Multiname default models under stochastic time-change

Michael Gordy and Pawel Szerszen

Federal Reserve Board

March 2015

The opinions expressed here are our own, and do not reflect the views of the Board of Governors or its staff.

Email: ⟨michael.gordy@frb.gov⟩

Spreads on five US dealer banks

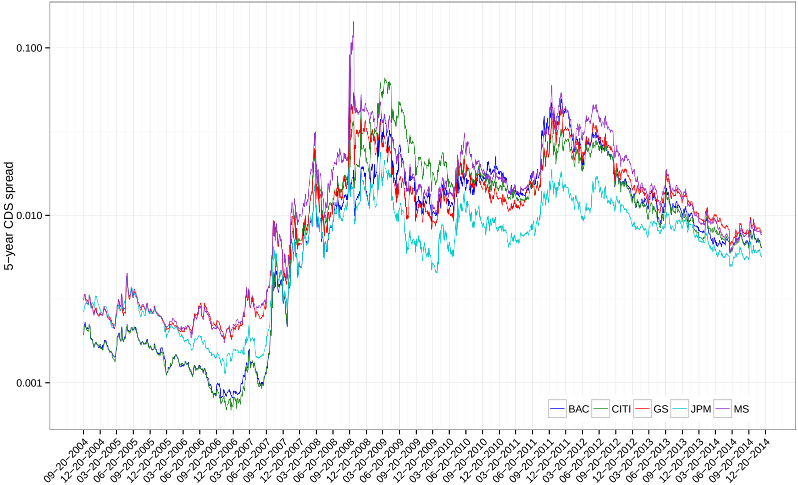

Daily changes in the log-spreads

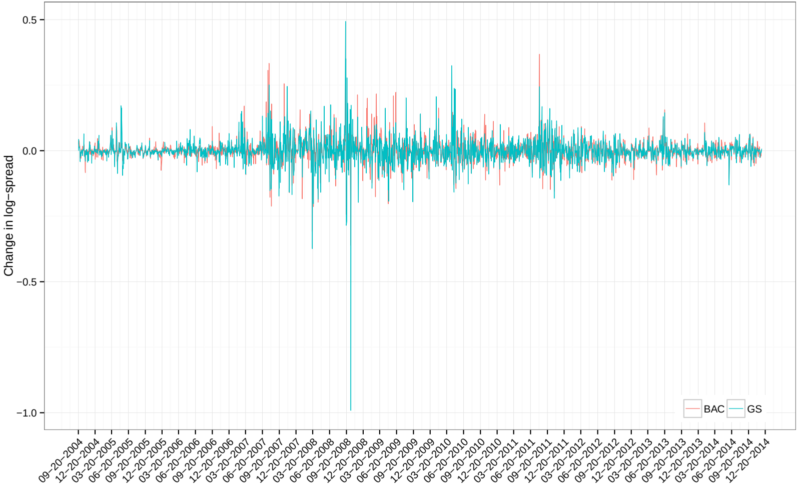

Objectives

- Specify and estimate a multiname model of stochastic volatility in default intensities and estimate on CDS spread data for a set of B large US dealer banks.
- Model should allow us to bound intensities nonnegative.
- Why not simply model joint process for CDS spreads? "No arbitrage" model is more demanding but has some advantages:
  - Distinguish effect of interest rate changes from changes in default intensity.
  - Consistent pricing of other instruments, e.g., corporate bonds, swaps at other maturities, swaptions, CDO tranches.
  - Risk-measures for portfolios of bank CDS at any horizon.
- Estimation may entail over 100 million pricing calls per reference entity, so the model must have a robust semi-analytic solution.

Research agenda

Third part of a sequence of papers on stochastic volatility in reduced-form credit risk models.

1. Methodology for pricing credit sensitive instruments under stochastic time-change.

Costin, Gordy, Huang and Szerszen (2015)

"Expectations of functions of stochastic time with application to credit risk modeling," forthcoming in Mathematical Finance.

2. Methodology for estimating single-name models.

Gordy and Szerszen (2015)

"Bayesian estimation of time-changed default intensity models," FEDS 2015-002.

