
Pricing and Hedging of Credit Derivatives via Nonlinear Filtering

Rüdiger Frey · Thorsten Schmidt ·
Abdelali Gabih

Version from October 16, 2007

Abstract In this paper a new, information-based approach for modelling the dynamic evolution of a portfolio of credit risky securities is proposed. In this context market prices of liquidly traded derivatives are given by the solution of a nonlinear filtering problem. This problem is solved via the innovations approach to nonlinear filtering. Moreover, we derive the ensuing asset price dynamics and compute risk-minimizing hedging strategies. In the last part of the paper we discuss a numerical approach - based on particle filtering - to some of the arising filtering problems.

Keywords Credit risk, incomplete information, nonlinear filtering, risk minimization

1 Introduction

Given the tremendous growth of the market for credit derivatives in recent years, the development of models for the pricing and the hedging of these products has become a major concern in credit-risk related research. In this paper we propose a new, information-based approach to modelling the dynamics of credit derivatives prices and discuss pricing and hedging issues in this framework.

We start with a brief description of our modelling strategy. We consider defaultable securities issued by m firms; the random time τ_i denotes the default time of firm i , $Y_{t,i} = \mathbb{1}_{\{\tau_i \leq t\}}$ is the corresponding default indicator, and $Y_t = (Y_{t,1}, \dots, Y_{t,m})$ gives the current default state of the portfolio. The default intensities (the intensities of the multivariate point process Y) are assumed to depend on some factor process X . In our analysis we distinguish three different layers of information: full information; information of informed market

Department of Mathematics, University of Leipzig, D-04081 Leipzig, Germany.
Emails: {frey,tschmidt,gabih}@math.uni-leipzig.de

participants (market-information); information of secondary-market investors (investor-information).

Full information. The full-information setup is a theoretical device used for the construction of the model. We work on some filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, Q)$ with Q being the risk-neutral measure and \mathbb{F} the full-information filtration. We assume that the τ_i are conditionally independent doubly stochastic random times with (Q, \mathbb{F}) -default intensity $\lambda_{t,i} = \lambda_i(X_t)$ and that X follows a finite-state Markov chain. This setup is akin to the model of di Graziano & Rogers (2006). We immediately work with discounted quantities and therefore take the risk-free short rate equal to zero. We define the *full-information value* of a \mathcal{F}_T^Y -measurable¹ claim P (such as a typical credit derivative) by $\mathbb{E}^Q(P | \mathcal{F}_t)$. By the Markov property of (X, Y) the full-information value is given by $p_t(X_t)$ for some \mathcal{F}_t^Y -measurable function p_t , see for details Equation (2.1).

Market-information. The prices of traded credit derivatives are determined by *informed market-participants*. These investors have access to so-called market information, given by the filtration $\mathbb{F}^M := \mathbb{F}^Y \vee \mathbb{F}^Z$. The stochastic process Z represents noisy observations of X and can be viewed as an abstract form of ‘insider information’. Mathematically, Z is given by $Z_t = \int_0^t \mathbf{a}(X_s) ds + dB_t$, B a standard \mathbb{F} -Brownian motion independent of X and Y . The market price of a traded security with payoff P is defined as

$$\hat{p}_t := \mathbb{E}^Q(P | \mathcal{F}_t^M) = \mathbb{E}^Q(p_t(X_t) | \mathcal{F}_t^M). \quad (1.1)$$

Since Y_t is known, in order to compute the market price \hat{p}_t , one needs to determine the conditional distribution of X_t given \mathcal{F}_t^M , given by the probability vector $\boldsymbol{\pi}_t = (\pi_t^1, \dots, \pi_t^K)$ with

$$\pi_t^k = Q(X_t = k | \mathcal{F}_t^M), \quad 1 \leq k \leq K$$

(for expository purposes we identify the state space of X with $\{1, \dots, K\}$). This is a typical nonlinear filtering problem which is solved in Section 3 using martingale representation results and the innovations approach to nonlinear filtering. By the same token we derive the dynamics of the market price of the traded credit derivatives.

Investor-information. Since the process Z is not directly related to observable economic quantities, Section 4 is devoted to the analysis of the pricing and the hedging of credit derivatives from the viewpoint of secondary-market investors with information set $\mathbb{F}^I \subset \mathbb{F}^M$. It is assumed that \mathbb{F}^I contains the default history \mathbb{F}^Y and noisy price observations of traded credit derivatives. As will be shown in Section 4, in this setup the computation of prices and risk-minimizing hedging strategies lead to a second filtering problem: one has to determine the conditional distribution of the probability vector $\boldsymbol{\pi}_t$ given investor information \mathcal{F}_t^I . General filtering equations for this problem and a numerical solution based on particle filtering are discussed in the last part of the paper.

¹ Throughout, we always denote the natural filtration of a stochastic process, say X , by $\mathbb{F}^X = (\mathcal{F}_t^X)_{t \geq 0}$.

The proposed modelling approach has a number of advantages. First, prices are weighted averages of full-information values (by (1.1)), so that hence actual computations are done mostly in the context of the full-information model. Since the latter has a very simple structure, computations become relatively straightforward; various approaches for computing full-information values are for instance discussed in di Graziano & Rogers (2006). Second, the fact that prices of traded securities are given by the projection of their full-information value on the market filtration $\mathbb{F}^{\mathbb{M}}$ leads to rich credit-spread dynamics: the proposed approach accommodates *spread risk* (as credit spreads fluctuate in response to fluctuations in Z) and *default contagion* (as defaults of firms in the portfolio lead to an update of the conditional distribution of X given $\mathcal{F}_t^{\mathbb{M}}$ and hence to a jump in the $(Q, \mathbb{F}^{\mathbb{M}})$ -default intensities. This feature is important in the derivation of robust dynamic hedging strategies. Third, the model has a natural factor structure with factors given by the conditional probabilities π_t^k , $1 \leq k \leq K$. Finally, the approach gives great flexibility in terms of calibration methodologies as will be discussed in detail in Section 4.

This paper is closely related to the companion paper Frey & Runggaldier (2006). The latter concentrates on the mathematical analysis of filtering problems in reduced-form credit risk models, whereas here we are interested in the dynamics of credit-derivative prices under incomplete information and in hedging issues. The idea of interpreting observed prices of derivatives as noisy observations of some factor process has been pursued previously by Gombani, Jaschke & Runggaldier (2005) in the context of default-free term-structure models. Reduced-form credit risk models under incomplete information are also discussed in Collin-Dufresne, Goldstein & Helwege (2003) and Schönbucher (2004), albeit on a simpler mathematical level. Moreover, non-linear filtering problems arise in a natural way in structural credit risk models with incomplete information about the current value of assets or liabilities such as Kusuoka (1999), Duffie & Lando (2001), Jarrow & Protter (2004), Coculescu, Geman & Jeanblanc (2006) or Frey & Schmidt (2007).

2 The Model

Our model is constructed on some filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, Q)$, with $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ satisfying the usual conditions; all processes considered are by assumption \mathbb{F} -adapted. Q is the risk-neutral martingale measure used for pricing. Throughout we consider a fixed credit portfolio consisting of a set of m firms. The current default state is $Y_t = (Y_{t,1}, \dots, Y_{t,m})$ with $Y_{t,i} = \mathbb{1}_{\{\tau_i \leq t\}}$, τ_i the default time of firm i ; note that $Y_t \in \{0, 1\}^m$. To avoid trivialities, we assume $Y_0 = 0$. We introduce the *ordered default times* $0 = T_0 < T_1 < \dots < T_m < T_{m+1} := \infty$. Denote the identity of the firm defaulting at T_n by ξ_n , $1 \leq n \leq m$. The sequence $(T_n, \xi_n)_{1 \leq n \leq m}$ gives a representation of Y as marked point process. The σ -field $\mathcal{F}_t^Y = \sigma(Y_{s,i}, s \leq t, 1 \leq i \leq m) = \sigma(\{(T_n, \xi_n) : T_n \leq t\})$ is the default history of the portfolio at time t . Note that every \mathcal{F}_t^Y -measurable

function $p_t : \Omega \times S^X \rightarrow \mathbb{R}$ admits the following representation

$$p_t(\omega; x) = \sum_{n=0}^m \mathbb{1}_{\{T_n(\omega) \leq t < T_{n+1}(\omega)\}} p^n(t, x; (T_i(\omega), \xi_i(\omega))_{1 \leq i \leq n}), \quad (2.1)$$

with functions $p^n : [0, \infty) \times S^X \times ((0, \infty) \times \{1, \dots, m\})^n \rightarrow \mathbb{R}$.

2.1 The Full-Information Model

The default intensities of the firms under consideration may depend on a process X , the so-called factor or state process. X is modelled as a finite-state Markov chain; in the sequel the state space S^X is usually identified with the set $\{1, \dots, K\}$. We make the following assumption.

A1 The default times are conditionally independent, doubly-stochastic random times with default intensity $\lambda_{t,i} := \lambda_i(X_t)$, i.e. there are functions $\lambda_i : S^X \rightarrow (0, \infty)$, such that the processes

$$Y_{t,i} - \int_0^{t \wedge \tau_i} \lambda_i(X_{s-}) ds \quad (2.2)$$

are \mathbb{F} -martingales, $1 \leq i \leq m$. Moreover, the random variables τ_1, \dots, τ_m are conditionally independent given $\mathcal{F}_\infty^X = \sigma(X_s : s \geq 0)$.

Note that conditional independence of default times excludes joint defaults, i.e. the processes Y_1, \dots, Y_m have a.s. no common jumps.

