%	7%–10%	7%–10%

Figure 8. *Comparison among delta and gamma hedging strategies.*

name CDS movements results in a consistent reduction of volatility across all tranches, while delta hedging index movements fails to do so for the three most senior tranches.

6.2. Does gamma hedging improve hedging performance? In order to examine whether gamma hedging can improve performance, we consider two different choices of tranches as the hedging instruments, as illustrated in Table 5.

Hedging tranches in Case 1 are chosen such that they give the best performance in our sample in terms of the relative hedging errors, and those in Case 2 are chosen as a comparison with Case 1. Figure 8 shows that gamma hedging can help reduce the hedging error for the $[0\%, 3\%]$, $[3\%, 7\%]$, and $[15\%, 30\%]$ tranches and reduce the P&L volatility for the tranches $[3\%, 7\%]$ and $[15\%, 30\%]$. However, gamma hedging worsens the hedging performance in all other cases. Moreover, its performance is very sensitive to the choice of hedging tranches.

In summary, we conclude that gamma hedging, while very sensitive to the choice of hedging instruments, does not necessarily perform well. Indeed, observations in section 3.1 suggest that the appropriate correction to delta hedging is not the second-order sensitivity with respect to the CDS spread movements but a correction taking into account jumps in spreads. This situation also arises in other contexts when jump risk is present [8].

6.3. Can hedging parallel shifts in correlations improve performance? As shown in Figure 9, immunizing the portfolio against parallel shifts in the (implied) correlation does not improve performance: for almost all tranches neither the hedging error nor the P&L volatility is reduced. This suggests that the (observed) changes in the compound correlations are typically not parallel. Hedging performance may be improved if we consider other scenarios for changes in the correlations.
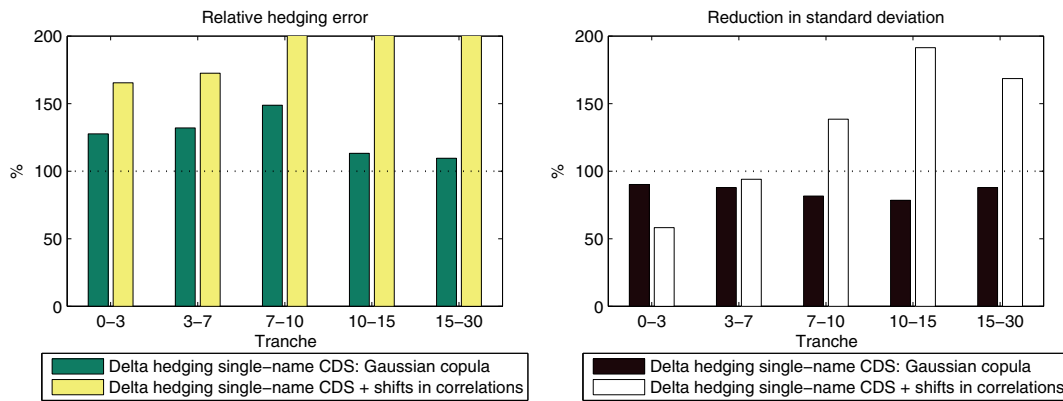
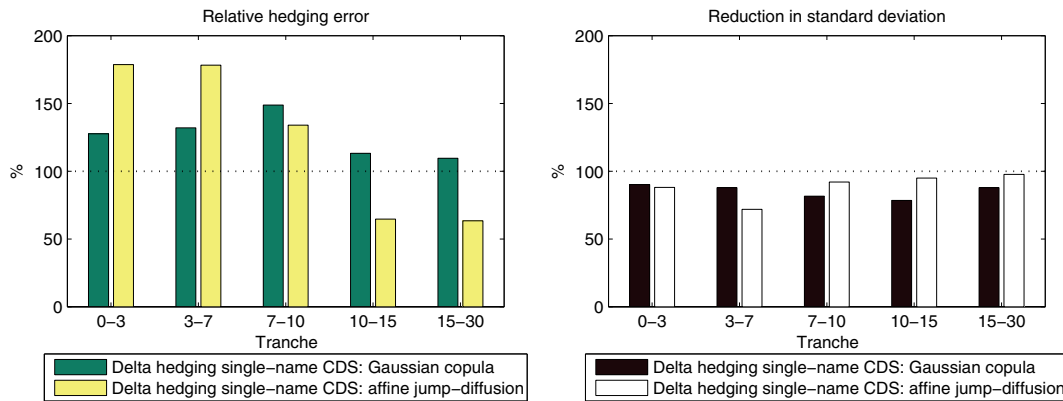**Figure 9.** *Comparison between delta hedging and the addition of hedging parallel shifts in compound correlations.*

Figure 10. *Comparison between delta hedging under the static Gaussian copula and the dynamic affine jump-diffusion models.*

6.4. Do dynamic models have better hedging performance than static models?. Copula-based factor models have faced a lot of criticism for their insufficiency for hedging, and one popular explanation is that this is due to their static nature. In order to verify this claim, we compare the hedging performance of delta hedging under the static Gaussian copula model and under the dynamic affine jump-diffusion model [9].