3. Extension to multiname setting (in progress).

Dependence but no contagion

- Dependence in our model is generated by common factors in the level of default intensities and in their volatility.
  - Extends Duffie and Gârleanu (FAJ 2001) model for pricing CDO tranches. We introduce stochastic volatility via a common stochastic time-change.
  - Stochastic time-change first used in CDO pricing by Joshi and Stacey (Risk 2006), but with deterministic background default intensity.
- No contagion in our framework.
  - Ding, Giesecke and Tomecek (OR 2009) suggest time-changed self-exciting process for aggregate number of defaults in a portfolio.
  - We are modeling portfolio from bottom-up.
  - Estimation of multiname models with contagion still quite difficult, e.g., Aït-Sahalia, Laeven and Pelizzon (J. Econometrics 2014).

Systemic risk?

- Whether this really constitutes a model of systemic risk is debatable.
- Applied to ordinary firms, it would properly be understood as a model of systematic risk.
- But joint distress of large dealer banks unavoidably has systemic consequences.
- Model outputs are coincident risk measures.
  - Most closely related to distress insurance premium of Huang, Zhou and Zhu (JBF 2009), which is a hypothetical insurance premium on a basket of banks.
  - Could develop analogs to SRISK measure of Brownlees and Engle, and other similar measures.
  - Central banks have taken great interest in such indicators as practical tools.

Stochastic volatility in credit risk models

Despite ubiquity of stochastic volatility (SV) in financial data, there has been limited exploration of SV in the credit risk literature.

- Spreads: Gordy and Willemann (MS 2012) find evidence of SV in CDS spreads.
- Structural models: Fouque, Sircar and Sølna (AMF 2006) add a volatility factor to Black-Cox; empirical evidence in Zhang, Zhou and Zhu (RFS 2009).
- Reduced-form models: Jacobs and Li (MS 2008) estimate Fong-Vasicek model for default intensities on bond prices.
    - For realistic parameter values, intensity often negative.

Stochastic time change

We employ stochastic time change as a parsimonious and tractable way to introduce SV.

- Associated with calendar time t is a stochastic business time T_t.
  - T_t is an increasing right-continuous process with left limits.
  - Intuition is that news flows to the markets at an uneven rate. The faster the rate of news flow, the quicker the passage of business time, and so the greater the likelihood of default events or large spread changes.
- Let $\tilde{\tau}$ be the calendar default time for a firm, and let $\tau = T_{\tilde{\tau}}$ be the business-clock default time.

Survival probabilities in business and calendar time

- Let λ_t denote the default intensity in business time.
- Assume that λ_t and T_t are independent.
  - We are ruling out any "leverage effect" in volatility.
  - Leverage effect appears significant for equity market.
  - Evidence less compelling for credit market. Jacobs and Li (2008) find median correlation of 1% between default intensity factor and the volatility factor.
- Fixing current time to 0, let $S(t;\ell) \equiv \Pr(\tau > t|\lambda_0 = \ell)$ be survival probability to business time t. By LIE, calendar-time survival probability is

$$\begin{aligned}\tilde{S}(t;\ell) \equiv \Pr(\tilde{\tau} > t|\lambda_0 = \ell) &= \Pr(T_{\tilde{\tau}} > T_t|\lambda_0 = \ell) \\ &= \mathrm{E}[\Pr(T_{\tilde{\tau}} > T_t|T_t, \lambda_0 = \ell)|\lambda_0 = \ell] = \mathrm{E}[S(T_t;\ell)]\end{aligned}$$

Expansion in exponential functions

- Assumption: S has a convergent series expansion of the form

$$S(t) = \exp(at) \sum_{n=0}^{\infty} \beta_n \exp(-n\gamma t)$$

  for constants a, γ with $\gamma > 0$, and $\sum_{n=0}^{\infty} |\beta_n| < \infty$.
- Let $M_t(u)$ denote the moment generating function for T_t.
  Assumption: $M_t(u)$ exists for $u < a$.
  - Many time-change processes of empirical interest have known moment generating functions.
- Then we have uniformly convergent series expansion

$$\boxed{\tilde{S}(t) = \mathrm{E}[S(T_t)] = \sum_{n=0}^{\infty} \beta_n \mathrm{E}[\exp((a - n\gamma)T_t)] = \sum_{n=0}^{\infty} \beta_n M_t(a - n\gamma)}$$

