Recovering portfolio default intensities implied by CDO quotes*

Rama CONT

Andreea MINCA

Revised version. First version appeared as: Columbia Financial
Engineering Report No. 2008-01, Jan 2008.

Abstract

We propose a stable non-parametric algorithm for the calibration of pricing models for portfolio credit derivatives: given a set of observations of market spreads for CDO tranches, we construct a risk-neutral default intensity process for the portfolio underlying the CDO which matches these observations, by looking for the risk neutral loss process 'closest' to a prior loss process, verifying the calibration constraints. We formalize the problem in terms of minimization of relative entropy with respect to the prior under calibration constraints and use convex duality methods to solve the problem: the dual problem is shown to be an intensity control problem, characterized in terms of a Hamilton–Jacobi system of differential equations, for which we present an analytical solution. Given a set of observed CDO tranche spreads, our method allows to construct a default intensity process which leads to tranche spreads consistent with the observations. We illustrate our method on ITRAXX index data: our results reveal strong evidence for the dependence of loss transitions rates on the past number of defaults, thus offering quantitative evidence for contagion effects in the risk–neutral loss process.

Keywords: collateralized debt obligation, duality, portfolio credit derivatives, reduced-form models, default risk, intensity control, top-down credit risk models, relative entropy, inverse problem, model calibration, stochastic control.

*This work was presented at the at the Workshop on PDEs in finance (Stockholm Aug 2007), the Stanford Financial mathematics seminar (Oct 07), Conference on Credit Risk (Chicago Univ, Oct 07), the CCCP 2007 Conference (Princeton Nov 07), the S&P Credit Risk Summit (New York Nov 07) the Conference on Credit Risk (Evry 2008) Bachelier Congress 2008 and the Chicago University Credit Risk conference (2008). We thank Marco Avellaneda, Damiano Brigo, Bruno Dupire, YuHang Kan, Johannes Ruf, Kay Giesecke and Igor Halperin for helpful remarks.

Contents

1	Introduction	3
2	Portfolio credit derivatives	4
2.1	Index default swaps	5
2.2	Collateralized Debt Obligations (CDOs)	5
2.3	Top-down models for CDO pricing	6
3	The information content of CDO tranches	9
3.1	Mimicking marked point processes with Markovian jump processes	9
3.2	Information content of portfolio credit derivatives	11
4	The calibration problem	12
4.1	Calibration as relative entropy minimization under constraints .	12
4.2	Computation of the relative entropy	13
4.3	Dual problem as an intensity control problem	15
4.4	Hamilton Jacobi equations	16
4.5	Handling payment dates	18
5	Recovering market-implied default rates	19
5.1	Calibration algorithm	19
5.2	Application to ITRAXX tranches	20
6	Conclusion	24

1 Introduction

Credit derivatives markets have witnessed an extraordinary activity in the last decade, especially with the development of a large market in portfolio credit derivatives of which collateralized debt obligations (CDOs) are the most well known example [8]. Yet, as illustrated in the recent market turmoil, commonly used modeling approaches –mostly static, copula-based pricing models such as the Gaussian copula model– appear to be insufficient for pricing and hedging these complex derivatives. One of the reasons has been the lack of transparency of such pricing methods in which non-intuitive and unobservable “default correlation” parameters are required as an input.

The Gaussian copula model, which has been widely used for the pricing of CDOs, has some well known shortcomings: its inability to reproduce market values of CDO tranche spreads, as exemplified by the base correlation skew, the instability of its “default correlation” parameters –as revealed by the GM/FORD crisis in May 2005 and the subprime crisis in 2007– and, most importantly, the lack of a well-defined dynamics for the risk factors which prevents any model-based assessment of hedging strategies. Other copula-based models may provide better fits to market quotes but share the other drawbacks of the Gaussian copula model, most notably its static character. These shortcomings have inspired a lot of research on alternative approaches to credit risk modeling [24]. On the other hand, a great advantage of static copula models is the ease with which the parameters can be calibrated to market data: this is a feature which many of the more complex, multi-name dynamic models such as Duffie & Garleanu [16], have lacked so far. The key challenge in improving on the Gaussian copula model lies therefore not so much in adding more realistic features to the model but in adding these features while maintaining analytical tractability, especially in regard to the calibration to market data.

To tackle some of these issues while allowing for a parsimonious parametrization of the model, several recent works [28, 21, 19, 1, 25] have proposed a “top-down” approach to the problem, in which one models in “reduced form” the dynamics of the portfolio loss, as a jump process whose intensity λ_t represents the (conditional) rate of occurrence of the next default and whose jump sizes represent the losses given default. Though top-down pricing models are typically much simpler to simulate or implement than high-dimensional reduced form models, numerical methods –Laplace transforms, numerical resolution of ODEs– are still required for the pricing of CDO tranches which makes parameter calibration computationally challenging. Existing studies of top-down pricing models [28, 21, 19, 1, 25] address model calibration by applying black box optimization procedures, whose convergence is not guaranteed, to the resulting high-dimensional nonlinear optimization problems. The lack of convexity of the optimization problems involved may lead to multiple solutions and numerical sensitivity of the results, making such results difficult to reproduce and rendering their interpretation delicate.

In this work we propose a rigorous nonparametric approach to the calibration of “top-down” pricing models for portfolio credit derivatives to a set of observed

CDO tranche spreads. First, we show a “mimicking theorem” for point processes which states that the marginal distributions of a loss process with arbitrary stochastic intensity can be matched using a *Markovian* point process. This result implies that, given any risk-neutral loss process with given default intensity we can construct a Markovian loss process which leads to the same prices. This observation allows to narrow down the calibration problem to the search for a Markovian loss process verifying a set of calibration constraints. We formalize this problem in terms of the minimization of relative entropy with respect to the law of a prior loss process under calibration constraints. We use convex duality techniques to solve the problem: the dual problem is shown to be an intensity control problem, characterized in terms of a Hamilton-Jacobi system of differential equations which can be analytically solved using a change of variable.

Given a set of observed CDO tranche spreads, our method allows to construct an implied intensity process λ_t which leads to tranche spreads consistent with the observations. The implied intensity $\lambda_t = f(t, L_t)$ depends on the defaults in the portfolio, which naturally leads to ‘contagion’ effects in the occurrence of defaults. The resulting model is parameterized by the probability (per unit time) of the next default in the portfolio, which allows for an intuitive check on parameter values.

The article is structured as follows. Section 2 describes the cash flow structure of a (static) CDO and present a brief review of the “top-down” modeling approach for portfolio credit derivatives. In section 3 we discuss the level of information about the risk-neutral loss process which can be extracted from CDO tranches: we state a “mimicking theorem” for point processes which implies that, in a general setting, the information content of CDO tranche quotations can be represented in the form of an *effective intensity function* allowing for dependence of the default rate on the current number of defaults in the portfolio and calendar time. The model calibration problem is defined in section 4 and formulated in terms of relative entropy minimization under constraints. In section 4.3 we show that, via convex duality, the calibration problem maps into an intensity control problem for a point process, which is then solved using dynamic programming. These results translate into a calibration algorithm which can be used to extract the risk-neutral default intensity from CDO tranche spreads: the algorithm is laid out in detail in section 5 and applied to ITRAXX index data. Section 6 discusses the implications of our results.

2 Portfolio credit derivatives

Let $(\Omega, (\mathcal{F}_t)_{t \leq T})$ be the set of market scenarios endowed with a filtration $(\mathcal{F}_t)_{0 \leq t \leq T}$ representing the flow of information with time. Consider a reference portfolio on which the credit derivatives we consider will be indexed. The main object of interest are the number of defaults N_t and the (cumulative) *default loss* L_t of this reference portfolio during a period $[0, t]$. Although the discussion below can be generalized to account for correlation between interest rates and default rates, for simplicity we shall assume independence between default risk and in-

terest rate risk. We denote by $B(t, T)$ the discount factor at date t for the maturity $T \geq t$.

A *portfolio credit derivative* can be modeled as a contingent claim whose payoff is a (possibly path-dependent) function of the portfolio loss process $(L_t)_{t \in [0, T]}$. The most important example of portfolio credit derivatives are *index default swaps* and *collateralized debt obligations* (CDO) [8].