Interestingly, the hedging error and the reduction in volatility ratios in Figure 10 do not provide any evidence that dynamic models perform better than this simple static model. Although the dynamic model successfully reduces the hedging error for tranches [10%, 15%] and [15%, 30%], it amplifies the hedging error significantly for the two most junior tranches. Moreover, the residual volatilities from the dynamic model are larger than those from the static model for tranches [7%, 10%], [10%, 15%], and [15%, 30%].

6.5. Do bottom-up models perform better than top-down models? Although top-down models are more flexible for calibration, it has been suggested that they may be inadequate for hedging [2]. However, this claim has not been backed by any empirical evidence: we
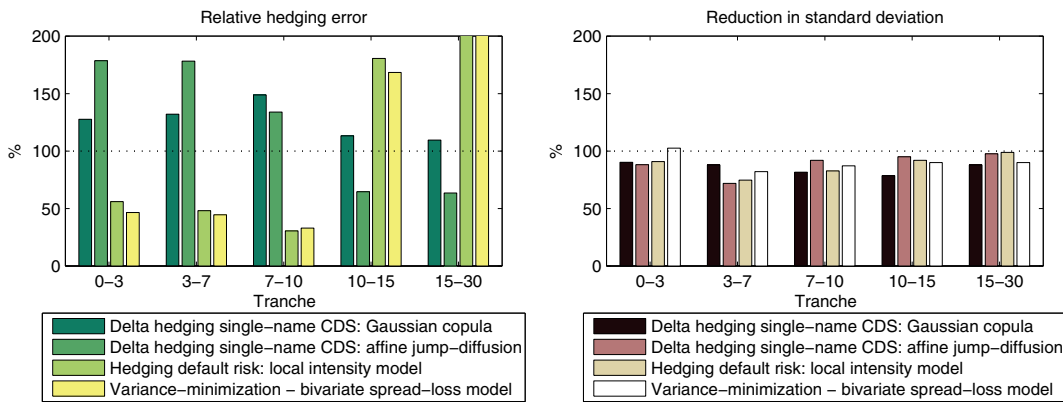
Figure 11. *Comparison of top-down and bottom-up hedging.*

will attempt to verify whether there is indeed an advantage in using bottom-up models for hedging.

Figure 11 compares the performance of various strategies based on bottom-up models, delta hedging based on the Gaussian copula model and the affine jump-diffusion model, versus the top-down strategies, hedging default risk based on the local intensity model and variance minimization based on the bivariate spread-loss model. First, we observe that the reduction of volatility is similar across the strategies, which does not provide much information to distinguish them. By comparing the hedging error, we find that the top-down models perform better than the bottom-up models for the three junior tranches, $[0\%, 3\%]$, $[3\%, 7\%]$, and $[7\%, 10\%]$, while bottom-up models fare better for the other two senior tranches.

Overall, there is no strong evidence that hedging based on the bottom-up models must perform better than that based on the top-down models. This observation contradicts the statements of Bielecki, Crépey, and Jeanblanc [2], who compare bottom-up and top-down hedging based on simulation. Although bottom-up models provide additional degrees of freedom, the effectiveness of a hedging strategy is not about goodness of fit but depends on how well the model can predict short-term comovements of the target instrument and the hedging instruments. From our results it appears that existing bottom-up models do a poor job at predicting such short-term comovements.

6.6. How does model-based hedging compare to regression-based hedging? Given the simplicity and intuitiveness of regression-based hedging, it is interesting to examine how well it performs relative to the model-based strategies. In Figure 12, we observe that regression-based hedging performs well across all tranches, consistently reducing both the hedging error and the daily P&L volatility. In particular, it reduces the volatilities for all tranches more than do the model-based strategies which are theoretically "optimal" in the respective models. This suggests that model misspecification is nonnegligible in all the models considered above.