Example 2.1 In order to illustrate the modelling possibilities under **A1**, we give two examples which will be taken up later. First, we consider a homogeneous model where the default intensities of all firms are identical. Such a model can alternatively be viewed as *total loss model* in the sense of Arnsdorf & Halperin (2007), i.e. as a model for the portfolio loss given by $\sum_{i=1}^m Y_{t,i}$ (neglecting recovery payments). In that case it is natural to model the default intensities by some increasing function $\lambda : \{1, \dots, K\} \rightarrow (0, \infty)$ of the states of the economy. The elements of S^X thus represent different states of the economy, 1 being the best state (lowest default intensity) and K the worst state.

As a second example we consider a portfolio with one-factor structure where the default intensity of firm i depends on a systematic factor as well as on a firm-specific factor: let $S^X = \{0, 1\}^m \times \{1, \dots, \kappa\}$ and write X^i , $i = 1, \dots, m+1$ for the i th component of X . The components X^1, \dots, X^{m+1} are assumed to be independent Markov chains; for $i = 1, \dots, m$, X^i refers to the individual state of company i , which is good ($X^i = 0$) or bad ($X^i = 1$), while X^{m+1} represents the systematic factor. The default intensities take the form $\lambda_i(x) = f_i(x^i) + g_i(x^{m+1})$ where f_i and g_i are increasing functions from $\{0, 1\}$ respectively $\{1, \dots, \kappa\}$ to $(0, \infty)$.

Models with several country- or industry factors can be constructed in an analogous fashion.

2.2 Market Information

Recall that in our setting the prices of traded credit derivatives are determined by informed market participants having access to the filtration $\mathbb{F}^{\mathbb{M}}$ (market information).

Market filtration. We assume that the market filtration $\mathbb{F}^{\mathbb{M}}$ is generated by the default history \mathbb{F}^Y and observations of functions of the state variable X in additive Gaussian noise:

A2 $\mathbb{F}^{\mathbb{M}} = \mathbb{F}^Y \vee \mathbb{F}^Z$, where the l -dimensional process Z is given by

$$dZ_t = \int_0^t \mathbf{a}(X_s) ds + dB_t. \quad (2.3)$$

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

Example 2.1 ctd. In the homogeneous-portfolio situation it is natural to assume that $a(\cdot) = c\lambda(\cdot)$. The constant $c \geq 0$ models the information-content of Z : for $c = 0$, Z carries no information, for c large the state can be observed with relatively high precision. Analogously, in the one-factor case one could take $l = m + 1$, $a_i(x) = c_i f_i(x^i)$, $1 \leq i \leq m$, and $a_{m+1}(x) = c_{m+1} g(x^{m+1})$ for some increasing function $g : \{1, \dots, \kappa\} \rightarrow \mathbb{R}$. This models a situation where the market has noisy observations regarding the current state of each company and moreover observes the current state of the economy in Gaussian noise.

Traded securities. We consider a market of N liquidly traded credit derivatives, with - for notational simplicity - common maturity T and square integrable \mathcal{F}_T^Y -measurable payoffs P_1, \dots, P_N . These securities could be corporate bonds, credit default swaps (CDSs) or portfolio products such as synthetic CDO-tranches. Recall that we work immediately with discounted quantities. Hence the prices of these securities are given by²

$$\mathbb{E}(P_i | \mathcal{F}_t^{\mathbb{M}}) =: \hat{p}_{t,i}. \quad (2.4)$$

3 Asset Price Dynamics

The main objective of this section is to derive the dynamics of the traded credit securities \hat{p}_i , $1 \leq i \leq n$. As explained in the introduction, \hat{p}_i is the

² For notational simplicity we abstract from intermediate cash flows such as spread- or default payments; such payments can be handled in our framework without additional difficulties. Moreover, the superscript Q will be omitted from the expectation operator if there are no ambiguities.

solution of a non-linear filtering problem. More precisely, we get by iterated conditional expectations

$$\widehat{p}_{t,i} = \mathbb{E}(P_i | \mathcal{F}_t^{\mathbb{M}}) = \mathbb{E}(\mathbb{E}(P_i | \mathcal{F}_t) | \mathcal{F}_t^{\mathbb{M}}) = \mathbb{E}(p_{t,i}(X_t) | \mathcal{F}_t^{\mathbb{M}}),$$

where the \mathcal{F}_t^Y -measurable function $p_{t,i}(\cdot)$ gives the full-information value of derivative i . Hence, in order to compute the market price $\widehat{p}_{t,i}$ it is sufficient to determine the conditional distribution of X_t given $\mathcal{F}_t^{\mathbb{M}}$.

3.1 Asset Price Dynamics under the Market Filtration

Following the innovations approach to nonlinear filtering we will derive a representation of the martingales \widehat{p}_i as a stochastic integral w.r.t. the *innovations processes*. The latter are defined as follows:

$$\begin{aligned} M_{t,j} &:= Y_{t,j} - \int_0^{t \wedge \tau_j} \widehat{\lambda}_j(X_{s-}) ds && \text{for } j = 1, \dots, m \\ \mu_{t,i} &:= Z_{t,i} - \int_0^t \widehat{a}_i(X_s) ds && \text{for } i = 1, \dots, l. \end{aligned}$$

By Theorem II.14 of Brémaud (1981) the process M_j is an $\mathbb{F}^{\mathbb{M}}$ -martingale. Moreover, μ is an $\mathbb{F}^{\mathbb{M}}$ -Brownian motion; see for instance Lemma 1 in Davis & Marcus (1981).

The following martingale representation result is a key tool in our analysis; its proof is relegated to the appendix.

Lemma 3.1 *For every $\mathbb{F}^{\mathbb{M}}$ -martingale $(U_t)_{0 \leq t \leq T}$ there is an \mathbb{R}^m -valued $\mathbb{F}^{\mathbb{M}}$ -predictable process γ and an \mathbb{R}^l -valued $\mathbb{F}^{\mathbb{M}}$ -adapted process α such that U has the representation*

$$\widehat{U}_t = \widehat{U}_0 + \int_0^t \gamma_s^\top dM_s + \int_0^t \alpha_s^\top d\mu_s, \quad 0 \leq t \leq T. \quad (3.1)$$

In the sequel we use the notation $\widehat{U}_t := \mathbb{E}(U_t | \mathcal{F}_t^{\mathbb{M}})$ for the *optional projection* of a generic process U w.r.t. the market filtration $\mathbb{F}^{\mathbb{M}}$. The following proposition extends the martingale representation result of Fujisaki, Kallianpur & Kunita (1972) to the case where the observations are of mixed type (point processes and diffusions).

Proposition 3.2 *Consider a real-valued \mathbb{F} -semimartingale*

$$J_t = J_0 + \int_0^t A_s ds + M_t^J, \quad t \leq T,$$

with $\mathbb{E}(|J_0|) < \infty$ and $\mathbb{E}(\int_0^T |A_s| ds) < \infty$. Assume that M^J is an \mathbb{F} -martingale with $[M^J, B] = 0$. Then the optional projection \widehat{J}_t has the representation

$$\widehat{J}_t = \widehat{J}_0 + \int_0^t \widehat{A}_s ds + \int_0^t \gamma_s^\top dM_s + \int_0^t \alpha_s^\top d\mu_s, \quad t \leq T, \quad (3.2)$$

where γ and α are given by

$$\alpha_t = \widehat{J}_t \widehat{\mathbf{a}}(X_t) - \widehat{J}_t \mathbf{a}(X_t), \quad (3.3)$$

$$\gamma_{t,j} = (1 - Y_{t-,j}) (\mathbb{E}(J_t | \mathcal{F}_{t-}^{\mathbb{M}} \vee \{\tau_j = t\}) - \mathbb{E}(J_t | \mathcal{F}_{t-}^{\mathbb{M}})), \quad j = 1, \dots, m. \quad (3.4)$$

Remark 3.3 The precise meaning of the expressions in (3.4) is as follows. Note that, in analogy to (2.1),

$$\mathbb{E}(J_t | \mathcal{F}_t^{\mathbb{M}}) = \sum_{n=0}^m \mathbb{1}_{\{T_n(\omega) \leq t < T_{n+1}(\omega)\}} F_n(t, Z^t(\omega); (T_i(\omega), \xi_i(\omega))_{1 \leq i \leq n}),$$

where $Z^t(\omega)$ denotes the stopped process $(Z_{s \wedge t})_{s \geq 0}$. Then we have on the predictable set $\{T_n < t, T_{n+1} \geq t, \tau_j \geq t\}$,

$$\mathbb{E}(J_t | \mathcal{F}_{t-}^{\mathbb{M}}) = F_n(t, Z^t(\omega); (T_i(\omega), \xi_i(\omega))_{1 \leq i \leq n}), \quad \text{and} \quad (3.5)$$

$$\mathbb{E}(J_t | \mathcal{F}_{t-}^{\mathbb{M}} \vee \{\tau_j = t\}) = F_{n+1}(t, Z^t(\omega); (T_i(\omega), \xi_i(\omega))_{1 \leq i \leq n}, (t, j)). \quad (3.6)$$

Proof The proof uses the following two well-known facts.

1. For every \mathbb{F} -martingale N , the optional projection \widehat{N} is an $\mathbb{F}^{\mathbb{M}}$ -martingale.
2. For any progressively measurable process ϕ with $\mathbb{E}(\int_0^T |\phi_s| ds) < \infty$ the process $\widehat{\int_0^t \phi_s ds} - \int_0^t \widehat{\phi}_s ds$, $s \leq T$, is an $\mathbb{F}^{\mathbb{M}}$ martingale.

The first fact is simply a consequence of iterated expectations, while the second follows from the Fubini theorem, see for instance Davis & Marcus (1981).

