Credit Innovation: Pricing and Hedging of Credit Derivatives via the Innovations Approach to Nonlinear Filtering

Rüdiger Frey
Universität Leipzig

Nice, September 2009

ruediger.frey@math.uni-leipzig.de
www.math.uni-leipzig.de/~frey

joined work with
T. Schmidt

1. Introduction

Development of a sound methodology for portfolio credit derivatives is a challenging problem: size of portfolios, scarcity of data, contagion- and network effects, spread dynamics . . .
Some progress in recent years. Nonetheless market practice mostly relies on the static Gauss copula model with its ad hoc techniques for calibration and risk management.

In this talk we present a new, information-based approach for constructing portfolio credit risk models. Key ideas/results:

- ⋆ Nonlinear filtering for deriving dynamics of traded credit derivatives; default contagion generated via updating of believes.
- ⋆ Dynamic version of the Hull-White implied copula model.
- ⋆ Consistent methodology for pricing exotic derivatives (e.g. credit index options).

Incomplete-information models: some literature

- Structural credit risk models: [Duffie and Lando, 2001], [Giesecke and Goldberg, 2004], [Jarrow and Protter, 2004], [Coculescu et al., 2006] or [Frey and Schmidt, 2006].
- Doubly-stochastic models with incomplete information such as [Collin-Dufresne et al., 2003], [Schönbucher, 2004], [Duffie et al., 2006] (empirical focus).
- [Frey and Runggaldier, 2008]. Relation between credit risk and nonlinear filtering and analysis of filtering problems in very general reduced-form model; dynamics of credit derivatives not studied.
- Default-free term-structure models: [Landen, 2001]: construction of short-rate model via nonlinear filtering.

2. The model

Throughout m firms with default times τ_i and default indicator $Y_{t,i} = 1_{\{\tau_i \leq t\}}$, $1 \leq i \leq m$; $Y_t = (Y_{t,1}, \ldots, Y_{t,m})$.

Several layers of information:

- Underlying factor model Default times τ_i are conditionally independent doubly-stochastic random times; intensities are driven by an unobservable factor X (a random variable or a finite-state Markov chain).

- Market information. Prices of traded assets are conditional expectation wrt market information $\mathbb{F}^M := \mathbb{F}^Y \vee \mathbb{F}^Z$. Z gives X in additive Gaussian noise. Filtering wrt $\mathbb{F}^M$ is used to obtain asset price dynamics and factor structure of asset prices.

- Z not directly observable $\Rightarrow$ study pricing, model calibration and hedging for investors who observe only default- and price history.

Advantages

- Prices are weighted averages of full-information values (the theoretical price wrt $\mathbb{F}^X \vee \mathbb{F}^Y$), so that most computations are done in underlying Markov model. $\Rightarrow$ numerics relatively easy.

- Rich credit-spread dynamics with spread risk (spreads fluctuate in response to fluctuations in Z) and default contagion.

  Note that dynamic models are necessary for for model-based hedging and for pricing certain exotic credit derivatives.

- Model has has a natural factor structure with factors given by the conditional probabilities $\pi_t^k = Q(X_t = k \mid \mathcal{F}_t^M)$, $1 \leq k \leq K$.

- Great flexibility for calibration.

