Credit Risk With Point Processes: An Introduction

Kay Giesecke
Management Science & Engineering
Stanford University
giesecke@stanford.edu
www.stanford.edu/∼giesecke

Credit risk

- The distribution of financial loss due to a **broken financial agreement**, for example
  - Failure to pay according to schedule
  - Chapter 11 bankruptcy filing
- Pervades **virtually all** financial transactions
- A credit derivative is a security that facilitates the distribution of credit risk among market participants

Outline

1. Motivating example: Credit Default Swap (CDS)
2. Point processes
   (a) Transform calculation
   (b) Simulation methods
   (c) Likelihood estimation and fitness testing
3. Applications
   (a) Corporate bonds and CDS
   (b) Index swaps and tranches
   (c) Portfolio credit risk

Credit default swap

Mechanics

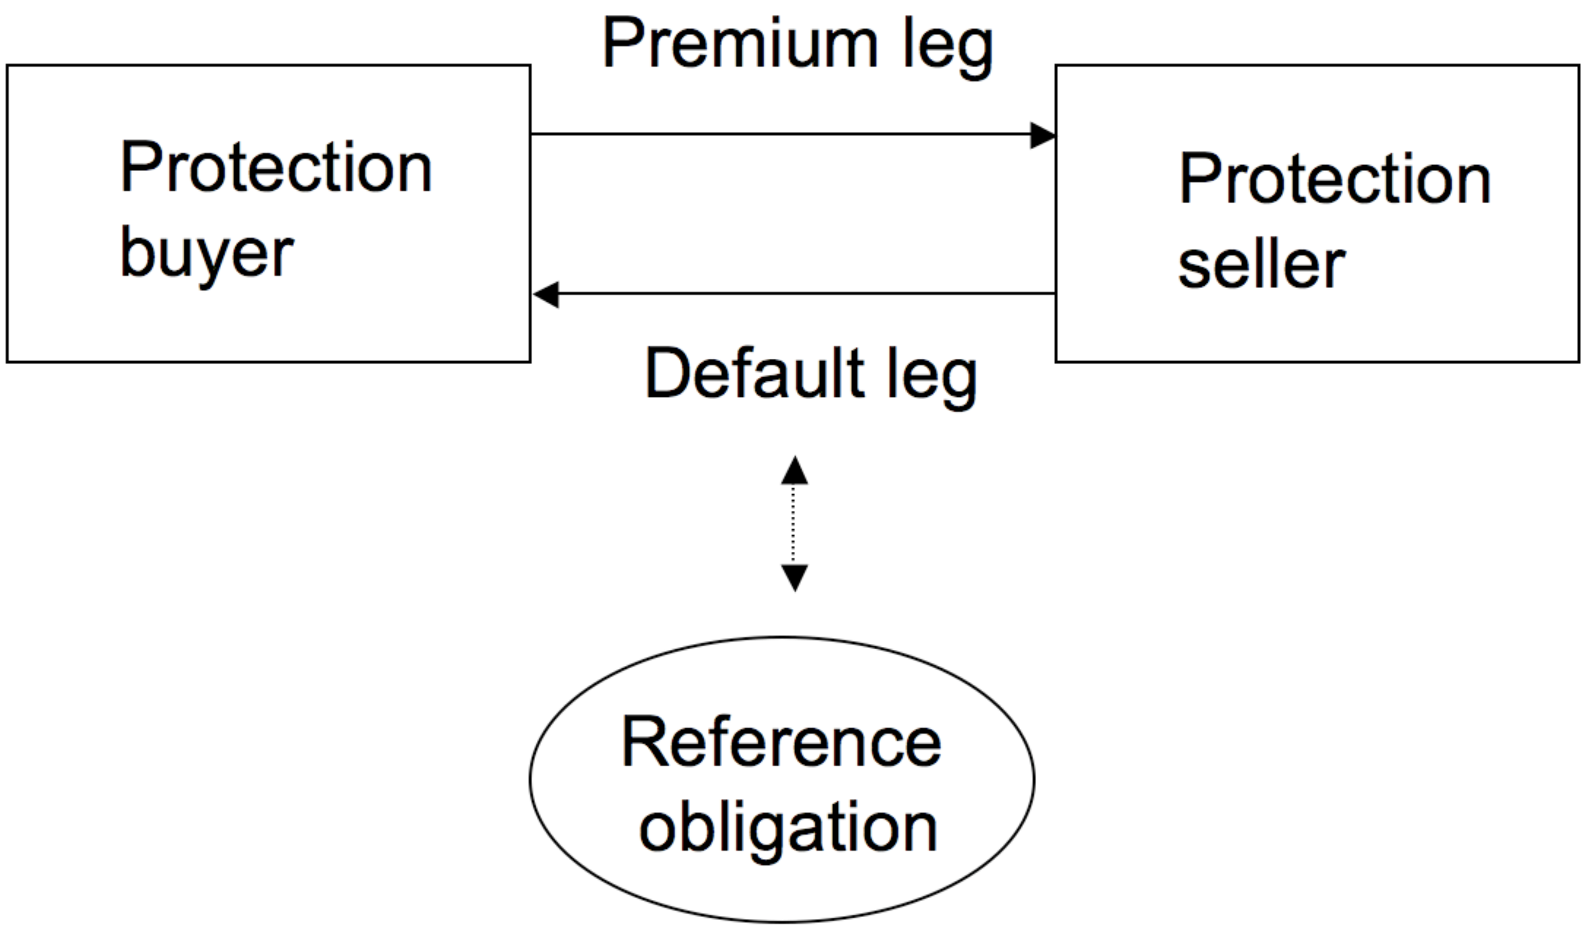

Credit default swap

A CDS is a bilateral over-the-counter transaction

- **Default leg**
  The protection seller compensates the protection buyer for the loss if the reference entity experiences a credit event before the maturity of the contract
- **Premium leg**
  The protection buyer pays either
  - An upfront fee at inception of the contract, or
  - A quarterly fee, called the CDS spread and stated as a fraction of the notional per annum, until the credit event or maturity, whichever is earlier

Credit default swap

Discussion

- Position of parties
  - The protection buyer shorts the reference credit without the risk of short squeezes or repo specials
  - The protection seller longs the reference credit
- Similar to an insurance contract, with the difference that the protection buyer often has no relation to the reference entity
  - Ban on naked positions in some markets (e.g. Europe)
- No exchange of notional takes place; the transaction is unfunded and therefore off-balance sheet
- Regulation being discussed: exchange trading, central clearing

Credit default swap

Applications

- Hedging
  - A bank may obtain regulatory capital relief by buying protection on a corporate loan it has made
- Trading ("speculation")
  - A fixed-income investor (e.g. insurance firm, hedge fund) can assume credit exposure by selling protection without having to fund a cash bond purchase; also liquidity is less of an issue
  - Express view on the behavior of the reference credit
    * Protection buyer (seller) can realize a mark-to-market profit if the reference credit deteriorates (improves)

Credit default swap

Credit event or **default**

- Events that trigger a payment are defined in the ISDA Master Agreement along with general terms and conditions governing a transaction, such as provisions relating to payment netting
- They include
  - Bankruptcy (Chapter 11, Chapter 7 etc.)
  - Obligation acceleration
  - Obligation default
  - Failure to pay
  - Repudiation/moratorium
  - Restructuring

Credit default swap

Alternative settlements at default

- Physical delivery
  - The protection buyer can choose to deliver to the seller any asset in a fixed pool of delivery assets, in exchange for the notional (cheapest-to-deliver option)
- Cash settlement
  - Unless the contract specifies a fixed recovery, market participants are polled to estimate the recovery value of the defaulted asset
  - The protection seller pays the protection buyer the notional minus the recovery

Credit default swap

Example

- A protection buyer purchases 5yr protection on a corporate name with notional \$10m at an annual spread of 300bp
- After 3 months, the protection buyer makes the first payment, roughly equal to $\$10m \times 0.03 \times 0.25 = \$75,000$
- The reference name defaults 1 month after the first payment and the reference obligation has a recovery rate of 45%
  - The seller pays $\$10m \times (100\% - 45\%) = \$5.5m$ at default
  - The buyer pays the accrued premium, roughly equal to $\$10m \times 0.03 \times 1/12 = \$25,000$, at default
- The contract expires

Credit default swap

Cash flows to the protection buyer

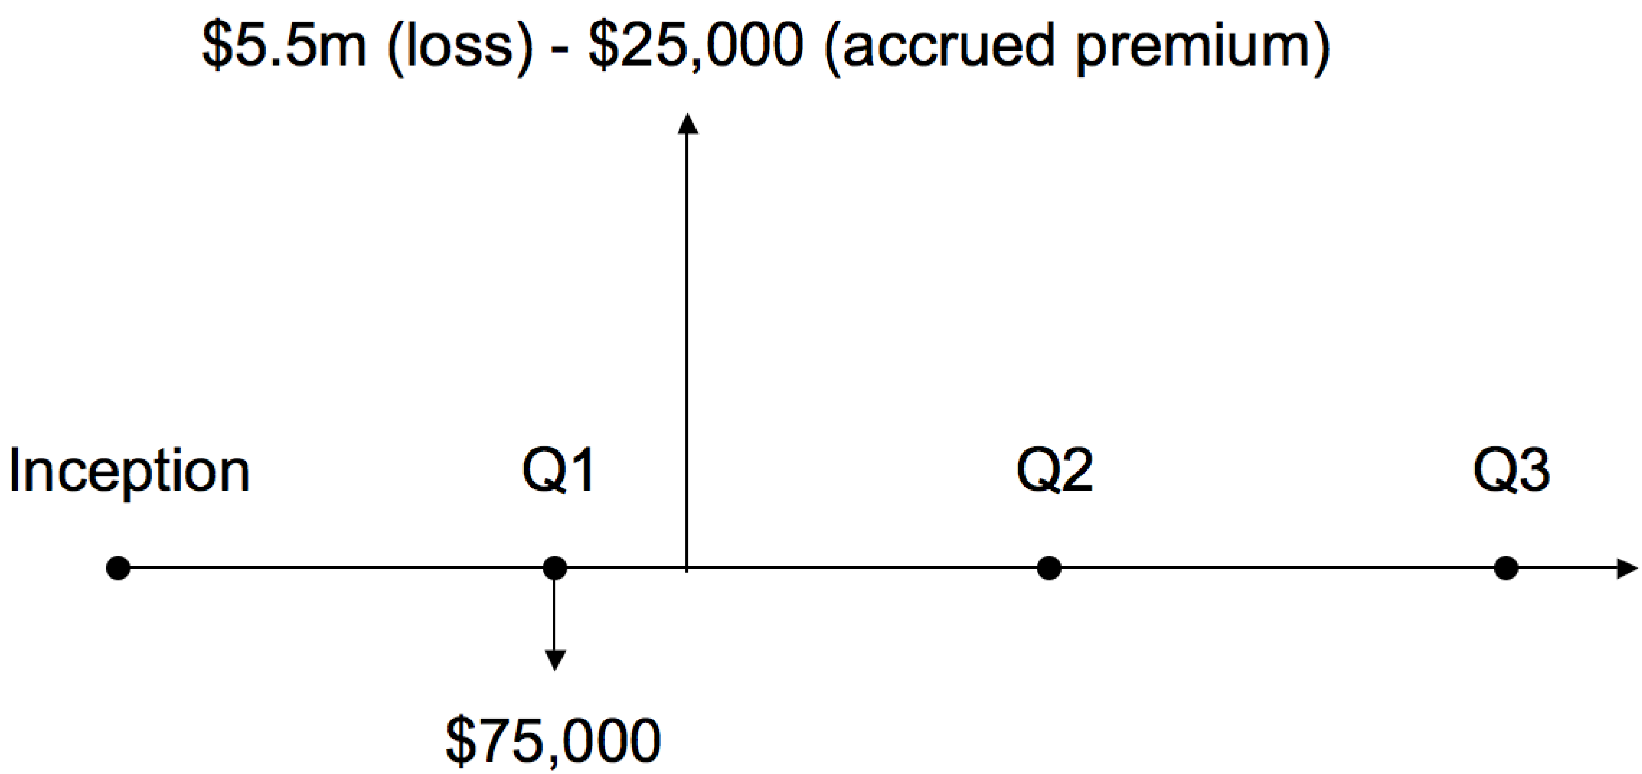

Credit default swap

Reverse engineering

- Credit swap spreads for various maturities between 0.5 and 10 years are quoted for a universe of reference names
- From market spreads we can infer a name's risk-neutral probability of default over future horizons
  - Mark-to-market of credit swap positions
  - Construction of forward CDS spread curve
  - Construction of intrinsic index swap spread
  - Design of credit trading strategies
- Need a model for default timing and the loss at default

Point process

Basic setting

- Fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with filtration $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ that is right-continuous and complete
- Consider a sequence of default stopping times $(T^k)_{k \geq 0}$
  - $0 = T^0 < T^1 < T^2 < \cdots$
  - $\lim_{k \to \infty} T^k = \infty$
- The stopping times generate the **default process** N by

$$N_t = \sum_{k \geq 1} 1_{\{T^k \leq t\}}$$

  - N is non-explosive: $N_t < \infty$
  - N has the same information as (T^k)

Point process

The sample paths of N are right continuous with left limits

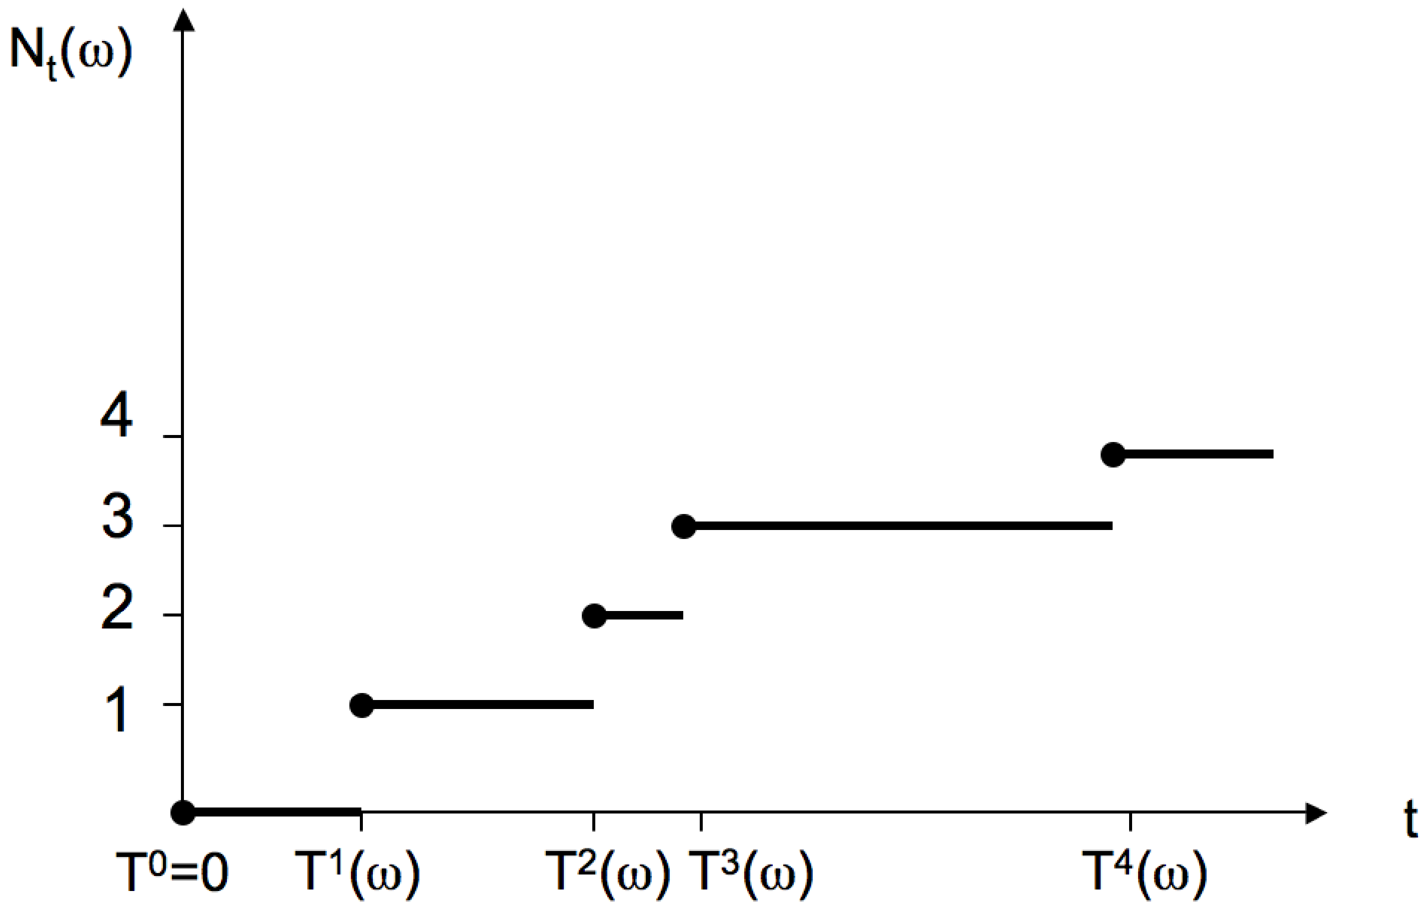

Point process

Compensator

- Since it is increasing, N is a submartingale:

$$N_t \leq \mathbb{E}(N_s \,|\, \mathcal{F}_t), \quad t \leq s$$

- The Doob-Meyer theorem guarantees that $N = A + M$, where A is an increasing, right continuous predictable process starting at 0 called the **compensator**, and M is a local martingale

- **Theorem.** The compensator A has continuous paths if and only if the T^k are totally inaccessible.
  - A stopping time τ is called predictable if there exists stopping times τ_n that increase to τ almost surely
  - A stopping time τ is called totally inaccessible if $\mathbb{P}(\tau = T < \infty) = 0$ for all predictable stopping times T

Point process

Intensity

- If A has paths that are absolutely continuous wrt. to the Lebesgue measure, i.e., if A takes the form

$$A_t = \int_0^t \lambda_s \, ds,$$

then the density λ is called the **intensity** of N

- The intensity represents the mean arrival rate: for "small" Δ

$$\lambda_t \Delta \approx \mathbb{E}(N_{t+\Delta} - N_t \,|\, \mathcal{F}_t)$$

- The existence of an intensity is convenient but not essential; for the analysis below we require A to have continuous paths only (i.e., we assume the T^k to be totally inaccessible)

Point process

Characteristic martingale

- Let $\psi(u) = 1 - e^{-u}$ and define

$$Z_t(u) = \exp(\psi(u)A_t - uN_t), \qquad u \geq 0$$

- Since the paths of N and A are right-continuous functions of finite variation, Stieltjes integration by parts (see Appendix) yields

$$Z_t(u) = 1 - \psi(u)\int_0^t Z_{s-}(u)dM_s$$

- We see that $Z(u)$ is a local martingale, since $M = N - A$ is one
- If $\mathbb{E}(\exp(A_T)) < \infty$, then $(Z_t(u))_{t\leq T}$ is a martingale

Point process

Poisson process

- Suppose A is deterministic
- By the martingale property of Z, we have

$$\mathbb{E}(e^{-u(N_T-N_t)} \,|\, \mathcal{F}_t) = e^{-\psi(u)(A_T-A_t)}$$

- We see that N has independent increments and that the increment $N_T - N_t \sim \text{Poi}(A_T - A_t)$, so N is a **Poisson process**
  - If $A_t = t$, then N is called a standard Poisson process
- **Watanabe's theorem.** A counting process is a Poisson process if and only if it has a deterministic compensator.
  - Analogous to Lévy's theorem for Brownian motion

Point process

Doubly-stochastic Poisson process (or Cox process)

- Suppose A is allowed to be random, but such that conditional on a path of A, N is a Poisson process with compensator A
  - The two-step randomization motivates the terminology **doubly-stochastic Poisson process**
- With $\mathbb{G} = (\mathcal{G}_t)_{t\geq 0} \subset \mathbb{F}$ the filtration generated by A, we have

$$\begin{aligned}\mathbb{E}(e^{-u(N_T-N_t)} \,|\, \mathcal{F}_t) &= \mathbb{E}(\mathbb{E}\{e^{-u(N_T-N_t)} \,|\, \mathcal{F}_t \vee \mathcal{G}_T\} \,|\, \mathcal{F}_t) \\ &= \mathbb{E}(e^{-\psi(u)(A_T-A_t)} \,|\, \mathcal{F}_t)\end{aligned}$$

- While generating tractability, the formulation is also restrictive: it does not permit feedback from N to A (for example, none of the T^k can be a $\mathbb{G}$-stopping time)

Point process transform

Measure change

- Since $Z_0(u) = 1$, can use $Z_T(u)$ to define $\mathbb{P}^u \approx \mathbb{P}$ on $\mathcal{F}_T$ via

$$\mathbb{P}^u(B) = \mathbb{E}(Z_T(u) 1_B), \quad B \in \mathcal{F}_T$$

- Each $\mathbb{P}^u$ corresponds to a Laplace transform of A:

$$\mathcal{L}^u(v, t, T) = \mathbb{E}^u(e^{-v(A_T - A_t)} \,|\, \mathcal{F}_t), \quad u, v \geq 0$$

- We compute the $\mathbb{P}$-Laplace transform of N as

$$\begin{aligned}
\mathbb{E}(e^{-u(N_T - N_t)} \,|\, \mathcal{F}_t) &= \mathbb{E}(e^{-\psi(u)(A_T - A_t)} e^{-u(N_T - N_t)} e^{\psi(u)(A_T - A_t)} \,|\, \mathcal{F}_t) \\
&= \mathbb{E}(e^{-\psi(u)(A_T - A_t)} Z_T(u)/Z_t(u) \,|\, \mathcal{F}_t) \\
&= \mathbb{E}^u(e^{-\psi(u)(A_T - A_t)} \,|\, \mathcal{F}_t) \\
&= \mathcal{L}^u(\psi(u), t, T)
\end{aligned}$$

Point process transform

Discussion

- The transform of N at u is given by the $\mathbb{P}^u$-transform of A evaluated at the characteristic exponent $\psi(u)$ of a Poisson process
  - Holds for any point process N with continuous compensator A satisfying $\mathbb{E}(\exp(A_T)) < \infty$
  - $\mathbb{P}^u$ absorbs any feedback from N to λ
  - Measure change is redundant in the doubly-stochastic case, where $\mathcal{L}^u = \mathcal{L}^0$ for all $0 \geq 0$
- Alternative derivation using Carr & Wu's (2004) complex-valued measure change for time-changed Lévy processes
  - By Meyer's (1971) theorem, N can be represented as a time-changed Poisson process N^0_A (see below)

Point process transform

Discussion

- Transform formula can be extended, see Giesecke & Zhu (2010)
  - Vector of interacting point processes
  - Random jump sizes
  - Include a stochastic discount factor and a random future cash flow (each possibly correlated with event times)

Point process transform

Transform of compensator

- The transform of A is analogous to the price at time t of a security that pays 1 at T when the risk-free interest rate is $v\lambda$:

$$\mathcal{L}^u(v,t,T) = \mathbb{E}^u(e^{-\int_t^T v\lambda_s ds} \mid \mathcal{F}_t)$$

- Adopt models for λ from default-free security valuation
  - Affine models, Duffie, Pan & Singleton (2000)
  - Quadratic models, Leippold & Wu (2002)
  - Linear-quadratic models, Cheng & Scaillet (2007)

Point process transform

Girsanov's theorem

- To calculate $\mathcal{L}^u(v,t,T)$, we need to understand how the dynamics of A are adjusted when the measure is changed to $\mathbb{P}^u$
  - Recall that A is initially specified under $\mathbb{P}$
- Let V be a local $\mathbb{P}$-martingale and $\langle V, N\rangle$ is the $\mathbb{P}$-conditional covariation, i.e. the $\mathbb{P}$-compensator of the quadratic variation $[V,N]$; by Girsanov's theorem the process

$$V + \psi(u)\langle V, N\rangle$$

is a $\mathbb{P}^u$-local martingale on $[0,T]$

Point process transform

Implications of Girsanov's theorem

- If V does not have jumps in common with N, then $[V, N] = 0$ and $\langle V, N\rangle = 0$. Thus, V remains a local martingale under $\mathbb{P}^u$.
  - A $\mathbb{P}$-Brownian motion is also a $\mathbb{P}^u$-Brownian motion
- The jumps of $V = N - A$ coincide with those of N, and $[V, N] = [N, N] = N$, whose compensator is $\langle V, N\rangle = A$. Thus,

$$V + \psi(u)\langle V, N\rangle = V + \psi(u)A = N - e^{-u}A$$

  is a local martingale under $\mathbb{P}^u$ on $[0, T]$
- If N has $\mathbb{P}$-intensity λ, then it has $\mathbb{P}^u$-intensity $e^{-u}\lambda$ on $[0, T]$: The measure change calls for a deterministic scaling of the intensity that depends on the variable u.