2.1 Index default swaps

Index default swaps are now commonly traded on various credit indices such as ITRAXX and CDX series, which are equally weighted indices of credit default swaps on European and US names [8]. In an index default swap transaction, a protection seller agrees to pay all default losses in the index (default leg) in return for a fixed periodic spread S paid on the total notional of obligors remaining in the index (premium leg). Denoting by $t_j, j = 1..J$ the payments dates,

- the default leg pays at t_j the losses $L(t_j) - L(t_{j-1})$ due to defaults in $]t_{j-1}, t_j]$
- the premium leg pays at t_j an interest (spread) S on the notional of the remaining obligors

$$(t_j - t_{j-1})S(1 - \frac{N_{t_j}}{n}). \quad (1)$$

In particular the cash flows of the index default swap only depend on the portfolio characteristics via N_t and L_t . The value at $t = 0$ of the default leg is therefore

$$\sum_{j=1}^J E^{\mathbb{Q}}[B(0, t_j)(L(t_j) - L(t_{j-1}))]$$

while the value at $t = 0$ of the premium leg is

$$S \sum_{j=1}^J E^{\mathbb{Q}}[B(0, t_j)(t_j - t_{j-1})(1 - \frac{N_{t_j}}{n})].$$

The index default swap spread at $t = 0$ is defined as the (fair) value of the spread which equalizes the two legs at inception:

$$S_{\text{index}} = \frac{E^{\mathbb{Q}}[\sum_{j=1}^J B(0, t_j)(L(t_j) - L(t_{j-1}))]}{\sum_{j=1}^J E^{\mathbb{Q}}[B(0, t_j)(t_j - t_{j-1})(1 - \frac{N_{t_j}}{n})]}. \quad (2)$$

2.2 Collateralized Debt Obligations (CDOs)

Consider a *tranche* defined by an interval $[a, b], 0 \leq a < b < 1$ for the loss process normalized by the total nominal. A CDO tranche swap (or simply CDO tranche)

is a bilateral contract in which an investor sells protection on all portfolio losses within the interval $[a, b]$ over some time period $[0, t_J]$ in return for a periodic spread $S(a, b)$ paid on the nominal remaining in the tranche after losses have been accounted for.

The loss of an investor exposed to the tranche $[a, b]$ is

$$L_{a,b}(t) = (L_t - a)_+ - (L_t - b)_+. \quad (3)$$

The premium leg is represented by the cash flow payed by the protection buyer to the protection seller. In case of a premium S , its value at time $t = 0$ is

$$P(a, b, t_J) = \sum_{j=1}^J S(t_j - t_{j-1}) E^{\mathbb{Q}}[B(0, t_j)((b - L(t_j))^+ - (a - L(t_j))^+)]$$

The default leg is represented by the cash payed by the protection seller to the protection buyer in case of default. Its value at time $t = 0$ is

$$D(a, b, t_J) = \sum_{j=1}^J E^{\mathbb{Q}}[B(0, t_j)(L_{a,b}(t_j) - L_{a,b}(t_{j-1}))].$$

The “fair spread” (or simply, the tranche spread) is the premium value $S_0(a, b, t_J)$ that equates the values of the two legs:

$$S_0(a, b, t_J) = \frac{E^{\mathbb{Q}} \sum_{j=1}^J B(0, t_j)[L_{a,b}(t_j) - L_{a,b}(t_{j-1})]}{E^{\mathbb{Q}} \sum_{j=1}^J B(0, t_j)(t_j - t_{j-1})[(b - L(t_j))^+ - (a - L(t_j))^+]}$$

Table 1 gives an example of such a tranche structure and the corresponding spreads for a standardized portfolio, the ITRAXX index. Note that these expressions for the tranche spreads depend on the portfolio loss process via the expected tranche notionals $C(t_j, K)$ where

$$C(t, K) = E^{\mathbb{Q}}[(K - L_t)^+]. \quad (4)$$

2.3 Top-down models for CDO pricing

It is immediately observed that the above expressions for the spread of a CDO tranche depend on the portfolio characteristics only through the (risk-neutral) law of the loss process L_t . The idea of “top-down” pricing models [1, 19, 21, 25, 28] is to model the risk neutral loss process, either by specifying the dynamics of the cumulative loss [1, 19, 21, 25] or by looking at the forward loss distribution [28]. We adopt here the former approach, which is simpler to implement.

The loss L_t is a piecewise constant process with upward jumps at each default event: its path is therefore completely characterized by the default times $(\tau_j)_{j \geq 1}$, representing default events and the jump sizes ΔL_j representing the loss given default. Here τ_j denotes the j -th default event observed in the portfolio: the

Maturity	Low	High	Bid\ Upfront	Mid\ Upfront	Ask\ Upfront
5Y	0%	3%	11.75%	11.88%	12.00%
	3%	6%	53.75	54.50	55.25
	6%	9%	14.00	14.75	15.50
	9%	12%	5.75	6.25	6.75
	12%	22%	2.13	2.50	2.88
	22%	100%	0.80	1.05	1.30
7Y	0%	3%	26.88%	27.00%	27.13%
	3%	6%	130	131.50	132
	6%	9%	36.75	37.00	38.25
	9%	12%	16.50	17.25	18.00
	12%	22%	5.50	6.00	6.50
	22%	100%	2.40	2.65	2.90
10Y	0%	3%	41.88%	42%	42.13%
	3%	6%	348	350.50	353
	6%	9%	93	94.00	95
	9%	12%	40	41.00	42
	12%	22%	13.25	13.75	14.25
	22%	100%	4.35	4.60	4.85

Table 1: CDO tranche spreads, in bp, for the ITRAXX index on March 15 2007. For the equity tranche the periodic spread is 500bp and figures represent upfront payments.

index j is *not* associated with the default of a given obligor but with the ordering in time of the events. The idea of aggregate loss models is to represent the rate of occurrence of defaults in the portfolio via the *portfolio default intensity* λ_t : we model the number of defaults $(N_t)_{t \in [0, T^*]}$ is a point process with $\mathcal{F}_t$ -intensity $(\lambda_t)_{t \in [0, T^*]}$ under $\mathbb{Q}$ i.e.

$$N_t - \int_0^t \lambda_t dt$$

is an $\mathcal{F}_t$ -local martingale under $\mathbb{Q}$ [6]. Intuitively, λ_t can be seen as probability per unit time of the next default conditional on current market information:

$$\lambda_t = \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \mathbb{Q}[N_{t+\Delta t} = N_t + 1 | \mathcal{F}_t]$$

Here $\mathcal{F}_t$ represents the coarse-grained information resulting from the observation of the aggregate loss process L_t of the portfolio and risk factors affecting it. In the simplest case it corresponds to the information (filtration) generated by the variables $\tau_j, \Delta L_j$ but it may also contain information on other market variables. This risk neutral intensity λ_t can be interpreted as the short term credit spread for protection against the first default in the portfolio [28].

λ_t can be modeled as a stochastic process that may depend on the loss process itself. The simplest specification is to model the loss L_t as a compound Poisson process [7], but since the intensity is constant and independent of the loss process, this does not enable to model features such as spread volatility or default contagion [14]. Spread volatility can be introduced by modeling λ_t as an autonomous jump-diffusion process and then constructing N_t as a Cox process: conditional on $(\lambda_t)_{t \in [0, T^*]}$, N has the law of a Poisson process with intensity $(\lambda_t)_{t \in [0, T^*]}$. This approach, common in the credit risk literature, has been used by Longstaff & Rajan [25] to model aggregate default rates in the CDX index. Default contagion can be incorporated in the model by introducing a dependence of the default intensity on the number of defaults. Ding et al. [15] construct the default process by starting from a linear birth process with immigration $\lambda_t = c + gN_t$ and applying a time change, while Arnsdorff & Halperin [1] use a two factor specification: $\lambda_t = \lambda_0(N_0 - N_t)Y_t$ where Y_t is a non-negative stochastic process (see also [26]). Finally, one can argue that not only the occurrence of defaults but also their timing and magnitude can affect the default intensity: this feature has been modeled using Hawkes or self-exciting processes [19, 21].