Fact 1 and 2 immediately yield that $\widehat{J}_t - \widehat{J}_0 - \int_0^t \widehat{A}_s ds$ is an $\mathbb{F}^{\mathbb{M}}$ -martingale. Lemma 3.1 thus gives the existence of the representation (3.2). It remains to identify γ and α . First, in order to identify γ , note that

$$\Delta \widehat{J}_t = \sum_{j=1}^m \mathbb{1}_{\{\Delta Y_{t,j} > 0\}} \left(\mathbb{E}(J_t | \mathcal{F}_{t-}^{\mathbb{M}} \vee \{\tau_j = t\}) - \mathbb{E}(J_t | \mathcal{F}_{t-}^{\mathbb{M}}) \right).$$

On the other hand, since μ is continuous we get from (3.2) that $\Delta \widehat{J}_t = \sum_{i=1}^m \gamma_{t,i} \Delta Y_{t,i}$. Equating these expressions and solving for γ_t gives the representation (3.4).

Second, in order to identify α , we use an argument similar to the derivation of the Fujisaki-Kallianpur-Kunita equation in standard filtering theory. The trick is to use the elementary identity

$$\widehat{JZ}_i = \widehat{J} Z_i, \quad 1 \leq i \leq m.$$

Each side of this equation gives rise to a different semimartingale decomposition of \widehat{JZ}_i ; comparing those one obtains α_i . On the one hand,

$$d(J_t Z_{t,i}) = J_{t-}(a_{t,i} dt + dB_{t,i}) + Z_{t,i}(A_t dt + dM_t^J) + d[Z_i, J]_t.$$

The covariation $[Z_i, J]$ vanishes as a consequence of $[M^J, B] = 0$. Applying Facts 1 and 2 above, we obtain the following form of the semimartingale decomposition of \widehat{JZ}_i

$$d(\widehat{J_t Z_{t,i}}) = ((\widehat{J_t a_{t,i}}) + Z_{t,i} \widehat{A_t}) dt + d\tilde{N}_t, \quad (3.7)$$

where \tilde{N} is a \mathbb{F}^M -martingale. On the other hand, since $dZ_{t,i} = d\mu_{t,i} + \widehat{a}_{t,i} dt$,

$$d(\widehat{J_t Z_{t,i}}) = d(\widehat{J_t} Z_{t,i}) = \widehat{J_t} \widehat{a}_{t,i} dt + Z_{t,i} \widehat{A_t} dt + d[\widehat{J}, Z_i]_t + \tilde{N}_t,$$

where \tilde{N} is some \mathbb{F}^M -martingale. Now, since μ is \mathbb{F}^M -standard Brownian motion, we have from (3.2) the relation $d[\widehat{J}, Z_i]_t = \sum_{j=1}^l \alpha_{t,j} d[\mu_j, \mu_i]_t = \alpha_{t,i} dt$, so that we obtain a second semimartingale decomposition of \widehat{JZ}_i :

$$d(\widehat{J_t Z_{t,i}}) = (\widehat{J_t} \widehat{a}_{t,i} + Z_{t,i} \widehat{A_t} + \alpha_{t,i}) dt + \tilde{N}_t \quad (3.8)$$

Uniqueness of the semimartingale decomposition of \widehat{JZ}_i implies that the finite-variation terms in (3.7) and (3.8) coincide. Hence $\alpha_{t,i} = \widehat{J_t} \widehat{a}_{t,i} - \widehat{J_t} \widehat{a}_{t,i}$, so that we obtain (3.3). \square

The following theorem is the main result of this section; it describes the dynamics of the discounted prices of traded securities and their instantaneous quadratic covariation.

Theorem 3.4 *Under A1 and A2 the (discounted) price processes of the traded securities have the martingale representation*

$$\widehat{p}_{t,i} = \widehat{p}_{0,i} + \int_0^t (\gamma_s^{\widehat{p}_i})^\top dM_s + \int_0^t (\alpha_s^{\widehat{p}_i})^\top d\mu_s, \quad \text{with} \quad (3.9)$$

$$\alpha_t^{\widehat{p}_i} = \widehat{p}_{t,i} \cdot \widehat{\mathbf{a}}_t - \widehat{p}_{t,i} \widehat{\mathbf{a}}_t \quad (3.10)$$

$$\gamma_{t,j}^{\widehat{p}_i} = (1 - Y_{t-,j}) \left(\mathbb{E}(p_{t,i}(X_t) | \mathcal{F}_{t-}^M \vee \{\tau_j = t\}) - \mathbb{E}(p_{t,i}(X_t) | \mathcal{F}_{t-}^M) \right). \quad (3.11)$$

The predictable quadratic variations of the asset prices with respect to the market information \mathbb{F}^M satisfy $d\langle \widehat{p}_i, \widehat{p}_j \rangle_t^M = v_t^{ij} dt$ with

$$v_t^{ij} := \sum_{n=1}^m \gamma_{t,n}^{\widehat{p}_i} \gamma_{t,n}^{\widehat{p}_j} \widehat{\lambda}_{t-,n} + \sum_{n=1}^l \alpha_{t-,n}^{\widehat{p}_i} \alpha_{t-,n}^{\widehat{p}_j}. \quad (3.12)$$

Proof Recall that $\widehat{p}_{t,i} = \mathbb{E}(p_{t,i}(X_t) | \mathcal{F}_t^{\mathbb{M}})$ and that the fundamental value process $p_i(\cdot)$ given by $p_{t,i}(X_t) = \mathbb{E}(P_i | \mathcal{F}_t)$ is an \mathbb{F} -martingale. Since $p_i(\cdot)$ is adapted to $\mathbb{F}^X \vee \mathbb{F}^Y$ whereas B is independent of $\mathbb{F}^X \vee \mathbb{F}^Y$, we clearly have $[B, p_i(\cdot)] = 0$. Hence Proposition 3.2 applies; the form of $\alpha^{\widehat{p}_i}$ and $\gamma^{\widehat{p}_i}$ above is a straightforward consequence of (3.3) and (3.4).

In order to compute the predictable quadratic variations note that we have $[M_i]_t = \sum_{s \leq t} (\Delta M_{s,i})^2 = Y_{t,i}$; hence the predictable quadratic variation $\langle M_i \rangle_t^{\mathbb{M}}$ is given by the compensator of Y_i w.r.t. $\mathbb{F}^{\mathbb{M}}$, i.e. $\langle M_i \rangle_t^{\mathbb{M}} = (\int_0^t \widehat{\lambda}_{s-,i} ds)$. Moreover, since there are no joint defaults, for $i \neq j$ we have $0 = [M_i, M_j] = \langle M_i, M_j \rangle_t^{\mathbb{M}}$. Finally $\langle \mu_i, \mu_j \rangle_t^{\mathbb{M}} = \delta_{ij}t$. Hence

$$d\langle \widehat{p}_i, \widehat{p}_j \rangle_t^{\mathbb{M}} = \sum_{n=1}^m \gamma_{t,n}^{\widehat{p}_i} \gamma_{t,n}^{\widehat{p}_j} d\langle M_n \rangle_t^{\mathbb{M}} + \sum_{n=1}^l \alpha_{t-,n}^{\widehat{p}_i} \alpha_{t-,n}^{\widehat{p}_j} d\langle \mu_n \rangle_t^{\mathbb{M}},$$

and the form of v_t^{ij} follows. \square

We denote the predictable instantaneous covariance matrix of the asset price processes by $\mathbf{v}_t = (v_t^{ij})_{1 \leq i, j \leq N}$.

Remark 3.5 Note that the assumption that X was a finite state Markov chain is not needed in Proposition 3.2 or in Theorem 3.4; modulo certain boundedness conditions on the functions $\mathbf{a}(\cdot)$ and $\lambda_i(\cdot)$, these results hold for any model where the τ_i are conditionally independent, doubly stochastic random times in the sense of Assumption A1, independent of the specific dynamics of X . The filtering results in Section 3.2 below on the other hand do exploit the specific structure of X .

3.2 Filtering and Factor Representation of Market Prices

Define the K -dimensional conditional probability vector $\boldsymbol{\pi}_t = (\pi_t^1, \dots, \pi_t^K)^\top$ with $\pi_t^k := Q(X_t = k | \mathcal{F}_t^{\mathbb{M}})$. $\boldsymbol{\pi}$ is a natural state variable process for the model in the market filtration; in particular, once $\boldsymbol{\pi}_t$ is at hand, we are able to compute $\mathbb{E}(f(X_t) | \mathcal{F}_t^{\mathbb{M}}) = \sum_{k=1}^K f(k) \pi_t^k$ for any function $f : S^X \rightarrow \mathbb{R}$.

Filtering. The following proposition shows that the process $\boldsymbol{\pi}$ is the solution of a K -dimensional SDE system driven by the innovations processes M and μ .