Example: Affine point process

- Take $\lambda = \Lambda(X)$ for an affine function Λ
- Under $\mathbb{P}$, the risk factor X solves the SDE

$$dX_t = \mu(X_t)\,dt + \sigma(X_t)\,dW_t + \delta\,dL_t$$

  where μ and σ^2 are affine functions

- $L_t = L^1 + \cdots + L^{N_t}$ is a (loss) point process whose jump times T^k are those of N; jump sizes L^k are iid with transform θ
  - Self-exciting model if $\delta > 0$
- From Duffie et al. (2000), know that for R affine

$$\mathbb{E}(e^{-\int_t^T R(X_s)ds + zX_T} \mid \mathcal{F}_t) = e^{a(t)+b(t)X_t}$$

  where $b(t) = b(z,t,T)$ and $a(t) = a(z,t,T)$ satisfy ODEs

Example: Affine point process

- Under $\mathbb{P}^u$, the intensity is Λe^{-u} and W is a Brownian motion
- Applying the previous formula under $\mathbb{P}^u$, we get

$$\mathcal{L}^u(v,t,T) = \mathbb{E}^u(e^{-v\int_t^T \Lambda(X_s)ds} \,|\, \mathcal{F}_t) = e^{\alpha(t)+\beta(t)X_t}$$

where $\beta(t) = \beta(u,v,t,T)$ and $\alpha(t) = \alpha(u,v,t,T)$ satisfy

$$\partial_t \beta(t) = v\Lambda_1 - K_1\beta(t) - \frac{1}{2}H_1\beta(t)^2 - e^{-u}\Lambda_1(\theta(\delta\beta(t)) - 1)$$
$$\partial_t \alpha(t) = v\Lambda_0 - K_0\beta(t) - \frac{1}{2}H_0\beta(t)^2 - e^{-u}\Lambda_0(\theta(\delta\beta(t)) - 1)$$

- The transform of N is given by

$$\begin{aligned}\mathbb{E}(e^{-u(N_T-N_t)} \,|\, \mathcal{F}_t) &= \mathcal{L}^u(\psi(u),t,T) \\ &= \exp(\alpha(u,\psi(u),t,T) + \beta(u,\psi(u),t,T)X_t)\end{aligned}$$

Example: Affine point process

A specific affine model

- The intensity has dynamics

$$d\lambda_t = \kappa(c - \lambda_t)\,dt + \sigma\sqrt{\lambda_t}dW_t + \delta\,dL_t, \quad \lambda_0 > 0 \tag{1}$$

- Intensity responds to events; if $\delta = \sigma = 0$: Poisson process
- Intensity decays after an event; if $\kappa = \sigma = 0$: birth process
- Intensity diffuses between events; if $\sigma = 0$: Hawkes process
- Positive correlation between λ and the L^ks (position losses)

Example: Affine point process

L^k is uniform on $\{0.4, 0.6, 0.8, 1\}$, $c = \lambda_0 = 0.7$, $\delta = 1$, $\kappa = 5$

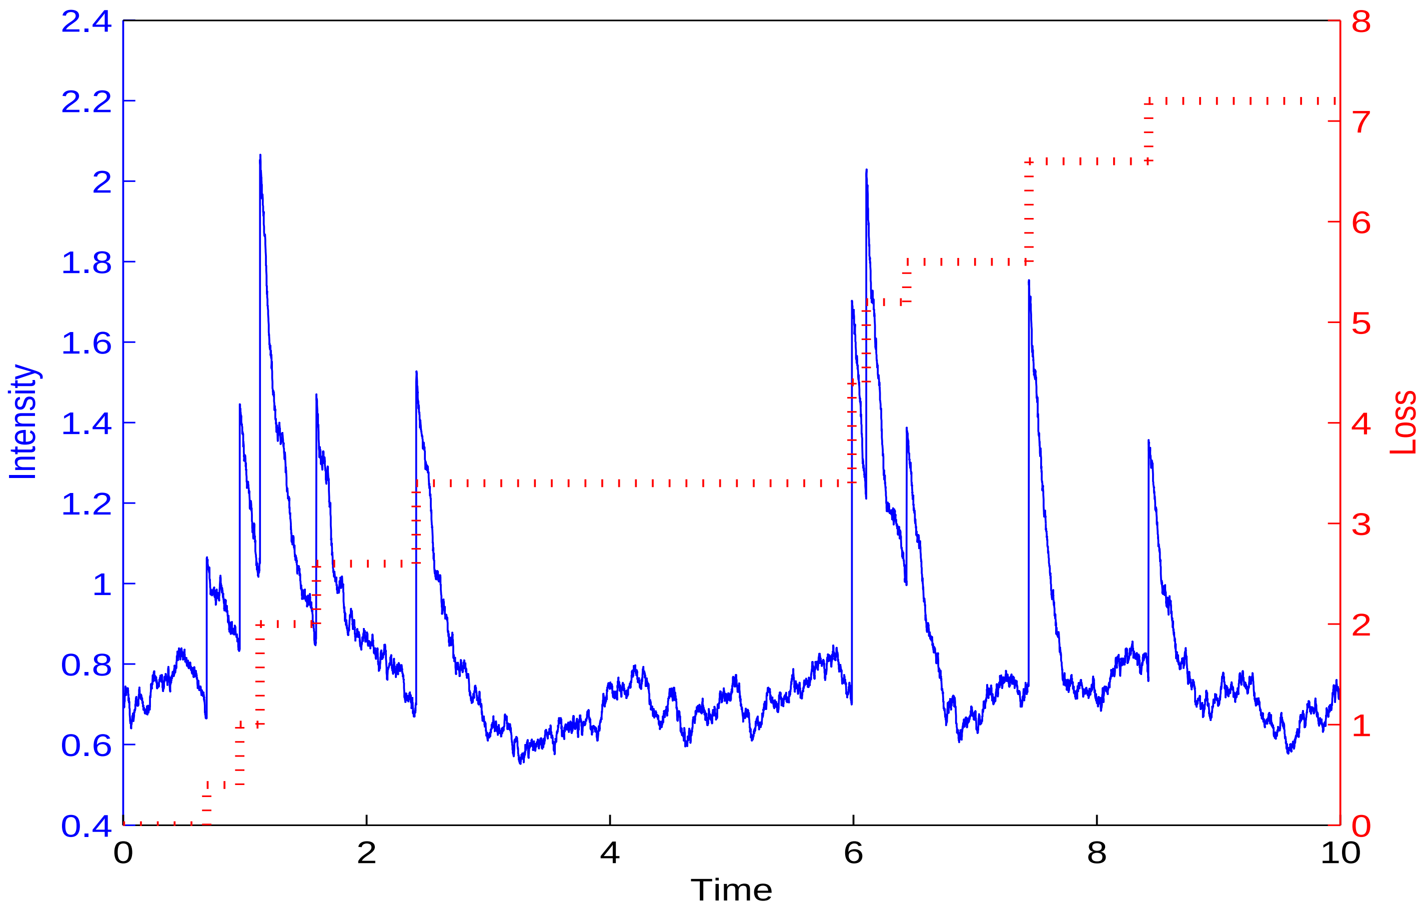

Example: Affine point process

Smoothed distribution of L_5, obtained by Fourier inversion

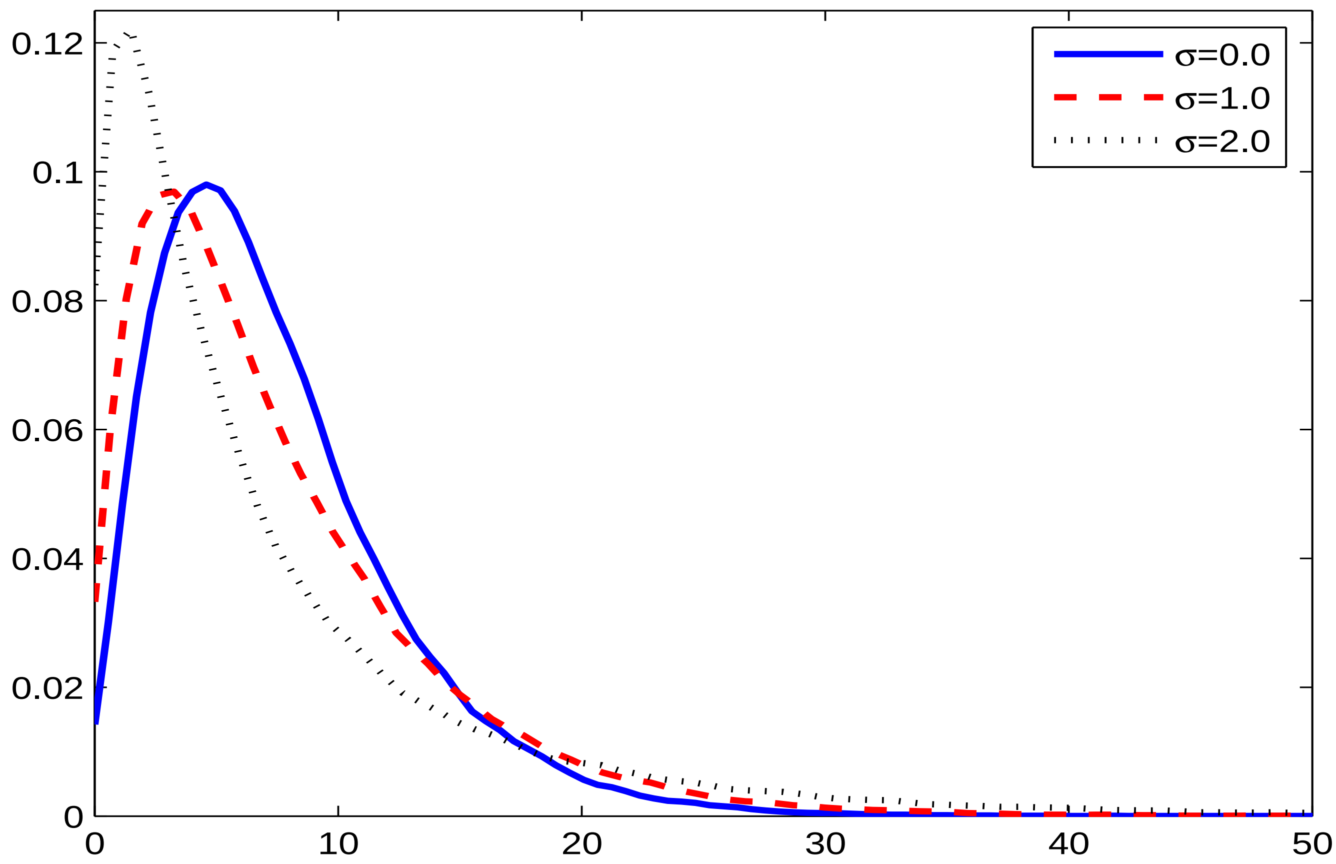

Point process simulation

Time change

- **Theorem.** (Meyer (1971)) If $A_t \to \infty$, then the variables

$$S^k = A_{T^k} = \int_0^{T^k} \lambda_s ds$$

  are the arrival times of a standard Poisson process in the time-scaled filtration defined by the stopping-time σ-fields $\mathcal{F}_{A_t}$

- Thus, any point process N can be represented as N^0_A, where N^0 is a standard Poisson process

- Analogous to the theorem of Dambis-Dubins-Schwartz, which states that any continuous local martingale V can be represented as $W_{[V,V]}$, where W is a standard Brownian motion

Point process simulation

Time change

- Thus, T^k is the hitting time of A to the random level $S^k = \mathcal{E}^1 + \cdots + \mathcal{E}^k$, where $\mathcal{E}^n \sim$ Exp(1):

$$T^k = A^{-1}_{S^k} = \inf\left\{t : \int_0^t \lambda_s ds \geq S^k\right\}$$

- To generate T^k using this representation, we need to approximate A on a discrete-time grid
  - Approximate λ on discrete-time grid, and then integrate
  - Exact sampling of λ at grid points is often possible (see Beskos & Roberts (2005), Chen (2009), Giesecke & Smelov (2010))

Point process simulation

Paths of λ_t, $A_t = \int_0^t \lambda_s ds$ and arrivals for the affine model (1). The values of λ_t are sampled exactly from the non-central chi-squared law.

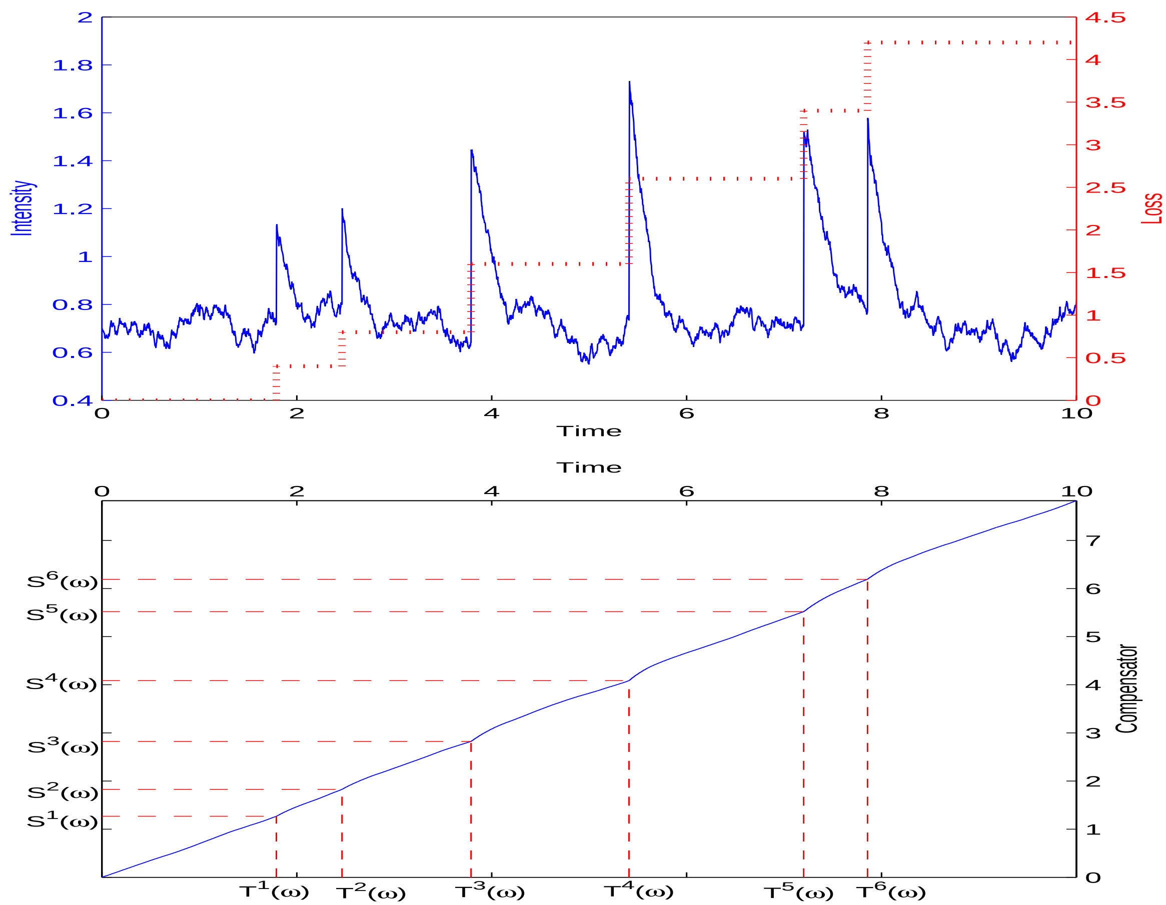

Point process simulation

- Due to approximation of A, simulation estimators are biased
  - Magnitude of bias is hard to quantify (work in progress)
  - Difficult to obtain valid confidence intervals
  - Difficult to determine optimal allocation of comp. budget
- **Exact methods** eliminate the need to discretize λ
  - Thinning scheme: Lewis & Shedler (1979), Glasserman & Merener (2003) for bounded λ and Giesecke, Kim & Zhu (2011) for general case
  - Inverse transform scheme: Giesecke & Kim (2007)
  - Mimicking Markov chain method: Giesecke, Kakavand, Mousavi & Takada (2010)
  - Projection method: Giesecke, Kakavand & Mousavi (2010)
  - Rejection method: Giesecke & Smelov (2010)

Point process simulation

Distribution of L_5 (smoothed) for the affine model (1), 1m paths, discretization 1 day, $\kappa = c = \sigma = \delta = \lambda_0 = 1$, $L^k \sim U\{0.4, 0.8\}$

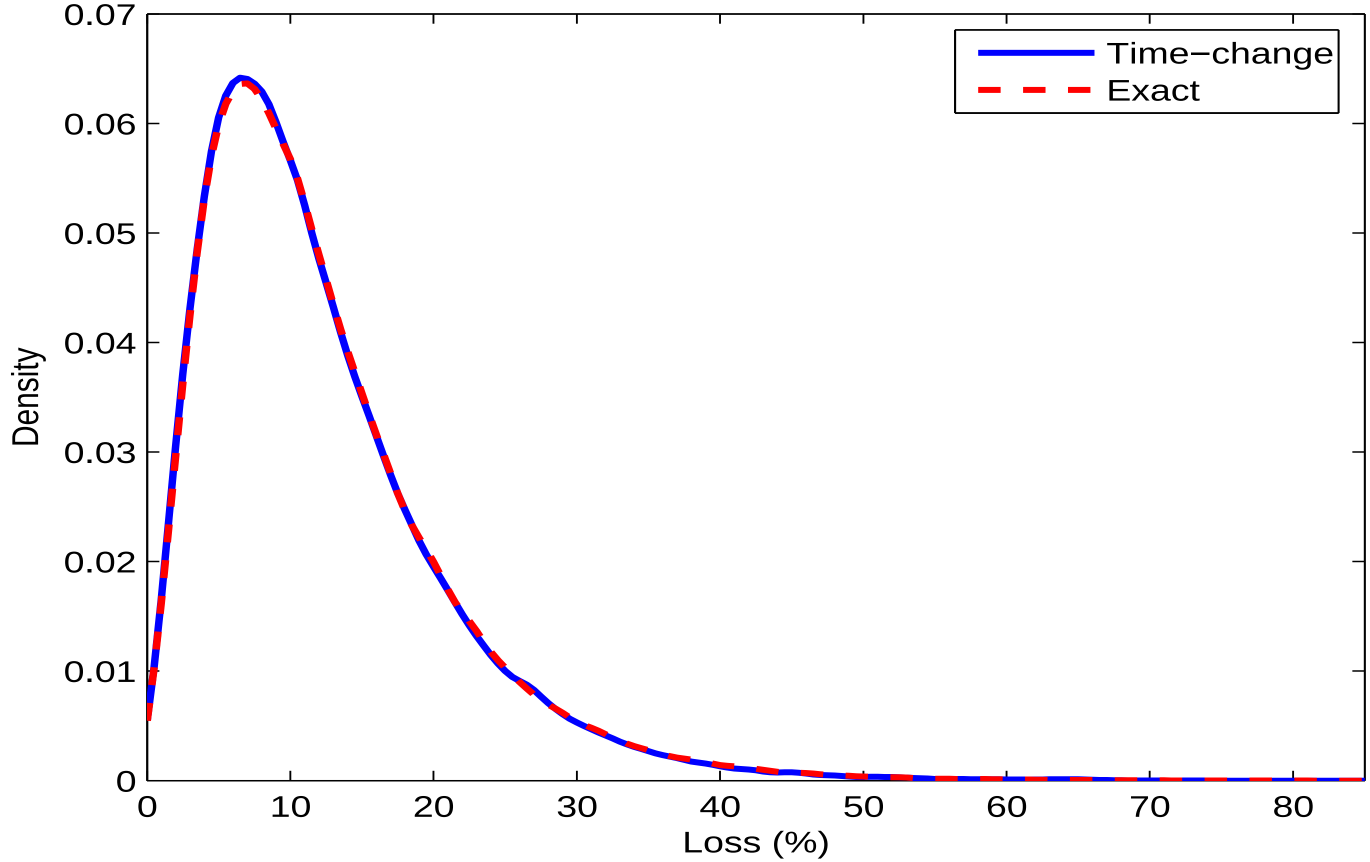

Point process simulation

Log-distribution of L_5 (smoothed) for the affine model (1), 1m paths, discretization 1 day, $\kappa = c = \sigma = \delta = \lambda_0 = 1$, $L^k \sim U\{0.4, 0.8\}$