Empirical specification: default intensity

- We adopt the conventional CIR specification for λ_t in business time.
- λ_t is the risk-neutral intensity, but we must specify dynamics under both $\mathbb{P}$ and $\mathbb{Q}$.
- Under $\mathbb{P}$,
$$d\lambda_t = (\mu - \kappa^{\mathbb{P}} \lambda_t)dt + \sigma \sqrt{\lambda_t}\, dW_t^{\mathbb{P}}$$
- As in Duffee (RFS 1999), SDE under $\mathbb{Q}$ has mean-reversion $\kappa^{\mathbb{Q}}$, but μ, σ invariant under change of measure.
- If $\mu > 0$ and $\lambda_0 \geq 0$ then $\lambda_t \geq 0$.
- Default intensity in calendar time is
$$\tilde{\lambda}_t = -\tilde{S}'(0; \lambda_{T(t)})$$

  Bounded nonnegative whenever λ_t bounded nonnegative.

Survival probabilities in business time

- Conventional solution takes form $S(t;\lambda_0) = \exp(A(t) + B(t)\lambda_0)$.
- We derive a series solution of the form

$$S(t;\lambda_0) = \exp(at)\sum_{n=0}^{\infty}\beta_n(\lambda_0)\exp(-n\gamma t)$$

  where the β_n are weighted Laguerre polynomials.
- Equivalent to eigenfunction expansion of Mendoza-Arriaga and Linetsky (AAP 2014).
- Series is uniformly convergent when $\kappa^{\mathbb{Q}} > 0$.

Big problem!

Empirically, $\kappa^{\mathbb{Q}}$ typically negative! Series solution divergent.

Generalized expansion

- Assumption: S has a convergent series expansion of the form

$$S(t) = \exp(at) \sum_{n=0}^{\infty} \beta_n \left(\omega(t) - \exp(-\gamma t)\right)^n$$

  for constants a, γ with $\gamma > 0$, and $\sum_{n=0}^{\infty} |\beta_n| < \infty$, and deterministic function $\omega(t)$ mapping to unit interval.
- Let $M_t(u)$ denote the moment generating function for T_t.
  Assumption: $M_t(u)$ exists for $u < a$.
- Then we have pointwise convergent series expansion

$$\boxed{\tilde{S}(t) = \mathrm{E}[S(T_t)] = \sum_{n=0}^{\infty} \beta_n \sum_{m=0}^{n} \binom{n}{m} \omega(t)^{n-m} (-1)^m M_t(a - m\gamma)}$$

Application to CIR model

- For CIR default intensity, we find

$$S(t;\lambda_0) = \exp(at)\sum_{n=0}^{\infty}\beta_n(\lambda_0;\omega(t))\cdot(\omega(t)-\exp(-\gamma t))^n$$

  for $\omega(t) = 1/2$ is uniformly convergent for all $\mu > 0$, $\sigma > 0$, $\kappa \in \mathbb{R}$.
- Can generalize to allow λ_t to follow basic affine process.
- As before, the β_n are weighted Laguerre polynomials.
  $\Rightarrow$ Two-term recurrence rule for β_n.
- Find series expansion for $\tilde{S}(t)$ converges in few terms if we set $\omega(t) = \mathrm{E}[\exp(-\gamma T_t)] = M_t(-\gamma)$.

Empirical specification: stochastic time change

T_t is an inverse Gaussian process with marginal distributions

$$T_t^{\mathbb{P}}, T_t^{\mathbb{Q}} \sim \mathrm{IG}(t, \alpha t^2) \Rightarrow \mathrm{E}^{\mathbb{P}}[T_t] = \mathrm{E}^{\mathbb{Q}}[T_t] = t, \ \mathrm{V}^{\mathbb{P}}[T_t] = \mathrm{V}^{\mathbb{Q}}[T_t] = t/\alpha$$

- Interpret α as a precision parameter.
  - As $\alpha \to \infty$, business time converges to calendar time, a.s.
- We assume no risk-premium on time-change. In our specification, risk-premium is very weakly identified.
- Lévy process. No serial dependence, so no volatility clustering.
- Moment generating function of T_t for $u < \alpha/2$:

$$M_t(u) = \exp\left(\alpha\left(1 - \sqrt{1 - 2u/\alpha}\right) t\right)$$

Pricing credit default swaps

- Assumption: Riskfree short rate r_t and default intensity $\tilde{\lambda}_t$ are independent.
  - Duffee (RFS 1999) finds that the correlation between riskfree rates and default intensity is negative but has second-order importance.