A simulated trajectory

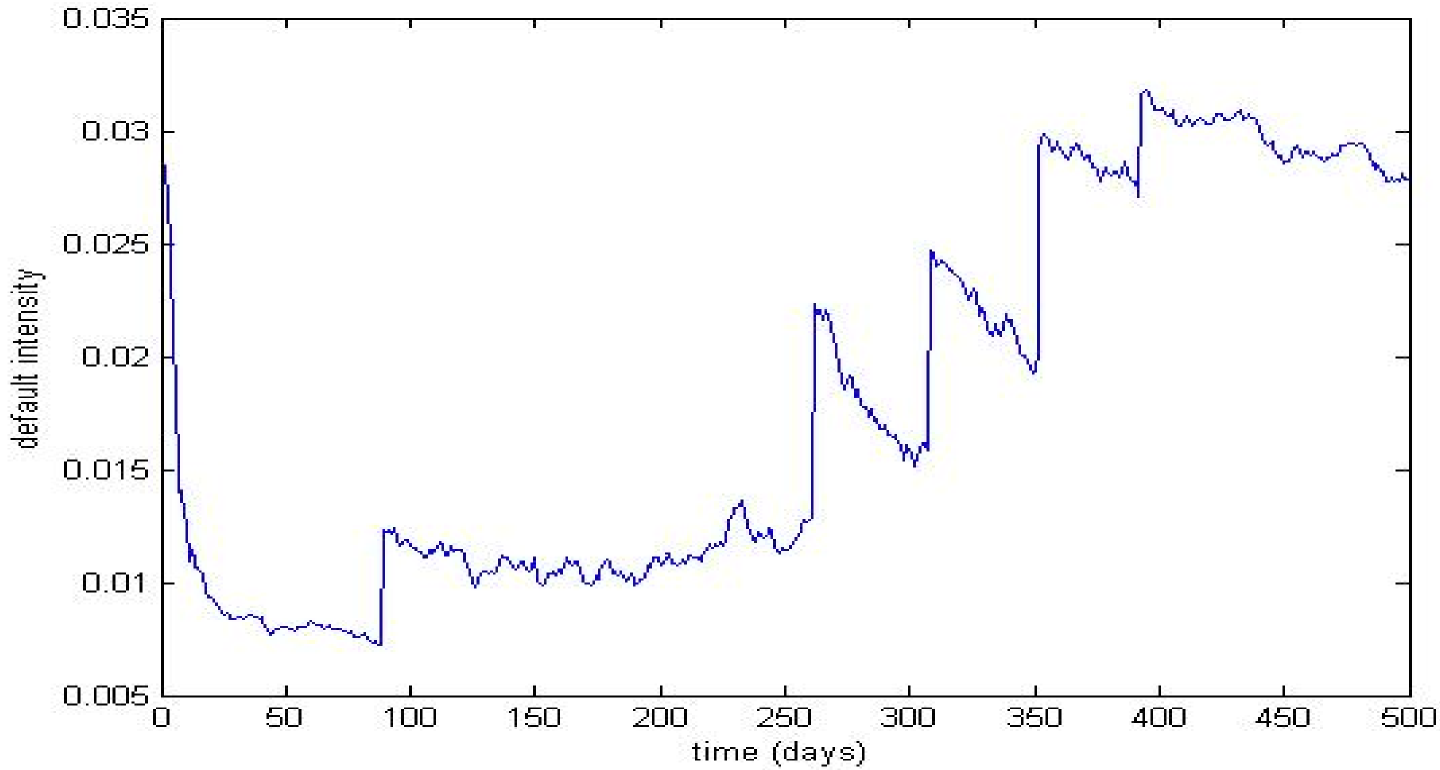

A simulated trajectory of the default intensity ($\approx$ short-term credit spread) generated within our framework.

The underlying Markov model

Consider a finite-state Markov chain X with $S^X := \{1, \ldots, K\}$ and generator Q^X on some $(\Omega, \mathcal{F}, \mathbb{F}, Q)$ (Q a risk neutral measure).

A1 The default times are conditionally independent, doubly stochastic random times with $(Q, \mathbb{F})$-default intensity $\lambda_i(X_t)$.

Implications.

- Recall that $Y_{t,j} := 1_{\{\tau_j \leq t\}}$. The processes $Y_{t,j} - \int_0^{t \wedge \tau_j} \lambda_j(X_{s-}) ds$, $1 \leq j \leq m$, are $\mathbb{F}$-martingales.
- $\tau_1, \ldots, \tau_m$ are conditionally independent given $\mathcal{F}_\infty^X$; in particular no joint defaults.
- The pair process (X, Y) is Markov wrt $\mathbb{F}$

Examples

Homogeneous model (default intensities of all firms identical). Intensities are modelled by some increasing function $\lambda : \{1, \ldots, K\} \to (0, \infty)$; elements of S^X thus represent different states of the economy (1 is the best state and K the worst.)

Global- and industry factors. Assume that we have $\bar{r}$ different industry groups. Let $S^X = \{1, \ldots, \kappa\} \times \{0, 1\}^r$; write X^0,. . . , $X^{\bar{r}}$ for the components of X, modelled as independent Markov chains. X^r is the state of industry r which is good $(X^r = 0)$ or bad $(X^r = 1)$; X^0 represents the global factor. Default intensity of firm i from industry group r takes the form $\lambda_i(x) = g_i(x^0) f_i(x^r)$ for increasing f_i and g_i.

Generator Q^X. Various possibilities; a simple but useful model takes X to be constant. We call this the dynamic implied copula model (dynamic extension of [Hull and White, 2006]) or Rosen and Saunders (2007).

Market information

Recall that information contained in prices of traded securities is modelled via observations of some process Z. Formally,

A2 $\mathbb{F}^M = \mathbb{F}^Y \vee \mathbb{F}^Z$, where the l-dim. process Z solves the SDE

$$dZ_t = \mathbf{a}(X_t)dt + dB_t.$$

Here, B is an l-dim standard $\mathbb{F}$-Brownian motion independent of X and Y, and $\mathbf{a}(\cdot)$ is a function from S^X to $\mathbb{R}^l$.

Notation. Given a generic RCLL process U, we denote by $\widehat{U}$ the optional projection of U w.r.t. the market filtration $\mathbb{F}^M$; recall that $\widehat{U}$ is a right continuous process with $\widehat{U}_t = \mathbb{E}(U_t | \mathcal{F}_t^M)$ for all $t \geq 0$.

Traded securities.