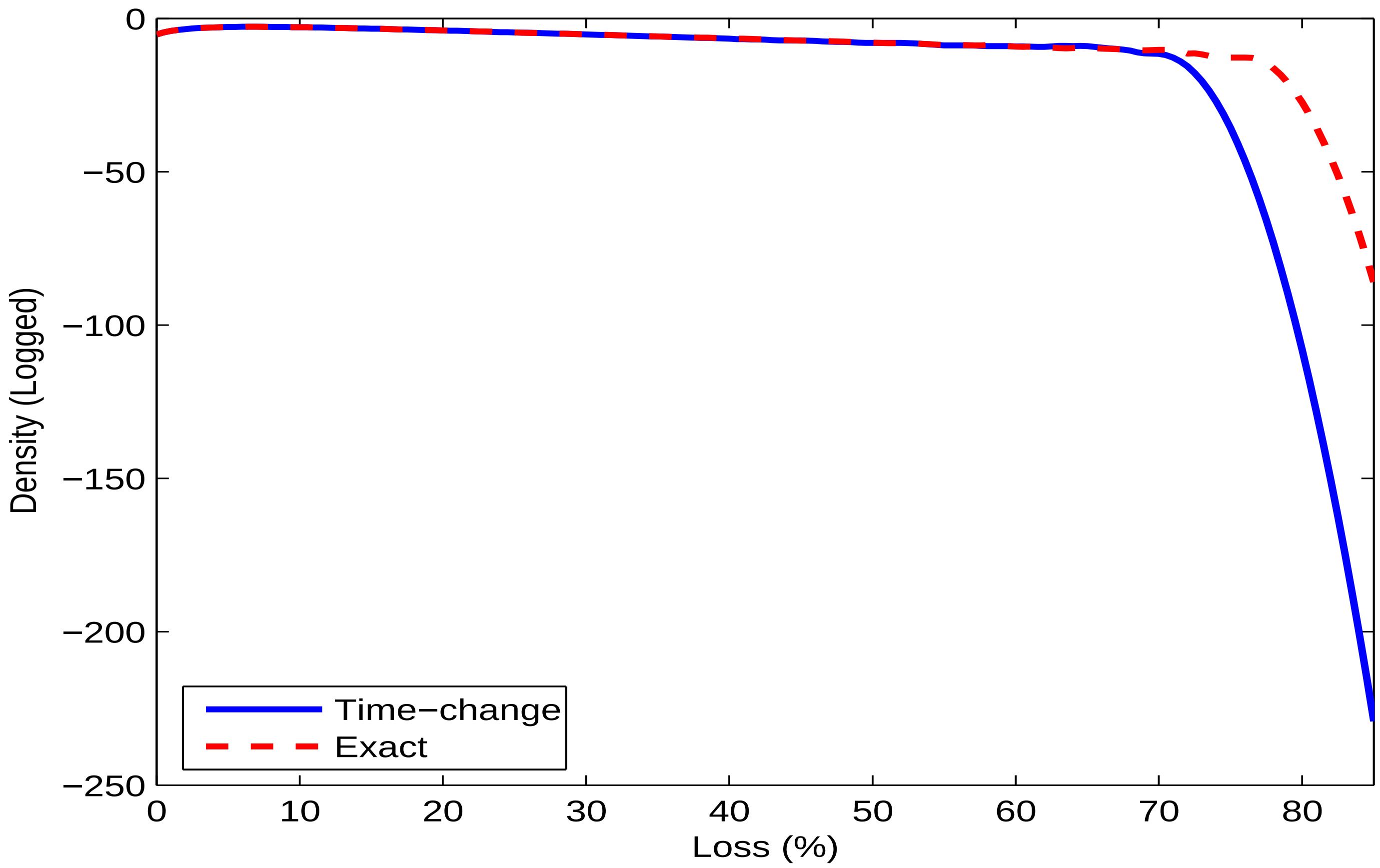

Point process simulation

Thinning

- **Proposition.** Let H be a predictable process with values in $[0, 1]$. Select an event time T^k of N with probability H_{T^k}. If N has intensity λ, then the counting process of the selected times (the thinned process) has intensity $H\lambda$.
- Justifies a scheme for generating N from a counting process $\overline{N}$ with intensity $\overline{\lambda}$ satisfying $\overline{\lambda}_t \geq \lambda_t$ almost surely
  - Sample event times of $\overline{N}$
  - Select an event time τ_k with probability $\lambda_{\tau_k-}/\overline{\lambda}_{\tau_k-}$

Point process simulation

Thinning

- Exact sampling of intensity values often feasible (see Beskos & Roberts (2005), Chen (2009))
- Tightness of bond determines efficiency
- Deterministic bound most convenient: then the dominating process $\overline{N}$ is a Poisson process which is easy to generate
  - Piece-wise deterministic bounds
- However, most standard models of λ are not almost surely bounded

Point process simulation

Mimicking Markov chain

- **Proposition.** Let π be a Markov chain on $[0, \infty)$ that takes values in $\mathbb{N}$, starts at 0, and has transition rate

$$\lambda(t, n) = \mathbb{E}(\lambda_t \mid N_t = n)$$

  at time t from state n to $n + 1$ and 0 for other transitions. Then $\pi_t = N_t$ in distribution, for each $t \geq 0$.

- Analogous to a result of Gyöngy (1986) for a continuous process with a general diffusion coefficient (Dupire's formula)

- **Proposition.** $\lambda(t, n) = \partial_t \mathbb{P}(N_t > n) / \mathbb{P}(N_t = n)$

- If λ can be computed and is bounded, can simulate π by thinning, and obtain unbiased estimator of $\mathbb{E}(f(\pi_t)) = \mathbb{E}(f(N_t))$ for any integrable function f on $\mathbb{N}$

Point process simulation

Projection onto a sub-filtration

- We project N onto its own filtration $\mathbb{G} = (\mathcal{G}_t)_{t\geq 0}$, the smallest sub-filtration of $\mathbb{F}$ that is compatible with N
- The $\mathbb{G}$-intensity h of N is given by the **optional projection** of λ onto $\mathbb{G}$; this is a unique $\mathbb{G}$-adapted process such that $\mathbb{E}(\lambda_T 1_{\{T<\infty\}} \,|\mathcal{F}_T) = h_T 1_{\{T<\infty\}}$ for every stopping time T
- Can show that h takes the form

$$h_t = \mathbb{E}(\lambda_t \,|\, \mathcal{G}_t) = \sum_{n\geq 0} h_n(t) 1_{\{N_t = n\}}$$

almost surely, where h_n is the $\mathcal{G}_{T_n}$-measurable function defined by

$$h_n(t) = \frac{\mathbb{E}(\lambda_t 1_{\{N_t=n\}} \,|\, \mathcal{G}_{T_n})}{\mathbb{P}(N_t = n \,|\, \mathcal{G}_{T_n})}, \quad t \geq T_n$$

Point process simulation

Projection onto a sub-filtration

- **Proposition.** If the intensity λ satisfies $\mathbb{E}(\int_0^t \lambda_s ds) < \infty$ for all $t \geq 0$, then, for $t \geq T_n$,

$$\mathbb{P}(T_{n+1} > t \,|\, \mathcal{G}_{T_n}) = \exp\left(- \int_{T_n}^t h_n(s) ds \right).$$

- Given $\mathcal{G}_{T_n}$, the waiting time to next event time T_{n+1} is equal in distribution to the first jump time of a $\mathbb{G}$-Poisson process started at T_n with intensity h_n

- The T^n can be generated sequentially in the filtration $\mathbb{G}$, using the thinning scheme or inverse transform method
  - The h_n can be computed for may standard models using filtering methods, see Kliemann, Koch & Marchetti (1990)

Point process simulation

Path of affine model (1): $d\lambda_t = \kappa(c - \lambda_t)dt + \sigma\sqrt{\lambda_t}dW_t + \delta dL_t$

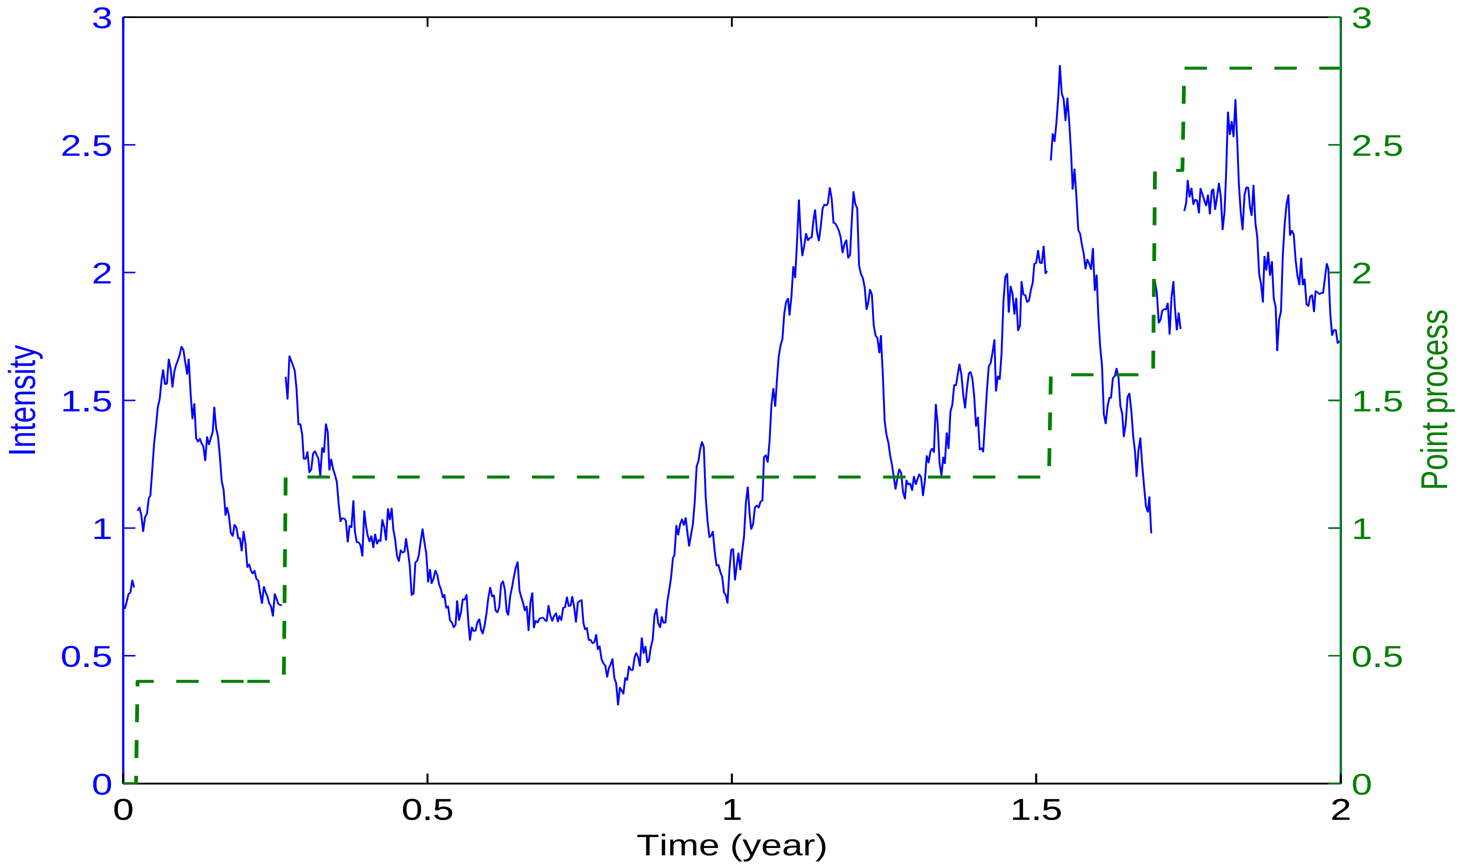

Point process simulation

Path of projected intensity h

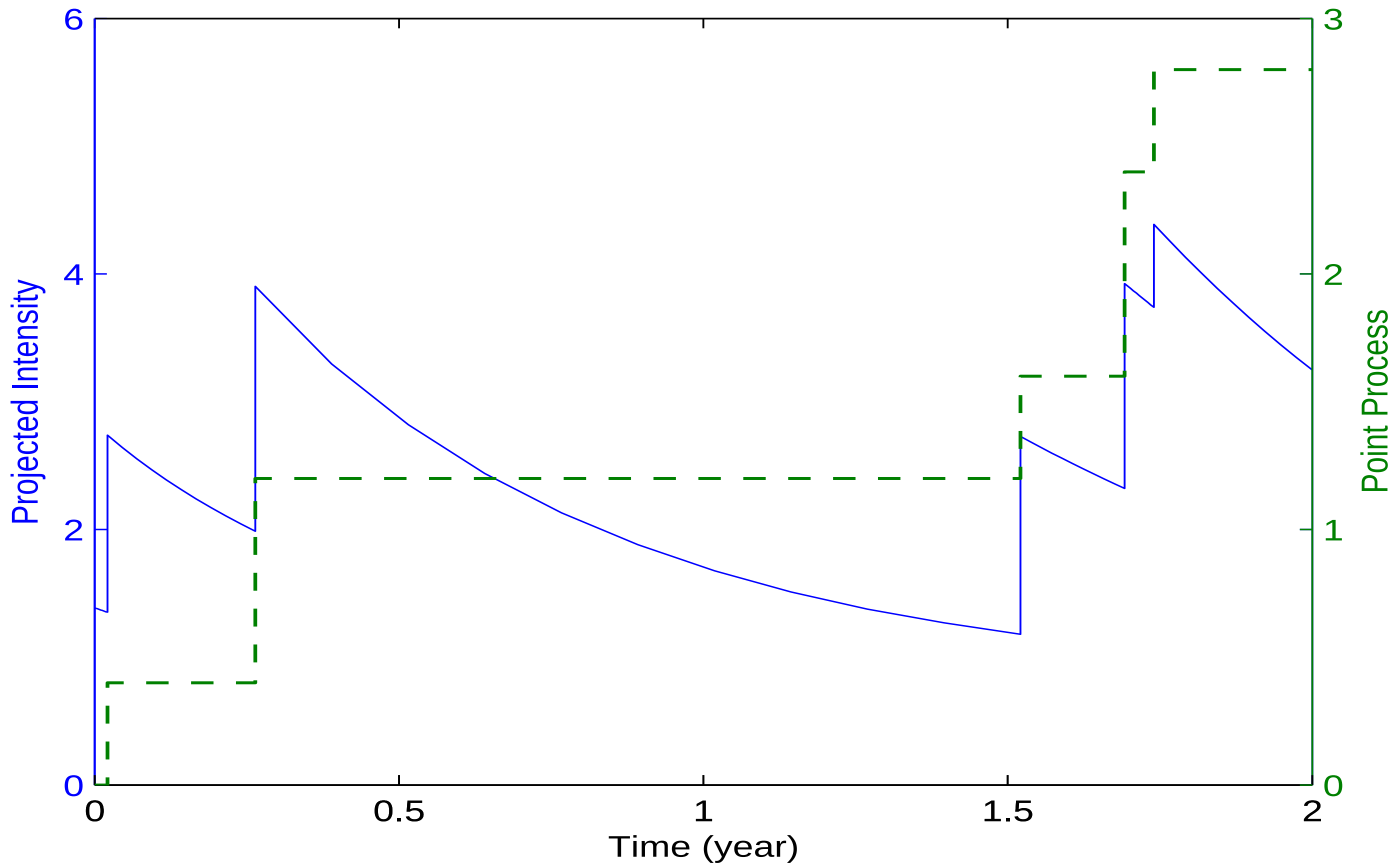

Parameter estimation

Setting

- Suppose $\lambda = \Lambda(N, X; \alpha)$ for some explanatory factor X and a parameter $\alpha \in \Theta^N \subset \mathbb{R}^p$ to be estimated
- Suppose X has transition density $p_t(\cdot; \gamma)$ for a parameter $\gamma \in \Theta^X \subset \mathbb{R}^q$ to be estimated
- The data available for the estimation of (α, γ) are a realization of the random variable $\mathcal{R} = (\mathcal{R}_X, \mathcal{R}_N)$ where

$$\mathcal{R}_N = (N_\tau, T_1, \ldots, T_{N_\tau}) \in R_N$$
$$\mathcal{R}_X = (X_{t_1}, \ldots, X_{t_m}) \in R_X$$

for $\tau > 0$ and $0 \leq t_1 < t_2 < \cdots t_m \leq \tau$

Parameter estimation

Point process likelihood

- $\mathbb{P}_{\mathcal{R}}$ is the law of $\mathcal{R}$ on $(\mathbb{R}^m \times R_N$, $\mathcal{B}^m \times \sigma(R_N))$
- $\mathbb{P}^*_N$ is the law on $(R_N, \sigma(R_N))$ of a std. Poisson process on $[0, \tau]$
- $\mathbb{L}^m$ is the Lebesgue measure on $(\mathbb{R}^m, \mathcal{B}^m)$
- The **likelihood function** $\mathcal{L}(\alpha, \gamma \,|\, \mathcal{R})$ is the Radon-Nikodym density of $\mathbb{P}_{\mathcal{R}}$ with respect to $\mathbb{P}^*_N \times \mathbb{L}$, evaluated at the data $\mathcal{R}$:

$$\mathcal{L}(\alpha, \gamma \,|\, \mathcal{R}) = \frac{d\mathbb{P}_{\mathcal{R}}}{d\mathbb{L}^m \times \mathbb{P}^*_N}$$

- The **maximum likelihood estimator** of (α, γ) is given by

$$(\hat{\alpha}, \hat{\gamma}) = \arg \max_{\alpha \in \Theta^N, \gamma \in \Theta^X} \mathcal{L}(\alpha, \gamma \,|\, \mathcal{R})$$

Parameter estimation

Point process likelihood

- Under technical conditions, $(\hat{\alpha}, \hat{\gamma})$ is consistent and asymptotically normal as $\tau \to \infty$, see Giesecke & Schwenkler (2011)

- To compute $\mathcal{L}(\alpha, \gamma \,|\, \mathcal{R})$, we introduce an auxiliary probability measure $\mathbb{P}^*$ on $\mathcal{F}_\tau$ with density (see Appendix)

$$\frac{d\mathbb{P}^*}{d\mathbb{P}} = Z_\tau = \exp\left(-\int_0^\tau \log(\lambda_{s-})dN_s - \int_0^\tau (1-\lambda_s)ds\right)$$

  N is a standard Poisson process on $[0, \tau]$ under $\mathbb{P}^*$ with law $\mathbb{P}^*_N$

- **Proposition.** Letting $\mathcal{L}^*_X(\gamma \,|\, \mathcal{R}_N)$ be the conditional $\mathbb{P}^*$-likelihood of $\mathcal{R}_X$ given $\mathbb{R}_N$, we have that

$$\mathcal{L}(\alpha, \gamma \,|\, \mathcal{R}) = \mathbb{E}^*(1/Z_\tau \,|\, \mathcal{R})(\alpha, \gamma) \times \mathcal{L}^*_X(\gamma \,|\, \mathcal{R}_N)$$

Parameter estimation

Discussion

- The ML problem can be decomposed into separate subproblems for α and γ only in special cases
  - If X is constant between observations (a std. assumption) or Λ does not depend on X, then Z_τ is measurable relative to $\sigma(\mathcal{R})$ and $\mathbb{E}^*(1/Z_\tau \,|\, \mathcal{R})(\alpha, \gamma) = 1/Z_\tau$ depends on α only
- In general, $\mathbb{E}^*(1/Z_\tau \,|\, \mathcal{R})$ is a **point process filter**
  - Filter is governed by Kushner-Stratonovic equation, see Kliemann et al. (1990)
  - Numerical methods: Giesecke & Schwenkler (2011)
- Extension to completely unobserved explanatory factors ("frailties"), and observation of mark variables

Goodness-of-fit testing

Time-scaling tests

- **Theorem.** (Meyer (1971)) If $A_t \to \infty$, then the variables

$$S^k = A_{T^k} = \int_0^{T^k} \lambda_s ds$$

  are the arrival times of a standard Poisson process in the time-scaled filtration defined by the stopping-time σ-fields $\mathcal{F}_{A_t}$

- Test whether the fitted intensity $\hat{\lambda}$ generates a time change that correctly transforms the observed T^k
  - Properties of time-scaled inter-arrival times
  - Properties of time-scaled event count

- In the presence of frailty, need fitted *filtered* intensity with respect to sub-filtration representing the observable information

Goodness-of-fit testing

Fitted compensator (time change)