Proposition 3.6 *Denote the generator matrix of X by $(q(\iota, k))_{1 \leq \iota, k \leq K}$. Then, for $k = 1, \dots, K$,*

$$d\pi_t^k = \sum_{\iota=1}^K q(\iota, k) \pi_t^\iota dt + (\boldsymbol{\gamma}^k(\boldsymbol{\pi}_{t-}))^\top dM_t + (\boldsymbol{\alpha}^k(\boldsymbol{\pi}_t))^\top d\mu_t, \quad (3.13)$$

where the coefficients are given by

$$\begin{aligned}\gamma_j^k(\boldsymbol{\pi}) &= \pi_k \left(\frac{\lambda_j(k)}{\sum_{n=1}^K \lambda_j(n) \pi_n} - 1 \right), \quad 1 \leq j \leq m, \\ \boldsymbol{\alpha}^k(\boldsymbol{\pi}) &= \pi_k \left(\mathbf{a}(k) - \sum_{n=1}^K \pi_n \mathbf{a}(n) \right).\end{aligned}\tag{3.14}$$

Proof Denote the generator of X by \mathcal{L} and set $f_k(x) = \mathbb{1}_{\{x=k\}}$. Then the \mathbb{F} -semimartingale decomposition of $(f_k(X_t))_{t \geq 0}$ is

$$f_k(X_t) = f_k(X_0) + \int_0^t \mathcal{L} f_k(X_s) ds + \left(f_k(X_t) - \int_0^t \mathcal{L} f_k(X_s) ds \right)$$

Note that $\pi^k = \widehat{f_k(X_t)}$ and that $\mathcal{L} f_k(X_t) = q(X_t, k)$. Moreover, $[f_k(\cdot), B] \equiv 0$. Hence Proposition 3.2 implies that π^k has the representation

$$d\pi_t^k = q(\widehat{X_t}, k) dt + (\boldsymbol{\gamma}_t^k)^\top dM_t + (\boldsymbol{\alpha}_t^k)^\top d\mu_t$$

with $q(\widehat{X_t}, k) = \sum_{l=1}^K q(l, k) \pi_t^l$ and

$$\begin{aligned}\boldsymbol{\alpha}_t^k &= f_k(\widehat{X_t}) \mathbf{a}(X_t) - \widehat{f_k(X_t)} \mathbf{a}(\widehat{X_t}) = \pi_t^k \mathbf{a}(k) - \pi_t^k \sum_{n=1}^K \pi_t^n \mathbf{a}(n), \\ \gamma_{t,j}^k &= (1 - Y_{t,j}) \left(Q(X_t = k | \mathcal{F}_{t-}^{\mathbb{M}} \vee \{\tau_j = t\}) - \pi_{t-}^k \right).\end{aligned}$$

Now the distribution of X conditional on a hypothetical default of firm j at t satisfies

$$Q(X_t = k | \mathcal{F}_{t-}^{\mathbb{M}} \vee \{\tau_j = t\}) = \frac{\lambda_j(k) \pi_{t-}^k}{\sum_{n=1}^K \lambda_j(n) \pi_{t-}^n}\tag{3.15}$$

see for instance Algorithm 5.3 in Frey & Runggaldier (2006) (Step 4). This gives (3.13) and (3.14). \square

Remark 3.7 Using the relations $dM_{t,i} = dY_{t,i} - (1 - Y_{t-,i}) \sum_{k=1}^K \lambda_i(k) \pi_{t-}^k dt$ as well as $d\mu_t = dZ_t - \sum_{k=1}^K \mathbf{a}(k) \pi_t^k dt$, the SDE-system (3.13) can alternatively be written with the observation processes Y and Z as drivers.

Related results have previously appeared in the filtering literature. For the case of diffusion observations, (3.13) is the Wonham filter, see Wonham (1965). For the case of marked-point-process observations we refer to Brémaud (1981) and further references therein.

Factor structure. All quantities of interest can be represented in terms of the process $\boldsymbol{\pi}$, as we now show. First, let $\mathbf{P}_t := (p_{t,i}(k))_{1 \leq i \leq N, 1 \leq k \leq K}$ be the matrix of the fundamental values of the traded securities. Then the vector of their prices satisfies $\widehat{\mathbf{p}}_t = \mathbf{P}_t \boldsymbol{\pi}_t$. Using Theorem 3.4 the dynamics of traded securities can be expressed in terms of $\boldsymbol{\pi}$ as well. We have

$$\begin{aligned} \boldsymbol{\alpha}_{\widehat{p}_t} &= \left(\sum_{k=1}^K p_{t,i}(k) \mathbf{a}(k) \pi_t^k - \left(\sum_{k=1}^K \pi_t^k p_{t,i}(k) \right) \left(\sum_{k=1}^K \pi_t^k \mathbf{a}(k) \right) \right) \\ \gamma_{\widehat{p}_t, j}^{\widehat{p}_t} &= \sum_{k=1}^K (1 - Y_{t-,j}) \left(\mathbb{E} \left(p_{t,i}(k) \mathbf{1}_{\{X_t=k\}} | \mathcal{F}_{t-}^{\mathbb{M}} \vee \{\tau_j = t\} \right) - p_{t,i}(k) \pi_{t-}^k \right). \end{aligned}$$

Using (2.1) and (3.15) we obtain the following representation for $\gamma_j^{\widehat{p}_t}$:

$$\begin{aligned} \gamma_{t,j}^{\widehat{p}_t} &= \sum_{n=0}^{m-1} \mathbf{1}_{(T_n(\omega), T_{n+1}(\omega)]}(t) \sum_{k=1}^K \left(p_{t,i}^{n+1}(k, (T_i(\omega), \xi_i(\omega))_{1 \leq i \leq n}, (t, j)) \frac{\lambda_j(k) \pi_{t-}^k}{\sum_{l=1}^K \lambda_j(l) \pi_{t-}^l} \right. \\ &\quad \left. - p_{t,i}^n(k, (T_i(\omega), \xi_i(\omega))_{1 \leq i \leq n}) \pi_{t-}^k \right). \end{aligned}$$

With $\boldsymbol{\gamma}$ and $\boldsymbol{\alpha}$ at hand the matrix \mathbf{v}_t of the instantaneous predictable quadratic covariations of the traded securities can also be expressed as function of $\boldsymbol{\pi}_{t-}$, see Equation (3.12).

Remark 3.8 The previous results permit us to give an explicit expression for the *contagion effects* induced by incomplete information (the fact that at the default of company j the default intensity of company $i \neq j$ - and hence also the credit spread of securities issued by that firm - is directly affected). We get from (3.15)

$$\widehat{\lambda}_{\tau_j, i} - \widehat{\lambda}_{\tau_j-, i} = \sum_{k=1}^K \lambda_i(k) \cdot \pi_{\tau_j-}^k \left(\frac{\lambda_j(k)}{\sum_{l=1}^K \lambda_j(l) \pi_{\tau_j-}^l} - 1 \right) = \frac{\text{cov}^{\boldsymbol{\pi}_{\tau_j-}}(\lambda_i, \lambda_j)}{\mathbb{E}^{\boldsymbol{\pi}_{\tau_j-}}(\lambda_j)}. \quad (3.16)$$

Recall that the probability vector $\boldsymbol{\pi}_{\tau_j-}$ gives the conditional distribution of X immediately prior to the default event. According to (3.16), default contagion increases (i) with increasing correlation of the random variables $\lambda_i(\cdot)$ and $\lambda_j(\cdot)$ under $\boldsymbol{\pi}_{\tau_j-}$, and (ii) with increasing variance of $\lambda_i(\cdot)$ or $\lambda_j(\cdot)$, i.e. with increasing dispersion of the measure $\boldsymbol{\pi}_{\tau_j-}$. Both effects are very intuitive.

4 Pricing and hedging for secondary-market investors

Now we consider the pricing and hedging w.r.t. the investor's filtration $\mathbb{F}^{\mathbb{I}}$. We start with general results on pricing and risk-minimizing hedging valid for an arbitrary investor filtration $\mathbb{F}^{\mathbb{I}}$ satisfying $\mathbb{F}^Y \subset \mathbb{F}^{\mathbb{I}} \subset \mathbb{F}^{\mathbb{M}}$. Specific examples and the corresponding calibration and filtering methodology are discussed in Sections 4.2 and 4.3.

4.1 Pricing and hedging: general results

The valuation of contingent claims. The secondary market value of a \mathcal{F}_T^Y -measurable, (non-traded) claim H is defined as $\mathbb{E}(H|\mathcal{F}_t^\mathbb{I})$. As previously, the full-information value of H is given by $h_t(X_t) := \mathbb{E}(H|\mathcal{F}_t)$, h_t a \mathcal{F}_t^Y -measurable function. Note that

$$\mathbb{E}(H|\mathcal{F}_t^\mathbb{I}) = \mathbb{E}\left(\mathbb{E}(H|\mathcal{F}_t^\mathbb{M})|\mathcal{F}_t^\mathbb{I}\right) = E\left(\sum_{k=1}^K h_t(k)\pi_t^k \mid \mathcal{F}_t^\mathbb{I}\right) = \sum_{k=1}^K h_t(k)\mathbb{E}(\pi_t^k \mid \mathcal{F}_t^\mathbb{I}). \quad (4.1)$$

Hence the computation of secondary-market values can be traced back to finding the mean of the conditional distribution of $\boldsymbol{\pi}_t$ given $\mathcal{F}_t^\mathbb{I}$; in the sequel this distribution is denoted $\nu_t = \nu_t(d\boldsymbol{\pi})$.

Risk-minimizing strategies under $\mathbb{F}^\mathbb{I}$. The concept of risk minimization under the risk neutral measure Q in the sense of Föllmer & Sondermann (1986) is well-suited for our setting: to begin with, it has a natural extension to problems under incomplete information, developed in Schweizer (1994). Moreover, working under the risk-neutral measure is natural for practitioners who usually prefer to set up pricing models directly under Q .

We begin by recalling the notion of a risk-minimizing strategy under restricted information as introduced in Schweizer (1994). Denote by $L^2(\widehat{\mathbf{p}}, \mathbb{F}^\mathbb{M})$ and $L^2(\widehat{\mathbf{p}}, \mathbb{F}^\mathbb{I})$ the space of $\mathbb{F}^\mathbb{M}$ respectively $\mathbb{F}^\mathbb{I}$ -predictable processes $\boldsymbol{\theta}$ such that $\mathbb{E}(\int_0^T \boldsymbol{\theta}_s^\top \mathbf{v}_s \boldsymbol{\theta}_s ds) < \infty$. Recall that we have assumed that $\mathbb{F}^Y \subset \mathbb{F}^\mathbb{I} \subset \mathbb{F}^\mathbb{M}$. This assumption on $\mathbb{F}^\mathbb{I}$ seems reasonable from an economic viewpoint since defaults are usually publicly observable. Note moreover, that the inclusion $\mathcal{F}_T^Y \subset \mathcal{F}_T^\mathbb{I}$ implies that at the maturity date T the price $\widehat{p}_{T,i} = P_i$ is observable for secondary-market investors.