We consider N liquidly traded credit derivatives with maturity T and $\mathbb{F}^Y$-adapted cumulative dividend processes $D_1, \dots, D_N$.

Examples.

- Defaultable zero-bond on firm i: $D_{t,i} = 0$, $t < T$; $D_{T,i} = 1 - Y_{T,i}$.
- CDS with fixed spread x: $D_t = \int_0^t dY_{s,i} - x \sum_{t_n \le t} \Delta t_n \left(1 - Y_{t_n,i}\right)$

We use martingale modelling to construct the model and let $r = 0$ for simplicity. Formally:

A3. Prices of traded credit derivatives are given by $\widehat{p}_{t,i} := E^Q\big(D_{T,i} - D_{t,i} \mid \mathcal{F}_t^M\big)$.

Market-pricing and nonlinear filtering.

For simplicity we assume deterministic recovery rates.

Define the full-information value of the traded securities by $E^Q\big(D_{T,i} - D_{t,i} \mid \mathcal{F}_t\big)$. Recall that (X, Y) is Markov w.r.t. $\mathbb{F} \Rightarrow$ for typical credit derivatives full information value is given by some function $p_i(t, X_t, Y_t)$.

We get from iterated conditional expectations

$$\widehat{p}_{t,i} = \mathbb{E}\big(\mathbb{E}(D_{T,i} - D_{t,i} | \mathcal{F}_t) \mid \mathcal{F}_t^M\big) = \mathbb{E}\big(p_i(t, X_t, Y_t) | \mathcal{F}_t^M\big). \qquad (1)$$

Evaluation of (1) is a typical nonlinear filtering problem: we need to determine the conditional probabilities $\pi_t^k = Q(X_t = k \mid \mathcal{F}_t^M)$ or, more generally, the conditional distribution of X_t given $\mathcal{F}_t^M$

Example: a CDS contract

Consider a CDS with fixed spread x^{CDS} and LGD δ. Full information value in t is given by $(1 - Y_{t,i})\big(V^{\mathsf{def}}(t, X_t) - x^{\mathsf{CDS}} V^{\mathsf{prem}}(t, X_t)\big)$, where

$$V^{\mathsf{def}}(t,k) = E\Big(\int_t^T \lambda(X_s)\delta e^{-\int_t^s \lambda(X_u)du}\, ds \mid X_t = k \Big)\,,$$

$$V^{\mathsf{prem}}(t,k) = \sum_{t_k \geq t} E\Big(\exp(-\int_t^{t_k} \lambda(X_s)\, ds) \mid X_t = k \Big)\,.$$

On $\{\tau > t\}$ the market value of the contract is thus given by

$$\sum_{k=1}^{K} \pi_t^k \, V^{\mathsf{def}}(t,k) - x^{\mathsf{CDS}} \sum_{k=1}^{K} \pi_t^k \, V^{\mathsf{prem}}(t,k)\,.$$

Computation of full-information values.

Many possibilities:

- Bond prices or legs of a CDS can be computed via Feynman-Kac
- For portfolio products such as CDOs we can use conditional independence and compute Laplace transform of portfolio loss, (as in [Graziano and Rogers, 2006]) or use Poisson- and normal approximations, combined with Monte Carlo.
- Often compact formulas can be given involving the matrix exponential of Q_X or of the generator matrix of (X, M) (M the number of defaults); joint work with A. Herbertsson

3. Dynamics of Security Prices

The following two processes will drive the model in the market filtration

$$M_{t,j} := Y_{t,j} - \int_0^{t\wedge\tau_j} \widehat{\lambda_j(X_{s-})} ds\,, \quad j = 1, \cdots, m$$
$$\mu_{t,i} := Z_{t,i} - \int_0^{t} \widehat{a_i(X_s)}\, ds\,, \qquad i = 1, \cdots, l.$$

Properties.

- M_j is an $\mathbb{F}^M$-martingale and μ is $\mathbb{F}^M$-Brownian motion.
- Every $\mathbb{F}^M$-martingale can be represented as stochastic integral wrt M and μ.

Filtering

Define the conditional probability vector $\boldsymbol{\pi}_t = (\pi_t^1, \ldots, \pi_t^K)^\top$ with $\pi_t^k := Q(X_t = k | \mathcal{F}_t^M)$. $\boldsymbol{\pi}_t$ is the natural state variable; in particular, prices of traded assets are linear functions of $\boldsymbol{\pi}_t$.