6.7. Performance on credit event dates. Of particular interest is the performance of the hedging strategies on the next business day after the Lehman Brothers and Fannie Mae/ Freddie Mac credit events. Figure 13 shows the hedging error on the next business days after the credit events under various hedging strategies. During the Lehman Brothers event, we
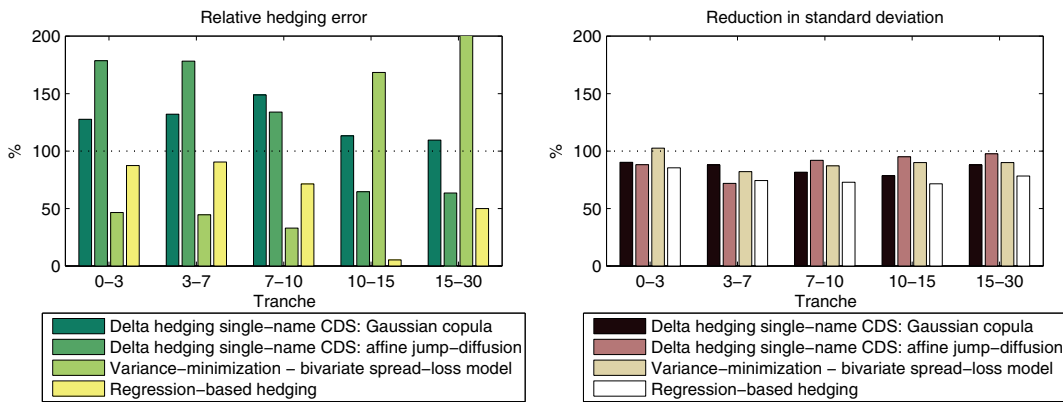
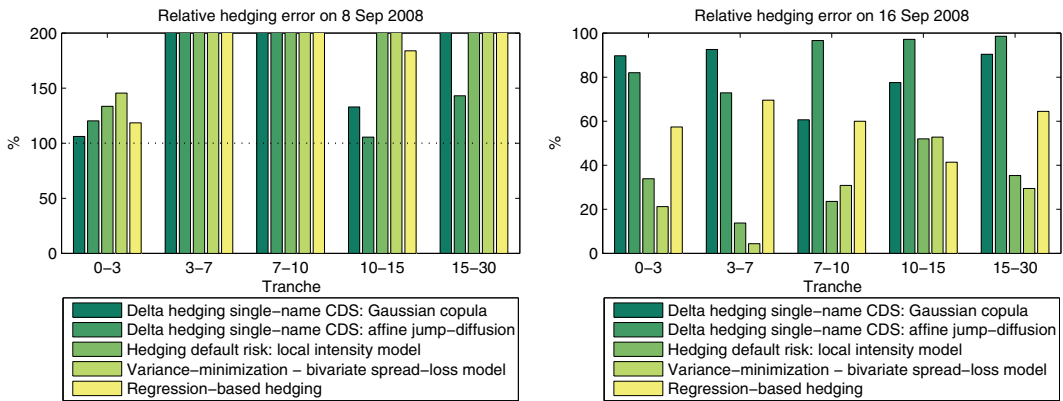**Figure 12.** *Comparison among top-down, bottom-up, and regression-based hedging.*

Figure 13. *Comparison of strategies on the next business day after Fannie Mae/Freddie Mac (8 September, 2008) and Lehman Brothers (16 September, 2008) credit events.*

observe that the top-down and regression-based hedging outperform the bottom-up hedging for all tranches. In particular, variance minimization based on the two-factor top-down model provides the best hedge for most tranches. This suggests that a macro event is better captured by top-down hedging.

On the other hand, all strategies failed to reduce the P&L during the Fannie Mae and Freddie Mac credit event, because, as we saw in section 3, the market happened to have anticipated the event, and there are no significant movements in the spreads. In particular, we observe a negative change in the CDS and index spreads, which leads to an increase in the overall exposure when we try to hedge the tranche positions with positive hedge ratios.

7. Conclusion. We have presented theoretical and empirical comparison of a wide range of hedging strategies for portfolio credit derivatives, with a detailed analysis of the hedging of index CDO tranches. By comparing the performance of these strategies in 2008, our analysis reveals several interesting features:

• Our analysis reveals a large proportion of unhedgeable risk in CDO tranches. This

suggests that market completeness is by no means an acceptable approximation, and toy models which assume a complete market may fail to provide useful insight for issues related to hedging of CDO tranches.

- Delta hedging of CDO tranche positions based on the Gaussian copula model is not effective.
- Although gamma hedging can improve the performance for certain tranches, its effectiveness is very sensitive to the choice of hedging instruments and is inconsistent across tranches.
- We do not find strong evidence that the Duffie and Gârleanu [9] bottom-up dynamic model performs better than the static Gaussian copula model when it comes to delta hedging with credit default swaps.
- Moreover, bottom-up models ([9] and [18]) do not appear to perform consistently better than top-down models ([1] and [16]), in contrast to what has been asserted (without justification) in the literature [2]. In fact, during the period around the Lehman Brothers default, hedging strategies based on top-down models performed substantially better than those based on bottom-up models. This leads us to question the need for computationally costly dynamic bottom-up models instead of the top-down models for hedging portfolio credit derivatives.
- Model-free regression-based hedging appears to be surprisingly effective when compared to other hedging strategies. This suggests—not surprisingly—that incorporating the statistical behavior of credit spreads is an important criterion for a successful hedging strategy.
- We find evidence for common jumps, or large comovements, in spreads. However, index and tranche spreads did not appear to have upward jumps at the default dates of index constituents. This observation goes against models, such as Markovian contagion models [1, 16] and self-exciting models [15], in which jumps in spreads occur only at constituent default dates. Jumps in spreads may also arise from unexpected events not necessarily related to defaults inside the portfolio.