Given the wide variety of models available for the default intensity, the choice of the model class among the above is not easy in practice. Indeed, even at the qualitative level it is not obvious which parametric specifications adequately reproduce observed features of market data. Also, once the class of models has been chosen, it is a nontrivial task to calibrate the model parameters in order to reproduce market spreads of index CDO tranches. In fact, in the models described above, numerical methods must be used to compute tranche spreads so the corresponding inverse problem of recovering parameters from market quotes is a computationally intensive one. Finally, these parameterizations mainly stem

from analytical convenience, more than from any fundamental economic considerations, so a nonparametric approach which makes fewer arbitrary assumptions on the form of the default intensity can provide some insight for model selection.

3 The information content of CDO tranches

One issue in the design and calibration of top-down models is how to parameterize the portfolio loss process in a general, yet parsimonious, way which can be flexible enough to accommodate market observations of tranche spreads and remain tractable. The main issue is how to specify the dependence of the default intensity λ_t with respect to other variables in the model: existing models range from a deterministic intensity $\lambda(t)$ to full path-dependence with respect to the loss process [19, 21].

While richer models might generate more realistic statistical features, an important issue in model calibration is the *identifiability* of such complex models. Given current prices of portfolio credit derivatives, what can be inferred from them in terms of the characteristics of the loss process? In this section we present a result which sheds light on this identifiability issue, showing that the marginal distributions of any marked point process with IID marks can be matched by a Markovian jump process. From this “mimicking theorem” we conclude that the retrievable information in the intensity process is exactly given by its conditional expectation given the loss process, which we call the *effective intensity*.

3.1 Mimicking marked point processes with Markovian jump processes

We first show a “mimicking theorem” which shows that the marginal distributions of any marked point process with IID marks can be matched by a Markovian jump process:

Proposition 1. *Consider any non-explosive jump process $(L_t)_{t \in [0, T^*]}$ with a intensity process $(\lambda_t(\omega))_{t \in [0, T^*]}$ and IID jumps with distribution F . Define $(\tilde{L}_t)_{t \in [0, T^*]}$ as the Markovian jump process with jump size distribution F and intensity*

$$\lambda_{\text{eff}}(t, l) = E^{\mathbb{Q}}[\lambda_t | L_{t-} = l, \mathcal{F}_0] \quad (5)$$

Then, for any $t \in [0, T^]$, L_t and $\tilde{L}_t$ have the same distribution conditional on $\mathcal{F}_0$. In particular, the flow of marginal distributions of $(L_t)_{t \in [0, T^*]}$ only depends on the intensity $(\lambda_t)_{t \in [0, T^*]}$ through its conditional expectation $\lambda_{\text{eff}}(., .)$.*

We call (5) the *effective intensity* associated to the process L . The relation between the intensity λ_t and the effective intensity $\lambda_{\text{eff}}(t, L_{t-})$ is analogous to the relation between instantaneous volatility and local volatility in diffusion models [11, 17].

Proof. Consider any bounded measurable function $f(\cdot)$. Using the pathwise decomposition of L_T into the sum of its jumps we can write

$$f(L_T) = f(L_0) + \sum_{0 < s \leq T} (f(L_{s-} + \Delta L_s) - f(L_{s-})) \quad (6)$$

so

$$\begin{aligned} E[f(L_T)|\mathcal{F}_0] &= f(L_0) + E\left[\sum_{0 < s \leq T} (f(L_{s-} + \Delta L_s) - f(L_{s-}))|\mathcal{F}_0\right] \\ &= f(L_0) + \int_0^\infty F(dy) \int_0^T dt \quad E[(f(L_{t-} + \Delta L_t) - f(L_{t-}))\lambda_t|\mathcal{F}_0] \end{aligned}$$

Denote

$$\mathcal{G}_t = \sigma(\mathcal{F}_0 \vee L_{t-})$$

the information set obtained by adding the knowledge of L_{t-} to the current information set $\mathcal{F}_0$. Define the *local intensity* function

$$\lambda_{\text{eff}}(t, l) = E^\mathbb{Q}[\lambda_t|\mathcal{F}_0, L_{t-} = l]. \quad (7)$$

Noting that $\mathcal{F}_0 \subset \mathcal{G}_t$ we have

$$\begin{aligned} E[(f(L_{t-} + y) - f(L_{t-}))\lambda_t|\mathcal{F}_0] &= E[E[(f(L_{t-} + y) - f(L_{t-}))\lambda_t|\mathcal{G}_t]|\mathcal{F}_0] = \\ E[(f(L_{t-} + y) - f(L_{t-}))E[\lambda_t|\mathcal{G}_t]|\mathcal{F}_0] &= E[\lambda_{\text{eff}}(t, L_{t-})(f(L_{t-} + y) - f(L_{t-}))|\mathcal{F}_0] \quad \text{so} \\ E[f(L_T)|\mathcal{F}_0] &= f(L_0) + E\left[\int_0^T dt \lambda_{\text{eff}}(t, L_{t-}) \int F(dy) (f(L_{t-} + y) - f(L_{t-}))\right]|\mathcal{F}_0 \end{aligned}$$

The above equality shows that $E[f(L_T)|\mathcal{F}_0]$ has the same value as $E[f(\tilde{L}_T)|\mathcal{F}_0]$ where $(\tilde{L}_t)_{0 \leq t \leq T}$ is the Markovian loss process with intensity $\gamma_t = \lambda_{\text{eff}}(t, \tilde{L}_{t-})$ and jump size distribution F , which shows the result. $\square$

Remark 1. *Proposition 1 can be viewed as a “mimicking theorem”: it states that the flow of marginal distributions of a (rather general) point process with (a possibly path-dependent) intensity λ_t can be matched by a Markovian jump process whose intensity is given by (7). In this sense, it is a discontinuous analogue of a similar result of Gyöngy [23] in the context of continuous martingales driven by Brownian motion.*

Note that this result also applies regardless of whether the filtration $\mathcal{F}_t$ is the natural filtration of L . In other words, the intensity (λ_t) can depend not only on the history of the (marked) point process itself but also on a richer information set as in the settings where λ_t is constructed through a stochastic differential equation involving an auxiliary Brownian motion W [1, 25, 20]. Even in these cases, however, the construction of $\tilde{L}_t$ does not involve any knowledge of the filtration of the Brownian motion: in other words, we can mimick the flow of marginal distributions of L by a process which is constructed on a smaller filtration.

3.2 Information content of portfolio credit derivatives

Consider now a portfolio loss model defined by a stochastic default intensity process (λ_t) and IID losses given default with distribution F . Applying the above result we obtain the following

Corollary 1. *The value $E^{\mathbb{Q}}[f(L_T)|\mathcal{F}_0]$ at $t = 0$ of any derivative whose payoff depends on the aggregate loss L_T of the portfolio on a fixed grid of dates, only depends on the default intensity $(\lambda_t)_{t \in [0, T^*]}$ through its risk-neutral conditional expectation with respect to the current loss level:*

$$\lambda_{\text{eff}}(t, l) = E^{\mathbb{Q}}[\lambda_t | L_{t-} = l, \mathcal{F}_0] \quad (8)$$

In particular, CDO tranche spreads and mark-to-market value of CDO tranches only depends on the transition rate $(\lambda_t)_{t \in [0, T^]}$ through the effective default intensity $\lambda_{\text{eff}}(\cdot, \cdot)$.*

Remark 2 (Case of index default swap). *The cash flows of an index default swap depend both on the loss L_t and on N_t , the number of defaults. In the case where recovery rates are constant $L_t = N_t \delta$ where $\delta = (1 - R)/n$ is the loss given a single default so the above results also apply to the index default swap rate, whose value only depends on the effective intensity (5). In the general case of random loss given default, N_t is a path-dependent functional of L but the above proof can be easily adapted to show that the index default swap only depends on $(\lambda_t)_{t \in [0, T^*]}$ through*

$$\bar{\lambda}(t, l, n) = E^{\mathbb{Q}}[\lambda_t | \mathcal{F}_0, L_{t-} = l, N_{t-} = n]$$

In the sequel we shall consider the more commonly used setting where the loss is proportional to the number of defaults.