Definition 4.1 (Risk-minimizing trading strategies)

1. An $\mathbb{F}^\mathbb{I}$ -admissible strategy is given by a pair $\varphi = (\boldsymbol{\theta}, \eta)$ where $\boldsymbol{\theta} \in L^2(\widehat{\mathbf{p}}, \mathbb{F}^\mathbb{I})$ and η is $\mathbb{F}^\mathbb{I}$ -adapted. Moreover the value process $V_t = V_t(\varphi) = \boldsymbol{\theta}_t^\top \widehat{\mathbf{p}}_t + \eta_t$ is RCLL and $\mathbb{E}(\sup_{0 \leq t \leq T} V_t^2) < \infty$. The cost process for trading strategy φ is defined by

$$C_t = C_t(\varphi) = V_t(\varphi) - \int_0^t \boldsymbol{\theta}_s^\top d\widehat{\mathbf{p}}_s,$$

and the risk process (w.r.t. $\mathbb{F}^\mathbb{I}$) is defined by $R_t^\mathbb{I} = R_t^\mathbb{I}(\varphi) = \mathbb{E}((C_T(\varphi) - C_t(\varphi))^2 | \mathcal{F}_t^\mathbb{I})$.

2. An $\mathbb{F}^\mathbb{I}$ -admissible strategy φ is called an $\mathbb{F}^\mathbb{I}$ -risk-minimizing hedging strategy for a claim $H \in L^2(\Omega, \mathcal{F}_T^Y, Q)$, if $V_T(\varphi) = H$ and if moreover for any $t \in [0, T]$ and any $\mathbb{F}^\mathbb{I}$ -admissible strategy $\tilde{\varphi}$ satisfying $V_T(\tilde{\varphi}) = H$, $\boldsymbol{\theta}_s = \tilde{\boldsymbol{\theta}}_s$ for $s \leq t$ and $\eta_s = \tilde{\eta}_s$ for $s < t$, we have $R_t^\mathbb{I}(\varphi) \leq R_t^\mathbb{I}(\tilde{\varphi})$.

Note that this definition allows for the case that the price processes \widehat{p}_i of the traded assets are not $\mathbb{F}^{\mathbb{I}}$ -adapted. However, since the $\widehat{p}_{T,i}$ is $\mathcal{F}_T^{\mathbb{I}}$ -measurable, the settlement value $V_T(\phi) = \theta_T \widehat{\mathbf{p}}_T + \eta_T$ of the strategy is $\mathcal{F}_T^{\mathbb{I}}$ -measurable, so that second part of the definition makes sense.

As a first step towards computing risk-minimizing hedging strategies we determine the *Galtchouk-Kunita-Watanabe* decomposition of a claim $H \in L^2(\Omega, \mathcal{F}_T^Y, Q)$ and the associated martingale $\widehat{h}_t = \mathbb{E}(H | \mathcal{F}_t^{\mathbb{M}})$ w.r.t. prices of traded securities:

$$\widehat{h}_t = \widehat{h}_0 + \sum_{i=1}^N \int_0^t \xi_{t,i}^H d\widehat{p}_{t,i} + L_t^H, \quad \xi_i^H \in L^2(\widehat{\mathbf{p}}, \mathbb{F}^{\mathbb{M}}), \quad \langle L^H, \widehat{\mathbf{p}} \rangle^{\mathbb{M}} \equiv 0. \quad (4.2)$$

It is well-known from Föllmer & Sondermann (1986) that the $\mathbb{F}^{\mathbb{M}}$ -risk-minimizing strategy has value process \widehat{h} and trading strategy ξ^H .

Lemma 4.2 (Galtchouk-Kunita-Watanabe decomposition) *A possible choice of ξ^H in (4.2) is given by*

$$\xi_t^H = \mathbf{v}_t^{\text{inv}} \frac{d\langle \widehat{h}, \mathbf{p} \rangle_t^{\mathbb{M}}}{dt}, \quad (4.3)$$

where \mathbf{v}_t denotes the instantaneous predictable covariations of the asset prices at time t (see (3.12)), where $\mathbf{v}_t^{\text{inv}}$ denotes the pseudo inverse of \mathbf{v}_t and where $d\langle \widehat{h}, \mathbf{p} \rangle_t^{\mathbb{M}}/dt$ is the predictable Lebesgue-density of $\langle \widehat{h}, \mathbf{p} \rangle_t^{\mathbb{M}}$.

Proof Since $\langle L, \widehat{\mathbf{p}} \rangle^{\mathbb{M}} \equiv 0$, we know that ξ_t^H solves the equation

$$\langle \widehat{h}, p_j \rangle_t^{\mathbb{M}} = \sum_{n=1}^N \int_0^t \xi_{s,j} d\langle p_n, p_j \rangle_s^{\mathbb{M}}. \quad (4.4)$$

Moreover, by Equation (3.12) we have $d\langle \mathbf{p} \rangle_t^{\mathbb{M}} = \mathbf{v}_t dt$. The pseudo inverse gives a vector which minimizes the distance between the range of \mathbf{v}_t and the ‘‘r.h.s.’’ $d\langle \widehat{h}, \mathbf{p} \rangle_t^{\mathbb{M}}/dt$. Hence a solution of (4.4) is given by (4.3). \square

Remark 4.3 By an analogous argument as in the proof of Theorem 3.4, \widehat{h}_t has a representation of the form (3.9) with integrands α^H and γ^H given by analogous expressions to (3.10) and (3.11). Then, $\langle \widehat{h}, \mathbf{p} \rangle^{\mathbb{M}}$ computes to

$$d\langle \widehat{h}, p_i \rangle_t^{\mathbb{M}} = \left(\sum_{j=1}^m (\gamma_{t,j}^H \gamma_{t,j}^i) \widehat{\lambda}_{t-,j} + \sum_{j=1}^l \alpha_{t-,j}^H \alpha_{t-,j}^{\widehat{p}_i} \right) dt, \quad 1 \leq i \leq N. \quad (4.5)$$

In particular, $\mathbf{v}_t \xi_t^H = d\langle \widehat{h}, \mathbf{p} \rangle_t^{\mathbb{M}}/dt$ can be written as a function of $\boldsymbol{\pi}_{t-}$.

The following assumption on the investor information enables us to compute the risk-minimizing hedging strategy w.r.t. $\mathbb{F}^{\mathbb{I}}$.

A3 Every $\mathbb{F}^{\mathbb{I}}$ -martingale $M^{\mathbb{I}}$ is quasi left continuous, i.e. $\Delta M_\tau^{\mathbb{I}} = 0$ for any $\mathbb{F}^{\mathbb{I}}$ -predictable stopping time τ .

For any measurable stochastic process U we denote by ${}^{o,\mathbb{F}^\mathbb{I}}U$ and ${}^{p,\mathbb{F}^\mathbb{I}}U$ the optional projection and the predictable projection, respectively, of U on the filtration $\mathbb{F}^\mathbb{I}$; see for instance Rogers & Williams (1994), Chapter VI.

Proposition 4.4 *If $\mathbb{F}^\mathbb{I}$ satisfies Assumption A3, the $\mathbb{F}^\mathbb{I}$ -risk-minimizing hedging strategy $\varphi^H = (\boldsymbol{\theta}^H, \eta^H)$ is given by*

$$\boldsymbol{\theta}_t^H := ({}^{o,\mathbb{F}^\mathbb{I}}\mathbf{v})_{t-}^{\text{inv}} ({}^{o,\mathbb{F}^\mathbb{I}}(\mathbf{v}\boldsymbol{\xi}^H))_{t-} \quad (4.6)$$

and $\eta_t^H = \mathbb{E}(H - (\boldsymbol{\theta}_t^H)^\top \widehat{\mathbf{p}}_t | \mathcal{F}_t^\mathbb{I})$.

Proof From Schweizer (1994), Theorem 3.1 respectively Relation (4.1) we obtain a $\mathbb{F}^\mathbb{I}$ -risk minimizing hedging strategy given by

$$\boldsymbol{\theta}_t^H := ({}^{p,\mathbb{F}^\mathbb{I}}\mathbf{v})_t^{\text{inv}} ({}^{p,\mathbb{F}^\mathbb{I}}(\mathbf{v}\boldsymbol{\xi}^H))_t$$

and $\eta_t^H = \mathbb{E}(H - (\boldsymbol{\theta}_t^H)^\top \widehat{\mathbf{p}}_t | \mathcal{F}_t^\mathbb{I})$. Denote for a generic RCLL process U by U^- the left continuous version given by $U_t^- = \lim_{s \rightarrow t-} U_s$. A similar reasoning as in the proof of Frey (2000), Proposition 3.1, yields that under Assumption A3 $({}^{p,\mathbb{F}^\mathbb{I}}\mathbf{v}) = ({}^{o,\mathbb{F}^\mathbb{I}}\mathbf{v})^-$ as well as $({}^{p,\mathbb{F}^\mathbb{I}}(\mathbf{v}\boldsymbol{\xi}^H)) = ({}^{o,\mathbb{F}^\mathbb{I}}(\mathbf{v}\boldsymbol{\xi}^H))^-$ and we obtain the representation for $\boldsymbol{\theta}^H$ as claimed. \square

Note that for $N = 1$ the result can be written more explicitly; in that case $\boldsymbol{\theta}_t$ is the left-continuous version of the process

$$\mathbb{E}(v_t \boldsymbol{\xi}_t^H | \mathcal{F}_t^\mathbb{I}) / \mathbb{E}(v_t | \mathcal{F}_t^\mathbb{I}), \quad 0 \leq t \leq T.$$

In order to determine $\boldsymbol{\theta}^H$ one has to compute the optional projections in (4.6). Since \mathbf{v}_t and $\mathbf{v}_t \boldsymbol{\xi}_t^H$ are nonlinear functions of $\boldsymbol{\pi}_{t-}$, the computation of hedging strategies for secondary-market investors leads to the problem of finding $\nu_t(d\boldsymbol{\pi})$ (the conditional distribution of $\boldsymbol{\pi}_t$ given $\mathcal{F}_t^\mathbb{I}$), a problem which we encountered already in the analysis of pricing problems with respect to $\mathbb{F}^\mathbb{I}$. Note that the determination of $\nu_t(d\boldsymbol{\pi})$ is a second nonlinear filtering problem, this time with state process $\boldsymbol{\pi}$ and observation filtration $\mathbb{F}^\mathbb{I}$.