We have left out many practically important considerations, such as liquidity, transaction costs, and computational issues, when assessing hedging performance. The *(il)liquidity* of CDS contracts leads to questions of feasibility of hedging strategies which require frequent rebalancing of positions in single-name CDS. *Transaction costs*—as reflected, for instance, in bid-ask spreads—are known to be higher for single-name CDS contracts than for the index, and taking them into account would favor top-down/index hedging strategies as opposed to hedging with single-name CDS, which requires rebalancing more than a hundred single-name CDS positions. Finally, computational costs are much lower for the top-down models, especially when it comes to calibration: various fast calibration methods have been proposed for top-down models [1, 6, 7, 21], whereas parameter calibration, especially if it needs to be done on a periodic basis, remains nontrivial for bottom-up models [11]. Therefore it should be clear that taking these aspects into account would tilt the comparison even more in favor of top-down/index hedging as opposed to hedging with single-name CDS.

Appendix A. Variance-minimizing hedge and Galtchouk–Kunita–Watanabe decomposition. Our work in this section is similar to the earlier work by Frey and Backhaus [14]. We

first define the *gain process* for the index and tranches. Let $P^{idx}(T_m)$ be the net payment received from an index default swap contract at time T_m which is bounded by definition. The mark-to-market value of an index default swap at time t is equal to

$$V_t^{idx} = \sum_{T_m > t} B(t, T_m) E^{\mathbb{Q}}[P^{idx}(T_m)|\mathcal{F}_t].$$

We define the discounted value process as

$$\tilde{V}_t^{idx} := B(0, t)V^{idx}(t).$$

Then, the gain process is defined as the present value of all cash-flows:

$$(21) \qquad G_t^{idx} = \sum_{T_m \le t} B(0, T_m) P^{idx}(T_m) + \tilde{V}_t^{idx} = E^{\mathbb{Q}}\left[\sum_{T_m > 0} B(0, T_m) P^{idx}(T_m) \Big| \mathcal{F}_t\right],$$

where (G_t^{idx}) is a square-integrable (bounded) martingale under $\mathbb{Q}$. Similarly, we can derive the expression for the gain process of an $[a, b]$ tranche $(G_t^{[a,b]})$ which is also a square-integrable (bounded) martingale under $\mathbb{Q}$.

Consider the variance minimization setting in Definition 5.1. Our goal is to hedge a tranche $[a, b]$ using the index default swap. Although we consider a terminal payoff H at maturity in Definition 5.1, the results in [23] allow us to replace the conditional expected value of the terminal payoff $E^{\mathbb{Q}}[H|\mathcal{F}_t]$ by the gain process value $G_t^{[a,b]}$. Let (ϕ_t) be the variance-minimizing hedging strategy which represents positions in the index default swap. Then, it can be shown that (see [13]) (ϕ_t) satisfies the Galtchouk–Kunita–Watanabe decomposition

$$(22) \qquad G_t^{[a,b]} = G_0^{[a,b]} + \int_0^t \phi_s dG_s^{idx} + Z_t,$$

where the process $(Z_t G_t^{idx})$ is a martingale under $\mathbb{Q}$. Therefore, the variance-minimizing strategy satisfies

$$(23) \qquad d\langle G^{[a,b]}, G^{idx}\rangle_t = \phi_t d\langle G^{idx}\rangle_t, \quad 0 \le t \le T,$$

where $(\langle G\rangle_t)$ denotes the unique predictable process with $\langle G\rangle_0 = 0$ and right-continuous increasing paths such that $(G_t^2 - \langle G\rangle_t)$ is a martingale under $\mathbb{Q}$.

Remark 1. Instead of variance minimization in Definition 5.1, Föllmer and Sondermann [13] introduce a stronger optimality condition which is known as *risk-minimization*. However, since the gain processes of the index and CDO tranches are bounded martingales under $\mathbb{Q}$, the two optimality criteria lead to the same optimal hedging strategies.

Appendix B. Derivation of the variance-minimizing hedge ratio. The key to computing the variance-minimizing hedge under a particular model is to express (G_t^{idx}) and $(G_t^{[a,b]})$ as stochastic integrals, use the Ito isometry formula for these stochastic integrals, then solve (23) (see, e.g., [8]). In particular, Frey and Backhaus [14] show the derivation for the *convex counterparty risk model*, and our following computations are similar to those in [14].