  Discount at observation date s for term t is written $D_s(t)$.
- Assume fixed recovery rate R.
- Let $\tilde{f}_s(t;\ell) = -\frac{\partial}{\partial t}\tilde{S}_s(t;\lambda_{T(s)} = \ell)$ be density of $\tilde{\tau}$ at term t as of observation date s.
- Protection leg value: $(1-R)\int_0^t D_s(z)\tilde{f}_s(z)dz$
- Premium leg value: $(c/4)\sum D_s(z_k)\tilde{S}_s(z_k)$ over settlement dates $z_k \leq t$.
- Par spread is coupon equating protection and premium legs.

Time-change affects kurtosis at short horizons

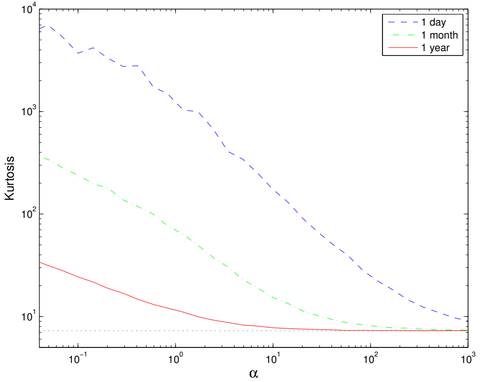

Kurtosis of stationary distribution of change in 5 year CDS spread over 1 day, 1 month or 1 year interval. Simulated distribution with 4 million trials.

Parameters: $\kappa^{\mathbb{P}} = \kappa^{\mathbb{Q}} = 0.2$, $\mu = 0.02\kappa^{\mathbb{Q}}$, $\sigma = 0.1$, riskless rate 3%, recovery 40%.

Swaption prices

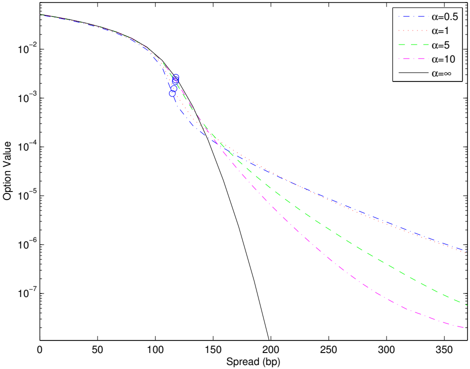

Value of one-month payer option on a five-year CDS as function of strike spread. Circles mark par spread at date 0. Parameters: $\lambda_0 = 0.01$, $\kappa^{\mathbb{Q}} = 0.2$, $\mu = 0.02\kappa^{\mathbb{Q}}$, $\sigma = 0.1$, riskless rate $r = 0.03$, recovery $R = 0.4$.

Data

- CDS spreads on five large US dealer banks (Markit).
  - Maturities 1, 2, 3, 5, 7, 10 years.
  - For observations since April 2009 change in market convention, upfront converted to par spread.
  - Estimation sample: Daily data 9/2004 to 10/2014.
  - Out-of-sample forecast period: 11/2014 to 2/2015.
- CDS spreads on 10 liquid non-financial corporate names (Markit).
  - Estimation sample: Daily data 1/2002 to 12/2009.
  - Out-of-sample forecast period: 1/2010 to 12/2011.
- Discount function $D_s(t)$ from fitting Svensson formula to Treasury yield curve on each observation date s.

MCMC estimation of single-name model

- Markov Chain Monte Carlo estimation
  - Smoothed estimates of latent (λ_s, T_s) as by-product of MCMC.
  - Primary challenge is in Metropolis-Hastings simulation of latent default intensities. Due to persistence in λ_s, states highly correlated, so single-move MH not a viable approach.