4.2 Modelling the investor filtration

In this section we discuss two possible specification for the investor filtration. First, we discuss a simple scenario with observable asset prices. Given enough independent price observations, the probability vector $\boldsymbol{\pi}_t$ can then be backed out with regression-type methods so that this scenario permits a simple, pragmatic calibration approach. The second scenario is more in line with the overall spirit of the paper. Here we consider the case where the price vector $\widehat{\mathbf{p}}$ is observed in Gaussian noise and show that in this case pricing and hedging for secondary-market investors leads to an interesting filtering problem with signal process $\boldsymbol{\pi}$ and observation process given by Y and by the noisily observed prices. A numerical solution of this problem via particle filtering is sketched in Section 4.3.

Implied state probabilities. Recall that the observed market prices of traded securities have the form $\widehat{\mathbf{p}}_t = \mathbf{P}_t \boldsymbol{\pi}_t$. Assume that the investor observes the N -dimensional price vector $\widehat{\mathbf{p}}$ directly, i.e. $\mathbb{F}^{\mathbb{I}} = \mathbb{F}^{\widehat{\mathbf{p}}} \vee \mathbb{F}^Y$. If moreover the rank of \mathbf{P}_t is equal to K or, equivalently, if $\mathbf{P}_t^\top \mathbf{P}_t$ is positive definite, the probability vector $\boldsymbol{\pi}_t$ can be backed out from the observed prices by solving the following regression-type minimization problem:³

$$\boldsymbol{\pi}_t = \operatorname{argmin}_{\{\boldsymbol{\pi} \geq 0, \sum_{k=1}^K \pi_k = 1\}} \|\widehat{\mathbf{p}}_t - \mathbf{P}_t \boldsymbol{\pi}\|^2. \quad (4.7)$$

With $\boldsymbol{\pi}_t$ at hand, we get from (4.1) for any \mathcal{F}_T^Y -measurable claim H that $\mathbb{E}(H \mid \mathcal{F}_t^{\mathbb{I}}) = \sum_{k=1}^K h_t(k) \pi_t^k = \mathbb{E}(H \mid \mathcal{F}_t^{\mathbb{M}})$. Moreover, we have

$$\left({}^{o, \mathbb{F}^{\mathbb{I}}} \mathbf{v} \right)_{t-}^{\text{inv}} \left({}^{o, \mathbb{F}^{\mathbb{I}}} (\mathbf{v} \boldsymbol{\xi}^H) \right)_{t-} = \mathbf{v}_{t-}^{\text{inv}} (\mathbf{v}_{t-} \boldsymbol{\xi}_{t-}^H) = \boldsymbol{\xi}_t^H.$$

Summarizing, in this case secondary-market values and $\mathbb{F}^{\mathbb{I}}$ -risk-minimizing hedging strategies coincide with prices and risk-minimizing hedging strategies as defined with respect to the market filtration $\mathbb{F}^{\mathbb{M}}$.

Prices observed with Gaussian noise. Here we assume that the investor information is given by $\mathbb{F}^{\mathbb{I}} = \mathbb{F}^Y \vee \mathbb{F}^{\mathbf{U}}$ where the N -dimensional process \mathbf{U} solves the SDE

$$d\mathbf{U}_t = \widehat{\mathbf{p}}_t dt + w d\mathbf{W}_t = \mathbf{P}_t \boldsymbol{\pi}_t dt + w d\mathbf{W}_t, \quad (4.8)$$

for an N -dimensional standard Brownian motion \mathbf{W} , independent of all other processes considered, and a deterministic $N \times N$ matrix w with full rank. We denote the instantaneous covariance matrix of \mathbf{U} by $\Sigma_{\mathbf{U}} := w w^\top$. Note that an argument analogous to the proof of Lemma 3.1 shows that every \mathbb{F}^I -adapted martingale can be written as stochastic integral with respect to the \mathbb{F}^I -martingales $Y_{t,i} - \int_0^{t \wedge \tau_i} \mathbb{E}(\widehat{\lambda}_{s,i} \mid \mathcal{F}_s^{\mathbb{I}}) ds$, $1 \leq i \leq m$, and $\mathbf{U}_t - \int_0^t \mathbb{E}(\widehat{\mathbf{p}}_s \mid \mathcal{F}_s^{\mathbb{I}}) ds$, so that Assumption A3 is satisfied.

\mathbf{U} can be viewed as noisy price information of the traded assets $\widehat{p}_1, \dots, \widehat{p}_N$. This interpretation is motivated by the following heuristic argument, borrowed from Frey & Runggaldier (2006). Suppose that secondary market investors observe market quotes at times $t_k = k\Delta$ for some small $\Delta > 0$ and that the prices actually quoted on the market are of the form $\mathbf{u}_{t_k} = \widehat{\mathbf{p}}_{t_k} + \boldsymbol{\epsilon}_k$ for an iid sequence of noise variables $(\boldsymbol{\epsilon}_k)_k$, independent of $\widehat{\mathbf{p}}$, with $\mathbb{E}(\boldsymbol{\epsilon}_1) = 0$ and positive definite covariance matrix $\operatorname{cov}(\boldsymbol{\epsilon}_1) = \tilde{w} \tilde{w}^\top$ for some deterministic and invertible $n \times N$ matrix \tilde{w} . The information can be represented in terms of the cumulative observation process $\mathbf{U}_t^\Delta := \Delta \sum_{t_k \leq t} \mathbf{u}_{t_k}$. Letting $w = \tilde{w} \sqrt{\Delta}$, we obtain, using Donsker's invariance principle,

$$\mathbf{U}_t^\Delta = \sum_{t_k \leq t} \widehat{\mathbf{p}}_{t_k} \Delta + \sum_{t_k \leq t} \Delta \boldsymbol{\epsilon}_k \approx \int_0^t \widehat{\mathbf{p}}_s ds + w \mathbf{W}_t.$$

³ Positive definiteness of $\mathbf{P}_t^\top \mathbf{P}_t$ implies that the problem (4.7) has a unique solution.

Note that $\Sigma_{\mathbf{U}} = \Delta \text{cov}(\boldsymbol{\epsilon}_1)$, so that the instantaneous covariance matrix of \mathbf{U} is proportional to the covariance matrix $\text{cov}(\boldsymbol{\epsilon}_1)$ of the observation noise and inversely proportional to the observation frequency $1/\Delta$. The noise $(\boldsymbol{\epsilon}_k)_k$ respectively $w\mathbf{W}$ - represents classical observation errors such as bid-ask spreads and transmission errors as well as model errors. In applications the error covariance matrix $\Sigma_{\mathbf{U}}$ will be chosen by the modeler in an attempt to balance calibration errors (resulting from error variances which are too large) against potential instabilities of the filter (resulting from error variances which are too low).

Recall that the conditional probability vector $\boldsymbol{\pi}_t$ solves the SDE (3.13), driven by μ and Y . On the other hand, $\mathbb{F}^{\mathbb{I}}$ is generated by Y and the process \mathbf{U} as defined above. In order to determine $\nu_t(d\boldsymbol{\pi})$ one therefore has to solve a challenging second filtering problem with usually high-dimensional signal process $\boldsymbol{\pi}$; with observations of mixed type (diffusion and marked point processes); and with common jumps of observation Y and signal $\boldsymbol{\pi}$. This problem is discussed in the next section.

4.3 Filtering for secondary-market investors

Filtering problems with observations of mixed type and with common jumps of signal process and observation process have been studied extensively in Frey & Runggaldier (2006) ([FR] in the sequel). In the first part of this section we apply the results provided therein to our context and obtain general filtering equations in weak form. In the second part a numerical solution of the filtering problem via particle-filtering in the sense of Crisan & Lyons (1999) is proposed; this method is well-suited for high-dimensional problems. In particular, we use the general filtering equations to extend the algorithm of Crisan & Lyons (1999) to our more complicated setting.