B.1. Proof of Proposition 5.2. Let (U_t) be a deterministic counting process which jumps by 1 at payment dates, i.e., $U_{T_m} - U_{T_{m-1}} = 1$ for all $m = 1, \ldots, M$. Using Ito's lemma, the dynamics of the discounted value process under $\mathbb{Q}$ for the index follows

$$d\tilde{V}_{lo}^{idx}(t, N_{t-}) = \frac{\partial}{\partial t}\tilde{V}_{lo}^{idx}(t, N_{t-})dt + \delta_N \tilde{V}_{lo}^{idx}(t, N_{t-})dN_t - B(0, t)P^{idx}(t)dU_t,$$

where $P^{idx}(t)$ is the net payment received at time t. Using Ito's lemma again and from (21), the dynamics of the gain process under $\mathbb{Q}$ for the index follows

$$\begin{aligned}
dG_{lo}^{idx}(t) &= B(0,t)P^{idx}(t)dU_t + d\tilde{V}_{lo}^{idx}(t, N_{t-}) \\
&= \frac{\partial}{\partial t}\tilde{V}_{lo}^{idx}(t, N_{t-})dt + \delta_N \tilde{V}_{lo}^{idx}(t, N_{t-})dN_t \\
&= \left(\frac{\partial}{\partial t}\tilde{V}_{lo}^{idx}(t, N_{t-}) + \delta_N \tilde{V}_{lo}^{idx}(t, N_{t-})\lambda_t \right) dt + \delta_N \tilde{V}_{lo}^{idx}(t, N_{t-})dN_t^c \\
&= \delta_N \tilde{V}_{lo}^{idx}(t, N_{t-})dN_t^c,
\end{aligned}$$

where $(N_t^c) = (N_t - \int_0^t \lambda_s ds)$ is the compensated version of (N_t). The last equality comes from the fact that $(G_{lo}^{idx}(t))$ is a martingale under $\mathbb{Q}$. Similarly, the gain process of a tranche $[a, b]$ under $\mathbb{Q}$ follows

$$dG_{lo}^{[a,b]}(t) = \delta_N \tilde{V}_{lo}^{[a,b]}(t, N_{t-})dN_t^c.$$

We now compute the compensators involved in (23) and obtain

$$\begin{aligned}
\phi_t &= \frac{\delta_N \tilde{V}_{lo}^{[a,b]}(t, N_{t-})\delta_N \tilde{V}_{lo}^{idx}(t, N_{t-})\lambda_t}{(\delta_N \tilde{V}_{lo}^{idx}(t, N_{t-}))^2 \lambda_t} \\
&= \frac{B(0,t)\delta_N V_{lo}^{[a,b]}(t, N_{t-})}{B(0,t)\delta_N V_{lo}^{idx}(t, N_{t-})} \\
&= \frac{V_{lo}^{[a,b]}(t, N_{t-} + 1) - V_{lo}^{[a,b]}(t, N_{t-})}{V_{lo}^{idx}(t, N_{t-} + 1) - V_{lo}^{idx}(t, N_{t-})}.
\end{aligned}$$

Then, we want to show that the hedging portfolio obtained by implementing the above strategy has finite variance, i.e., $E^{\mathbb{Q}}\left[\int_0^T \phi_t dG_{lo}^{idx}(t)\right]^2 < \infty$, and it is sufficient to show that $E^{\mathbb{Q}}\left[\int_0^T \phi_t^2 d\langle G_{lo}^{idx}\rangle_t\right] < \infty$. Since all the cash flows of tranche $[a, b]$ are bounded, there exists a $K > 0$ such that $|V_{lo}^{[a,b]}(t, N)| \le K/2$ for all $t \in [0, T]$, $N = 0, \ldots, n$, which implies that $|\delta_N V_{lo}^{[a,b]}(t, N)| \le K$ for all $t \in [0, T]$, $N = 0, \ldots, n$. Consider

$$\begin{aligned}
E^{\mathbb{Q}}\left[\int_0^T \phi_t^2 d\langle G_{lo}^{idx}\rangle_t\right] &= E^{\mathbb{Q}}\left[\int_0^T \left(\frac{\delta_N V_{lo}^{[a,b]}(t, N_{t-})}{\delta_N V_{lo}^{idx}(t, N_{t-})}\right)^2 \left(\delta_N \tilde{V}_{lo}^{idx}(t, N_{t-})\right)^2 \lambda_t dt\right] \\
&\le E^{\mathbb{Q}}\left[\int_0^T \left(\delta_N V_{lo}^{[a,b]}(t, N_{t-})\right)^2 \lambda_t dt\right] \\
&\le K^2 E^{\mathbb{Q}}\left[\int_0^T \lambda_t dt\right] \\
&\le K^2 n < \infty,
\end{aligned}$$

which gives the desired result.