- Adapt method of Stroud, Müller and Polson (JASA 2003)
  - Auxilliary model approximates true model as a locally linearized mixture model with K nodes. New latent state k_s is mixing variable.
  - Auxilliary model is linear and Gaussian within blocks, so draw all states jointly via forward filtering backward sampling FFBS algorithm of Carter and Kohn (Biometrika 94), and Frühwirth-Schnatter (JTSA 94).
  - MH accept/reject determined by true model. Acceptance probability high if auxilliary model is a good approximation.

State-space representation in discrete time

$$
\begin{aligned}
y_{s,t} &= F_{s,t}(h_s, \mu, \kappa^{\mathbb{Q}}, \sigma, \alpha) + \zeta\ \varepsilon^{(y)}_{s,\,t} \\
h_{s+\Delta} &= h_s + (\mu - \kappa^{\mathbb{P}} h_s)\chi_{s+\Delta} + \sigma\ \sqrt{h_s \chi_{s+\Delta}}\ \varepsilon^{(\lambda)}_{s+\Delta} \\
\chi_{s+\Delta} &= T_{s+\Delta} - T_s \overset{\text{iid}}{\sim} \ IG(\Delta, \alpha\Delta^2)
\end{aligned}
$$

- Calendar time increment is $\Delta = 1/250$ (daily data).
- $y_{s,t}$: observed log-spread for maturity t on observation date s.
- $h_s = \lambda(T_s)$ is the business-time default intensity at the business time associated with calendar time s.
- $F_{s,t}$: model implied log-spreads, depend on $\mathbb{Q}$-measure parameters.
- ζ^2 is variance of noise in the measurement equation.
- State equations depend on $\mathbb{P}$-measure parameters.
- $\varepsilon^{(y)}, \varepsilon^{(\lambda)}$ are iid standard normal.
- CIR special case obtained by fixing $\chi_s = \Delta$ and $\alpha = \infty$.

Parameter estimates: Across models

$\kappa^{\mathbb{P}}$	σ	100μ	$\kappa^{\mathbb{Q}}$	ζ	α
Bank of America: CIR Model					
0.4502	0.1665	0.0298	-0.3643	0.2262	
(0.3397)	(0.0016)	(0.0005)	(0.0051)	(0.0013)	
Bank of America: CIR + IG Time Change Model					
0.6986	0.1729	0.0212	-0.5063	0.2186	2.0763
(0.5406)	(0.0014)	(0.0005)	(0.0074)	(0.0013)	(0.0733)
Goldman Sachs: CIR Model					
0.4516	0.1847	0.0933	-0.2504	0.2164	
(0.3367)	(0.0030)	(0.0014)	(0.0070)	(0.0013)	
Goldman Sachs: CIR + IG Time Change Model					
0.6985	0.2216	0.0699	-0.4261	0.2100	2.7927
(0.5422)	(0.0019)	(0.0014)	(0.0100)	(0.0012)	(0.1705)

