International Journal of Theoretical and Applied Finance
Vol. 17, No. 3 (2014) 1450017 (14 pages)

DOI: 10.1142/S0219024914500174

A SPREAD-RETURN MEAN-REVERTING MODEL FOR CREDIT SPREAD DYNAMICS

BRENDAN O'DONOGHUE[*,¶], MATTHEW PEACOCK[†,‖],
JACKY LEE[‡,**] and LUCA CAPRIOTTI[§,††]

[*]*Electrical Engineering Department*
Stanford University

[†]*White Oak Asset Management SA*

[‡]*Quantitative Strategies*
Investment Banking Division
Credit Suisse Group

[§]*Quantitative Strategies*
Investment Banking Division
Credit Suisse Group
[¶]*bodonoghue85@gmail.com*
[‖]*matthew.peacock@whiteoakam.com*
[**]*shinghoi.lee@credit-suisse.com*
[††]*luca.capriotti@credit-suisse.com*

Received 6 August 2013
Accepted 20 January 2014
Published 12 May 2014

In this paper, we propose a novel, analytically tractable, one-factor stochastic model for the dynamics of credit default swap (CDS) spreads and their returns, which we refer to as the spread-return mean-reverting (SRMR) model. The SRMR model can be seen as a hybrid of the Black–Karasinski model on spreads and the Ornstein–Uhlenbeck model on spread returns, and is able to capture empirically observed properties of CDS spreads and returns, including spread mean-reversion, heavy tails of the return distribution, and return autocorrelations. Although developed for modeling CDS spreads, the SRMR model has applications for many other stochastic processes with similar empirical properties, including more general rate processes.

Keywords: Credit default swaps; credit risk; risk management; return autocorrelation; heavy tails; model fitting.

1. Introduction

Much of the literature on portfolio credit risk has been concentrated on building realistic models of default risk and correlation among defaults (O'Kane 2008). This was, in part, due to the popularity enjoyed by default correlation products, such as basket default swaps and collateralized debt obligations (CDOs), in the years immediately preceding the 2008 subprime crisis. Nonetheless, investors in bonds, loans, credit default swaps (CDS) and other simple credit derivatives are exposed to price fluctuations even in the absence of defaults. This is because changes in market perceptions of the credit worthiness of an obligor affect the risk premium demanded by investors, and therefore the obligor's credit spread.

Successful management of spread risk, either in derivatives pricing, portfolio optimization or Value at Risk (VaR) applications, relies on a realistic description of the credit spread dynamics. However, most of the models for credit spreads proposed in the past, like the celebrated Cox–Ingersoll–Ross (CIR) model (Cox *et al.* 1985 and Duffie & Singleton 1999) and its generalizations (Brigo & Alfonsi 2005), prioritized analytic tractability over an accurate description of the empirical properties of credit spreads. A notable exception is the recent work of Cont & Kan (2011) which undertook a systematic study of the empirical properties of CDS spreads. This work highlighted several important properties of the dynamics of CDS spreads, including stationarity of returns, positive autocorrelations, and two-sided heavy tailed distributions. In order to fit these properties, these authors proposed a discrete-time AR(1)-GARCH model.

In this paper, we continue this line of research and propose a simple and analytically tractable model that is able to capture the principal features of the CDS spread dynamics observed in practice. We refer to this model as the *spread-return mean-reverting* (SRMR) model. The SRMR model can be seen as a hybrid of the Black–Karasinski (BK) model for spreads and the Ornstein–Uhlenbeck (OU) model for spread returns; it shares their analytical tractability while overcoming some of their weaknesses.

There are many competing frameworks in the literature that can be used to model spreads. Most of these were first used to model interest rates but have since been applied, in the literature or in practice, to CDS spreads. The first important model on interest rates is the Vasicek model (1977), which is an OU model for the instantaneous rate. If the parameters that define the model are allowed to vary with time then the model is sometimes referred to as the Hull–White model (1990) or the extended Vasicek model. Ho & Lee (1986) modeled interest rates as a time varying drift-diffusion process. Rendleman & Bartter (1980) proposed a model on interest rates using a geometric Brownian motion process. The CIR (Cox *et al.* 1985, Duffie & Singleton 1999, and Brigo & Alfonsi 2005) model, which is popular among practitioners, is closely related to the Vasicek model family. The difference between the CIR model and others is that the volatility depends on the square-root of the rate, which precludes negative interest rates. Black & Karasinski (1991), proposed a model on the log-rates, which can be interpreted as a log-normal application of the Hull–White model. For (much) more details on these and other models see, e.g., Brigo & Mercurio (2006), Hull (2012), and James & Webber (2000).