Note that if we implement this hedging strategy, we have

$$dG_{lo}^{[a,b]}(t) = \phi_t dG_{lo}^{idx}(t).$$

The tranche is perfectly hedged in this case, which is consistent with the results of [17].

B.2. Proof of Proposition 5.3. Defining the deterministic counting process (U_t) as in Appendix B.1 and applying Ito's lemma, the dynamic of the discounted value process for the index becomes

$$d\tilde{V}_{bi}^{idx}(t, Y_{t-}, N_{t-}) = \frac{\partial}{\partial t}\tilde{V}_{bi}^{idx}(t, Y_{t-}, N_{t-})dt + \frac{\partial}{\partial Y}\tilde{V}_{bi}^{idx}(t, Y_{t-}, N_{t-})dY_t$$
$$+ \frac{1}{2}\frac{\partial^2}{\partial Y^2}\tilde{V}_{bi}^{idx}(t, Y_{t-}, N_{t-})\sigma^2 dt + \delta_N\tilde{V}_{bi}^{idx}(t, Y_{t-}, N_{t-})dN_t - B(0,t)P^{idx}(t)dU_t,$$

where $P^{idx}(t)$ is the net payment received at time t. Then, using Ito's lemma again and from (21), the dynamics of the gain process for the index follows

$$dG_{bi}^{idx}(t) = B(0,t)P^{idx}(t)dU_t + d\tilde{V}_{bi}^{idx}(t, Y_{t-}, N_{t-})$$
$$= \frac{\partial}{\partial t}\tilde{V}_{bi}^{idx}(t, Y_{t-}, N_{t-})dt + \frac{\partial}{\partial Y}\tilde{V}_{bi}^{idx}(t, Y_{t-}, N_{t-})dY_t$$
$$+ \frac{1}{2}\frac{\partial^2}{\partial Y^2}\tilde{V}_{bi}^{idx}(t, Y_{t-}, N_{t-})\sigma^2 dt + \delta_N\tilde{V}_{bi}^{idx}(t, Y_{t-}, N_{t-})dN_t$$
$$= \left(\frac{\partial}{\partial t}\tilde{V}_{bi}^{idx}(t, Y_{t-}, N_{t-}) - \frac{\partial}{\partial Y}\tilde{V}_{bi}^{idx}(t, Y_{t-}, N_{t-})\kappa Y_t\right.$$
$$\left. + \frac{1}{2}\frac{\partial^2}{\partial Y^2}\tilde{V}_{bi}^{idx}(t, Y_{t-}, N_{t-})\sigma^2 + \delta_N\tilde{V}_{bi}^{idx}(t, Y_{t-}, N_{t-})\lambda_t\right)dt$$
$$+ \frac{\partial}{\partial Y}\tilde{V}_{bi}^{idx}(t, Y_{t-}, N_{t-})\sigma dW_t + \delta_N\tilde{V}_{bi}^{idx}(t, Y_{t-}, N_{t-})dN_t^c$$
$$= \frac{\partial}{\partial Y}\tilde{V}_{bi}^{idx}(t, Y_{t-}, N_{t-})\sigma dW_t + \delta_N\tilde{V}_{bi}^{idx}(t, Y_{t-}, N_{t-})dN_t^c.$$

Similarly, the gain process of a tranche $[a,b]$ follows

$$dG_{bi}^{[a,b]}(t) = \frac{\partial}{\partial Y}\tilde{V}_{bi}^{[a,b]}(t, Y_{t-}, N_{t-})\sigma dW_t + \delta_N\tilde{V}_{bi}^{[a,b]}(t, Y_{t-}, N_{t-})dN_t^c.$$

We can now compute the compensators of $(G_{bi}^{[a,b]}(t))$ and $(G_{bi}^{idx}(t))$ and use (23) to compute (ϕ_t) which gives (19).

Let us now show that the hedging portfolio based on (19) has finite variance. Consider the following:

$$E^{\mathbb{Q}}\left[\int_0^T \phi_t d\langle G_{bi}^{idx}\rangle_t\right]$$

$$= E^{\mathbb{Q}}\left[\int_0^T \frac{\left[\partial_Y \tilde{V}_{bi}^{[a,b]}(t,Y_{t-},N_{t-})\,\partial_Y \tilde{V}_{bi}^{idx}(t,Y_{t-},N_{t-})\,\sigma^2 + \delta_N \tilde{V}_{bi}^{[a,b]}(t,Y_{t-},N_{t-})\delta \tilde{V}_{bi}^{idx}(t,Y_{t-},N_{t-})\lambda_t\right]^2}{(\partial_Y \tilde{V}_{bi}^{idx}(t,Y_{t-},N_{t-})\,\sigma)^2 + (\delta_N \tilde{V}_{bi}^{idx}(t,Y_{t-},N_{t-}))^2\,\lambda_t}\,dt\right]$$