MCMC smoothed state estimates: Goldman Sachs

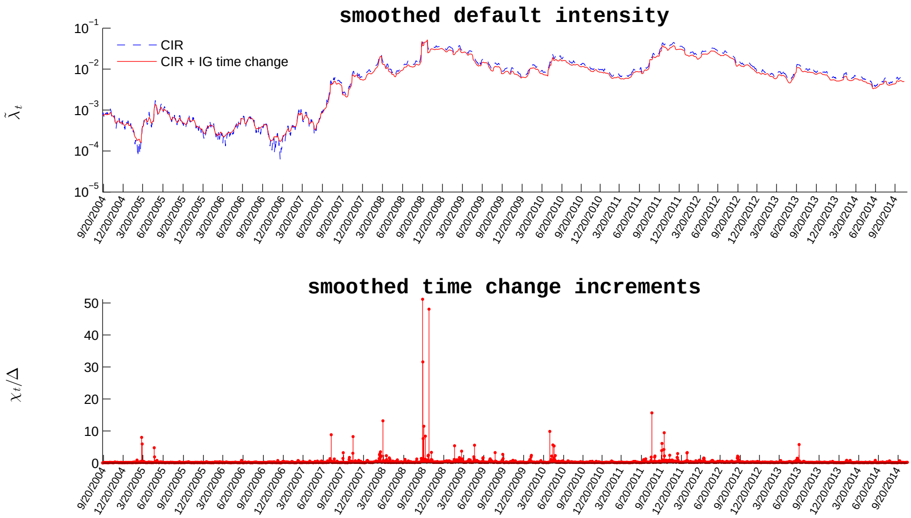

Residual diagnostics: measurement equation

Moments of measurement equation residuals $\varepsilon^{(y)}$

Mean	Std Dev	Skewness	Kurtosis
Bank of America: CIR Model			
-.0195	.9996	-.4809	4.3952
Bank of America: CIR + IG Time Change			
-.0042	1.0000	-.2526	4.5424
Goldman Sachs: CIR Model			
-.0140	.9997	-.4940	4.2928
Goldman Sachs: CIR + IG Time Change			
-.0031	1.0000	-.2313	4.2111

Residual diagnostics: state equation

Moments of state equation residuals $\varepsilon^{(\lambda)}$

Mean	Std Dev	Skewness	Kurtosis
Bank of America: CIR Model			
.0351	.5191	.1529	3.8284
Bank of America: CIR + IG Time Change			
.0019	.9325	.0077	3.0746
Goldman Sachs: CIR Model			
.0408	.6207	.2000	4.2336
Goldman Sachs: CIR + IG Time Change			
-.0019	.9348	.0142	3.0672

Testing out-of-sample performance

- Diebold and Mariano (JBES 1995) and Amisano and Giacomini (JBES 2007) test statistic based on predictive likelihood p_M evaluated at realized return realizations:

$$D_M = \frac{1}{T-1} \sum_{t=1}^{T-1} \log(p_M(y_{t+1}|\mathcal{F}_t, \Theta)),$$

  where $\mathcal{F}_t = \{y_1, y_2, \ldots, y_t\}$ is the information set available at time t.
- Predictive density p_M estimated by auxiliary particle filter (Pitt and Shephard, JTSA 1999).
  - Fix parameter Θ to mean of (in-sample) posterior distribution.
  - Propagation of default intensity h_t adapted to y_t using auxiliary linearized model, similar to our MCMC strategy.
  - Sampling of time-change increments χ_t non-adapted.
- For models A and B statistic $\widehat{D}_A - \widehat{D}_B$ is asymptotically normal.

Diebold-Mariano test statistics

Test favors IG time-change model over standard CIR default intensity model for all five banks, and for 9 of 10 non-financials.

Banks	U	Non-financials	U
Bank of America	27.95	Alcoa	10.98
Citi	37.13	Anadarko	6.47
Goldman Sachs	24.01	Century Link	17.65
JP Morgan	22.82	Ford	19.44
Morgan Stanley	34.62	Lennar	-14.73

U-statistic is standard normal under H_0 of equal performance.

Cross-firm correlations in smoothed states

Correlation in log intensities: Above the diagonal,
Correlation in log time increments: *Below the diagonal*.

		BAC	GS	Ford
Pre-crisis	BAC		0.6086	-0.3268
09/2004 – 06/2007	GS	*0.3112*		-0.1637
	Ford	*0.1407*	*0.1169*	

		BAC	GS	Ford
Crisis	BAC		0.9024	0.8232
07/2007 – 06/2009	GS	*0.6022*		0.8866
	Ford	*0.2683*	*0.1933*	

		BAC	GS	Ford
Post-crisis	BAC		0.9447	0.6769
07/2009 – 10/2014	GS	*0.7387*		0.5995
	Ford	*0.3476*	*0.4362*	