Being able to mimick the marginal distribution of the loss processes using a Markovian model allows for considerable simplifications in pricing algorithms. First, it is well known that for a Markovian jump process the transition probabilities can be computed by solving a Fokker Planck equation. Combined with Proposition 1, this shows that the transition probabilities $q_j(t, T) = \mathbb{Q}(N_T = j | \mathcal{F}_t)$ also solve the Fokker-Planck equation corresponding to the effective intensity:

$$\frac{dq_0}{dT}(t, T) = -\lambda_{\text{eff}}(T, 0)q_0(t, T) \quad (9)$$

$$\frac{dq_j}{dT}(t, T) = -\lambda_{\text{eff}}(T, j)q_j(t, T) + \lambda_{\text{eff}}(T, j-1)q_{j-1}(t, T) \quad (10)$$

$$\frac{dq_n}{dT}(t, T) = \lambda_{\text{eff}}(T, n-1)q_{n-1}(t, T) \quad (11)$$

Moreover, by analogy with the Dupire equation for diffusion models [17], one can show that the expected tranche notional $C(T, K)$ can be obtained by solving

a (single) Dupire-type forward equation [11]:

$$\begin{aligned} \frac{\partial C(T, K)}{\partial T} - C(T, K - \delta) \lambda_k(T) + \lambda_{k-1}(T) C(T, K) \\ + \sum_{j=1}^{k-2} [\lambda_{j+1}(T) - 2\lambda_j(T) + \lambda_{j-1}(T)] C(T, j\delta) = 0 \end{aligned} \quad (12)$$

where $\lambda_k(T) = \lambda_{\text{eff}}(T, k\delta)$. This is a bidiagonal system of ODEs which can be solved efficiently in order to compute the expected tranche notionals (and thus the values of CDO tranches) given the local intensity function $\lambda_{\text{eff}}(\cdot, \cdot)$ without Monte Carlo simulation.

4 The calibration problem

The model calibration problem for CDO pricing models can be defined as the problem of recovering the law of the portfolio default intensity $(\lambda_t)_{t \in [0, T]}$ from market observations, which consist of spreads for (a small number of) CDO tranches.

Denote by $T_1 < \dots < T_m$ the maturities of the observed CDO tranches (usually $m = 3$ or 4) with $T = T_m$ being the largest maturity and $0, K_1, \dots, K_I$ the attachment points. We shall use the notations of section 2: the payment dates are denoted $(t_j, j = 1..J)$. At $t = 0$ we observe the tranche spreads $(S_0(K_i, K_{i+1}, T_k), i = 1..I - 1, k = 1..m)$. The calibration problem for a “top-down” CDO pricing model can be formulated as follows:

Problem 1 (Calibration problem). *Given a set of observed CDO tranche spreads $(S_0(K_i, K_{i+1}, T_k), i = 1..I - 1, k = 1..m)$ for a reference portfolio, construct a (risk-neutral) default rate/ loss intensity $\lambda = (\lambda_t)_{t \in [0, T]}$ such that the spreads computed under the model $\mathbb{Q}^\lambda$ match the market observations:*

$$S_0(K_i, K_{i+1}, T_k) = \frac{\sum_{t_j \leq T_k} B(0, t_j) E^{\mathbb{Q}^\lambda} [L_{K_i, K_{i+1}}(t_j) - L_{K_i, K_{i+1}}(t_{j-1})]}{\sum_{t_j \leq T_k} B(0, t_j) (t_j - t_{j-1}) E^{\mathbb{Q}^\lambda} [(K_{i+1} - L(t_j))^+ - (K_i - L(t_j))^+]} \quad (13)$$

Proposition 1 tells us that if the data can be calibrated using a default intensity process λ they can also be calibrated using its Markovian projection λ_{eff} defined by (5). Using this fact we can restrict $\mathbb{Q}^\lambda$ to the set Λ of *Markovian* loss processes.

Stated in this form, the calibration problem is an ill-posed problem: it requires to reconstitute the law of an unknown stochastic process (the portfolio loss) given a finite (and typically, small) number of observations.

4.1 Calibration as relative entropy minimization under constraints

Problem 1 is an ill-posed inverse problem, similar to the one which arises in the calibration of pricing models for equity and index derivatives, where one

attempts to recover a risk-neutral probability measure from a finite set of option prices: there is little hope to obtain a unique solution, let alone to compute it in a stable manner. One solution strategy is to restore uniqueness and stability by adding some information into the problem in the form of a prior model $\mathbb{Q}_0$ and looking for the risk-neutral loss process verifying the calibration constraints (13) which is the “closest” to $\mathbb{Q}_0$ in some sense. The *relative entropy* of $\mathbb{Q}$ with respect to $\mathbb{Q}_0$, defined as

$$E^{\mathbb{Q}_0} \left[\frac{d\mathbb{Q}^\lambda}{d\mathbb{Q}_0} \ln \frac{d\mathbb{Q}^\lambda}{d\mathbb{Q}_0} \right]$$

may be used to quantify the closeness of a given loss process with law $\mathbb{Q}$ to the prior model; it also allows for an information-theoretic interpretation [13]. This approach has been successfully applied to calibration of option pricing models in various static [3, 30] and dynamic [2, 4, 9] settings. Finally, this approach to model calibration is linked via duality to exponential utility maximization problems [22].

For a Markovian loss process $\mathbb{Q}^\lambda \in \Lambda$ denote by $(\lambda_t)_{t \in [0, T^*]}$ its $\mathcal{F}_t$ -predictable default intensity. Given market tranche spreads $S(K_i, K_{i+1}, T_k)$ and a prior guess γ_t for the loss intensity, the calibration problem for the default intensity can be formalized as:

Problem 2 (Calibration via relative entropy minimization). *Given a prior loss process with law $\mathbb{Q}_0$, find a loss process with law $\mathbb{Q}^\lambda$ and default intensity $(\lambda_t)_{t \in [0, T^*]}$ which minimizes*

$$\inf_{\mathbb{Q}^\lambda \in \Lambda} E^{\mathbb{Q}_0} \left[\frac{d\mathbb{Q}^\lambda}{d\mathbb{Q}_0} \ln \frac{d\mathbb{Q}^\lambda}{d\mathbb{Q}_0} \right] \quad \text{under} \quad E^{\mathbb{Q}^\lambda} [H_{i,k}] = 0 \quad (14)$$

where

$$\begin{aligned} H_{ik} = & S_0(K_i, K_{i+1}, T_k) \sum_{t_j \leq T_k} B(0, t_j) (t_j - t_{j-1}) [(K_{i+1} - L(t_j))^+ - (K_i - L(t_j))^+] \\ & + \sum_{t_j \leq T_k} B(0, t_j) [(K_{i+1} - L(t_j))^+ - (K_i - L(t_j))^+ - (K_{i+1} - L(t_{j-1}))^+ + (K_i - L(t_{j-1}))^+] \end{aligned} \quad (15)$$

and $\mathbb{Q}^\lambda$ denotes the law of the point process with intensity $(\lambda_t)_{t \in [0, T^*]}$ and $\mathbb{Q}_0$ is the law of the point process with intensity $(\gamma_t)_{t \in [0, T^*]}$.

The rest of the paper is devoted to the solution of this problem. We will see that the choice of relative entropy as calibration criterion makes the problem both well-posed and tractable: we will exhibit an efficient numerical method for solving the problem and apply this method to data sets of index CDOs to extract implied default intensities from index CDO tranche spreads.