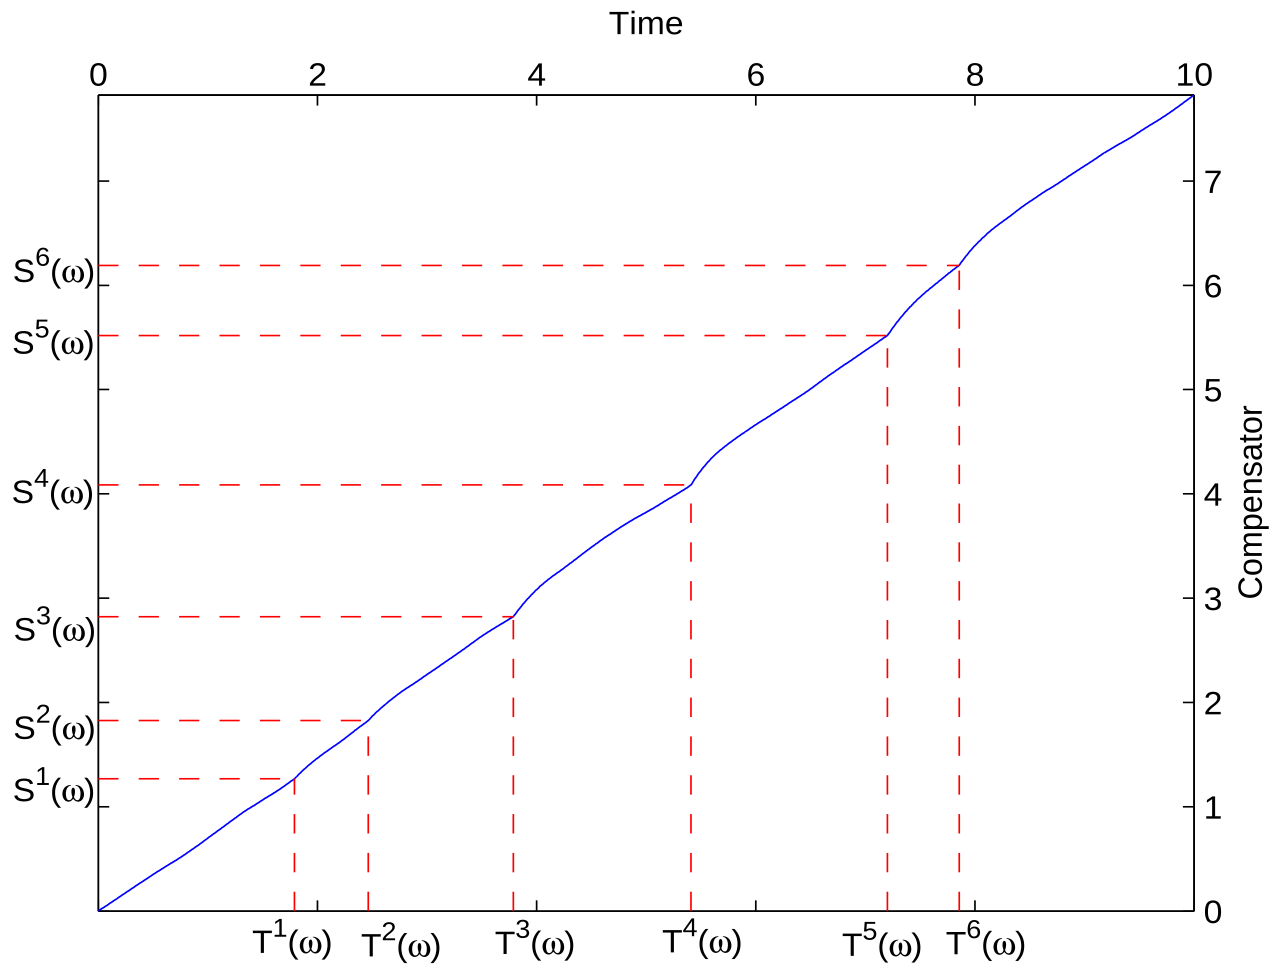

Applications

- Valuation of securities exposed to (correlated) default risk
  - Corporate bonds, CDS, forwards and options on CDS
  - Index and tranche swaps ("CDOs")
- Counterparty valuation adjustment (CVA)
- Risk premia extraction
- Estimation of portfolio credit risk

Valuation

Alternative formulations of the valuation problem

- Model N under the **actual measure** that represents the empirical likelihood of events, and then specify an equivalent change of measure to a risk-neutral measure (see Appendix)
  - Pricing and empirical time series applications that require risk premia specifications (as in Berndt, Douglas, Duffie, Ferguson & Schranz (2005), Eckner (2007), Azizpour, Giesecke & Kim (2011), and others)
- Model the point process N under a **risk-neutral measure**
  - Pricing applications

Single-name valuation

- Let $\mathbb{P}$ be a pricing measure relative to a constant risk-free rate r
- The issuer's default time is the first jump time T^1 of N; the corresponding default process is $N^1 = \min(N, 1)$
- The financial loss at default is modeled by an $\mathcal{F}_{T^1}$-measurable random variable ℓ^1, which is independent of T^1 and has $\ell = \mathbb{E}(\ell^1)$
  - Independence assumption can be relaxed
- At any time $t < \min(T, T^1)$, the firm's risk-neutral conditional survival probability satisfies

$$\begin{aligned}\mathbb{P}(T^1 > T \,|\, \mathcal{F}_t) &= \mathbb{P}(N_T = 0 \,|\, \mathcal{F}_t) \\ &= \lim_{u \uparrow \infty} \mathcal{L}^u(\psi(u), t, T) \\ &= \mathbb{E}^\infty(e^{-(A_T - A_t)})\end{aligned}$$

Single-name valuation

The limiting measure $\mathbb{P}^\infty$

- Since $\lim_{u\to\infty}\psi(u) = 1$, we have for the limiting density

$$Z_T^\infty = \lim_{u\to\infty} Z_T(u) = 1_{\{N_T=0\}} \exp(A_T)$$

- It follows that $\mathbb{P}(Z_T^\infty = 0) = \mathbb{P}(N_T > 0) > 0$, and therefore, the limiting measure $\mathbb{P}^\infty$ defined by Z_T^∞ is only absolutely continuous with respect to $\mathbb{P}$, but not equivalent

- Under $\mathbb{P}^\infty$, the compensator of N is null on $[0,T]$, so

$$\mathbb{P}^\infty(T^1 > T) = 1;$$

  $\mathbb{P}^\infty$ puts 0 measure on the event of default by T

Single-name valuation

The limiting measure $\mathbb{P}^\infty$

- To provide an example, suppose N is an affine point process driven by the SDE (1) for X, as introduced above
- Taking the limit, for $t < T^1 \wedge T$ we get the formula

$$\mathbb{P}(T^1 > T \,|\, \mathcal{F}_t) = \exp(a(t,T) + b(t,T)X_t)$$

  where $b(t) = b(t,T)$ and $a(t) = a(t,T)$ satisfy the ODEs

$$\partial_t b(t) = \Lambda_1 - K_1 b(t) - \frac{1}{2} H_1 b(t)^2$$
$$\partial_t a(t) = \Lambda_0 - K_0 b(t) - \frac{1}{2} H_0 b(t)^2$$

  with boundary conditions $b(T) = a(T) = 0$

- The feedback from N to λ, which comes from the jump term L in (1), is irrelevant for the survival probability

Corporate zero bond

A zero coupon bond with unit face value and maturity T pays

- The **face value** 1 at $T < T^1$
  This has value $F(t,T)\mathbb{P}(T^1 > T \mid \mathcal{F}_t)$, where $F(t,T)$ is the price of a unit face value, T-maturity zero coupon government bond
- The **recovery** $(1-\ell^1)$ of face value at $T^1 \leq T$
  This has value

$$\mathbb{E}(F(t,T^1)(1-\ell^1)N^1_T \mid \mathcal{F}_t) = (1-\ell)R_t(T)$$

where $R_t(T)$ is the pre-default value of a unit recovery payment at $T^1 \leq T$. By Stieltjes integration by parts

$$R_t(T) = \mathbb{E}\left(\int_t^T F(t,s)dN^1_s \,\middle|\, \mathcal{F}_t \right)$$
$$= F(t,T)\mathbb{P}(T^1 \leq T \mid \mathcal{F}_t) + r\int_t^T F(t,s)\mathbb{P}(T^1 \leq s \mid \mathcal{F}_t)ds$$

Corporate coupon bond

A corporate coupon bond with unit face value, annualized coupon rate c, coupon dates (t_m) and maturity T pays

- The **coupon** cC_m at each $t_m < T^1$ where C_m is the day count fraction for period m ($=$ portfolio of zero-recovery zero bonds)
- The **face value** 1 at $T < T^1$ ($=$ zero-recovery zero bond)
- The **recovery** $(1-\ell^1)$ of face value at $T^1 \leq T$
- The **accrued coupon** $\frac{T^1 - t_{m-1}}{\Delta_m} cC_m$ at T^1 if $t_{m-1} < T^1 \leq t_m$, where $\Delta_m = t_m - t_{m-1}$

Its pre-default value is given by

$$F(t,T)\mathbb{P}(T^1 > T \,|\, \mathcal{F}_t) + cV_t(T) + (1-\ell)R_t(T)$$

where $V_t(T)$ is the *risky DV01*, the pre-default value of a unit stream at coupon times (t_m) until $\min(T^1, T)$ plus any accruals

Credit swap

A credit swap with unit notional, annualized swap spread S, premium payment dates (t_m) and maturity T is a bilateral contract in which

- The **protection seller** pays the default loss ℓ^1 at $T^1 \leq T$
  This has pre-default value $D_t = \ell R_t(T)$

- The **protection buyer** pays the swap spread SC_m at each $t_m < T^1$ plus any accruals (assuming 0 points upfront)
  This has pre-default value $P_t(S) = SV_t(T)$

The fair spread S equates the values of the default and premium legs. Since there is no cash flow at inception, the fair swap spread at inception date t is the solution $S = S_t(T)$ to the equation $D_t = P_t(S)$:

$$S_t(T) = \ell R_t(T)/V_t(T)$$

Credit swap

Discussion

- $S_t(T)$ can be expressed in terms of $\mathbb{P}(T^1 > s \,|\, \mathcal{F}_t)$ for various $s \leq T$, and the expected loss ℓ
- The formula for $S_t(T)$ ignores **counterparty risk**, the risk that the protection seller fails with the reference entity
  - Need model for correlated default to address this issue
- The formula can be extended to incorporate a stochastic interest rate, and a loss at default correlated with the default time

Credit swap

Market spreads on 9/8/2008 for Comcast and Lehman Brothers

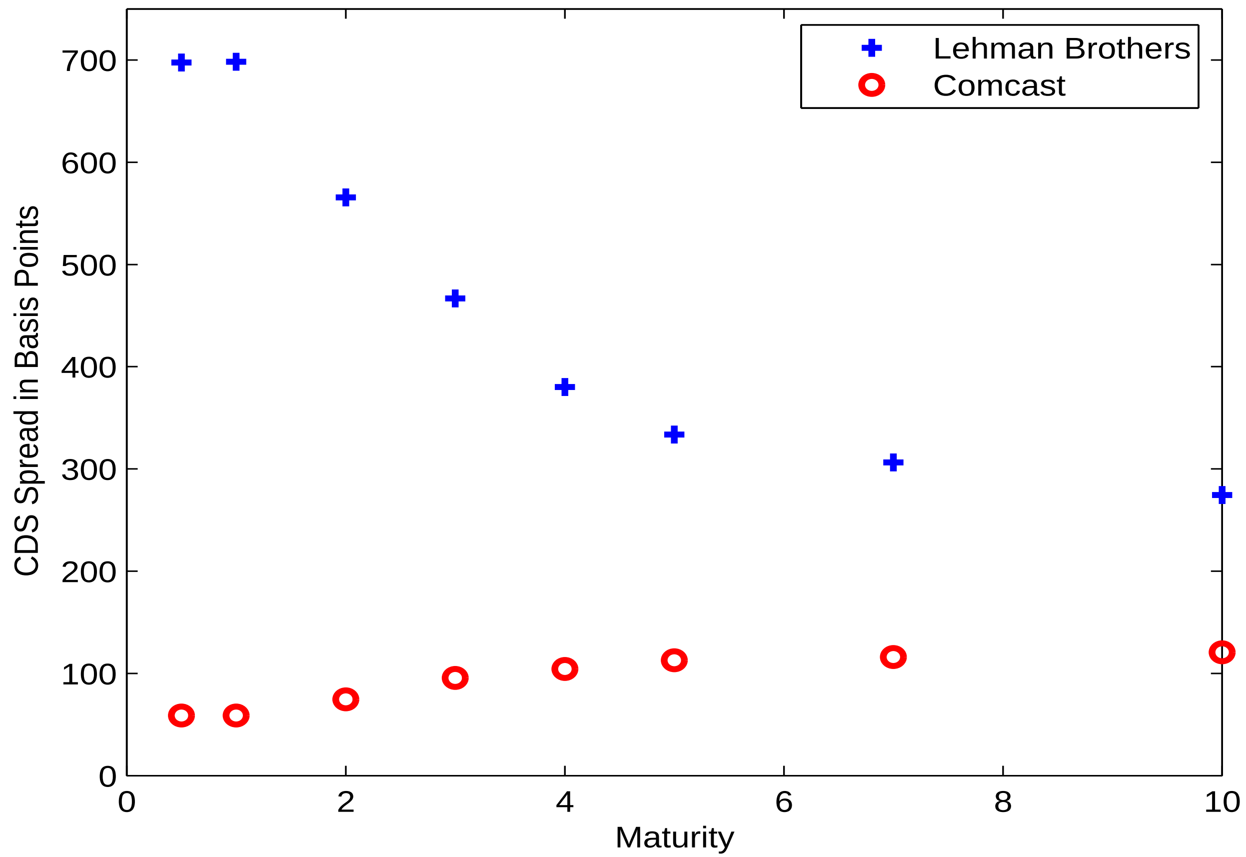

Credit swap

Market-implied, piece-wise constant intensity λ, assuming $\ell = 0.6$

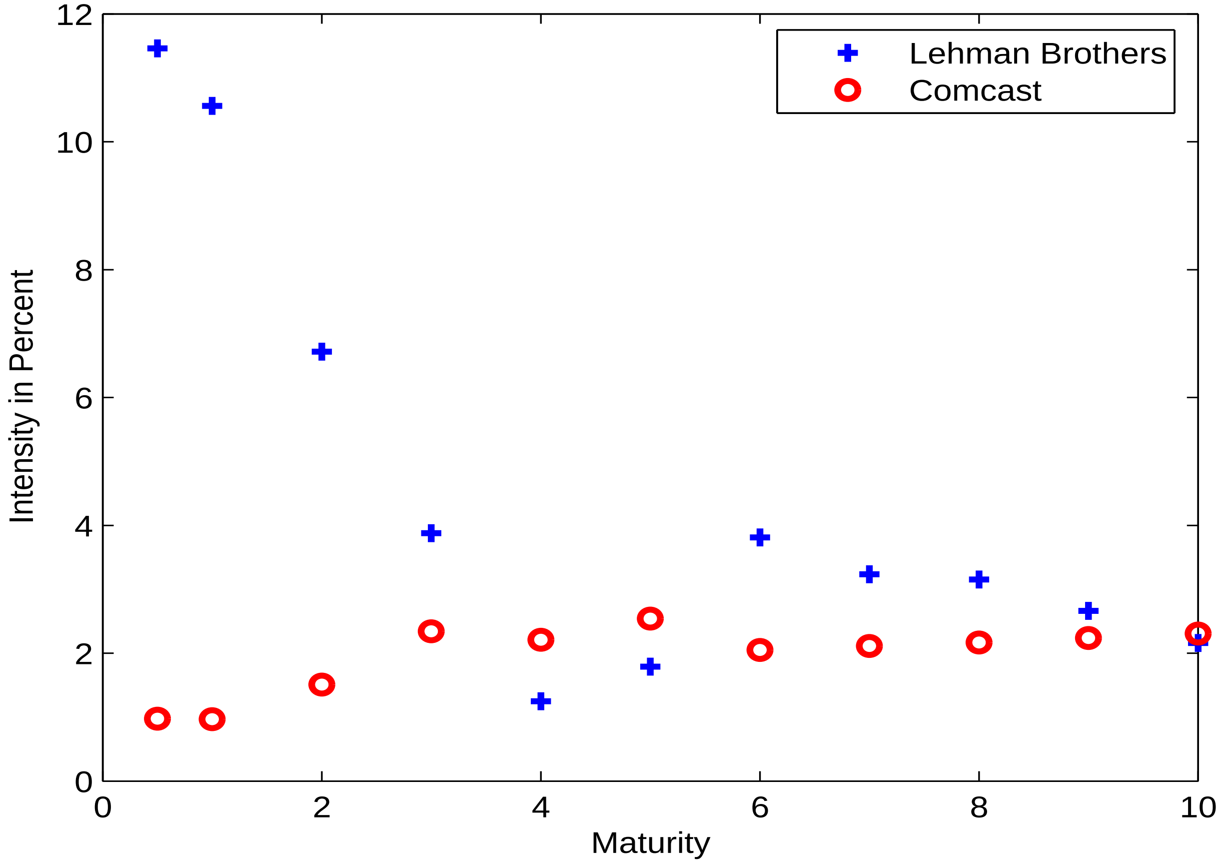

Credit swap

Market-implied survival probability $\mathbb{P}(T^1 > T)$, assuming a piece-wise constant intensity λ and $\ell = 0.6$

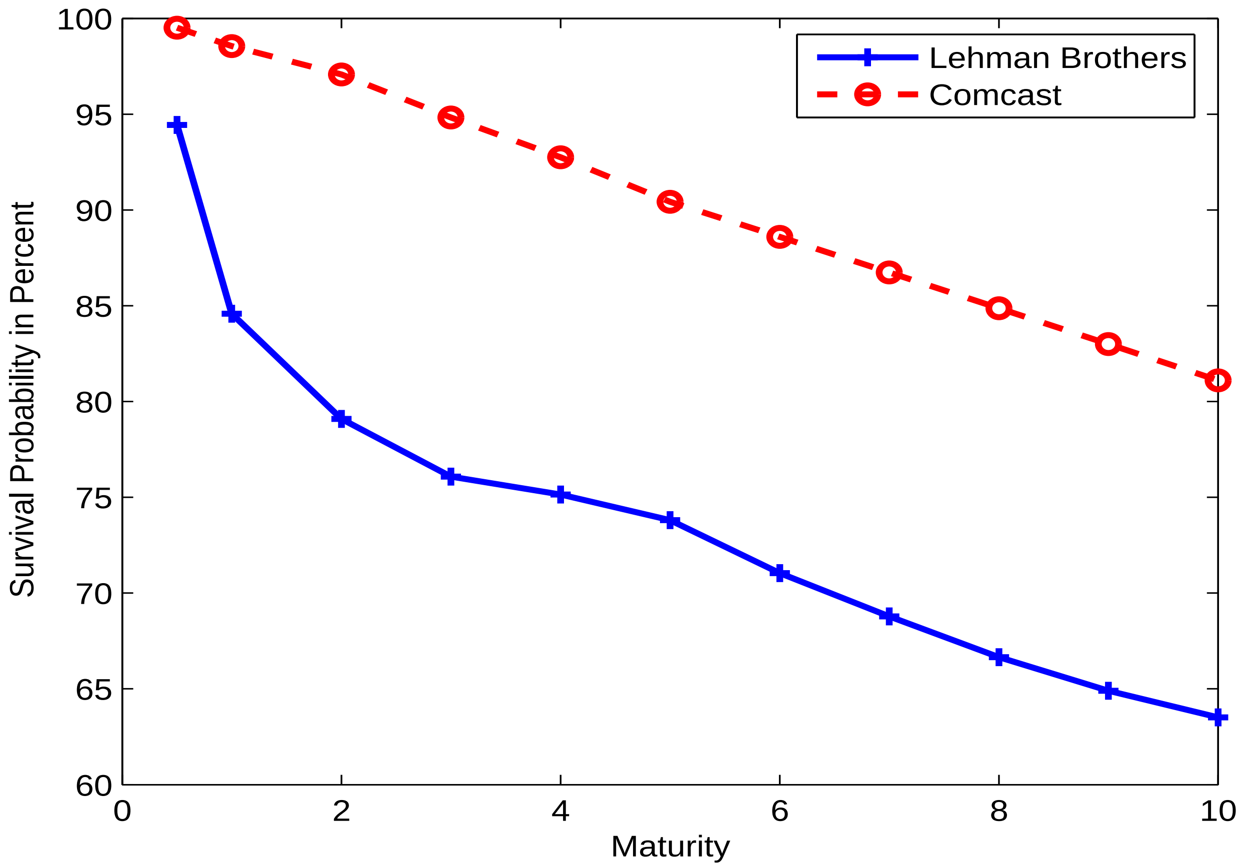

Credit swap

Bloomberg CDSW

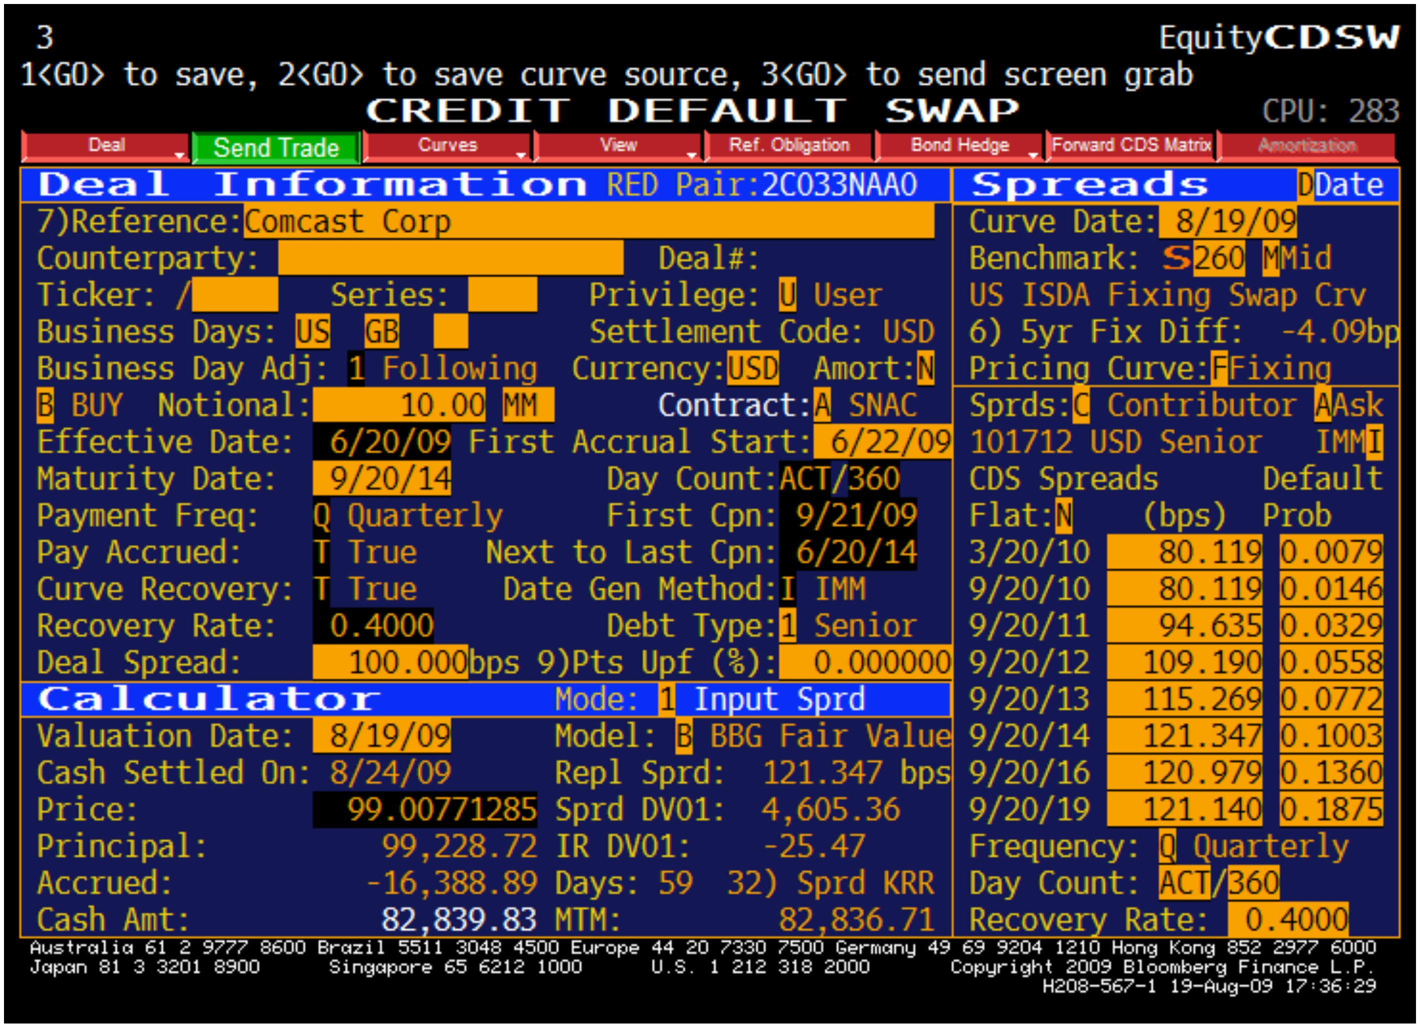

Credit swap

CDS mark-to-market

- Consider an investor who buys protection at $t = 0$ for the period $[0, T]$ at a swap spread of $S_0(T)$. The **mark-to-market value** of the position at time $t \geq 0$, when the market spread is $S_t(T)$, is

$$V_t(T)(S_t(T) - S_0(T)) = D_t - P_t(S_0(T))$$

- Important for CDS trading strategies
- To estimate the distribution of future mark-to-market of a swap position, we need a model of the evolution of spreads
  - Maximum likelihood estimation of λ using spread time series
  - Need to change the measure
- Forward and option on CDS: contracts on the mark-to-market