Common factor in default intensities

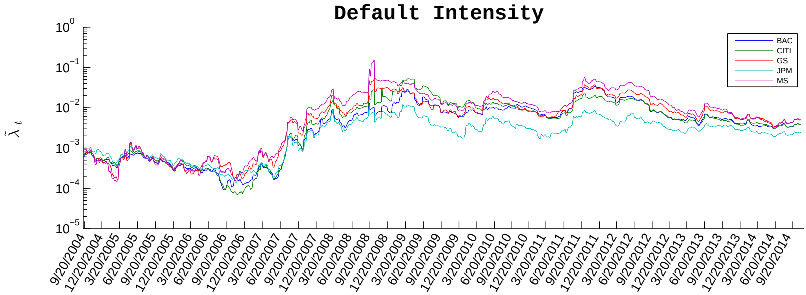

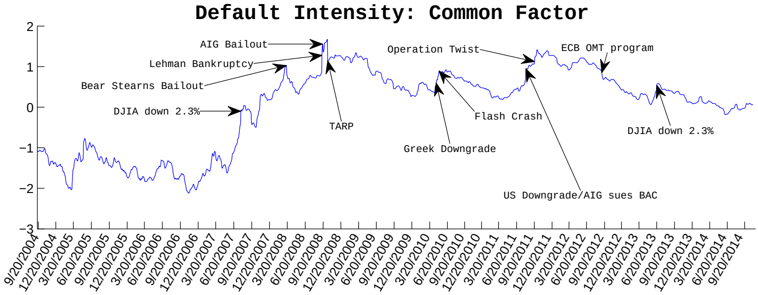

Common factor in stochastic time increments

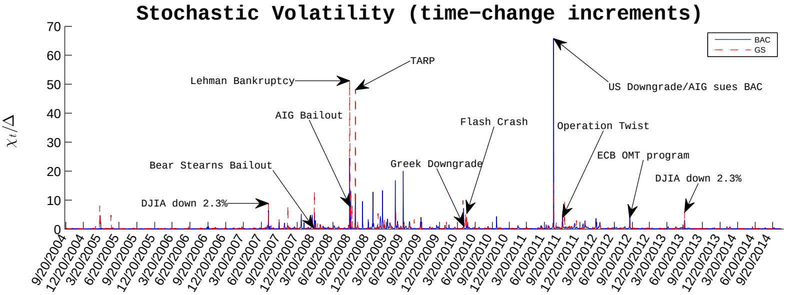

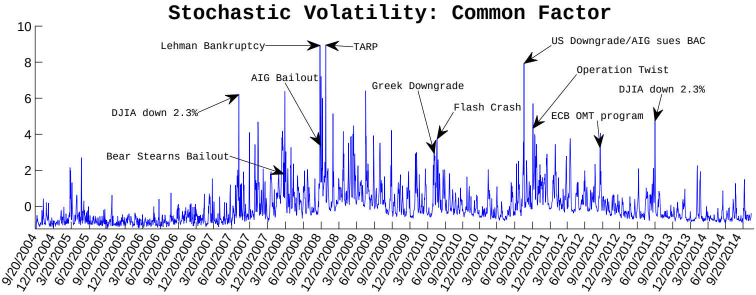

Total communality

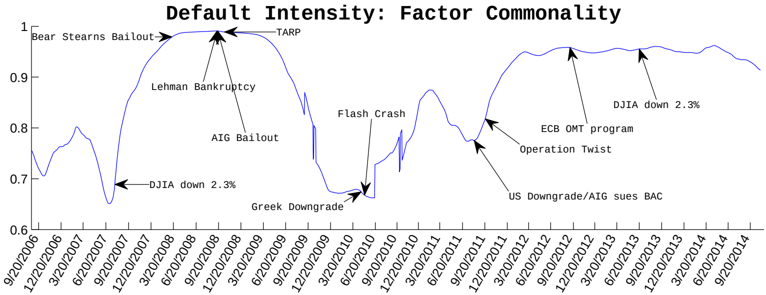

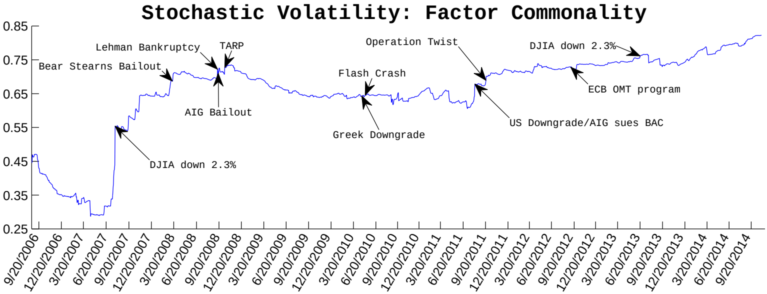

Stylized facts for a multiname model

Features for a new multiname model:

- Common factors in business-time default intensity and in stochastic time-change.
- Time-varying correlations.
- Persistence in time-change increments, i.e., volatility clustering.