4.2 Computation of the relative entropy

Other than its information-theoretic interpretation and its convexity, another advantage of using the relative entropy as a calibration criterion is that it can

be computed efficiently in the case of point processes. This is due to the fact that the law of a point process with a given intensity can be obtained from, say, the law of a Poisson process by an *exponential* change of measure [5, 6], while the computation of entropy involves computing the logarithm of the corresponding density. This feature makes the relative entropy explicitly computable for many useful classes of loss processes. We shall use the following result [5, 6] on equivalent changes of measure for point processes:

Proposition 2. *Let N_t be a point process with intensity γ_t on $(\Omega, \mathcal{F}_t, \mathbb{Q}_0)$. Let $\lambda = (\lambda_t)_{t \in [0, T]}$ be a nonnegative, $\mathcal{F}_t$ -predictable process such that*

$$\int_0^t \lambda_s ds < \infty \quad \mathbb{Q}_0 - a.s. \quad (16)$$

and $\{t \geq 0, \lambda_t > 0\} = \{t \geq 0, \gamma_t > 0\}$ a.s. Define the process

$$Z_t = \left(\prod_{\tau_j \leq t} \frac{\lambda_{\tau_j}}{\gamma_{\tau_j}} \right) \exp \left\{ \int_0^t (\gamma_s - \lambda_s) ds \right\} \quad (17)$$

where $\tau_1 \leq \tau_2 \leq \tau_3 \leq \dots$ are the jump times of N . Suppose moreover that $E^{\mathbb{Q}_0}[Z_T] = 1$ and define the probability measure $\mathbb{Q}^\lambda$ on $\mathcal{F}_t$ by

$$\frac{d\mathbb{Q}^\lambda}{d\mathbb{Q}_0} = Z_T \quad (18)$$

Then N_t is a point process with $\mathcal{F}_t$ intensity $(\lambda_t)_{t \in [0, T]}$ under $\mathbb{Q}^\lambda$.

Taking $\gamma_0 = 1$ (Poisson process) this result can be used to construct (via change of measure) the law of a process with a given intensity $(\lambda_t)_{t \in [0, T]}$. The above result allows to compute the entropy of $\mathbb{Q}^\lambda$ relative to $\mathbb{Q}_0$, the law of the point process with intensity $(\gamma_t)_{t \in [0, T]}$:

Proposition 3 (Computation of relative entropy). *Denote by*

- $\mathbb{Q}_0$ *the law on $[0, T]$ of a point process with intensity $(\gamma_t)_{t \in [0, T]}$ and*
- $\mathbb{Q}^\lambda$ *the law on $[0, T]$ of the point process with intensity $(\lambda_t)_{t \in [0, T]}$ verifying the assumptions of proposition 2.*

The relative entropy of $\mathbb{Q}^\lambda$ with respect to $\mathbb{Q}_0$ is given by:

$$E^{\mathbb{Q}_0} \left[\frac{d\mathbb{Q}^\lambda}{d\mathbb{Q}_0} \ln \frac{d\mathbb{Q}^\lambda}{d\mathbb{Q}_0} \right] = E^{\mathbb{Q}^\lambda} \left[\int_0^T \left(\lambda_t \ln \frac{\lambda_t}{\gamma_t} dt - \lambda_t + \gamma_t \right) dt \right] \quad (19)$$

Proof. It is a straightforward application of Proposition 2.

$$E^{\mathbb{Q}_0} \left[\frac{d\mathbb{Q}^\lambda}{d\mathbb{Q}_0} \ln \frac{d\mathbb{Q}^\lambda}{d\mathbb{Q}_0} \right] = E^{\mathbb{Q}^\lambda} \left[\sum_{\tau_i \leq T} \ln \frac{\lambda_{\tau_i}}{\gamma_{\tau_i}} + \int_0^T (\gamma_t - \lambda_t) dt \right].$$

The intensity $(\lambda_t)_{t \in [0, T]}$ of the loss L under $\mathbb{Q}^\lambda$ is characterized [6] by the property that for any $\mathcal{F}_t$ -predictable process $C(t)$,

$$E^{\mathbb{Q}^\lambda} \left[\sum_{0 < \tau_i \leq T} C(\tau_i) \right] = E^{\mathbb{Q}^\lambda} \left[\int_0^T \lambda_t C(t) dt \right]$$

It follows that

$$E^{\mathbb{Q}^\lambda} \left(\sum_{0 < \tau_i \leq T} \ln \frac{\lambda_{\tau_i}}{\gamma_{\tau_i}} \right) = E^{\mathbb{Q}^\lambda} \left(\int_0^T \ln \frac{\lambda_s}{\gamma_s} dN_s \right) = E^{\mathbb{Q}^\lambda} \left(\int_0^T \lambda_s \ln \frac{\lambda_s}{\gamma_s} ds \right) \quad (20)$$

□

4.3 Dual problem as an intensity control problem

To solve the constrained optimization problem (14) by introducing Lagrange multipliers and using convex duality methods [18]. Define the Lagrangian

$$\mathcal{L}(\lambda, \mu) = E^{\mathbb{Q}^\lambda} \left[\int_0^T \left(\lambda_s \ln \frac{\lambda_s}{\gamma_s} + \gamma_s - \lambda_s \right) ds - \sum_{i=1}^I \sum_{k=1}^m \mu_{i,k} H_{ik} \right] \quad (21)$$

where μ_{ik} is the Lagrange multiplier for the inequality constraints in (15). The (primal) problem (14) is equivalent to

$$\inf_{\lambda \in \Lambda} \sup_{\mu \in \mathbb{R}^{m \cdot I}} E^{\mathbb{Q}^\lambda} \left[\int_0^T \left(\lambda_s \ln \frac{\lambda_s}{\gamma_s} + \gamma_s - \lambda_s \right) ds - \sum_{i=1}^I \sum_{k=1}^m \mu_{i,k} H_{ik} \right] \quad (22)$$

We assume that the calibration problem (14) has at least one solution i.e. the primal problem (22) is finite-valued. Then the strict convexity of the objective function implies that the above problem has the same value function (and solution) as the associated dual problem [18] given by

$$\sup_{\mu \in \mathbb{R}^{m \cdot I}} \inf_{\lambda \in \Lambda} E^{\mathbb{Q}^\lambda} \left[\int_0^T \left(\lambda_s \ln \frac{\lambda_s}{\gamma_s} + \gamma_s - \lambda_s \right) ds - \sum_{i=1}^I \sum_{k=1}^m \mu_{i,k} H_{ik} \right] \quad (23)$$

The inner optimization problem

$$J(\mu) = \mathcal{L}(\lambda^*(\mu), \mu) = \inf_{\mathbb{Q}^\lambda \in \Lambda} \mathcal{L}(\lambda, \mu)$$

is an example of an *intensity control* problem [5, 6]: the optimal choice of the intensity of a jump process in order to minimize a criterion of the type

$$E^{\mathbb{Q}^\lambda} \left[\int_0^T \varphi(t, \lambda_t) dt + \sum_{j=1}^J \Phi_j(t_j, L_{t_j}) \right], \quad (24)$$

where $\varphi(t, \lambda_t)$ is a *running cost* and $\Phi_j(t_j, L_{t_j})$ represents a “terminal” cost. In our case

$$\varphi(t, x) = x \ln \frac{x}{\gamma_t} + \gamma_t - x \quad \text{and} \quad \Phi_j(t_j, L_{t_j}) = \sum_{i=1}^I M_{ij} (K_i - L_{t_j})^+ \quad (25)$$

$$\text{where} \quad M_{ij} = B(0, t_{j+1}) \sum_{T_k \geq t_{j+1}} (\mu_{ik} - \mu_{i-1,k}) +$$

$$B(0, t_j) \sum_{T_k \geq t_j} [\mu_{ik}(1 - \Delta S(K_i, K_{i+1}, T_k)) - \mu_{i-1,k}(1 - \Delta S(K_{i-1}, K_i, T_k))] \quad (26)$$

where $\Delta = t_j - t_{j-1}$ is the interval between payment dates.

The solution of an intensity control problem can be obtained using a dynamic programming principle and is characterized in terms of a system of Hamilton-Jacobi equations [6, Ch. VII]. We will now use these properties to solve (24).

Once the inner optimization/ intensity control problem has been solved we have to solve the outer problem by optimizing $J(\mu)$ over the Lagrange multipliers $\mu \in \mathbb{R}^{mI}$: the corresponding optimal control λ^* then yields precisely the default intensity which calibrates the observations. The problem setting is similar to the one formulated by Avellaneda et al. [2] in the context of diffusion models. We will observe however that, unlike the setting of [2], we are able to solve the stochastic control problem in (24) analytically thereby greatly simplifying the algorithm.