$$\leq E^{\mathbb{Q}}\left[\int_0^T \frac{\left[\partial_Y V_{bi}^{[a,b]}(t,Y_{t-},N_{t-})\,\partial_Y V_{bi}^{idx}(t,Y_{t-},N_{t-})\,\sigma^2 + \delta_N V_{bi}^{[a,b]}(t,Y_{t-},N_{t-})\delta_N V_{bi}^{idx}(t,Y_{t-},N_{t-})\lambda_t\right]^2}{(\partial_Y V_{bi}^{idx}(t,Y_{t-},N_{t-})\,\sigma)^2 + (\delta_N V_{bi}^{idx}(t,Y_{t-},N_{t-}))^2\,\lambda_t}\,dt\right]$$

$$\leq 2E^{\mathbb{Q}}\left[\int_0^T \frac{\left[\partial_Y V_{bi}^{[a,b]}(t,Y_{t-},N_{t-})\,\partial_Y V_{bi}^{idx}(t,Y_{t-},N_{t-})\,\sigma^2\right]^2}{(\partial_Y V_{bi}^{idx}(t,Y_{t-},N_{t-})\,\sigma)^2 + (\delta_N V_{bi}^{idx}(t,Y_{t-},N_{t-}))^2\,\lambda_t}\,dt\right]$$

$$+ 2E^{\mathbb{Q}}\left[\int_0^T \frac{\left[\delta_N V_{bi}^{[a,b]}(t,Y_{t-},N_{t-})\delta_N V_{bi}^{idx}(t,Y_{t-},N_{t-})\lambda_t\right]^2}{(\partial_Y V_{bi}^{idx}(t,Y_{t-},N_{t-})\,\sigma)^2 + (\delta_N V_{bi}^{idx}(t,Y_{t-},N_{t-}))^2\,\lambda_t}\,dt\right].$$

Consider the first term. For a fixed time t, the integrand is equal to zero if $\partial_Y V_{bi}^{idx}(t,Y_{t-},N_{t-}) = 0$. If $\partial_Y V_{bi}^{idx}(t,Y_{t-},N_{t-}) \neq 0$, we have

$$\frac{\left[\partial_Y V_{bi}^{[a,b]}(t,Y_{t-},N_{t-})\,\partial_Y V_{bi}^{idx}(t,Y_{t-},N_{t-})\,\sigma^2\right]^2}{(\partial_Y V_{bi}^{idx}(t,Y_{t-},N_{t-})\,\sigma)^2 + (\delta_N V_{bi}^{idx}(t,Y_{t-},N_{t-}))^2\,\lambda_t} \leq \left[\partial_Y V_{bi}^{[a,b]}(t,Y_{t-},N_{t-})\,\sigma\right]^2.$$

For the second expectation, for a fixed time t, the integrand is equal to zero if $\delta_N V_{bi}^{idx}(t,Y_{t-},N_{t-})\lambda_t = 0$. Otherwise, we have

$$\frac{\left[\delta_N V_{bi}^{[a,b]}(t,Y_{t-},N_{t-})\delta_N V_{bi}^{idx}(t,Y_{t-},N_{t-})\lambda_t\right]^2}{(\partial_Y V_{bi}^{idx}(t,Y_{t-},N_{t-})\,\sigma)^2 + (\delta_N V_{bi}^{idx}(t,Y_{t-},N_{t-}))^2\,\lambda_t} \leq \left(\delta_N V_{bi}^{[a,b]}(t,Y_{t-},N_{t-})\right)^2\lambda_t.$$

Using the fact that $\delta_N V_{bi}^{[a,b]}(t,Y,N)$ is bounded, we obtain

$$E^{\mathbb{Q}}\left[\int_0^T \phi_t d\langle G_{bi}^{idx}\rangle_t\right] \leq 2E^{\mathbb{Q}}\left[\int_0^T \left[\partial_Y V_{bi}^{[a,b]}(t,Y_{t-},N_{t-})\,\sigma\right]^2 + \left(\delta_N V_{bi}^{[a,b]}(t,Y_{t-},N_{t-})\right)^2\lambda_t dt\right]$$

$$\leq 2\sigma^2 E^{\mathbb{Q}}\left[\int_0^T \left(\partial_Y V_{bi}^{[a,b]}(t,Y_{t-},N_{t-})\right)^2 dt\right] + 2K^2 n < \infty.$$

Therefore, $E^{\mathbb{Q}}\left[\int_0^T \phi_t dG_{bi}^{idx}(t)\right]^2 < \infty$.

REFERENCES

[1] M. ARNSDORF AND I. HALPERIN, *BSLP: Markovian bivariate spread-loss model for portfolio credit deriv-atives*, J. Comput. Finance, 12 (2008), pp. 77–107.