Forward credit swap

CDS forward

- In a forward CDS with maturity T, a party agrees to buy protection at a future time $U < T$ at the **forward spread** S_f
- The forward spread equals the spot spread S for $U = 0$
- The contract is knocked out at a default before U

Forward credit swap

CDS forward

- For a unit notional, the **premium leg** consists of a spread stream S_f during $[\min(\tau,U),\min(\tau,T)]$ plus any accruals so this leg has value

$$P_f(S_f) = S_f\big(V(T) - V(U)\big)$$

- The **default leg** covers the loss during $[U,T]$
  - It is equal to the difference between the default leg of a spot CDS for T and the default leg of a spot CDS for U
  - Its value must be equal to the difference between the values of the corresponding premium legs:

$$D_f = S(T)V(T) - S(U)V(U)$$

Forward credit swap

CDS forward

- The fair **forward spread** is the solution $S_f = S_f(U,T)$ to the equation

$$D_f = P_f(S_f)$$

so

$$S_f(U,T) = \frac{S(T)V(T) - S(U)V(U)}{V(T) - V(U)}$$

- We can estimate the forward CDS curve from the market spot CDS curve and the risky DV01s implied by that curve
  - Note that estimates depend on default timing model used

CDS option

CDS option

- A CDS option or swaption with maturity T imparts the right to enter into a CDS at a future time $U < T$ at a **strike spread** S
- A **payer swaption** gives the right to buy protection for S
  - Can include a knock-out provision: if the reference name defaults between the trade date and the expiry date U the option is canceled
  - Targets widening spreads: exercise if $S_U(T) > S$
- A **receiver swaption** gives the right to sell protection for S
  - Targets tightening spreads: exercise if $S_U(T) < S$

CDS option

Cancelable CDS

- This is a CDS with an embedded receiver swaption
- The protection buyer
  - Buys regular spot protection with maturity T
  - Buys a receiver swaption expiring at $U < T$ on the reference name, imparting the right to sell protection over $[U, T]$ at U for a fixed strike spread S (often equal to spot spread for T)
- The protection buyer exercises the option and thereby effectively closes out the spot CDS position at U if $S_U(T) < S$
  - Can then buy protection for $S_U(T)$ for the term $[U, T]$
- The protection seller receives the option premium in addition to the spot CDS spread

CDS option

Valuation

- In any knock-out forward transaction with maturity T, conditional on survival to time U, the payoff $Y_U(T)$ to a party at U is a function of the mark-to-market (**??**) of the position to the party

- In a forward CDS, the survival payoff to the protection buyer is

$$Y_U(T) = V_U(T)(S_U(T) - S_t(U, T))$$

while that to the protection seller is the negative of this value

- In a knock-out payer swaption, the survival payoff to the protection buyer is the positive part of the mark-to-market

$$Y_U(T) = V_U(T)(S_U(T) - S)^+$$

while for the protection seller it is $V_U(T)(S - S_U(T))^+$

Corporate defaults cluster

Value-weighted default rate 1865–2008, US nonfinancial (Source: Giesecke, Longstaff, Strebulaev & Schaefer (2011))

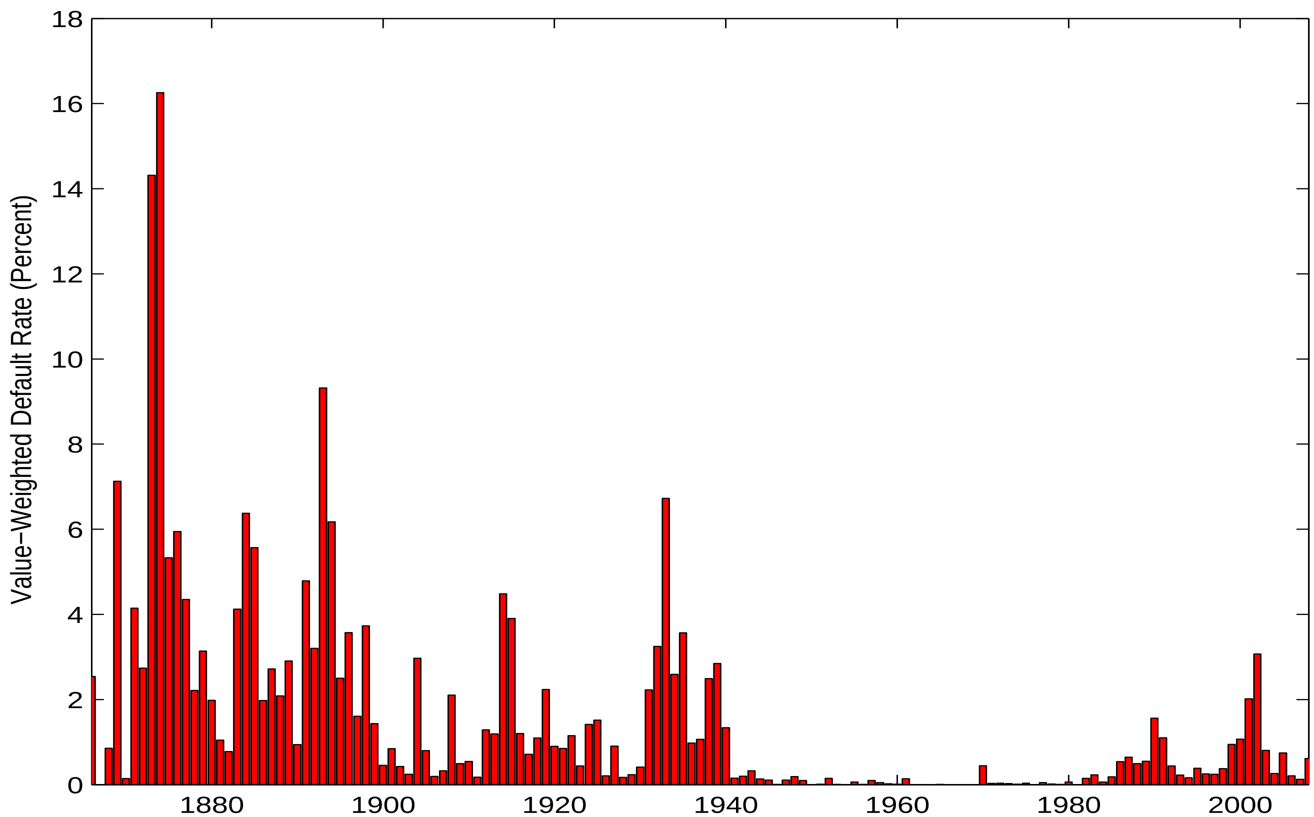

Corporate defaults cluster

Sources of clustering

- First, firms are exposed to **common or correlated risk factors**. The movements of these factors cause correlated changes in firms' conditional default rates (Duffie, Saita & Wang (2006))
- Second, some of the risk factors may be unobservable **frailties**. The uncertainty regarding the values of these factors has an influence on the conditional default rates of the firms that depend on the same frailties (Duffie, Eckner, Horel & Saita (2009))
- Third, a default may be **contagious**, and have a direct impact on the conditional default rates of other firms (Azizpour, Giesecke & Schwenkler (2009))

Portfolio derivatives

- They facilitate the transfer of the correlated default risk in a portfolio of reference names
  - There are multiple families of standard reference portfolios, called indices, including the CDX family (North American issuers) and the iTraxx family (European and Asian names)
- They are contingent claims on the portfolio loss point process $L = \sum_{k=1}^{N} \ell^k$ where ℓ^k is the loss at the kth default
- The reference pool often consists of single-name CDS with common notional that we normalize to 1, common maturity date T and common premium payment dates (t_m).

Portfolio derivatives are claims on pool loss

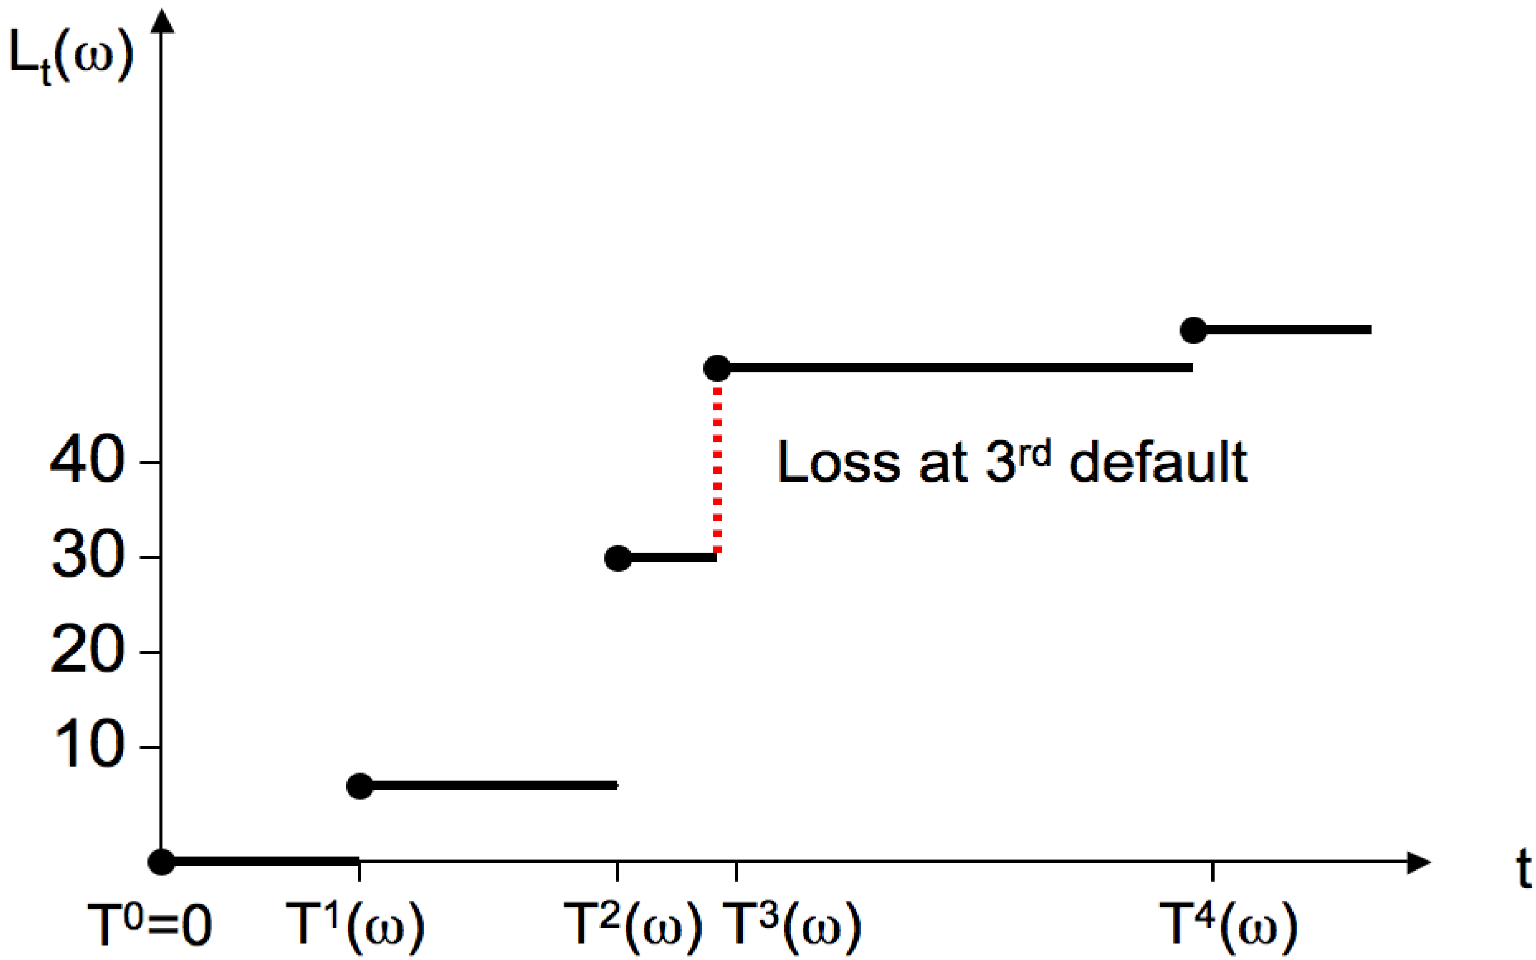

Index swap

In an index swap with swap spread S and maturity T,

- The **protection seller** covers portfolio losses as they occur, i.e. the increments of L; this leg has value

$$D_t = \mathbb{E}\left(\int_t^T F(t,s) dL_s \,\middle|\, \mathcal{F}_t \right)$$

$$= F(t,T)\mathbb{E}(L_T \,|\, \mathcal{F}_t) - L_t + r \int_t^T F(t,s)\mathbb{E}(L_s \,|\, \mathcal{F}_t)\, ds$$

- The **protection buyer** pays $SC_m(n - N_{t_m})$ at each date t_m; this leg has value

$$P_t(S) = S \sum_{t_m \geq t} F(t,t_m) C_m (n - \mathbb{E}(N_{t_m} \,|\, \mathcal{F}_t))$$

The spread at t is the solution $S = S_t(T)$ to the equation $D_t = P_t(S)$

Index swap default leg

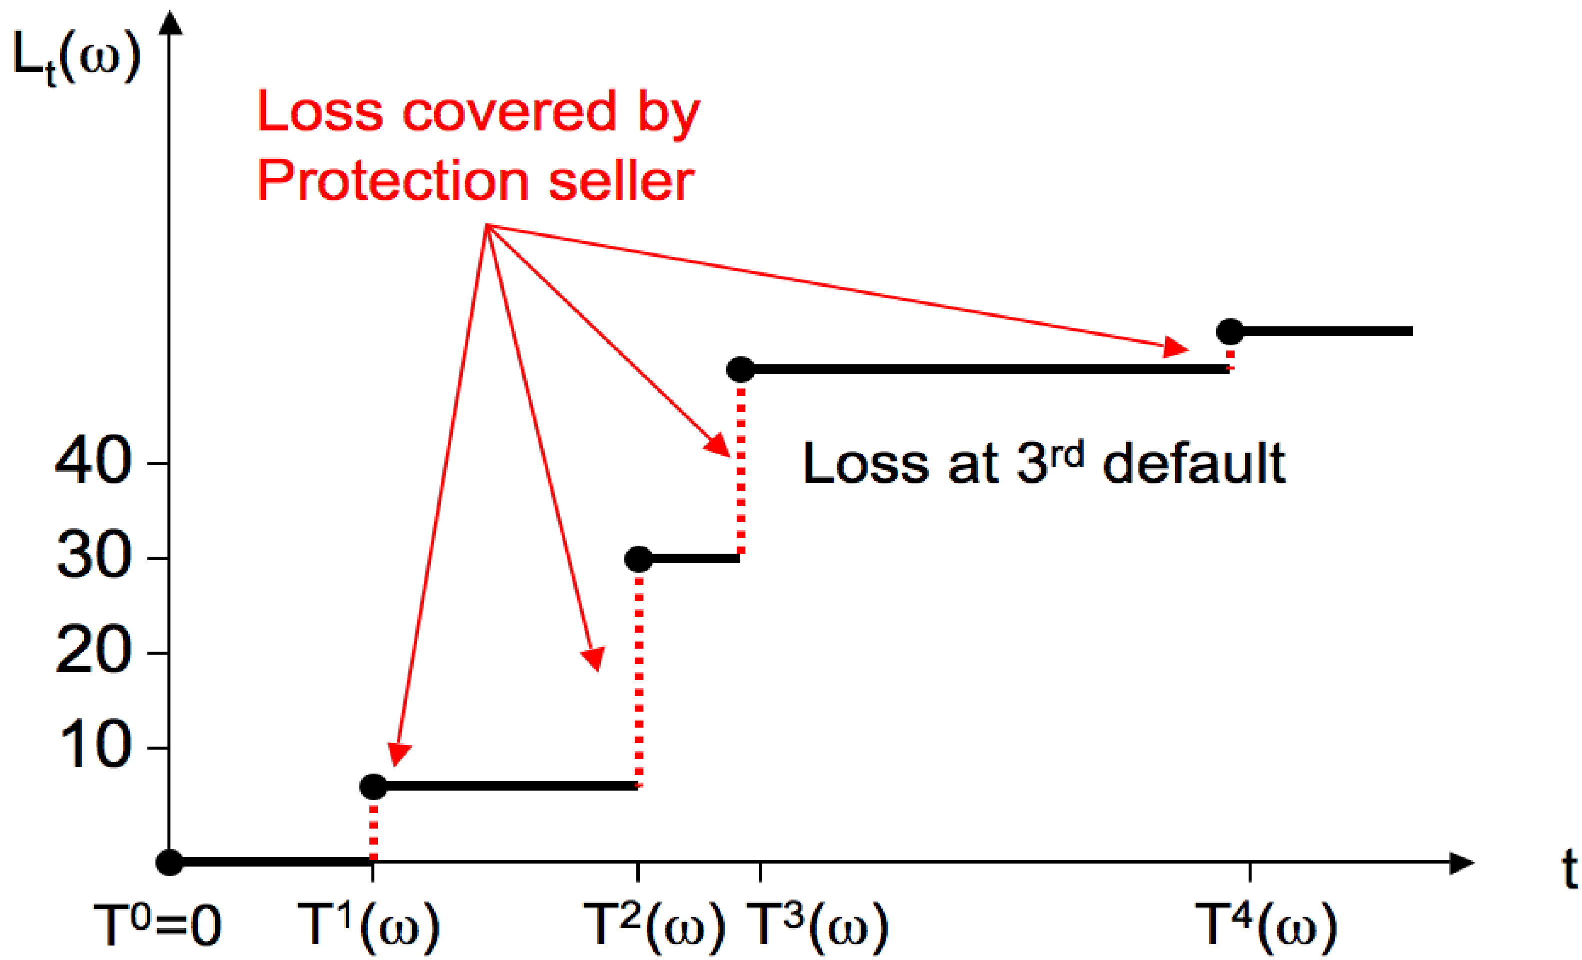

Index swap and single-name swaps

- The index is simply a portfolio of single name swaps with common maturity, notional $1/n$ and premium payment dates (t_m)
- The index default leg must equal the sum over index constituents of the single name premium legs $\frac{1}{n}S_t^i(T)V_t^i(T)$, where $S_t^i(T)$ is the single name spread and $V_t^i(T)$ is the single name DV01
- The index premium leg is the sum over constituents of $\frac{1}{n}V_t^i(T)$, i.e. the index DV01, times $S_t(T)$
- This gives the **intrinsic index spread** formula