Standard formulations of intensity control problems involve a single horizon ($J = 1$); we will first examine this case in the next section and then discuss how to extend the analysis to the case of several maturities in section 4.5.

4.4 Hamilton Jacobi equations

Let us consider first the case where $J = 1$ i.e a single time horizon is involved. The dual problem is then to minimize

$$E^{\mathbb{Q}^\lambda} \left[\int_0^T \varphi(t, \lambda_t) dt + \Phi(T, L_T) \right] \quad (27)$$

where $\Phi(\cdot)$ is of the form (25) (and thus depends on the Lagrange multipliers μ). The solution of the stochastic control problem (23) can be obtained using dynamic programming methods [5, 6]. The idea is to define a family of optimization problems indexed by the initial condition (t, n) ,

$$V(t, N_t) = \inf_{\mathbb{Q}^\lambda \in \Lambda([t, T])} E^{\mathbb{Q}^\lambda} \left[\int_t^T (\lambda_s \ln \frac{\lambda_s}{\gamma_s} + \gamma_s - \lambda_s) ds + \Phi(T, \delta N_T) \right] | \mathcal{F}_t \quad (28)$$

where $\delta = (1 - R)/n$ is the loss given default and $\Lambda([t, T])$ is the set of laws of point processes on $[t, T]$ parameterized by their intensity λ as in (17). The

value function $V(t, k)$ then solves the dynamic programming equation [6]:

$$\frac{\partial V}{\partial t}(t, k) + \inf_{\lambda \in [0, \infty[} \{ \lambda(t, k)[V(t, k+1) - V(t, k)] + \lambda(t, k) \ln \frac{\lambda(t, k)}{\gamma(t, k)} - \lambda(t, k) + \gamma(t, k) \} = 0 \quad (29)$$

$$\text{for } t \in [0, T] \text{ and } V(T, k) = \Phi(T, k\delta) \quad (30)$$

The value function of (27) is then given by $V(0, 0)$ and the optimal intensity control is obtained by maximizing over λ in the nonlinear term [6]:

Proposition 4 (Verification theorem). *If $V : [0, T] \times \mathbb{N}$ is a bounded solution of (29)–(30), differentiable in t then $\mathcal{L}(\lambda_\mu^*, \mu) = V(0, 0)$ and the optimal control λ_μ^* is given by the minimizer of*

$$\lambda_\mu^*(t, k) = \arg \min_{\lambda > 0} \lambda[V(t, k+1) - V(t, k)] + (\lambda \ln \frac{\lambda}{\gamma_t} + \gamma_t - \lambda),$$

for each t and $0 \leq k \leq n$.

In this case the maximum in the nonlinear term can be explicitly computed:

$$\lambda_\mu^*(t, k) = \gamma(t, k) e^{-[V(t, k+1) - V(t, k)]} \quad (31)$$

$$\frac{\partial V}{\partial t}(t, k) + \gamma(t, k)(1 - e^{-[V(t, k+1) - V(t, k)]}) = 0 \quad (32)$$

To solve the dual problem we need to solve the Hamilton–Jacobi equations (29)–(30). This is a system of n nonlinear ODEs which may seem daunting at first glance. Remarkably, in this case a logarithmic change of variable yields an explicit solution:

Proposition 5 (Value function). *Consider a function Φ such that $\Phi(x) = 0$ for $x \geq n\delta$. The solution of (29)–30 has the probabilistic representation*

$$V(t, k) = -\ln[1 + \sum_{j=0}^{n-k} \mathbb{Q}_0(N_T = k+j | N_t = k)(e^{-\Phi(T, (k+j)\delta)} - 1)] \quad (33)$$

Corollary 2 (Case of Poisson prior). *If the prior process is a Poisson process with intensity γ_0 stopped at n , then the value function V is given by*

$$V(t, k) = \Phi(T, n\delta) - \ln[1 + \sum_{j=0}^{n-k-1} \frac{\gamma_0^j (T-t)^j e^{-\gamma_0(T-t)}}{j!} (e^{\Phi(T, n\delta) - \Phi(T, (k+j)\delta)} - 1)] \quad (34)$$

Proof. If we consider $u(t, k) = e^{-V(t, k)}$ then u solves a linear equation

$$\frac{\partial u(t, k)}{\partial t} + \gamma(t, k)(u(t, k+1) - u(t, k)) = 0 \quad \text{with } u(T, k) = \exp(-\Phi(T, k\delta))$$

which is recognized as the backward Kolmogorov equation associated with the Markovian point process with intensity function $\gamma(t, k)$ (i.e. the prior process, with law $\mathbb{Q}_0$). The solution is thus given by the Feynman-Kac formula

$$u(t, k) = E^{\mathbb{Q}_0}[e^{-\Phi(T, \delta N_T)} | N_t = k]$$

The expectation is easily computed using the transition probabilities of the prior process, where the sum over jumps can be truncated using the fact that $\Phi(x) = 0$ for $x \geq n\delta$:

$$\begin{aligned} u(t, k) &= \sum_{j=0}^{n-k} \mathbb{Q}_0(N_T = k+j | N_t = k) e^{-\Phi(T, (k+j)\delta)} + \sum_{j>n-k} \mathbb{Q}_0(N_T = k+j | N_t = k) \\ &= \sum_{j=0}^{n-k} \mathbb{Q}_0(N_T = k+j | N_t = k) e^{-\Phi(T, (k+j)\delta)} + 1 - \sum_{j=0}^{n-k} \mathbb{Q}_0(N_T = k+j | N_t = k) \\ &= 1 + \sum_{j=0}^{n-k} \mathbb{Q}_0(N_T = k+j | N_t = k) [e^{-\Phi(T, (k+j)\delta)} - 1] \end{aligned}$$

which leads to (33). These transitions probabilities can be explicitly computed for a (stopped) Poisson process which then leads to (34). $\square$

The fact that a logarithmic change of variable linearizes the Hamilton Jacobi equation is not a coincidence: this is a common feature of stochastic control problems related to exponential utility maximization [31]. This result can also be derived using the dual representation of the entropic risk measure as in [27].

4.5 Handling payment dates

In the (realistic) case where several payment dates $0 \leq t_1 \leq t_2 \leq \dots \leq t_J$ are involved, the criterion to be optimized in the dual problem is of the form

$$E^{\mathbb{Q}^\lambda} \left[\int_0^{t_J} \varphi(t, \lambda_t) dt + \Phi_1(t_1, L_{t_1}) + \Phi_2(t_2, L_{t_2}) + \dots \Phi_J(t_J, L_{t_J}) \right].$$

We will now show that this problem can be treated as a sequence of single-horizon intensity control problems in a recursive manner using a dynamic programming principle. Denote by $\Lambda([t_j, t_{j+1}])$ the restriction of loss processes in Λ to $t \in [t_j, t_{j+1}]$. Consider the value function:

$$V(t, k; \mu) = \inf_{\Lambda([t, t_J])} E^{\mathbb{Q}^\lambda} \left[\int_t^{t_J} \varphi(t, \lambda_t) dt + \sum_{t_j > t} \Phi_j(t_j, L_{t_j}) | N_t = k \right]$$

We will compute V going backwards from t_J . First, we note that $V(t_{J-1}, k; \mu)$ is of the form (27) and can be computed using the formula (33) with $\Phi = \Phi_J$. Assume now we have computed $V(t, k; \mu)$ for $t \geq t_{j+1}$. Then

$$\begin{aligned} V(t_j, k; \mu) &= \inf_{\Lambda([t_j, t_J])} E^{\mathbb{Q}^\lambda} \left[\int_{t_j}^{t_{j+1}} \varphi(t, \lambda_t) dt + \Phi_{j+1}(t_{j+1}, L_{t_{j+1}}) \right. \\ &\quad \left. + \int_{t_{j+1}}^{t_J} \varphi(t, \lambda_t) dt + \sum_{i=j+2}^J \Phi_i(t_i, L_{t_i}) | N_{t_j} = k \right] \end{aligned} \quad (35)$$