General filter equations. In using the results from [FR] we make the following identifications: the state process is given by $\boldsymbol{\pi}$ and the observation filtration is $\mathbb{F}^{\mathbb{I}}$; the full-information-filtration of [FR] is given by $\mathbb{F}^{\mathbb{M}}$; the intensity of Y_i w.r.t. $\mathbb{F}^{\mathbb{M}}$ is

$$\widehat{\lambda}_i(\boldsymbol{\pi}_{t-}) := \sum_{k=1}^K \lambda_i(k) \pi_{t-}^k, \quad 1 \leq i \leq m; \quad (4.9)$$

and the drift of the continuous observation process is given by $\mathbf{P}_t \boldsymbol{\pi}_t$. Denote for given $\boldsymbol{\pi} \in [0, 1]^K$ with $\sum_{k=1}^K \pi^k = 1$ and for given $y \in \{0, 1\}^m$ by $\bar{Q}_{\boldsymbol{\pi}, y}$ the law of the solution of the K -dimensional SDE

$$d\bar{\pi}_t^k = \sum_{\iota=1}^K q(\iota, k) \bar{\pi}_{t-}^{\iota} dt - \sum_{i=1}^m (1 - y_i) \gamma_i^k(\bar{\boldsymbol{\pi}}_{t-}) \widehat{\lambda}_i(\bar{\boldsymbol{\pi}}_{t-}) dt + \boldsymbol{\alpha}^k(\bar{\boldsymbol{\pi}}_{t-})^{\top} d\mu_t, \quad (4.10)$$

with initial condition $\bar{\boldsymbol{\pi}}_0 = \boldsymbol{\pi}$; expectations with respect to $\bar{Q}_{\boldsymbol{\pi}, y}$ are denoted with $\mathbb{E}_{\boldsymbol{\pi}, y}$. A comparison with (3.13) shows that for $\boldsymbol{\pi} = \boldsymbol{\pi}_{T_{n-1}}$ and $y = Y_{T_{n-1}}$,

(4.10) describes the dynamics of the process $\boldsymbol{\pi}$ for $t \in [T_{n-1}, T_n)$ (i.e. between default times). Finally, define for $\mathbf{u} \in \mathbb{R}^N$ the quantity

$$\|\mathbf{u}\|_{\Sigma_{\mathbf{U}}^{-1}} = (\mathbf{u}^\top \Sigma_{\mathbf{U}}^{-1} \mathbf{u})^{1/2}$$

and put for $f : [0, 1]^K \rightarrow \mathbb{R}$ bounded and measurable $\nu_t f := \mathbb{E}(f(\boldsymbol{\pi}_t) \mid \mathcal{F}_t^{\mathbb{I}})$.

With this notation at hand, the filtering results from [FR] - which take the form of a recursion over the successive default times T_n - are readily applied. We start with the problem of *filtering between defaults*, i.e. we consider $\nu_t f$ for $t \in [T_{n-1}, T_n)$. From Theorem 4.1 in [FR] we obtain

$$\begin{aligned} \nu_t f \propto & \int_{[0,1]^K} \mathbb{E} \boldsymbol{\pi}_{Y_{T_{n-1}}} \left(f(\bar{\boldsymbol{\pi}}_{t-T_{n-1}}) \exp \left(- \int_0^{t-T_{n-1}} \sum_{i=1}^m (1 - Y_{T_{n-1}, i}) \widehat{\lambda}_i(\bar{\boldsymbol{\pi}}_{s-}) ds \right) \right. \\ & \cdot \exp \left(- \int_{T_{n-1}}^t (\mathbf{P}_s \bar{\boldsymbol{\pi}}(s - T_{n-1}))^\top \Sigma_{\mathbf{U}}^{-1} d\mathbf{U}_s - \frac{1}{2} \int_{T_{n-1}}^t \|\mathbf{P}_s \bar{\boldsymbol{\pi}}(s - T_{n-1})\|_{\Sigma_{\mathbf{U}}^{-1}}^2 ds \right) \left. \nu_{T_{n-1}}(d\boldsymbol{\pi}) \right). \end{aligned} \quad (4.11)$$

Next we consider the filtering *at a default time*. We have to take two effects into account: first, at the default time T_n the vector $\boldsymbol{\pi}$ jumps with $\Delta \pi_{T_n}^k = \gamma_{\xi_n}^k(\boldsymbol{\pi}_{T_n^-})$; second, the new information (T_n, ξ_n) leads to an update of the conditional distribution of $\boldsymbol{\pi}_{T_n^-}$ given $\mathcal{F}_{T_n}^{\mathbb{I}}$. From Theorem 4.3 of [FR] we therefore get with $S_n := T_n - \bar{T}_{n-1}$

$$\begin{aligned} \nu_{T_n} f \propto & \int_{[0,1]^K} \mathbb{E} \boldsymbol{\pi}_{Y_{T_{n-1}}} \left[f(\bar{\boldsymbol{\pi}}_{S_n} + \gamma_{\xi_n}(\bar{\boldsymbol{\pi}}_{S_n})) \cdot \widehat{\lambda}_{\xi_n}(\bar{\boldsymbol{\pi}}_{S_n}) \right. \\ & \cdot \exp \left(- \int_0^{S_n} \sum_{i=1}^m (1 - Y_{T_{n-1}, i}) \widehat{\lambda}_i(\bar{\boldsymbol{\pi}}_{s-}) ds \right) \\ & \cdot \exp \left(- \int_{T_{n-1}}^{T_n} (\mathbf{P}_s \bar{\boldsymbol{\pi}}(s - T_{n-1}))^\top \Sigma_{\mathbf{U}}^{-1} d\mathbf{U}_s - \frac{1}{2} \int_{T_{n-1}}^{T_n} \|\mathbf{P}_s \bar{\boldsymbol{\pi}}(s - T_{n-1})\|_{\Sigma_{\mathbf{U}}^{-1}}^2 ds \right) \left. \nu_{T_{n-1}}(d\boldsymbol{\pi}) \right]. \end{aligned} \quad (4.12)$$

Particle filtering. In particle filtering the conditional distribution ν_t is approximated by the occupation measure $\tilde{\nu}_t$ of a branching particle system. This branching system is constructed by a recursion over discrete time steps $t_k = k\Delta$, $k = 0, 1, \dots$ for Δ small. $\tilde{\nu}_{t_{k+1}}$ is constructed from $\tilde{\nu}_{t_k}$ in a two-stage procedure. In the *prediction step* one generates for each particle $\varpi^i(t_k)$ in the system at time t_k a trajectory of the SDE (4.10) of length Δ with initial value $\varpi^i(t_k)$. In the *updating step*, $\tilde{\nu}_{t_{k+1}}$ is generated from the simulated trajectories using the new default- and price information over $(t_k, t_{k+1}]$ by letting each particle branch into a random number of offsprings; the mean number of offsprings is taken proportional to the likelihood of the new observation as given in the filtering equations (4.11) or (4.12). Let $\beta(t_k)$ denote the number of particles at time t_k . In algorithmic form the evolution of the particle system can be described as follows.

1. The initial state $\tilde{\nu}_0$ of the system is given by the occupation measure of $\beta(0)$ particles of mass $1/\beta(0)$, i.e. $\tilde{\nu}_0 = \beta(0)^{-1} \sum_{i=1}^{\beta(0)} \delta_{\varpi^i(0)}$; here $\varpi^i(0)$, $1 \leq i \leq \beta(0)$ represent $\beta(0)$ independent draws from the initial distribution ν_0 .
2. (The prediction step) Given the particles $\varpi^1(t_k), \dots, \varpi^{\beta(t_k)}(t_k)$ in the system at time t_k , generate for $i = 1, \dots, \beta(t_k)$ independent trajectories $\bar{\pi}^i = (\bar{\pi}^i(s))_{0 \leq s \leq \Delta}$ of the SDE (4.10) with $y = Y_{t_k}$ and initial value $\bar{\pi}^i(0) = \varpi^i(t_k)$; the occupation measure

$$\tilde{\eta}_{t_{k+1}} := \beta(t_k)^{-1} \sum_{i=0}^{\beta(t_k)} \delta_{\bar{\pi}^i(\Delta)}$$

of the endpoints is then an approximation to the *predictive distribution* at t_{k+1} (the conditional distribution of $\pi_{t_{k+1}}$ given $\mathcal{F}_{t_k}^{\Pi}$).

3. (The updating step) We start with the case when there is no default in $(t_k, t_{k+1}]$. Given the new noisy price observation $(\mathbf{U}_t)_{t_k \leq s \leq t_{k+1}}$, in accordance with (4.11) we define for each trajectory $\bar{\pi}^i$, $1 \leq i \leq \beta(t_k)$, the likelihood

$$\begin{aligned} L^i := & \exp\left(-\int_0^\Delta \sum_{j=1}^m (1 - Y_{t_k, j}) \hat{\lambda}_j(\bar{\pi}^i(s-)) ds\right) \\ & \cdot \exp\left(-\int_{t_k}^{t_k+\Delta} (P_s \bar{\pi}^i(s - t_k))^\top \Sigma_{\mathbf{U}}^{-1} d\mathbf{U}_s - \frac{1}{2} \int_{t_k}^{t_k+\Delta} \|P_s \bar{\pi}^i(s - t_k)\|_{\Sigma_{\mathbf{U}}^{-1}}^2 ds\right). \end{aligned} \quad (4.13)$$

In a numerical implementation the stochastic integral in (4.13) would typically be computed by Euler approximation. Define

$$\mu^i := \frac{\beta(t_k) L^i}{\sum_{j=1}^{\beta(t_k)} L^j}$$

and denote by $[\mu^i]$ the integer part of μ^i . At t_{k+1} each particle $\varpi^i(t_k)$ produces independently a random number $m(i)$ of offsprings where $m(i)$ has support $\{[\mu^i], [\mu^i]+1\}$ and mean μ^i . The positions of the $m(i)$ offsprings of particle i are given by $\bar{\pi}^i(\Delta)$ (the endpoint of the trajectory with initial value $\varpi^i(t_k)$). We set $\beta(t_{k+1}) := \sum_{i=1}^{\beta(t_k)} m(i)$ and denote the new particles at time t_{k+1} by $\varpi^i(t_{k+1})$, $1 \leq i \leq \beta(t_{k+1})$. The approximation to the filter distribution at time t_{k+1} is then given by

$$\tilde{\nu}_{t_{k+1}} = \beta(t_{k+1})^{-1} \sum_{i=1}^{\beta(t_{k+1})} \delta_{\varpi^i(t_{k+1})}. \quad (4.14)$$

If there is a default event in $(t_k, t_{k+1}]$, we use (4.12) and proceed as follows⁴. Denote by $\xi \in \{1, \dots, m\}$ the identity of the defaulting firm and put $\tilde{L}^i := \hat{\lambda}_\xi(\bar{\pi}^i(\Delta)) L^i$ with L^i as in (4.13). The number of offsprings $m(i)$ is determined by the same mechanism as before but with $\tilde{L}^1, \dots, \tilde{L}^{\beta(t_k)}$ instead of $L^1, \dots, L^{\beta(t_k)}$; the position of the offsprings of particle i is given by $\bar{\pi}^i(\Delta) + \gamma_\xi \bar{\pi}^i(\Delta)$. The measure $\tilde{\nu}_{t_{k+1}}$ is then again given by (4.14).