$$S_t(T) = \frac{\sum_{i=1}^n S_t^i(T)V_t^i(T)}{\sum_{i=1}^n V_t^i(T)}$$

Tranche swap default leg

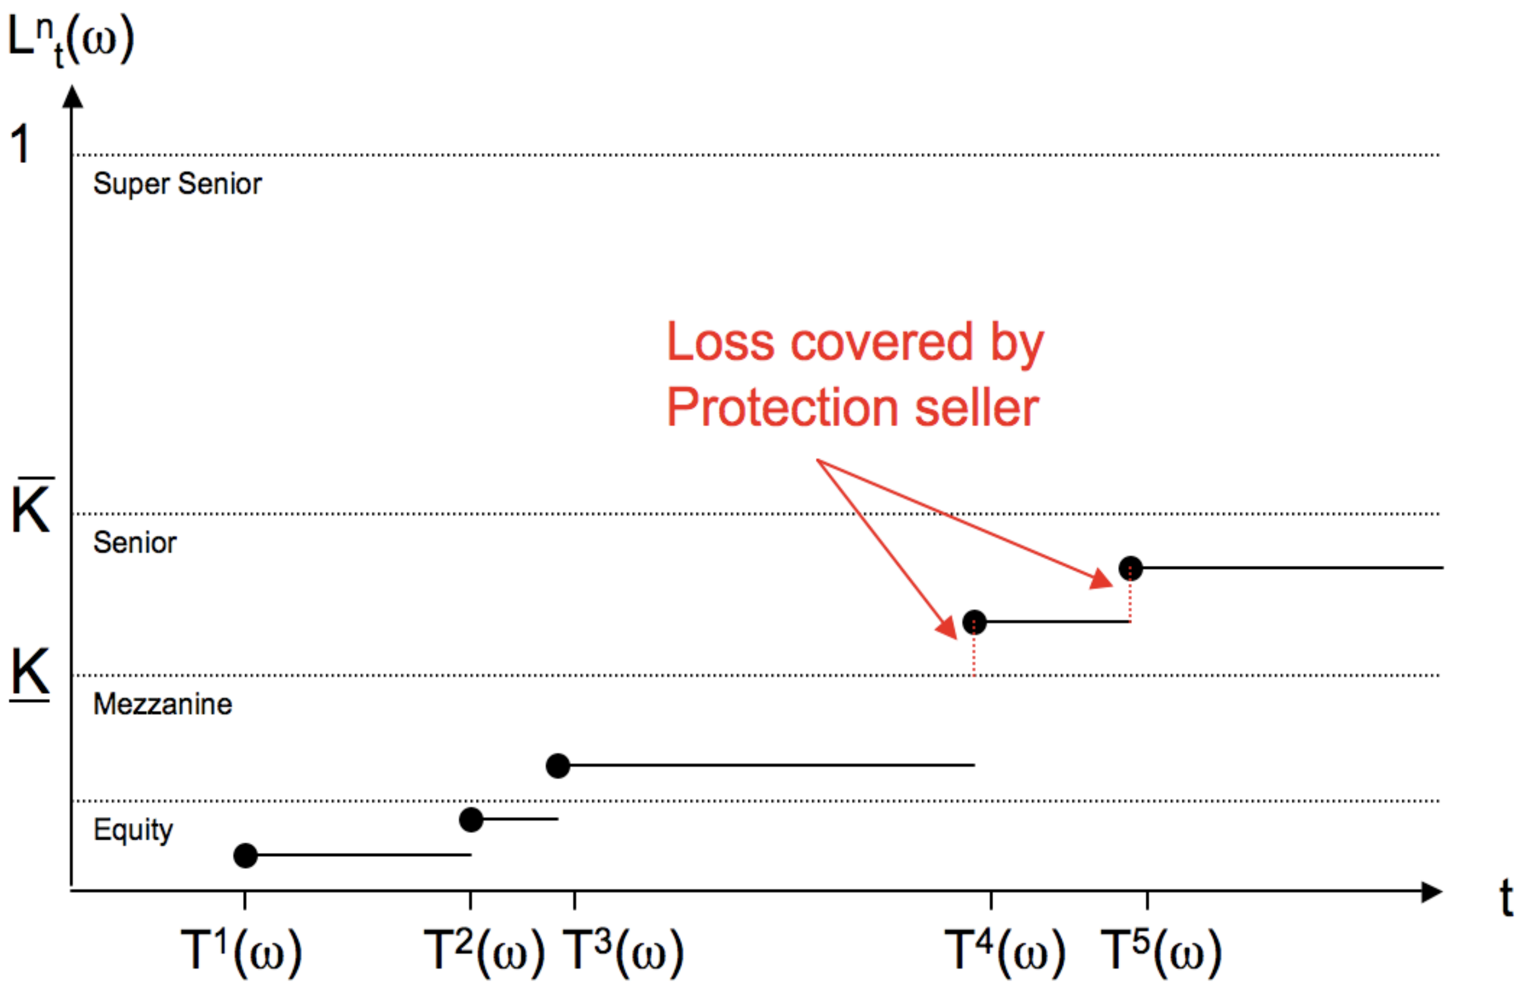

Tranche loss = call spread on portfolio loss

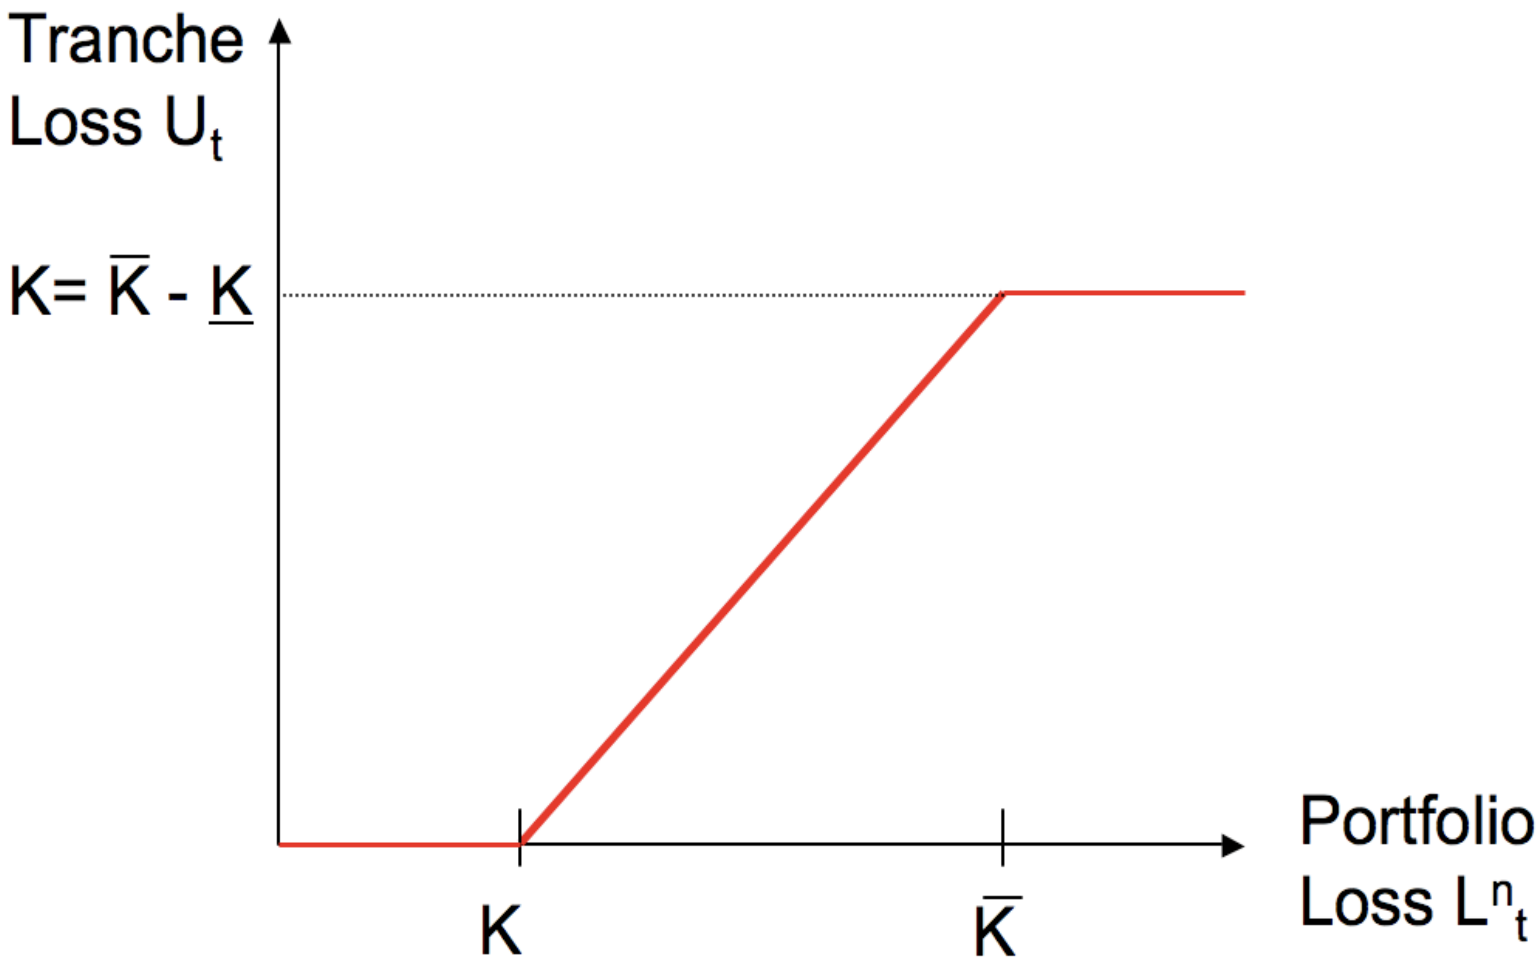

Tranche swap

In a tranche with lower attachment point $\underline{K} \in [0,1]$, upper attachment point $\overline{K} \in (\underline{K}, 1]$, upfront rate G, and swap spread S,

- The **protection seller** pays $U_t = (L_t - \underline{K}n)^+ - (L_t - \overline{K}n)^+$ at t; this has value

$$D_t(\underline{K}, \overline{K}) = F(t,T)\mathbb{E}(U_T \,|\, \mathcal{F}_t) - U_t + r\int_t^T F(t,s)\mathbb{E}(U_s \,|\, \mathcal{F}_t)\, ds$$

- The **protection buyer** pays GKn at inception t and $SC_m(Kn - U_{t_m})$ at each date t_m (assuming $\overline{K} < 1$)

$$P_t(\underline{K}, \overline{K}, G, S) = GKn + S\sum_{t_m \geq t} F(t, t_m)C_m(Kn - \mathbb{E}(U_{t_m} \,|\, \mathcal{F}_t))$$

For fixed G, the fair spread S is the solution $S = S_t(\underline{K}, \overline{K}, G, T)$ to $D_t(\underline{K}, \overline{K}) = P_t(\underline{K}, \overline{K}, G, S)$. For fixed S, the fair upfront rate G is the solution $G = G_t(\underline{K}, \overline{K}, S, T)$ to $D_t(\underline{K}, \overline{K}) = P_t(\underline{K}, \overline{K}, G, S)$.

Market calibration

- We calibrate the affine point process model (1) to index and tranche spreads/upfront rates on the CDX.NA.HY index observed on 5/11/2007, which has attachment points $0, 10, 15, 25, 35, 100\%$
- We assume the risk-neutral distribution of loss at default is uniform on $\{\ell_1, \ell_2\}$ with $0 < \ell_1 < \ell_2 < 1$, and we set the expected loss at default $\mathbb{E}(\ell^k) = \int z d\nu(z) = 0.6$
- The risk-neutral intensity λ is specified by the parameters $\lambda_0, c, \kappa, \sigma$ and δ
- The risk-free rate $r = 0.05$

Market calibration

- We fit the parameter vector $\theta = (\lambda_0, c, \kappa, \sigma, \delta, \ell_1)$ by numerically solving the optimization problem

$$\min_{\theta \in \Theta} \sum_i \left(\frac{\text{MarketMid(i)} - \text{Model(i,}\theta\text{)}}{\text{MarketAsk(i)} - \text{MarketBid(i)}} \right)^2 \tag{2}$$

  where $\Theta = [0,5]^3 \times [0,1] \times [0,5] \times [0.2, 0.6]$ and the sum ranges over the spread/upfront rate data points

- We use adaptive simulated annealing

- We analyze two model specifications:
  - Mod 1 is unrestricted
  - Mod 2 is restricted: the diffusive volatility $\sigma = 0$

Market calibration

5Y maturity: market data and fitting results

- Both models fit the data well; the basic affine model does better than the Hawkes model

	MarketBid	MarketAsk	Mod 1	Mod 2
0-10	70.50%	70.75%	71.11%	71.48%
10-15	34.25%	34.50%	32.85%	32.74%
15-25	316.00	319.00	316.80	311.43
25-35	79.00	81.00	81.47	77.34
Index	262.85	263.10	263.46	262.97
MinObj			41.63	60.41
AAPE			1.47%	2.24%

Market calibration

5Y maturity: initial and calibrated parameter values

- We ran several calibrations with different initial values; the values reported below generated the lowest objective function value
- The calibrated parameter values are very similar for the two models

	λ_0	c	κ	σ	δ	ℓ_1
Initial	2.50	2.50	2.50	0.50	2.50	0.40
Mod 1	0.70	1.61	2.62	0.62	2.99	0.24
Mod 2	0.75	1.60	2.58	0.00	2.94	0.24

Market calibration

5Y maturity: loss distribution implied by the calibrated basic affine model Mod 1

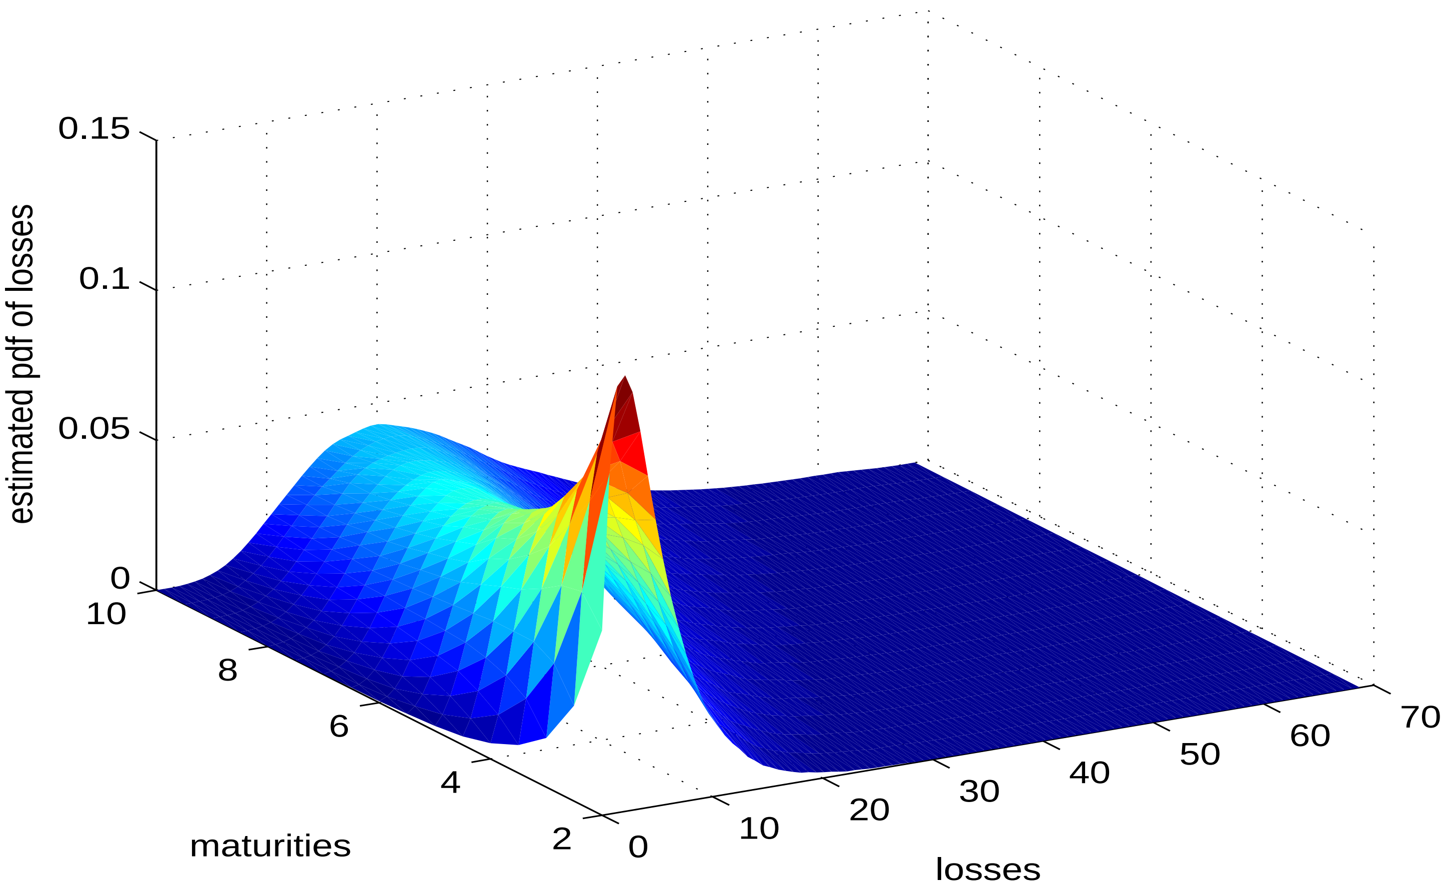

Market calibration

5Y maturity: parameter stability

- We re-calibrate each model specification at different dates starting with 5/11/07
- The initial values at a date after 5/11 are set to the optimal values from the previous date

Mod 1	Date	05/11	05/14	05/15	05/16	05/17
	MinObj	41.63	58.66	46.78	46.24	58.24
	AAPE	1.47%	1.70%	1.26%	1.32%	1.35%

Mod 2	Date	05/11	05/14	05/15	05/16	05/17
	MinObj	60.41	73.43	54.29	43.88	65.23
	AAPE	2.24%	2.04%	1.57%	1.45%	1.92%

Market calibration

5Y maturity: parameter stability Mod 1 (left) better than Mod 2

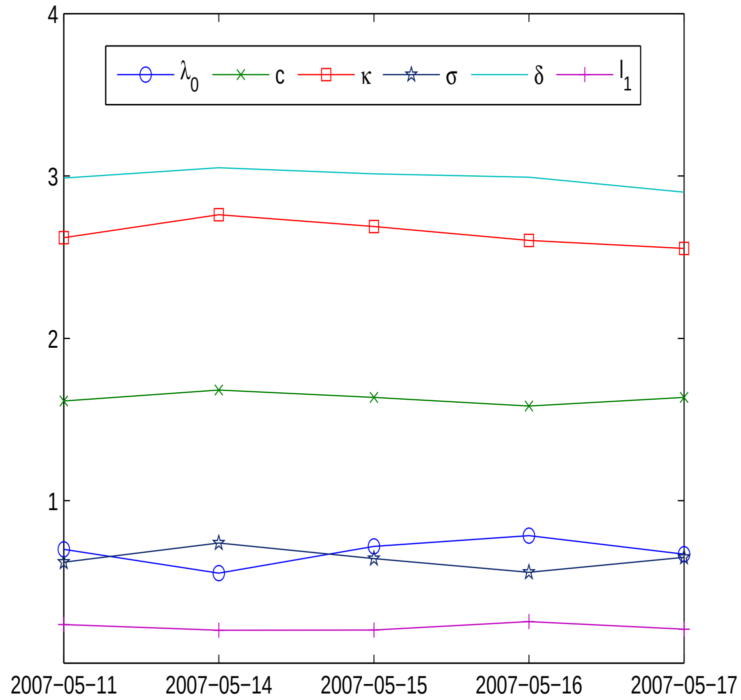

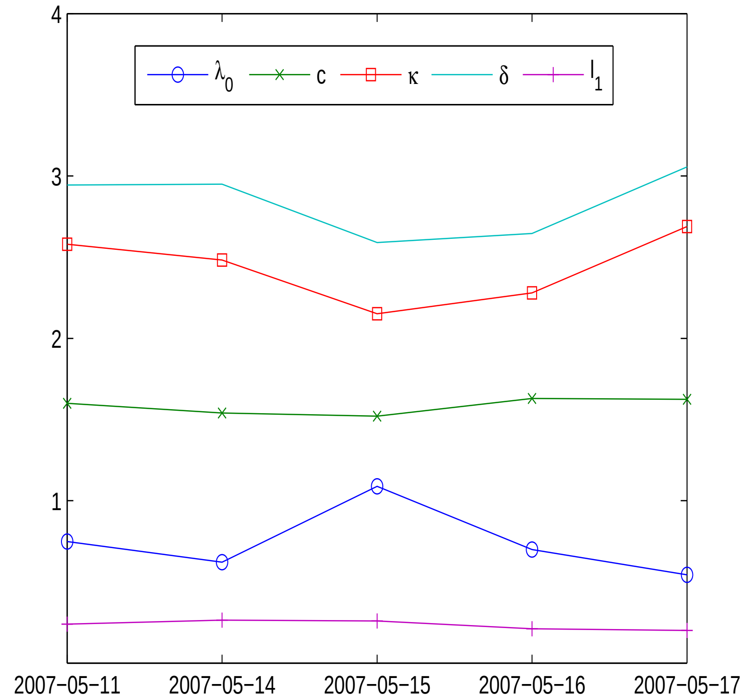

Market calibration

5 and 7Y maturities: market data and fitting results

		MarketBid	MarketAsk	Mod 1	Mod 2
5Y	0-10	70.50%	70.75%	71.69%	72.00%
	10-15	34.25%	34.50%	33.44%	33.47%
	15-25	316.00	319.00	316.13	309.93
	25-35	79.00	81.00	78.45	72.79
	Index	262.85	263.10	263.94	262.75
7Y	0-10	80.13%	80.38%	81.23%	81.49%
	10-15	55.50%	55.75%	55.17%	55.40%
	15-25	582.00	587.00	580.25	584.63
	25-35	180.00	183.00	206.41	207.98
	Index	307.50	307.75	307.53	308.91
	MinObj			136.12	192.82
	AAPE			2.35%	3.30%

Market calibration

5Y maturity: Mod 1 out-of-sample forecast for 7Y

		MarketBid	MarketAsk	5+7Y	5Y
5Y	0-10	70.50%	70.75%	71.69%	71.11%
	10-15	34.25%	34.50%	33.44%	32.85%
	15-25	316.00	319.00	316.13	316.80
	25-35	79.00	81.00	78.45	81.47
	Index	262.85	263.10	263.94	263.46
7Y	0-10	80.13%	80.38%	81.23%	81.79%
	10-15	55.50%	55.75%	55.17%	52.49%
	15-25	582.00	587.00	580.25	531.61
	25-35	180.00	183.00	206.41	186.26
	Index	307.50	307.75	307.53	296.00
	MinObj			136.12	
	AAPE			2.35%	3.04%

Market calibration

3, 5 and 7Y maturities: Mod 1

		MarketBid	MarketAsk	Mod 1
3Y	0-10	45.63%	45.88%	45.03%
	10-15	7.75%	8.00%	7.51%
	15-25	51.00	55.00	54.94
	25-35	12.00	14.00	7.77
	Index	184.70	184.95	189.21
5Y	0-10	70.50%	70.75%	70.25%
	10-15	34.25%	34.50%	32.10%
	15-25	316.00	319.00	310.74
	25-35	79.00	81.00	79.26
	Index	262.85	263.10	258.69
7Y	0-10	80.13%	80.38%	80.25%
	10-15	55.50%	55.75%	53.56%
	15-25	582.00	587.00	580.92
	25-35	180.00	183.00	218.18
	Index	307.50	307.75	308.04
	MinObj			930.19
	AAPE			5.93%