Kushner-Stratonovich equation. (K-dim SDE-system for $\boldsymbol{\pi}$) Let $q(\iota, k)$, $1 \le \iota, k \le K$ denote generator matrix of X. Then

$$d\pi_t^k = \sum_{\iota=1}^{K} q(\iota, k)\pi_t^\iota dt + (\boldsymbol{\gamma}^k(\boldsymbol{\pi}_{t-}))^\top dM_t + (\boldsymbol{\alpha}^k(\boldsymbol{\pi}_t))^\top d\mu_t\,, \text{ with} \tag{2}$$

$$\gamma_j^k(\boldsymbol{\pi}) = \pi_k \Big(\frac{\lambda_j(k)}{\sum_{n=1}^K \lambda_j(n)\pi_n} - 1 \Big), \quad 1 \le j \le m, \tag{3}$$

$$\boldsymbol{\alpha}^k(\boldsymbol{\pi}) = \pi_k \Big(\mathbf{a}(k) - \sum_{n=1}^{K} \pi_n \mathbf{a}(n) \Big)\,. \tag{4}$$

Default contagion

At τ_j the default intensity ($\approx$ short-term credit spread) of surviving firm i jumps by

$$\Delta\widehat{\lambda}_i(\tau_j) = \sum_{k=1}^{K} \lambda_i(k) \cdot \pi^k_{\tau_j -} \left(\frac{\lambda_j(k)}{\sum_{l=1}^{K} \lambda_j(l) \pi^l_{\tau_j -}} - 1 \right) \tag{5}$$

$$= \frac{\text{cov}^{\boldsymbol{\pi}_{\tau_j -}}(\lambda_i, \lambda_j)}{\mathbb{E}^{\boldsymbol{\pi}_{\tau_j -}}(\lambda_j)}. \tag{6}$$

Note that strength of contagion is greatest

- for firms with similar characteristics (high correlation of λ_i and λ_j)
- for a-priori distribution $\boldsymbol{\pi}_{\tau_j -}$ with a large variance (large incertitude about true state).

Security-price dynamics

Theorem 1. Under **A1** - **A3** the discounted cum-dividend price process $\widehat{g}_t = \widehat{p}_t + D_t$ of the traded assets has the martingale representation

$$\begin{aligned}
\widehat{g}_{t,i} &= \widehat{g}_{0,i} + \int_0^t \boldsymbol{\gamma}_s^{\widehat{g}_i,\top} dM_s + \int_0^t \boldsymbol{\alpha}_s^{\widehat{g}_i,\top} d\mu_s, \quad \text{with} \\
\boldsymbol{\alpha}_t^{\widehat{g}_i} &= \widehat{p_{t,i} \cdot \mathbf{a}}_t - \widehat{p}_{t,i}\, \widehat{\mathbf{a}}_t, \\
\gamma_{t,j}^{\widehat{g}_i} &= \frac{1}{(\widehat{\lambda}_j)_{t-}} \Big((\widehat{p_i \lambda_j})_{t-} - \widehat{p}_{t-} (\widehat{\lambda}_j)_{t-} + (\widehat{R^{g_i,j} \lambda_j})_{t-} \Big) \quad \text{and} \\
R_t^{g_i,j} &= p_i(t, X_t, Y_t^j) - p(t, X_t, Y_t) + \Delta D_{\tau_j,i}.
\end{aligned}$$

Predictable quadratic variations of the asset prices $\mathbb{F}^M$ satisfy $d\langle \widehat{g}_i, \widehat{g}_j \rangle_t^{M} = v_t^{ij}\, dt$ with

$$v_t^{ij} = \sum_{n=1}^{m} \gamma_{t,n}^{\widehat{g}_i}\, \gamma_{t,n}^{\widehat{g}_j}\, \widehat{\lambda}_{t-,n} + \sum_{n=1}^{l} \alpha_{t-,n}^{\widehat{g}_i} \alpha_{t-,n}^{\widehat{g}_j}. \tag{7}$$

4. Pricing and Calibration

Define price of a nontraded claim H as $H_t := \mathbb{E}^Q(H|\mathcal{F}_t^M)$. We distinguish two types of claims.

Options on the loss state. Here H is given by a function of the default state at maturity (eg. basket swaps or bespoke CDOs.) Let $h(t, X_t, Y_t) = E(H \mid \mathcal{F}_t)$. We get from iterated conditional expectations

$$H_t = \sum_{k=1}^{K} \pi_t^k \, h(t, k, Y_t),$$

i.e. the price depends only on $\boldsymbol{\pi}_t$ and on hypothetical value $h(\cdot)$.

Options on traded assets. Here H is of the form $\tilde{h}(Y_{\tilde{T}}, \widehat{p}_{\tilde{T},1}, \ldots, \widehat{p}_{\tilde{T},N})$ at $\tilde{T} < T$. (eg. CDS index options). Since $(Y, \boldsymbol{\pi})$ is $\mathbb{F}^M$-Markov,

$$H_t = \mathbb{E}\left(h(Y_{\tilde{T}}, \widehat{p}_{1,\tilde{T}}, \ldots, \widehat{p}_{N,\tilde{T}})|\mathcal{F}_t^M\right) = h(t, Y_t, \boldsymbol{\pi}_t),$$

but now price depends on dynamics of $\boldsymbol{\pi}$ as well.