⁴ Since the interval length is short we can neglect the possibility of more than one default per time step. Moreover, we may assume that the default happens exactly at $t = t_{k+1}$.

This algorithm has a number of advantages, in particular for high-dimensional problems. First, particles with small weights, corresponding to a-posteriori unlikely trajectories of $\boldsymbol{\pi}$, are not carried forward so that the particles concentrate mostly in the more probable regions of the state space. Moreover, the computational effort increases only linearly in the dimensionality K of the problem.

Remark 4.5 In Crisan & Lyons (1999) it is shown that for standard filtering problems $\tilde{\nu}_t$ does in fact converge in measure to ν_t if the initial number of particles $\beta(0)$ tends to infinity; see also Xiong & Zeng (2006) for a similar assertion in the context of marked-point-process observations. We are confident that a similar result could be established in our setting. However, a formal analysis is postponed, since it is not central to the issues discussed in this paper.

A Proofs

Proof of Lemma 3.1. The proof goes in three steps. First, we introduce a new measure Q^* , so that under Q^* the compensator of Y is independent of X and Z is a Q^* -Brownian motion. Next, we use available martingale representation results under Q^* and finally we change back to the original measure Q .

In the following, we simply write $\mathbf{a}_s := \mathbf{a}(X_s)$. Define the density martingale

$$L_t := \prod_{T_n \leq t} (\hat{\lambda}_{T_n-, \xi_n})^{-1} \exp \left(\int_0^t \sum_{i=1}^m (1 - Y_{s,i}) (\hat{\lambda}_{s-, i} - 1) ds - \int_0^t \hat{\mathbf{a}}_s^\top d\mu_s - \frac{1}{2} \int_0^t \|\hat{\mathbf{a}}_s\|^2 ds \right), \quad t \leq T,$$

and note that the dynamics of L is

$$dL_t = L_{t-} \left(\sum_{i=1}^m ((\hat{\lambda}_{t-, i})^{-1} - 1) dM_{t,i} - \hat{\mathbf{a}}_t^\top d\mu_t \right).$$

By Assumption A1, $\lambda_j > 0$. As S^X is finite, the functions $\lambda, \hat{\lambda}, \hat{\lambda}^{-1}$, and $\hat{\mathbf{a}}$ are bounded, hence L is a true martingale; see for instance Protter (2004), Exercise V.14. Define a measure Q^* by $dQ^*/dQ|_{\mathcal{F}_T^{\mathbb{M}}} = L_T$. Then, by the Girsanov theorem, Z is a Q^* -Brownian motion and

$$M_{t,i}^* := Y_{t,i} - \int_0^t (1 - Y_{s,i}) ds, \quad 1 \leq i \leq m$$

are martingales under Q^* .

Consider now the $(Q, \mathbb{F}^{\mathbb{M}})$ -martingale U and define the Q^* -integrable random variable $N_T := U_T L_T^{-1}$ and the associated martingale $N_t = E^{Q^*}(N_T | \mathcal{F}_t^{\mathbb{M}})$. Note that by the Bayes formula,

$$N_t = \frac{1}{L_t} E^Q(N_T L_T | \mathcal{F}_t^{\mathbb{M}}) = \frac{1}{L_t} E^Q(U_T | \mathcal{F}_t^{\mathbb{M}}) = \frac{U_t}{L_t}.$$

Due to the independence of Z and $\mathbf{M}^* = M_1^*, \dots, M_m^*$ we have a martingale representation of the \mathcal{Q}^* -martingale $(N_t)_{0 \leq t \leq T}$ using standard representation results such as Jacod & Shiryaev (1987), Theorem III.4.34:

$$N_t = \mathbb{E}(U_T) + \int_0^t \tilde{\alpha}_s^\top dZ_s + \int_0^t \tilde{\gamma}_s^\top d\mathbf{M}_s^*.$$

The final step is to compute the differential of $U_t = L_t N_t$. We have

$$\begin{aligned} dU_t &= d(L_t N_t) = L_{t-} dN_t + N_{t-} dL_t + d[L, N]_t \\ &= L_{t-} \tilde{\alpha}_t^\top (d\mu_t + \hat{\mathbf{a}}_t dt) + \sum_{i=1}^m L_{t-} \tilde{\gamma}_{t,i} (dY_{t,i} - (1 - Y_{t,i}) dt) \\ &\quad + \sum_{i=1}^m N_{t-} L_{t-} ((\hat{\lambda}_{t-,i})^{-1} - 1) (dY_{t,i} - \hat{\lambda}_{t,i} (1 - Y_{t,i}) dt) - L_{t-} N_{t-} \hat{\mathbf{a}}_t^\top d\mu_t - L_{t-} \hat{\mathbf{a}}_t^\top \tilde{\alpha}_t dt \\ &\quad + \sum_{i=1}^m L_{t-} ((\hat{\lambda}_{t-,i})^{-1} - 1) \tilde{\gamma}_{t,i} dY_{t,i}. \end{aligned}$$

Rearranging terms gives

$$dU_t = L_{t-} (\tilde{\alpha}_t - N_{t-} \hat{\mathbf{a}}_t) d\hat{\mu}_t + \sum_{i=1}^m L_{t-} \left(N_{t-} ((\hat{\lambda}_{t-,i})^{-1} - 1) + (\hat{\lambda}_{t-,i})^{-1} \tilde{\gamma}_{t,i} \right) dM_{t,i},$$

which is the desired martingale representation for U . \square

References

- Arnsdorf, M. & Halperin, I. (2007), ‘BSLP: Markovian bivariate spread-loss model for portfolio credit derivatives’, working paper, JP Morgan.
- Brémaud, P. (1981), *Point Processes and Queues*, Springer Verlag, Berlin Heidelberg New York.
- Coculescu, D., Geman, H. & Jeanblanc, M. (2006), ‘Valuation of default sensitive claims under imperfect information’, *Working paper*.
- Collin-Dufresne, P., Goldstein, R. & Helwege, J. (2003), ‘Is credit event risk priced? Modeling contagion via the updating of beliefs’, Preprint, Carnegie Mellon University.
- Crisan, D. & Lyons, T. (1999), ‘A particle approximation of the solution of the Kushner-Stratonovitch equation’, *Probability Theory and Related Fields* **115**, 549 – 578.
- Davis, M. H. A. & Marcus, S. I. (1981), An introduction to nonlinear filtering, in M. Hazewinkel & J. C. Willems, eds, ‘Stochastic Systems: The Mathematics of Filtering and Identifications and Applications’, Reidel Publishing Company, pp. 53–75.
- di Graziano, G. & Rogers, L. C. G. (2006), ‘A new approach to the modeling and pricing of correlation credit derivatives’, *Working paper*.
- Duffie, D. & Lando, D. (2001), ‘Term structures of credit spreads with incomplete accounting information’, *Econometrica* **69**, 633–664.
- Föllmer, H. & Sondermann, D. (1986), Hedging of non-redundant contingent-claims, in W. Hildenbrand & A. Mas-Colell, eds, ‘Contributions to Mathematical Economics’, North Holland, pp. 147–160.
- Frey, R. (2000), ‘Risk minimization with incomplete information in a model for high frequency data’, *Mathematical Finance* **10**(2), 215–225.
- Frey, R. & Runggaldier, W. (2006), ‘Credit risk and incomplete information: a nonlinear filtering approach’, *preprint, University of Leipzig*. Available from www.math.uni-leipzig.de/~frey.

-
- Frey, R. & Schmidt, T. (2006), 'Pricing corporate securities under noisy asset information', *forthcoming in Mathematical Finance*.
- Fujisaki, M., Kallianpur, G. & Kunita, H. (1972), 'Stochastic differential equations for the non-linear filtering problem', *Osaka Journal of Mathematics* **9**, 19 – 40.
- Gombani, A., Jaschke, S. & Runggaldier, W. (2005), 'A filtered no arbitrage model for term structures with noisy data', *Stochastic Processes and Applications* **115**, 381–400.
- Jacod, J. & Shiryaev, A. (1987), *Limit Theorems for Stochastic Processes*, 2nd edn, Springer Verlag, Berlin.
- Jarrow, R. & Protter, P. (2004), 'Structural versus reduced-form models: a new, information-based perspective', *Journal of Investment management* **2**, 1–10.
- Kusuoka, S. (1999), 'A remark on default risk models', *Advances in Mathematical Economics* **1**, 69–81.
- Protter, P. (2004), *Stochastic Integration and Differential Equations*, 2nd edn, Springer Verlag, Berlin Heidelberg New York.
- Rogers, L. C. G. & Williams, D. (1994), *Diffusions, Markov processes and martingales*, Wiley.
- Schönbucher, P. (2004), 'Information-driven default contagion', Preprint, Department of Mathematics, ETH Zürich.
- Schweizer, M. (1994), 'Risk-minimizing hedging strategies under restricted information', *Mathematical Finance* **4**(4), 327–342.
- Wonham, W. M. (1965), 'Some applications of stochastic differential equations to optimal non-linear filtering', *SIAM Journal of Control* **2**, 347 – 369.
- Xiong, J. & Zeng, Y. (2006), 'A branching particle approximation to the filtering problem with counting process observations', *submitted*.