Market calibration

Calibrating through September 2008

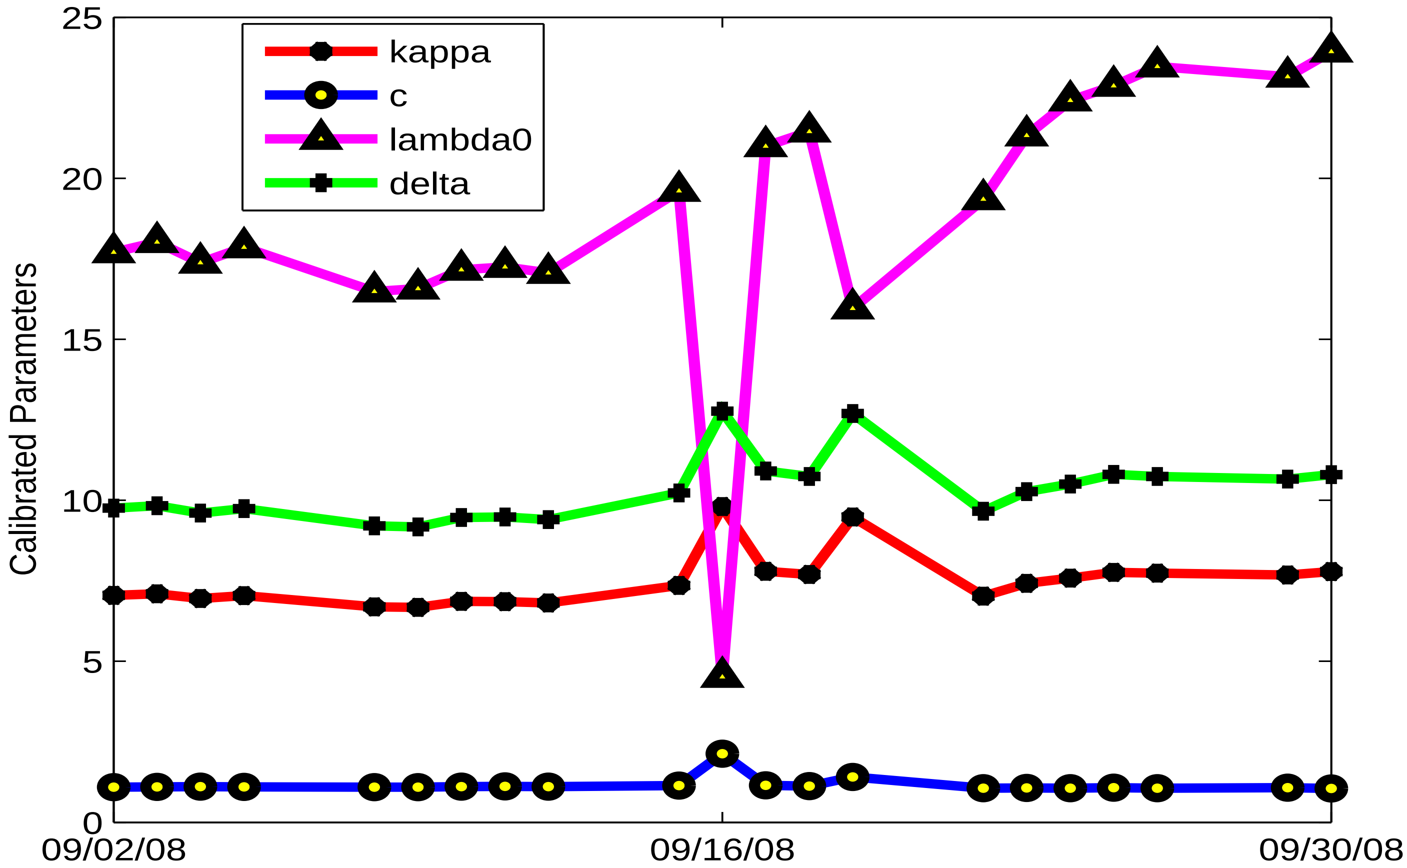

Market calibration

AAPE for September 2008

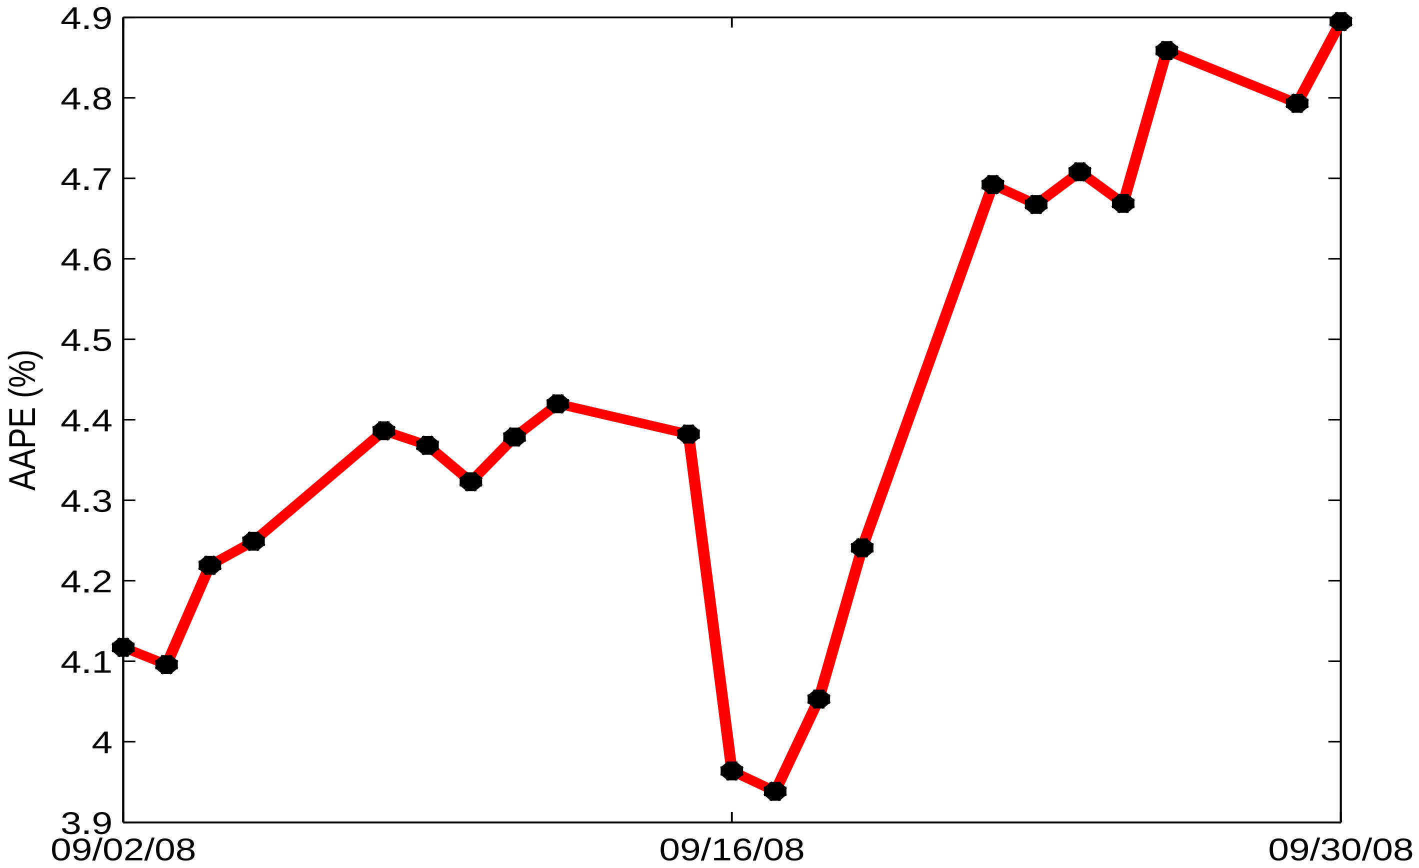

Market calibration

Discussion

- The parsimonious model (1) fits the entire index/tranche market across attachment points and maturities
  - The calibrated parameter values remain remarkably stable over different calibration dates
  - If required, a perfect fit can be achieved by introducing time-dependent parameters
- Maximum likelihood estimation from time series of tranche and index swap spreads to extract **premia for correlated default risk**, see Azizpour et al. (2011)
  - Requires specification of measure change (see Appendix)

Portfolio credit risk

Overview

- We estimate a parametric model for the **actual measure intensity** from historical default experience (i.e., a path of N) using maximum likelihood (see above)
  - For notational simplicity, we let $\mathbb{P}$ denote the actual measure
- We test the goodness-of-fit using the time-scaling test
- We compute the conditional distribution of N at future times, and the associated risk statistics such as VaR

Portfolio credit risk

Data and intensity model

- Our sample period is 1/1/1970 to 1/1/2010
- Data on economy-wide industrial and financial default timing are from Moody's Default Risk Service
- The data is a realization of a marked point process $(T_n, D_n)_{n=1}^{1109}$
  - T_n represents a date with at least one default incidence
  - D_n is the number of defaults at T_n, and $\sum_{n=1}^{1109} D_n = 1667$
- Our intensity model is

$$\lambda_t = \exp(\beta \cdot X_t) + \delta \sum_{n \leq N_t} \exp(-\eta(t - T_n))(D_n + wD_n^2)$$

where X is a vector of observable explanatory covariates

Portfolio credit risk

Number of defaults per year

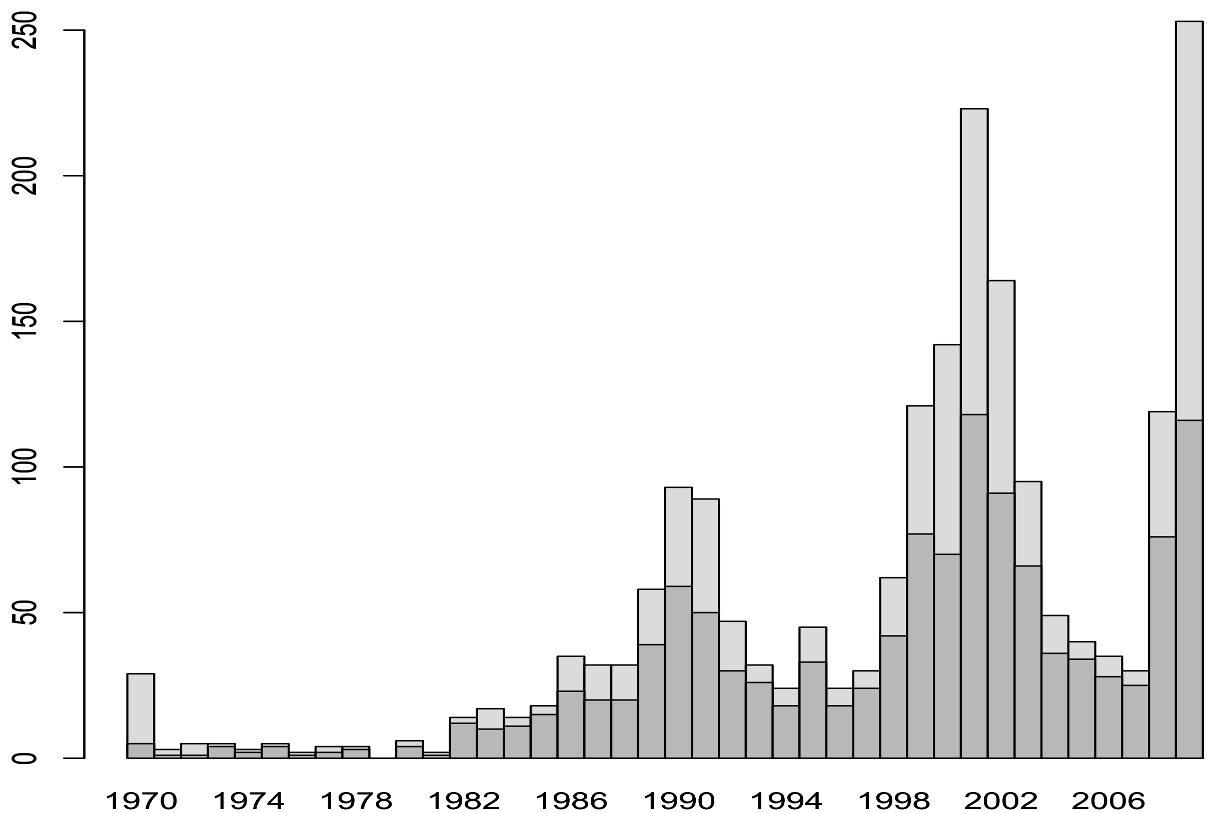

Portfolio credit risk

Number of defaults per day

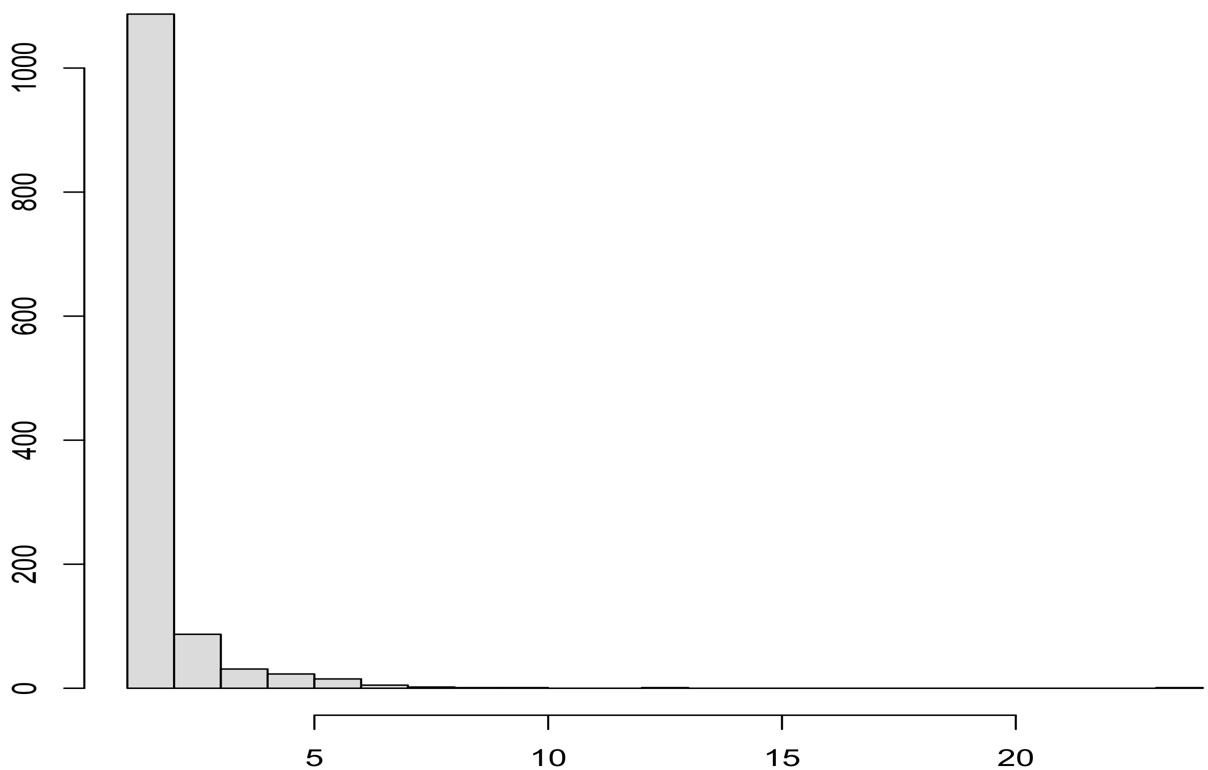

Portfolio credit risk

Vector X of observable covariates

- A constant
- The trailing 1-year return on the S&P 500
- The 3-month Treasury bill rate
- The spread between the 10 and 1-year Treasury rates
- The 1-year percentage growth of the US industrial production
- In different econometric settings, Duffie et al. (2006), Duffie et al. (2009), and others found these variables to be significant predictors of US industrial defaults

Portfolio credit risk

Fitted intensity vs. default times

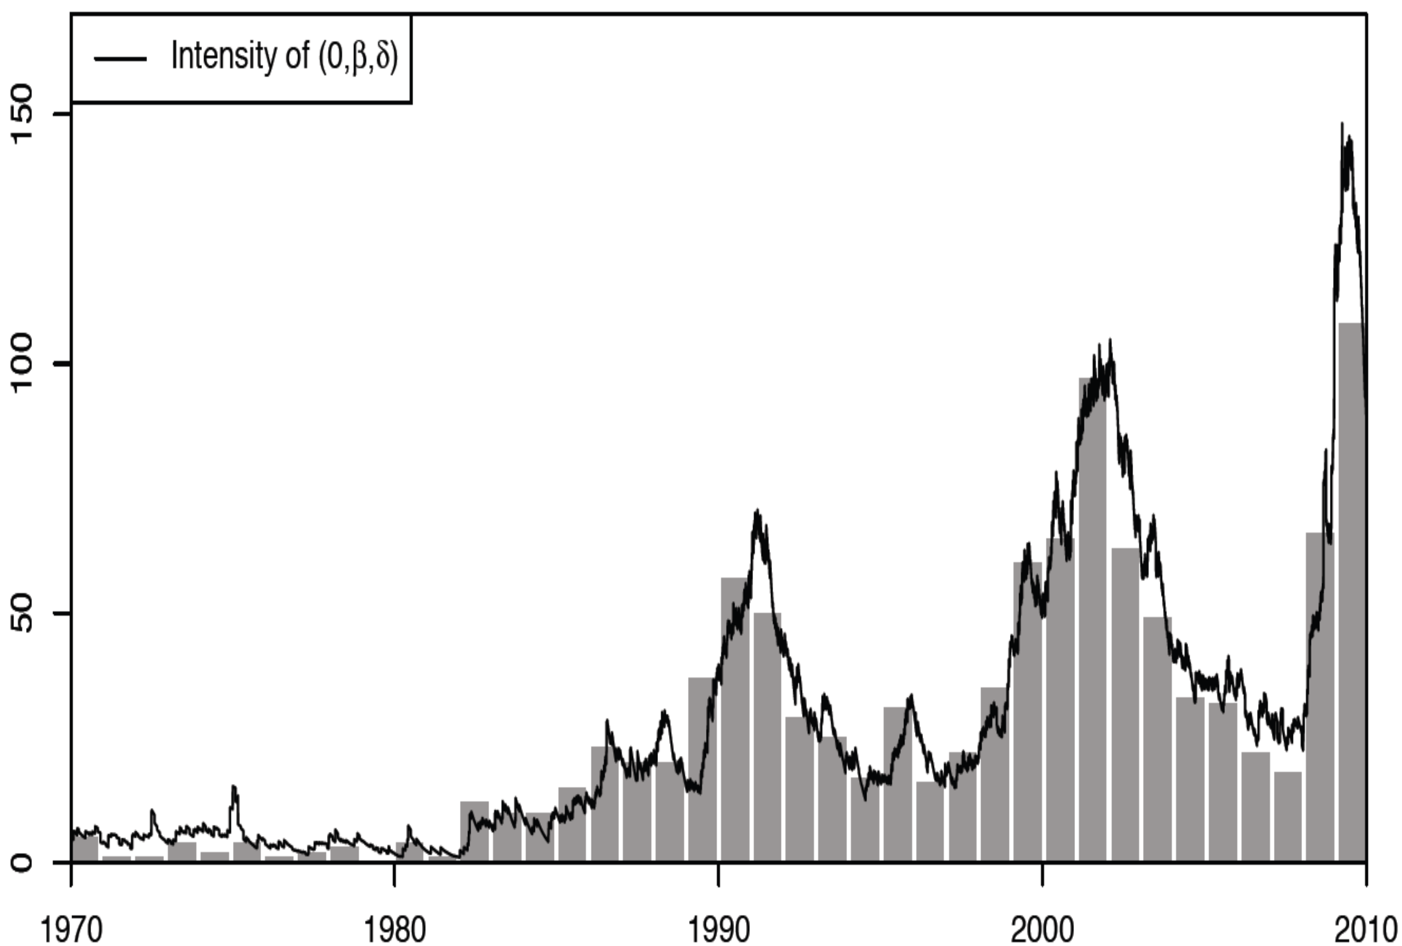

Portfolio credit risk

Time change tests

- We test whether the time-changed event counting process

$$N_{A^{-1}}, \qquad A_t = \int_0^t \hat{\lambda}_s ds$$

generated by the fitted filtered intensity $\hat{\lambda}$ is a realization of a standard Poisson process relative to its own filtration

- If λ is correctly specified, then
  - The inter-arrival times $W_n = A_{T_n} - A_{T_{n-1}}$ of $N_{A^{-1}}$ are independent samples from a standard exponential distribution
  - The binned event counts $P_i^b = \sum_{n=1}^{N_\tau} 1_{(b(i-1),\, bi]}(A_{T_n})$, defined for bins of size $b > 0$, are independent samples from a Poisson distribution with parameter b