Calibration for secondary market investors

Two separate tasks

- In order to use the pricing formulas investors need to determine current value of unobservable factor $\boldsymbol{\pi}_t$ by "matching" market and model prices. Two approaches:
  - ⋆ Standard (pragmatic) calibration using linear or convex programming
  - ⋆ Calibration via filtering [Frey and Runggaldier, 2008]
- Determine the drift $a(\cdot)$ of Z (and generator Q^X). $a(\cdot)$ largely governs dynamics of $\boldsymbol{\pi}_t$ and hence of asset prices. Largely an econometric problem; possible approach: EM-algorithm

Application to itraxx-tranches

We concentrate on models where X is constant $\Rightarrow$ Pricing and calibration of CDOs similar as in Hull-White (2006) or Rosen Saunders (2007) (but hedging is different!)

Example 1. Calibrate homogeneous version to itraxx data from various years (pre-crisis and during credit crisis). We obtained very good fit for all data sets. Note the increase in the implied probability of the extreme scenario $\lambda = 70\%$ (5 year PD $\approx$ 96 %).

λ (in %)	0.01	0.3	0.6	1.2	2.5	4.0	8.0	20	70
$\boldsymbol{\pi}^*$, data from 2004	12.6	22.9	42.0	17.6	2.5	1.45	0.54	0.13	0.03
$\boldsymbol{\pi}^*$, data from 2006	22.2	29.9	39.0	7.6	1.2	0.16	0.03	0.03	0.05
$\boldsymbol{\pi}^*$, data from 2008	1.1	7.9	57.6	10.8	11.7	4.9	1.26	1.79	2.60
$\boldsymbol{\pi}^*$, data from 2009	0.0	13.6	6.35	42.2	22.3	12.5	0.0	0.00	3.06

Components of $\boldsymbol{\pi}^*$ are given in percentage points.