[2] T. BIELECKI, S. CRÉPEY, AND M. JEANBLANC, *Up and down credit risk*, Quant. Finance, 10 (2010), pp. 1137–1151.

[3] T. BIELECKI, M. JEANBLANC, AND M. RUTKOWSKI, *Pricing and trading credit default swaps in a hazard process model*, Ann. Appl. Probab., 18 (2008), pp. 2495–2529.

[4] D. BRIGO AND M. MORINI, *Arbitrage-free Pricing of Credit Index Options. The No-Armageddon-Pricing Measure and the Role of Correlation after the Subprime Crisis*, SSRN working paper, 2008.

[5] R. CONT, *Model uncertainty and its impact on the pricing of derivative instruments*, Math. Finance, 16 (2006), pp. 519–547.

[6] R. CONT, R. DEGUEST, AND Y. H. KAN, *Default intensities implied by CDO spreads: Inversion formula and model calibration*, SIAM J. Financial Math., 1 (2010), pp. 555–585.

[7] R. CONT AND A. MINCA, *Recovering Portfolio Default Intensities Implied by CDO Tranches*, Financial Engineering Report 2008-01, Columbia University, New York, 2008.

[8] R. CONT, P. TANKOV, AND E. VOLTCHKOVA, *Hedging with options in presence of jumps*, in Stochastic Analysis and Applications: The Abel Symposium 2005 in Honor of Kiyosi Ito, F. Benth, G. Di Nunno, T. Lindstrom, B. Oksendal, and T. Zhang, eds., Springer, New York, 2007, pp. 197–218.

[9] D. DUFFIE AND N. GÂRLEANU, *Risk and valuation of collateralized debt obligations*, Financial Anal. J., 57 (2001), pp. 41–59.

[10] P. DURAND AND J.-F. JOUANIN, *Some short elements on hedging credit derivatives*, ESAIM Probab. Statist., 11 (2007), pp. 23–34.

[11] A. ECKNER, *Computational techniques for basic affine models of portfolio credit risk*, J. Comput. Finance, 13 (2009), pp. 63–97.

[12] E. ERRAIS, K. GIESECKE, AND L. R. GOLDBERG, *Affine point processes and portfolio credit risk*, SIAM J. Finan. Math., 1 (2010), pp. 642–665.

[13] H. FÖLLMER AND D. SONDERMANN, *Hedging of nonredundant contingent claims*, in Contributions to Mathematical Economics, North–Holland, Amsterdam, 1986, pp. 205–223.

[14] R. FREY AND J. BACKHAUS, *Dynamic hedging of synthetic CDO tranches with spread and contagion risk*, J. Econom. Dynam. Control, 34 (2010), pp. 710–724.

[15] K. GIESECKE, L. GOLDBERG, AND X. DING, *A top-down approach to multi-name credit*, Oper. Res., 59 (2011), to appear.

[16] A. HERBERTSSON, *Pricing synthetic CDO tranches in a model with default contagion using the matrix-analytic approach*, J. Credit Risk, 4 (2008), pp. 3–35.

[17] J.-P. LAURENT, A. COUSIN, AND J.-D. FERMANIAN, *Hedging Default Risks of CDOs in Markovian Contagion Models*, working paper, 2008.

[18] D. LI, *On default correlation: A copula function approach*, J. Fixed Income, 9 (2000), pp. 43–54.

[19] D. LI, *Base correlation*, in Encyclopedia of Quantitative Finance, R. Cont, ed., Wiley, New York, 2010.

[20] A. LIPTON AND A. RENNIE, EDS., *The Oxford Handbook of Credit Derivatives*, Oxford University Press, London, 2011.

[21] A. V. LOPATIN AND T. MISIRPASHAEV, *Two-dimensional Markovian model for dynamics of aggregate credit loss*, in Econometrics and Risk Management, Adv. Econom. 22, Emerald/JAI, Amsterdam, 2008, pp. 243–274.

[22] S. MORGAN AND A. MORTENSEN, *CDO Hedging Anomalies in the Base Correlation Approach*, Technical report, 2007.

[23] T. MØLLER, *Risk-minimizing hedging strategies for insurance payment process*, Finance Stoch., 5 (2001), pp. 419–446.

[24] M. NEUGEBAUER, J. CARTER, R. HRVATIN, T. CUNNINGHAM, AND R. HARDEE, *Understanding and Hedging Risks in Synthetic CDO Tranches*, Technical report, Fitch Ratings, New York, 2006.

[25] P. SCHÖNBUCHER, *Portfolio Losses and the Term Structure of Loss Transition Rates: A New Methodology for the Pricing of Portfolio Credit Derivatives*, working paper, 2005.