Duffie-Gârleanu specification

Duffie and Gârleanu (FAJ 2001) suggest a parsimonious specification for the background default intensities. For each bank $j = 1, \ldots, B$:

$$\lambda_{j,t} = Z_t^* + Z_t^j$$

- Common factor Z_t^* and idiosyncratic Z_t^j are independent CIR processes.
- If $\kappa_*^{\mathbb{Q}} = \kappa_j^{\mathbb{Q}} = \kappa^{\mathbb{Q}}$ and $\sigma_* = \sigma_j = \sigma$, then $\lambda_{j,t}$ is CIR with $\mathbb{Q}$-parameters $\mu_* + \mu_j$, $\kappa^{\mathbb{Q}}$, σ.
- Flexibility with respect to μ_*, μ_j allows heterogeneity in long-run mean CDS spreads.
- Tractability of estimation does not require that $\lambda_{j,t}$ be CIR under $\mathbb{P}$, so no need to restrict $\kappa_*^{\mathbb{P}} = \kappa_j^{\mathbb{P}}$. Due to weak identification of $\kappa^{\mathbb{P}}$, this is harmless to impose in practice.
- Correlations depend on $Z_t^*/(Z_t^* + Z_t^j)$, so vary over time and across banks.

Parameter estimates: Across dealer banks

	$\kappa^{\mathbb{P}}$	σ	100μ	$\kappa^{\mathbb{Q}}$	ζ	α
BAC	0.6986	0.1729	0.0212	-0.5063	0.2186	2.0763
	(0.5406)	(0.0014)	(0.0005)	(0.0074)	(0.0013)	(0.0733)
CITI	0.4852	0.1663	0.0201	-0.4720	0.2243	2.2752
	(0.3749)	(0.0015)	(0.0005)	(0.0072)	(0.0013)	(0.1114)
GS	0.6985	0.2216	0.0699	-0.4261	0.2100	2.7927
	(0.5422)	(0.0019)	(0.0014)	(0.0100)	(0.0012)	(0.1705)
JPM	1.3076	0.1797	0.0237	-0.6558	0.2008	1.7736
	(1.0083)	(0.0012)	(0.0006)	(0.0081)	(0.0012)	(0.0380)
MS	0.3847	0.2209	0.0827	-0.3598	0.2213	3.5631
	(0.2982)	(0.0025)	(0.0016)	(0.0104)	(0.0013)	(0.2466)

Common activity rate

- For parsimony, we assume a single common time-change process T_t for all banks.
- To get persistence, we model T_t as time-integral of an activity rate ν_t.

$$T_t = \int_0^t \nu_s \, ds$$

- ν_t is the speed of business time with respect to calendar time. Assume CIR specification.
- Series expansion for $\tilde{S}(t)$ works as before, but simply with a different $M_t(u)$ transform.
- This model won't accommodate "spikes" as easily as IG time-change.

Estimation strategy

- Introducing persistence in all factors makes the model well-suited to particle filtering, less amenable to our SMP MCMC estimator.
- We propose to use particle marginal Metropolis-Hastings (PMMH) of Andrieu, Doucet and Holenstein (JRSS 2010).
  - Special case of particle Markov chain Monte Carlo.
  - Unlike the SMP estimator, needs little customization and tuning.
  - Can propagate λ_t under exact transition density.
- In a nutshell, PMMH has
  - Outer MCMC loop over parameter vector Θ.
  - Inner PF over states to evaluate likelihood, which drives accept/reject of the MH proposal in outer loop.
- We have implemented estimator for model with no time-change, tested on simulated data. Computationally demanding, but works beautifully!