3d-Representation

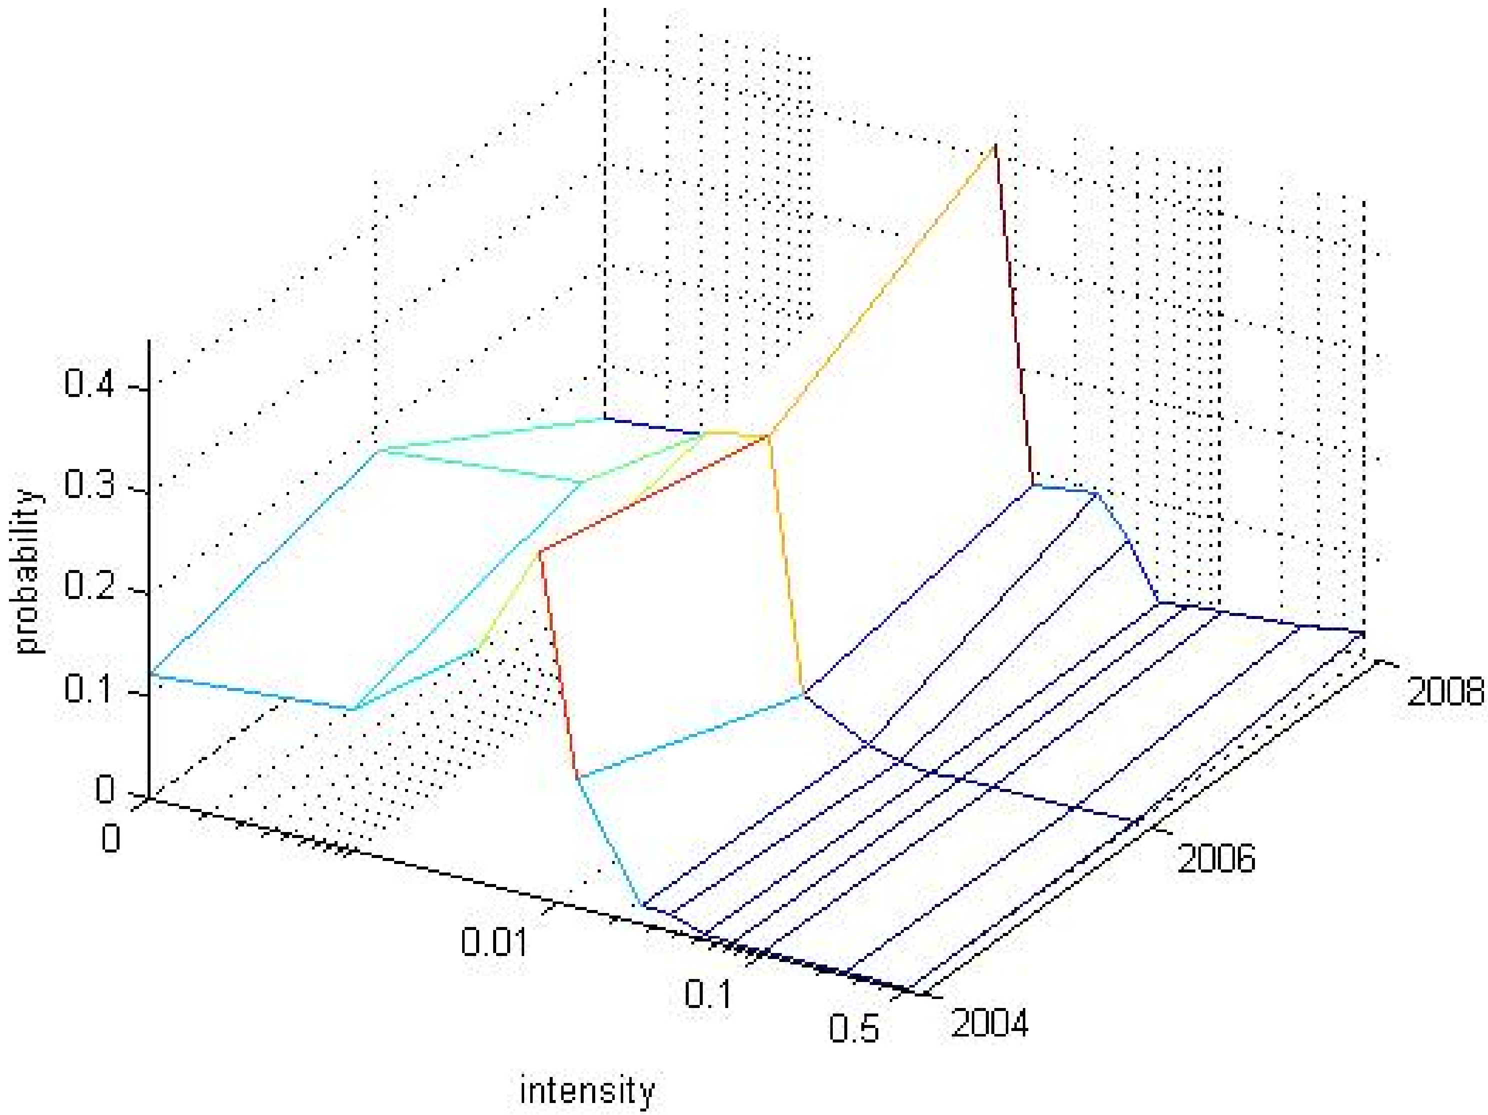

Simulated trajectory of π_t

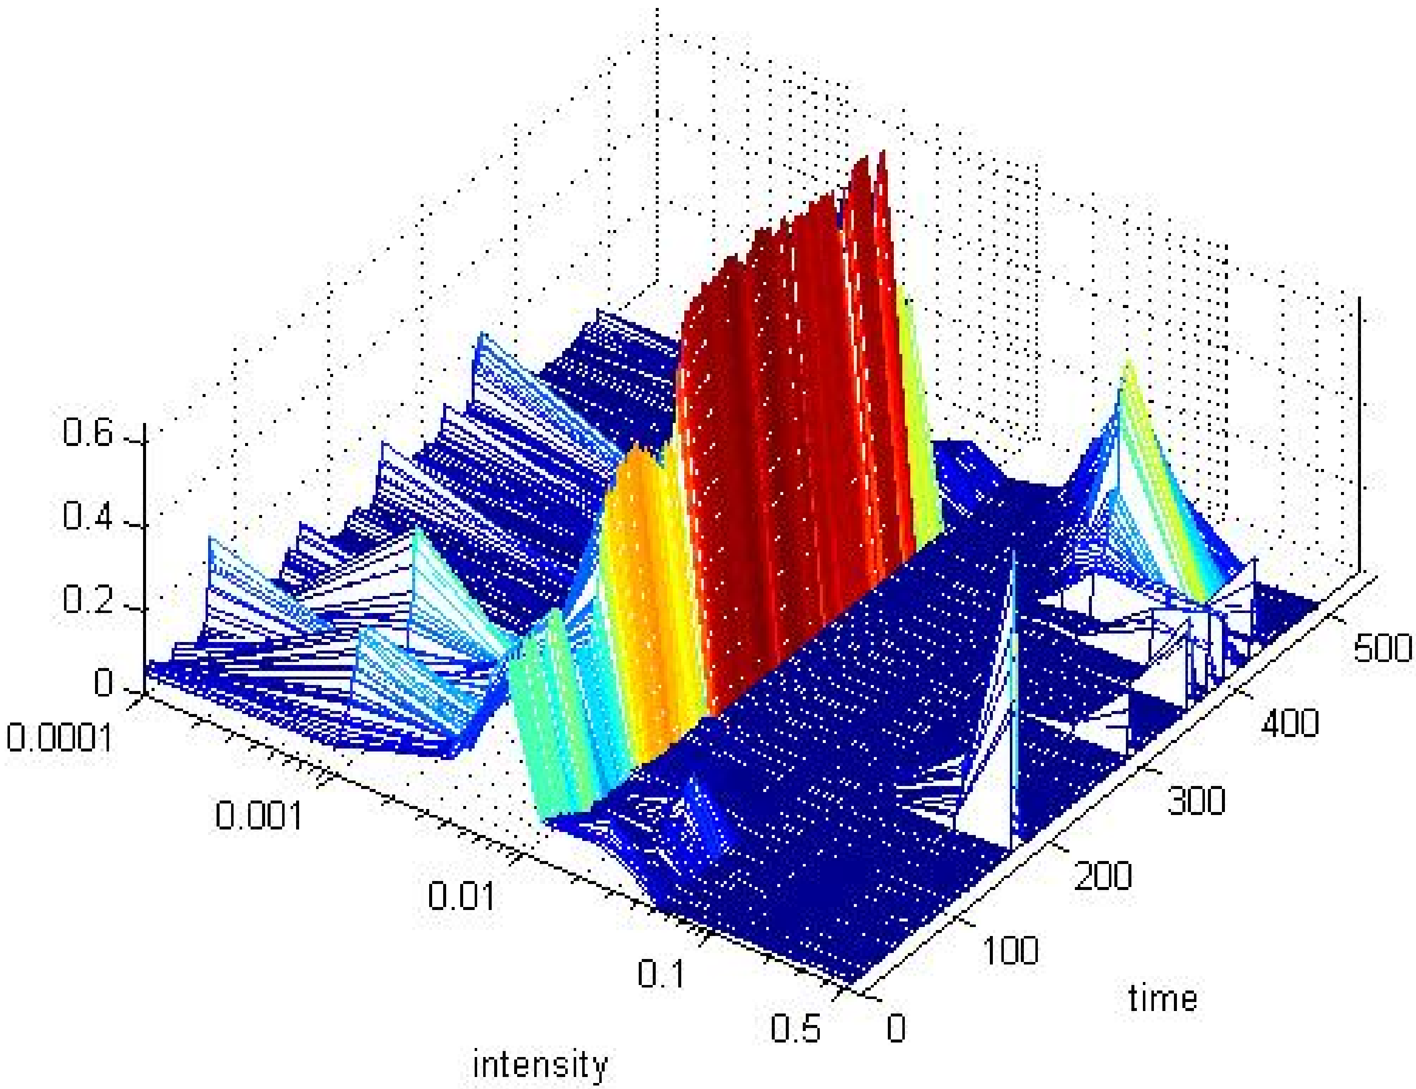

Credit Index Options

Payoff. A payer credit index option gives the right to enter into a CDS index as protection buyer at $\tilde{T}$ for a predetermined spread K (the strike). Moreover, there is front-end protection: upon exercise the holder receives the losses in the portfolio between inception and maturity $\tilde{T}$.

- Payoff depends on CDS-index spread at $\tilde{T}$ and hence on price of traded security
- Market pricing approach: assume (adjusted) spreads are lognormal after clever change of numeraire and apply Black formula. (Pedersen(2003), Brigo-Morini(2007)). The portfolio loss is not explicitly modelled.

Credit Index Options

Our approach provides a model for the joint evolution of portfolio losses and (index) spreads. Prices are computed via Monte carlo simulation.

Numerical results. We used 4 states, $\lambda \in \{0.01, 0.02, 0.04, 0.1\}$. Model was calibrated to CDS index spread S^*. Prices are quoted as implied volatility computed via the Pedersen(2003) approach.