The dynamic programming principle can be stated by saying that the cost functional is a martingale when computed at the optimal policy λ^* , hence:

$$\begin{aligned} V(t_j, k; \mu) &= \\ E^{\mathbb{Q}^*} \left[\int_{t_j}^{t_{j+1}} \varphi(t, \lambda_t^*) dt + \Phi_{j+1}(t_{j+1}, L_{t_{j+1}}) + \int_{t_{j+1}}^{t_J} \varphi(t, \lambda_t^*) dt + \sum_{i=j+2}^J \Phi_i(t_i, L_{t_i}) | N_{t_j} = k \right] \\ &= \inf_{\Lambda([t_j, t_J])} E^{\mathbb{Q}^\lambda} \left[\int_{t_j}^{t_{j+1}} \varphi(t, \lambda_t) dt + \Phi_{j+1}(t_{j+1}, L_{t_{j+1}}) + V(t_{j+1}, k; \mu) | N_{t_j} = k \right] \end{aligned}$$

Therefore on $[t_j, t_{j+1}[$ we also have a problem of the form (27) with $\Phi = F_{j+1} = \Phi_{j+1} + V(t_{j+1}, \cdot)$: $V(t_j, k; \mu)$ can therefore be computed using the formula (33) with $\Phi = F_{j+1}$. This results in the following method for computing recursively the value function $V(t, k; \mu)$:

1. Start from the last payment date $j = J$ and set $F_J(k) = \Phi_J(t_J, \delta k)$.
2. Solve the Hamilton–Jacobi equations (29) on $]t_{j-1}, t_j]$ backwards starting from the terminal condition

$$V(t_j, k, \mu) = F_j(k) \tag{36}$$

$V(t_j, k, \mu)$ can be explicitly computed for $t \in]t_{j-1}, t_j]$ using (33) with $\Phi = F_j$.

3. Set $F_{j-1}(k) = V(t_{j-1}, k) + \Phi_{j-1}(t_{j-1}, k\delta)$
4. Go to step 2 and repeat until $j = 0$ is reached.

The value function of the dual problem is then given by $V(0, 0, \mu)$. This procedure yields an explicit (although lengthy) formula for $V(0, 0, \mu)$, which is obtained by nesting J times the expression (33). In particular this formula can be used to compute $\nabla V(0, 0, \mu)$ and to use gradient-based methods to minimize $V(0, 0, \mu)$ with respect to μ in the last step of the algorithm.

5 Recovering market-implied default rates

5.1 Calibration algorithm

The above results lead to a non-parametric algorithm for recovering a market-implied portfolio default intensity from CDO spreads. The algorithm consists of the following steps:

1. Solve the dynamic programming equations (29)–(30) for $\mu \in \mathbb{R}^{m.I}$ to compute $V(0, 0, \mu)$.
2. Optimize $V(0, 0, \mu)$ over $\mu \in \mathbb{R}^{m.I}$ using a gradient-based method:

$$\sup_{\mu \in \mathbb{R}^{m.I}} V(0, 0, \mu) = V(0, 0, \mu^*) = V^*(0, 0)$$

3. Compute the calibrated default intensity (optimal control) as follows:

$$\lambda^*(t, k) = \gamma(t, k)e^{V^*(t, k) - V^*(t, k+1)} \quad (37)$$

4. Compute the term structure of loss probabilities by solving the Fokker-Planck equations (10).
5. The calibrated default intensity $\lambda^*(., .)$ can then be used to compute CDO spreads for different tranches, forward tranches,...: first we compute the expected tranche notional $C(T, K)$ by solving the forward equation (12) and then use the expected tranche notional to evaluate CDO tranche spreads, mark to market value, etc. In particular the calibrated default intensity can be used to “fill the gaps” in the base correlation surface in an arbitrage-free manner, by first computing the expected tranche loss for all strikes and then computing the base correlation for that strike.

5.2 Application to ITRAXX tranches

We have applied the above methodology to several data sets of CDO quotes; we present here only the results for three data sets, consisting of ITRAXX Europe IG tranche quotes on Sept 26, 2005, March 15, 2007 and March 25, 2008.

Maturity	Low	High	Bid \ Upfront	September 26, 2005	March 25, 2008
5Y	0%	3%	29.50%	29.875%	38.67 %
	3%	6%	96	98	454.08
	6%	9%	33	34.5	280.22
	9%	12%	13	14	189.40
	12%	22%	7.50	8.125	110.74
	22%	100%	2.25	3.125	46.87
7Y	0%	3%	47.1%	47.55 %	43.97%
	3%	6%	193	196.5	514.76
	6%	9%	52	54.5	312.50
	9%	12%	29	31.5	206.53
	12%	22%	12	13.5	115.47
	22%	100%	5.25	6.25	48.55
10Y	0%	3%	58.25%	58.75%	48.43%
	3%	6%	505	512.5	633.16
	6%	9%	100	103	362.40
	9%	12%	48	51.5	238.54
	12%	22%	22	23.5	25
	22%	100%	8.25	9.5	10.75

Table 2: ITRAXX IG Europe tranche spreads (mid), September 26, 2005 vs March 25, 2008.

Figure 1: Model vs market spreads: ITRAXX September 26, 2005 (left) Sept 2008 (right).

Figure 2: Implied ITRAXX default intensity functions: September 2005 (left) vs Sept 2008 (right).

Figure 2 displays the local intensity function $\lambda(t, l)$ as a function of time t and the number of defaults l .

Several features deserve to be commented. First, we note the strong dependence of the default intensity on the portfolio loss level: as noted in section 3, this dependence is the signature of 'contagion effects' in CDO tranches. Figure 2 shows the dependence of the default intensity with respect to the number of defaults at two different dates (in 2005 and 2008). We observe a similar pattern in both cases: while the initial default rate is quite low (less than 0.5, which means on average less one default every two years), this default intensity quickly increases as defaults occur in the portfolio, which is a clear signature of default contagion. Such contagion effects, which leads to the clustering of defaults, have been observed in historical time series [14]: our results indicated that their effect is also detectable in the *implied* default intensity, i.e. that contagion risk is effectively priced into market quotes of CDO tranches. In pricing terms, this steep initial slope means that in this period (2005-2007) equity tranches are priced

Figure 3: Implied default intensity as a function of number of defaults at a 2 year horizon: Sept 2005 (left) and Sept 2008 (right).

Figure 4: Left: term structure of loss distributions implied by ITRAXX Europe Series 6, March 15 2007. Right: loss distributions at various maturities.

relatively cheaply with respect to mezzanine or senior tranches. The values of $\lambda(t, k)$ for small k also give interesting insights for the pricing of first-to-default and k -th to default swaps.

Once the equity tranche of the portfolio is wiped out by defaults, we observe in figure 2 a plateau where the default intensity remains relatively insensitive to the number of defaults: in this regime, in fact, a constant, Poisson-type approximation seems to work well. This regime corresponds to the bulk of the portfolio, composed of obligors whose default risk is well represented by the average spread of the portfolio. From a pricing perspective, this flat region implied that, in these examples, apart from the equity tranche, the other tranches were priced assuming a constant (and high) value of the default intensity once the equity tranche has been wiped out. The steep decline of the $\lambda(t, k)$ for large k can be understood as corresponding to the group of obligors in the index with the lowest spreads/ default risk and which are the least exposed to systemic risk: they are the last to default, with a very low probability.

Finally we note that, as illustrated in Figures 1, 2 and 3, both the precision of the calibration the qualitative features of the default intensity function remain the same throughout the period 2005-2008, a particularly turbulent period during which base correlations computed using Gaussian copula models have been notoriously unstable and sometimes impossible to calibrate to market spreads. This shows that the instability of such “default correlations” parameters is linked more to model mis-specification than to genuine non-stationarity: using a richer model structure along with a stable calibration algorithm restores a greater degree of parameter stability. This aspect is of course essential if the model is to be used for hedging [10].