Portfolio credit risk

QQ-plot of time-changed inter-arrival times W_n

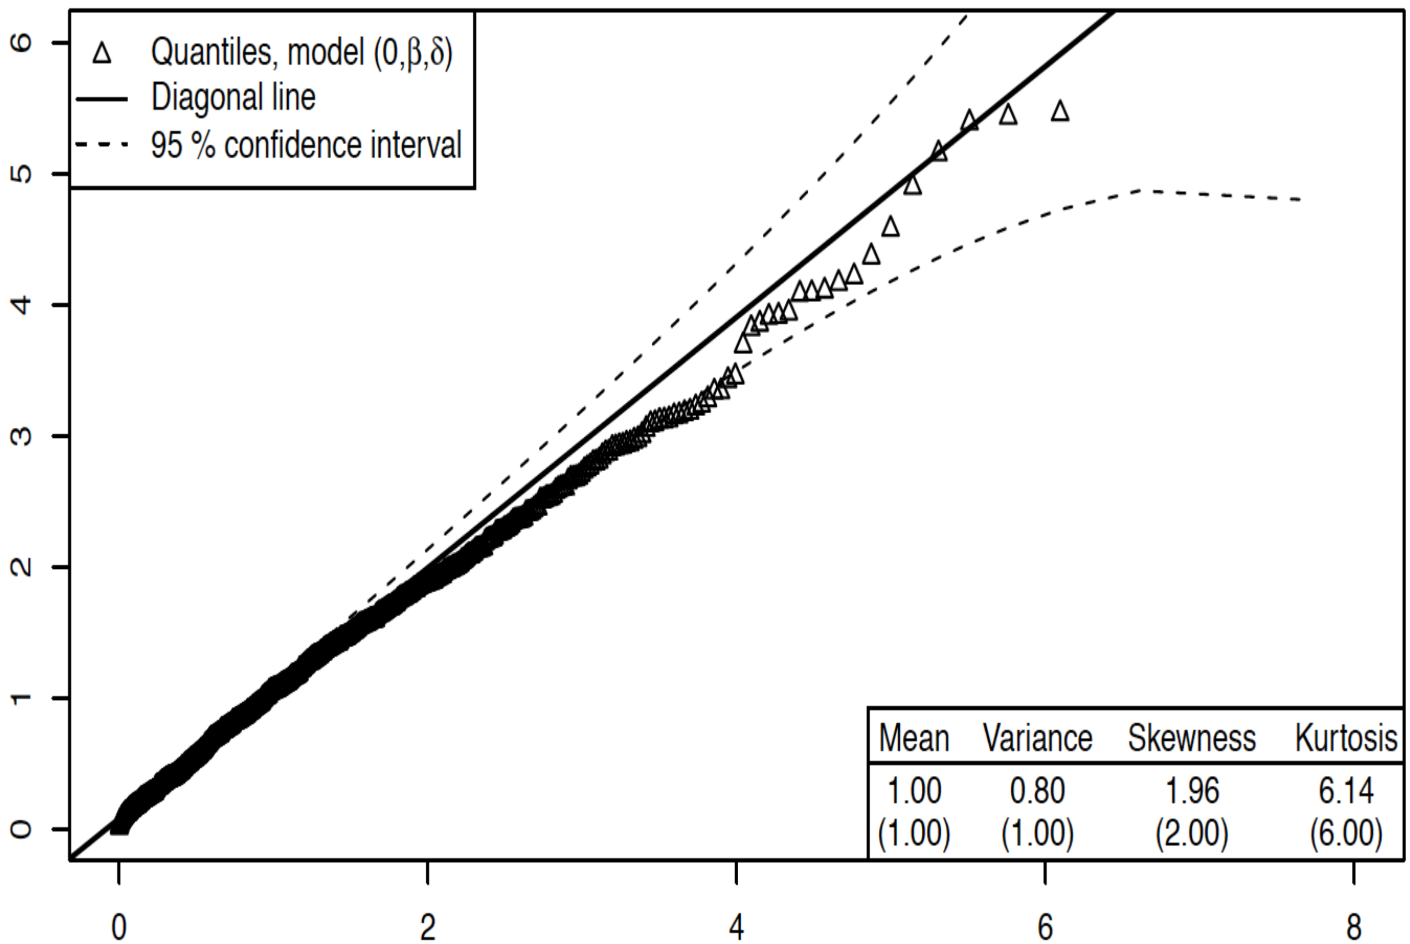

Portfolio credit risk

Out-of-sample forecasts vs. realized events

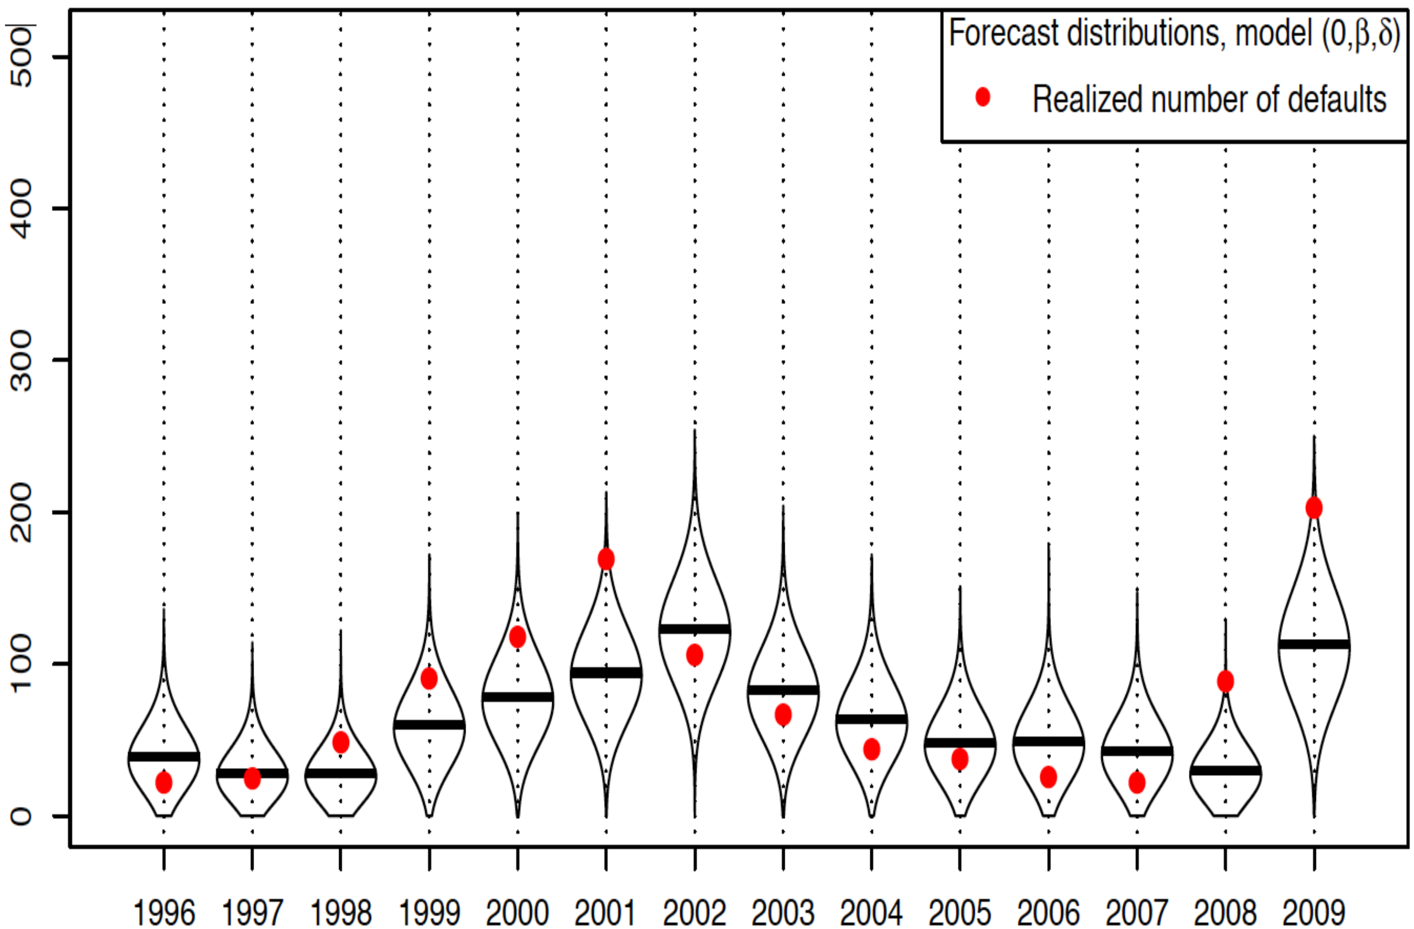

Portfolio credit risk

Discussion

- With two default correlation channels addressed in our model, the forecast performance is quite reasonable
  - Dependence of λ on common covariate vector X
  - Feedback (contagion)
- In- and out-of-sample fit can be improved somewhat by including an unobserved frailty process to address the third channel of default correlation
  - Computationally challenging filtering problem must be addressed, see Giesecke & Schwenkler (2011)

Portfolio credit risk

Out-of-sample forecasts when including an OU frailty

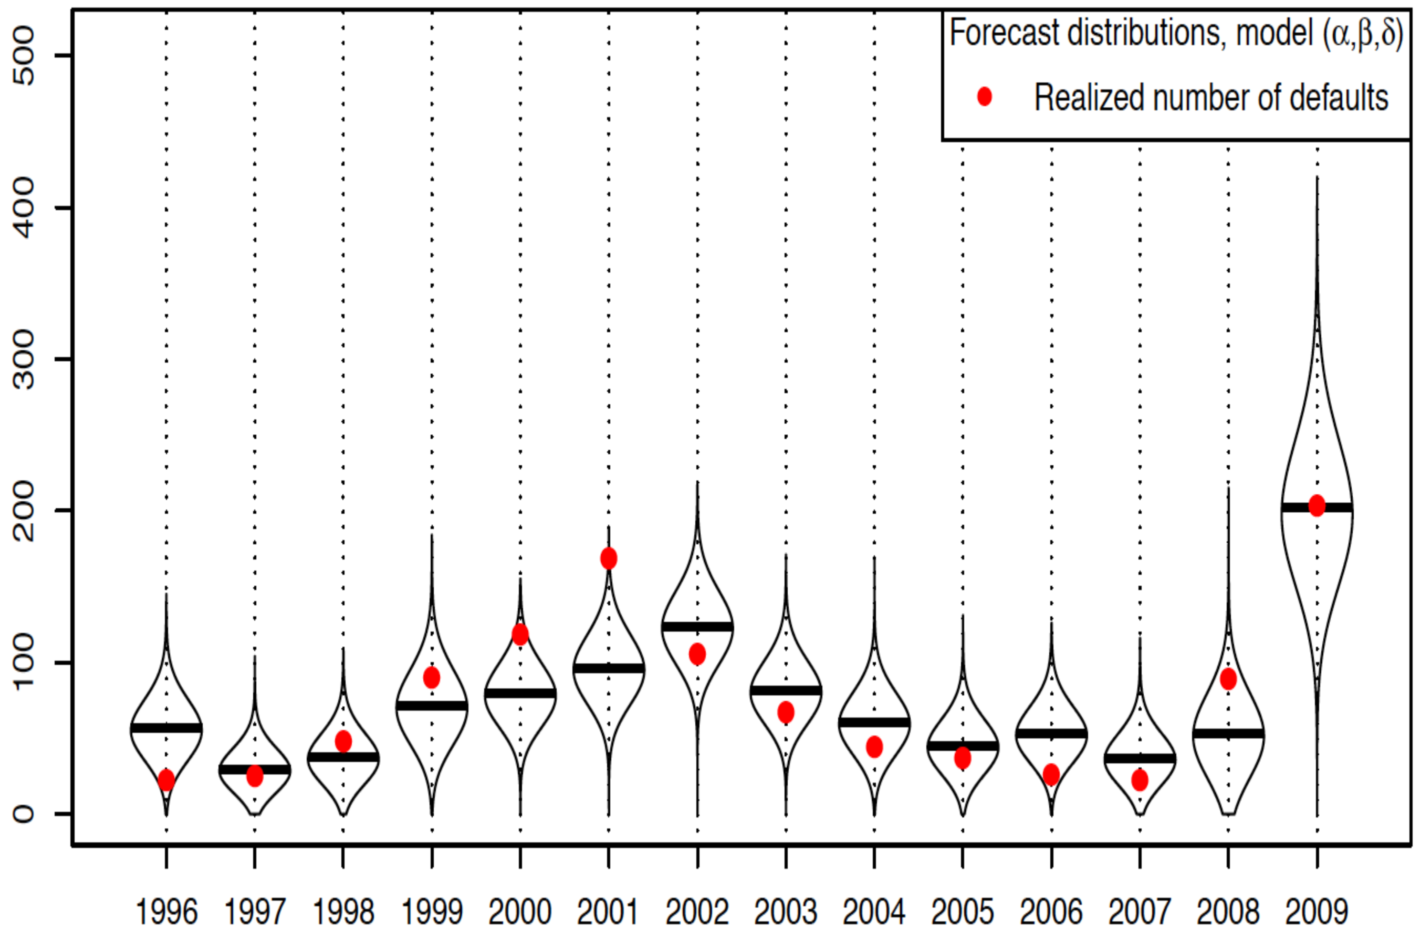

Conclusion

- Casting credit risk applications as point process problems is natural and effective
- Many other important topics we did not treat, e.g.
  - Structural models of default (Merton (1974), Black & Cox (1976), and many others)
  - Counterparty valuation adjustment (CVA)
  - Bottom-up models of correlated default (Duffie & Garleanu (2001), Eckner (2009), and many others)

Appendix

Stieltjes integrals

- Consider a real-valued, right continuous function Y on $[0, \infty)$
  - For fixed $\omega \in \Omega$, the paths $N_t(\omega)$ and $A_t(\omega)$ are real-valued, right continuous, non-decreasing and finite on $[0, t]$
- The variation V_t of Y on $[0, t]$ is defined as

$$V_t = \sup_{\Delta} \sum_i |Y_{t_{i+1}} - Y_{t_i}|$$

  where Δ is a subdivision of $[0, t]$ with $0 = t_0 < t_1 < \ldots < t_n = t$

- The function Y is of finite variation if $V_t < \infty$ for every t
- Monotone finite functions such as the paths $N_t(\omega)$ and $A_t(\omega)$ are of finite variation on any finite interval, and so is the path of the compensated point process martingale $N_t(\omega) - A_t(\omega)$

Appendix

Stieltjes integrals

- A right continuous function Y_t of finite variation corresponds to a measure μ on $[0, \infty)$ via

$$Y_t = \mu([0, t])$$

- If X_t is a locally bounded Borel function on $\mathbb{R}_+$, then its **Stieltjes integral** with respect to Y, denoted

$$I_t = \int_0^t X_s dY_s$$

  is the Lebesgue integral of X with respect to μ on $(0, t]$

- The function I is again right continuous and of finite variation

Appendix

Change of variables for Stieltjes integrals

- Let X be a right continuous function of finite variation
- Let $f : \mathbb{R} \to \mathbb{R}$ be continuously differentiable
- We have the change of variables formula for Stieltjes integrals

$$f(X_t) - f(X_0) = \int_0^t f'(X_{s-})dX_s + \sum_{0<s\leq t} \Big(f(X_s) - f(X_{s-}) - f'(X_{s-})\Delta X_s \Big)$$

- Here, $\Delta X_s = X_s - X_{s-}$ is the jump of X at s and

$$X_{s-} = \lim_{u \uparrow s} X_u, \quad X_{0-} = X_0$$

Appendix

Compare with Itô's formula

- Let X be a semimartingale, not necessarily of finite variation
- Let $f : \mathbb{R} \times \mathbb{R}_+ \to \mathbb{R}$ be twice continuously differentiable
- We have the change of variables formula

$$\begin{aligned} f(X_t,t) - f(X_0,0) &= \int_0^t f_x(X_{s-},s-)dX_s + \int_0^t f_t(X_s,s)ds \\ &+ \frac{1}{2}\int_0^t f_{xx}(X_{s-},s-)d[X,X]_s^c \\ &+ \sum_{0<s\leq t} \big(f(X_s,s) - f(X_{s-},s-) - f_x(X_{s-},s-)\Delta X_s\big) \end{aligned}$$

  with $[X,X]_t = [X,X]_t^c + \sum_{0\leq s\leq t}(\Delta X_s)^2$

- If X is right continuous and of finite variation, then $[X,X]^c = 0$ and we get the change of variables formula for Stieltjes integrals

Appendix

Stieltjes integration by parts

- Let X and Y be right continuous functions of finite variation
- We have the integration by parts formula

$$X_t Y_t = X_0 Y_0 + \int_0^t X_{s-} dY_s + \int_0^t Y_s dX_s$$

- We can write this formula equivalently as

$$X_t Y_t = X_0 Y_0 + \int_0^t X_{s-} dY_s + \int_0^t Y_{s-} dX_s + \sum_{s \leq t} \Delta X_s \Delta Y_s$$

since, using the decomposition $X_t = X_t^c + \sum_{s \leq t} \Delta X_s$ and the fact that Y has only a countable number of jumps,

$$\int_0^t \Delta Y_s dX_s = \int_0^t \Delta Y_s dX_s^c + \sum_{s \leq t} \Delta X_s \Delta Y_s = \sum_{s \leq t} \Delta X_s \Delta Y_s$$

Appendix

Equivalent change of measure

- Consider a probability measure $\mathbb{Q}$ on $\mathcal{F}_T$ equivalent to $\mathbb{P}$, meaning that $\mathbb{P}$ and $\mathbb{Q}$ have the same events of measure 0
- The Radon-Nikodym theorem asserts that there exists a random variable $Z = \frac{d\mathbb{Q}}{d\mathbb{P}} > 0$ with $\mathbb{E}(Z) = 1$ such that

$$\mathbb{Q}(U) = \mathbb{E}(Z 1_U) = \int_U Z(\omega) d\mathbb{P}(\omega), \quad U \in \mathcal{F}_T$$

- The $\mathbb{P}$-martingale Z defined by $Z_t = \mathbb{E}(Z \,|\, \mathcal{F}_t)$ for $t \leq T$ is called the density process

Appendix

Form of density

- Suppose we start with a Brownian motion W and a point process N on $(\Omega, \mathcal{F}, \mathbb{P})$; W may influence the intensity of N
- The martingale representation theorem implies that there are predictable processes a and $b > 0$ such that

$$Z_t = \exp\Big(-\int_0^t a_s dW_s - \frac{1}{2}\int_0^t a_s^2 ds + \int_0^t \log(b_s) dN_s - \int_0^t (b_s - 1) dA_s \Big), \quad t \leq T$$

- Conditions are required

Appendix

Girsanov's theorem

- If N has an intensity λ under $\mathbb{P}$, then it also admits an intensity $\lambda^{\mathbb{Q}}$ under $\mathbb{Q}$ (see Artzner & Delbaen (1995))
- The $\mathbb{Q}$-intensity is given by

$$\lambda_t^{\mathbb{Q}} = \lambda_t b_t, \quad t \leq T$$

- A $\mathbb{Q}$-Brownian motion $W^{\mathbb{Q}}$ is given by

$$W_t^{\mathbb{Q}} = W_t + \int_0^t a_s ds, \quad t \leq T$$

Appendix

Example

- Suppose N has $\mathbb{P}$-intensity $\lambda > 0$ and we want to transform N into a standard Poisson process under $\mathbb{Q}$
- We choose $b = 1/\lambda$ so N has $\mathbb{Q}$-intensity 1, and therefore is standard Poisson under $\mathbb{Q}$, by Watanabe's theorem
- Note that N and $W^{\mathbb{Q}}$ are $\mathbb{Q}$-independent (even though W may influence λ, and hence not be $\mathbb{P}$-independent of N)

References

Artzner, Philippe & Freddy Delbaen (1995), 'Default risk insurance and incomplete markets', *Mathematical Finance* **5**, 187–195.

Azizpour, Shahriar, Kay Giesecke & Baeho Kim (2011), 'Premia for correlated default risk', *Journal of Economic Dynamics and Control* **35**(8), 1340–1357.

Azizpour, Shahriar, Kay Giesecke & Gustavo Schwenkler (2009), Exploring the sources of default clustering. Working Paper, Stanford University.

Berndt, Antje, Rohan Douglas, Darrell Duffie, Mark Ferguson & David Schranz (2005), Measuring default risk premia from default swap rates and EDF's. Working Paper, Stanford University.

Beskos, Alexander & Gareth Roberts (2005), 'Exact simulation of diffusions', *Annals of Applied Probability* **15**(4), 2422–2444.

Black, Fischer & John C. Cox (1976), 'Valuing corporate securities: Some effects of bond indenture provisions', *Journal of Finance* **31**, 351–367.

Carr, Peter & Liuren Wu (2004), 'Time-changed Lévy processes and option pricing', *Journal of Financial Economics* **71**, 113–141.

Chen, Nan (2009), Localization and exact simulation of brownian motion driven stochastic differential equations. Working Paper, Chinese University of Hong Kong.

Cheng, Peng & Olivier Scaillet (2007), 'Linear-quadratic jump-diffusion modeling', *Mathematical Finance* **17**, 575–598.

Duffie, Darrell, Andreas Eckner, Guillaume Horel & Leandro Saita (2009), 'Frailty correlated default', *Journal of Finance* **64**, 2089–2123.

Duffie, Darrell, Jun Pan & Kenneth Singleton (2000), 'Transform analysis and asset pricing for affine jump-diffusions', *Econometrica* **68**, 1343–1376.

Duffie, Darrell, Leandro Saita & Ke Wang (2006), 'Multi-period corporate default prediction with stochastic covariates', *Journal of Financial Economics* **83**(3), 635–665.

Duffie, Darrell & Nicolae Garleanu (2001), 'Risk and valuation of collateralized debt obligations', *Financial Analysts Journal* **57**(1), 41–59.

Eckner, Andreas (2007), Risk premia in structured credit derivatives. Working Paper, Stanford University.

Eckner, Andreas (2009), 'Computational techniques for basic affine models of portfolio credit risk', *Journal of Computational Finance* **15**, 63–97.

Giesecke, Kay & Baeho Kim (2007), Estimating tranche spreads by loss process simulation, *in* S. G.Henderson, B.Biller, M.-H.Hsieh, J.Shortle, J. D.Tew & R. R.Barton, eds, 'Proceedings of the 2007

Winter Simulation Conference', IEEE Press, pp. 967–975.

Giesecke, Kay, Baeho Kim & Shilin Zhu (2011), 'Monte Carlo algorithms for default timing problems', *Management Science* **57**, 2115–2129.

Giesecke, Kay & Dmitry Smelov (2010), Exact sampling of jump-diffusions. Working Paper, Stanford University.

Giesecke, Kay, Francis Longstaff, Ilya Strebulaev & Stephen Schaefer (2011), 'Corporate bond default risk: A 150-year perspective', *Journal of Financial Economics* **102**(2), 233–250.

Giesecke, Kay & Gustavo Schwenkler (2011), Filtered likelihood for point processes. *Annals of Statistics*, forthcoming.

Giesecke, Kay, Hossein Kakavand & Mohammad Mousavi (2010), Exact simulation of point processes with stochastic intensities. *Operations Research*, forthcoming.

Giesecke, Kay, Hossein Kakavand, Mohammad Mousavi & Hideyuki Takada (2010), 'Exact and efficient simulation of correlated defaults', *SIAM Journal on Financial Mathematics* **1**, 868–896.

Giesecke, Kay & Shilin Zhu (2010), Transform analysis for point processes and applications in credit risk. *Mathematical Finance*, forthcoming.

Glasserman, Paul & Nicolas Merener (2003), 'Numerical solution of jump-diffusion LIBOR market models', *Finance and Stochastics* **7**, 1–27.

Gyöngy, Imre (1986), 'Mimicking the one-dimensional marginal distributions of processes having an Itô differential', *Probability Theory and Related Fields* **71**, 501–516.

Kliemann, W.H., G. Koch & F. Marchetti (1990), 'On the unnormalized solution of the filtering problem with counting process observations', *IEEE Transactions on Information Theory*

36(6), 1415 –1425.

Leippold, Markus & Liuren Wu (2002), 'Asset pricing under the quadratic class', *Journal of Financial and Quantitative Analysis* **37**, 271–295.

Lewis, Peter & George Shedler (1979), 'Simulation of nonhomogeneous Poisson processes by thinning', *Naval Logistics Quarterly* **26**, 403–413.

Merton, Robert C. (1974), 'On the pricing of corporate debt: The risk structure of interest rates', *Journal of Finance* **29**, 449–470.

Meyer, Paul-André (1971), Démonstration simplifée d'un théorème de Knight, *in* 'Séminaire de Probabilités V, Lecture Note in Mathematics 191', Springer-Verlag Berlin, pp. 191–195.