moneyness K/S^*	0.8	1.0	1.4
implied vol, $\boldsymbol{\pi} = (0.25, 0.53, 0.07, 0.15)$	1.13	1.31	1.49
implied vol, $\boldsymbol{\pi} = (0.20, 0.24, 0.56, 0.001)$	0.55	0.56	0.46

Wide range of levels and smile patterns can be generated by varying $\boldsymbol{\pi}$ and the values of λ.

5. Dynamic Hedging

Consider some claim H with price process $\hat{h}_t$. Look for risk-minimizing strategies as in [Föllmer and Sondermann, 1986].

- Allows to address potential incompleteness of the market.
- Tractable criterion (related to Kunita-Watanabe decomposition).

We seek a representation $\hat{h}_t - \hat{h}_0 = \sum_{j=1}^{n} \int_0^t \theta^H_{s,j}\, d\hat{p}_{s,j} + L_t$ such that the remaining risk (conditional error variance) $E\big((L_T - L_t)^2 \mid \mathcal{F}^M_t\big)$ is minimized simultaneously for all t.

Proposition 2. We have $\boldsymbol{\theta}^H_t = \mathbf{v}_t^{-1} \frac{d}{dt} \langle \widehat{h}, \widehat{\mathbf{p}} \rangle^M_t$, $\mathbf{v}_t$ the instantaneous predictable quadratic variation of the traded assets.