We also note that there is a discontinuity in the dependence on t at each observed maturity: this discontinuity is a structural feature related to properties of the dynamic programming equation and does not have any informational content. Such discontinuities are not present in quantities such as default probabilities (Figure 4).

The above approach can be used to construct an arbitrage-free interpolation/extrapolation of ‘base correlations’, by first calibrating the local intensity function to the observed tranche spreads then computing expected tranche losses for a fine grid of detachment points/maturities and converting them into a base correlation figure. Note that, unlike the usual linear interpolation of base correlations, this method also provides an arbitrage-free extrapolation of base correlations beyond the largest detachment point and below the smallest attachment point. Figure 5 shows the result of such an interpolation for the ITRAXX data, compared with the linear interpolation method used by many market participants. The difference between the two methods is striking, especially for senior tranches.

Figure 5: Base correlation surface generated by the calibrated model: ITRAXX Europe Series 6, March 15 2007.

6 Conclusion

We have proposed a rigorous methodology for calibrating a CDO pricing model to market data, by formulating the calibration problem as a relative entropy minimization problem under constraints and mapping it into an intensity control problem for a point process, which can be solved analytically.

By contrast with other calibration methods proposed for top-down CDO pricing models in the literature, our method proposed is *nonparametric*: it does not assume any arbitrary functional form for the default intensity. Another feature of algorithm proposed is that it does not require preliminary interpolation or smoothing of CDO data in maturity or strike (which may violate arbitrage constraints), nor does it require a preliminary (model-dependent) “stripping” of CDO spreads into expected tranche notionals. In particular, our algorithm yields meaningful and stable results even for sparse data sets such as the ones available in CDS index markets.

Our method allows to compute portfolio default rates implied by index CDO quotes. Results obtained on ITRAXX tranche spreads point to default contagion effects in the riskneutral loss process and also illustrate that the implied default intensity corresponding to the first few defaults are very different from those of the bulk of the portfolio.

The model obtained from our calibration is a Markovian loss process where the default intensity depends on the current loss level and time. When compared with other possible specifications of top-down pricing models, the Markovian loss process considered here is of course quite simple in structure. Though it does account for clustering of defaults, it does not include, for instance, spread risk and the influence of other factors such as interest rates. Although more complex specifications are possible, as shown in Proposition 1, the information content

of CDO spreads does not allow to identify such models uniquely. Recently, Lopatin & Misirpashayev [26] have suggested to use a Markovian loss model as an intermediate step in the calibration of a two-factor model with richer dynamics using a relation such as (5) to link the parameters of the. In this context our algorithm can be used as the first phase of a calibration algorithm for a more complex model.

References

- [1] Arnsdorff, M. and Halperin, I. (2008) BSLP: Markovian Bivariate Spread-Loss Model for Portfolio Credit Derivatives, *Journal of Computational Finance*, 12, No. 2.
- [2] Avellaneda, M. Friedman, C. Holmes, R and Samperi D. (1997) Calibrating the volatility surfaces via relative entropy minimization, *Applied Mathematical Finance* Vol. 4, 37–64.
- [3] Avellaneda, M. (1998) The minimum-entropy algorithm and related methods for calibrating asset-pricing models, in: *Proceedings of the International Congress of Mathematicians*, (Berlin, 1998), Vol. III, Documenta Mathematica, 1998, pp. 545–563.
- [4] Avellaneda, M., Buff, R., Friedman, C., Grandchamp, N., Kruk, L., and Newman, J. (2001) Weighted Monte Carlo: a new technique for calibrating asset-pricing models, *Int. J. Theor. Appl. Finance*, 4, pp. 91–119.
- [5] Bismut, J.M. (1975) Contrôle de processus de sauts, *C.R. Académie des Sciences Paris* A281, 767–770.
- [6] Brémaud, P. (1981) *Point Processes and Queues*, New York: Springer.
- [7] Brigo D., Pallavicini A., Torresetti, R (2007) Calibration of CDO tranches with the dynamic generalized Poisson loss model, working paper.
- [8] R. Bruyère, R. Cont, R. Copinot, L. Fery, C. Jaeck, T. Spitz (2005) *Credit derivatives and structured credit*, Chichester: Wiley.
- [9] Cont, R. and Tankov, P. (2004) Non-parametric calibration of jump-diffusion option pricing models, *Journal of Computational Finance*, 7, pp. 1–49.
- [10] Cont, R. and Kan, Y. (2008) Dynamic hedging of portfolio credit derivatives, Financial Engineering Report No. 2008-08, Columbia University.
- [11] Cont, R. & Savescu, I. (2008) Forward equations for portfolio credit derivatives, Chapter 11, in: Cont, R. (ed.): *Frontiers in quantitative finance: credit risk and volatility modeling*, Wiley.

- [12] Cousin, A. & Laurent, J.P. (2008) An overview of factor models for pricing CDO tranches, in: Cont, R. (ed.): *Frontiers in quantitative finance: credit risk and volatility modeling*, Wiley.
- [13] I. Csiszar (1967) Information-type measures of difference of probability distributions and indirect observation, *Studia Sci. Math. Hungar*, 2, 299-318.
- [14] Das S., Duffie, D. and Kapadia, N. (2007) Common failings: how corporate defaults are correlated, *Journal of Finance*, Vol. 62, No. 1, (February 2007), pp 93-117.
- [15] Ding X., K. Giesecke, and P. Tomecek (2006) Time-changed birth processes and multi-name credit, working paper.
- [16] Duffie, D. and Garleanu, N. (2001) Risk and valuation of collateralized debt obligations, *Financial Analysts Journal*, Vol. 57, No. 1, (January/February 2001), pp. 41-59.
- [17] B. Dupire (1994) Pricing with a smile, *RISK*, 7, pp. 18-20.
- [18] Ekeland, I. and Temam, R. (1987) *Convex analysis and variational problems*, SIAM.
- [19] Errais, E., Giesecke, K., Goldberg, L. (2006) Pricing credit from the top down with affine point processes, Working Paper.
- [20] Giesecke, K. and Kim, B. (2007) Estimating tranche spreads by loss process simulation, in: Hendersen et al. (eds.): *Proceedings of the 2007 Winter Simulation Conference*.
- [21] Giesecke, K. (2008) Portfolio credit risk: top-down vs bottom-up, in: Cont, R. (ed.): *Frontiers in Quantitative Finance: credit risk and volatility modeling*, Wiley.
- [22] Goll, T. and Rüschendorf, L. (2002) Minimal distance martingale measures and optimal portfolios consistent with observed market prices, in: *Stochastic processes and related topics*, 141-154, Stochastics Monographs, 12, Taylor & Francis, 2002.
- [23] Gyöngy, I. (1986) Mimicking the one dimensional distributions of processes having an Ito differential, *Probability theory and related fields*, Vol. 71, No. 4, pp. 501-516.
- [24] Lipton, A. and Rennie, A. (eds.): *Credit correlation: life after copulas*, World Scientific.
- [25] Longstaff, F. and Rajan, A. (2008) An empirical analysis of the pricing of collateralized debt obligations, *Journal of Finance*, 63, 509-563.
- [26] Misirpashayev, T. and Lopatin, A. (2007) Two-dimensional Markovian model for dynamics of aggregate credit loss, www.defaultrisk.com.

- [27] R. Rouge and N. El Karoui (2000) Pricing via utility maximization and entropy, *Mathematical Finance*, 10, no. 2, 259-276.
- [28] Schönbucher, P. (2005) Portfolio losses and the term structure of loss transition rates: a new methodology for the pricing of portfolio credit derivatives, Working paper.
- [29] Sidenius J., Piterbarg V. and Andersen, L. (2008) A new framework for dynamic credit portfolio loss modeling, *International Journal of Theoretical and Applied Finance*, 11 (2), 163 – 197.
- [30] Stutzer, M (1996) A simple nonparametric approach to derivative security valuation, *Journal of Finance*, 51, pp. 1633–1652.
- [31] Zariphopoulou, Th. (2001) A solution approach to valuation with unhedgeable risks, *Finance and Stochastics*, 5, No. 1, 61–82.