All ingredients are readily computed.

Example: hedging CDO-tranches with the index

Tranche	[0-3]	[3-6]	[6-9]	[9-12]	[12-22]
low spread volatility					
π calibrated to 2004 data	0.3249	0.1097	0.0749	0.0614	0.1462
π calibrated to 2006 data	0.2404	0.0684	0.0427	0.0340	0.0973
π calibrated to 2008 data	0.0674	0.0376	0.0359	0.0342	0.1073
high spread volatility					
π calibrated to 2004 data	0.6592	0.1471	0.0842	0.0604	0.1144
π calibrated to 2006 data	0.6799	0.0958	0.0418	0.0243	0.0555
π calibrated to 2008 data	0.0948	0.0516	0.0436	0.0370	0.1055

Risk-minimizing hedge ratio θ for hedging a CDO tranche with the underlying CDS index

6. Calibration via filtering

Here we assume that $\mathbb{F}^I = \mathbb{F}^Y \vee \mathbb{F}^U$ where U solves the SDE

$$dU_t = \widehat{p}_t dt + dW_t = \mathbf{p}(t, X_t, Y_t)\boldsymbol{\pi}_t dt + dW_t$$

for a Brownian motion W independent of X, Y, Z. U can be viewed as cumulative noisy price information of traded assets $\widehat{p}_1, \ldots, \widehat{p}_N$ (noise reflects observation- and model errors.)

Recall that $\boldsymbol{\pi}$ solves the KS-equation (2). $\Rightarrow$ finding conditional distribution of $\boldsymbol{\pi}_t$ given $\mathcal{F}_t^I$ is a nonlinear filtering problem with signal process $\boldsymbol{\pi}$ and observation processes U and Y.

Analysis of filtering problem. Challenges: observations of mixed type; high dimension of state process; joint jumps of $\boldsymbol{\pi}$ and Y. Numerical treatment via particle filtering, see [Frey and Runggaldier, 2008].

A. Outlook

- Practical issues: further numerical work on hedging and on pricing of exotic credit derivatives; extension to inhomogeneous portfolios and to models with $Q^X \neq 0$; performance of hedging strategies and model risk.

- Consider models where X has a continuous state space and study other finite-dimensional approximations to the filtering problem

- Filtering methods/EM algorithm for calibration and estimation of model parameters Q^X and in particular $a(\cdot)$.

- Extension of previous methodology to other markets, in particular markets for corporate securities or default-free term-structure models.

References

[Coculescu et al., 2006] Coculescu, D., Geman, H., , and Jeanblanc, M. (2006). Valuation of default sensitive claims under imperfect information. working paper, Université d' Evry.

[Collin-Dufresne et al., 2003] Collin-Dufresne, P., Goldstein, R., and Helwege, J. (2003). Is credit event risk priced? modeling contagion via the updating of beliefs. Preprint, Carnegie Mellon University.

[Duffie et al., 2006] Duffie, D., Eckner, A., Horel, G., and Saita, L. (2006). Frailty correlated defaullt. preprint, Stanford University.

[Duffie and Lando, 2001] Duffie, D. and Lando, D. (2001). Term structure of credit risk with incomplete accounting observations. *Econometrica*, 69:633–664.

[Föllmer and Sondermann, 1986] Föllmer, H. and Sondermann, D. (1986). Hedging of non-redundant contingent-claims. In Hildenbrand, W. and Mas-Colell, A., editors, *Contributions to Mathematical Economics*, pages 147–160. North Holland.

[Frey and Runggaldier, 2008] Frey, R. and Runggaldier, W. (2008). Pricing credit derivatives under incomplete information: a nonlinear filtering approach. preprint, Universität Leipzig, to appear in Finance and Stochastics.

[Frey and Schmidt, 2006] Frey, R. and Schmidt, T. (2006). Pricing corporate securities under noisy asset information. preprint, Universität Leipzig,forthcoming in Mathematical Finance.

[Giesecke and Goldberg, 2004] Giesecke, K. and Goldberg, L. (2004). Sequential defaults and incomplete information. *Journal of Risk*, 7:1–26.

[Graziano and Rogers, 2006] Graziano, G. and Rogers, C. (2006). A dynamic approach to the modelling of correlation credit derivatives using Markov chains. working paper, Statistical Laboratory, University of Cambridge.

[Hull and White, 2006] Hull, J. and White, A. (2006). The implied copula model. *The Journal of Derivatives*.

[Jarrow and Protter, 2004] Jarrow, R. and Protter, P. (2004). Structural versus reduced-form models: a new information based perspective. *Journal of Investment management*, 2:1–10.

[Landen, 2001] Landen, C. (2001). Bond pricing in a hidden markov model of the short rate. *Finance and Stochastics*, 4:371–389.

[Schönbucher, 2004] Schönbucher, P. (2004). Information-driven default contagion. Preprint, Department of Mathematics, ETH Zürich.