Throughout the paper we shall denote the CDS spread at time t as s_t, and the (log-)return of the spread at time t as r_t, i.e. $r_t = d\log(s_t)/dt$ in a continuous time setting and $r_t = \log(s_t/s_{t-1})$ in a discrete-time setting. We use log-returns so as to preclude the possibility of negative spreads, as we shall see. We consider daily returns, so throughout the paper t will have units of days.

1.1. *Outline*

The paper is structured as follows. In Sec. 2, we introduce the SRMR model. In Sec. 3, we derive closed-form expressions for the value of the spread and return at any time, under the model. In Sec. 4, we discuss the connections of the SRMR model to other well-known models in the literature. In Sec. 5, we discuss ways of fitting the parameters in the SRMR model to data. In Sec. 6, we present a real-world example of using the model to calculate VaR.

2. The SRMR Model

Cont & Kan (2011) show that modeling CDS spreads directly does not capture the full statistical properties of the empirically observed process. Instead they propose modeling the spread *returns*. Here, we continue this line of research and we propose a simple analytically tractable continuous time model for spread returns that captures many of the empirical properties of both spreads and returns. The starting point is a simplified continuous time, diffusive version of the model proposed in Cont & Kan (2011), namely an OU process on returns of the form

$$dr_t = (\gamma - \alpha r_t)dt + \sigma dW_t, \tag{1}$$

where γ, α and σ are constant parameters and W_t is a standard Wiener process. While this model is able to capture autocorrelation of returns, it results in undesirable credit spread dynamics. As illustrated in Fig. 1 and shown analytically below, the process (1) is characterized by a variance that is unbounded with time. This is because there is nothing in the model specification that forces the *spread* to revert to a long term mean. In order to address this problem, we modify the process in Eq. (1) as follows:

$$dr_t = \left(\gamma - (\alpha + \beta)r_t - \alpha\beta \int_0^t r_\tau d\tau\right) dt + \sigma dW_t, \tag{2}$$

where again γ, α, β and σ are constant parameters and W_t is a standard Wiener process. We refer to model in Eq. (2) as the SRMR model. The SRMR model is an extension of the OU model on returns (1) that incorporates a mean-reversion term on the integral of the log-return, which is simply the change in the log-rate from time $t = 0$, since by our definition of return we have

$$s_t = s_0 \exp\left(\int_0^t r_t\right).$$

This ensures that the spreads are always nonnegative under the model. Thus *both* the spreads and returns are explicitly taken into account by the model.

The SRMR model has a seemingly complicated form; the return at time t depends on the entire history of the return from 0 to t through the integral term. However, the SRMR model is essentially just as tractable as the BK or OU models. First, if we consider the state of the process to be $y_t = (r_t, \int_0^t r_\tau d\tau)^T \in \mathbf{R}^2$, then

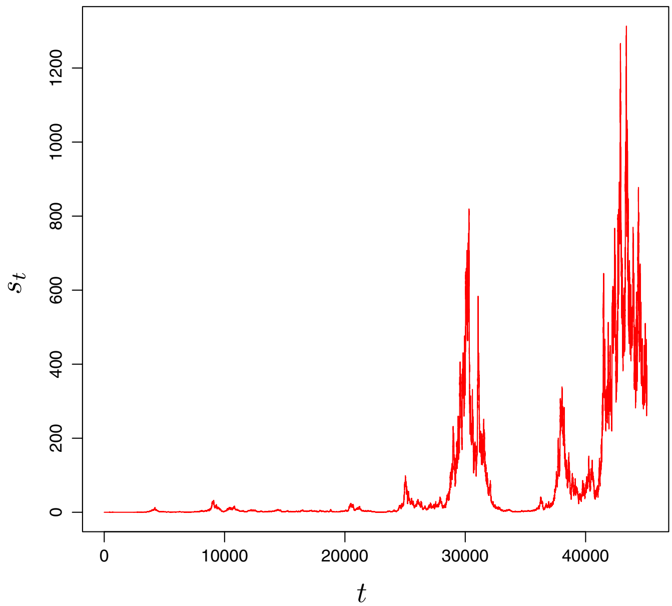

Fig. 1. A sample path for the spread as implied by the OU model for the log-returns, equation (1), illustrating unbounded variance.

we can write a stochastic differential equation governing y_t that has the Markov property. Indeed, we can rewrite the process as

$$dy_t = (\Gamma - Ay_t)dt + \Sigma dW_t,$$

where

$$\Gamma = \begin{bmatrix} \gamma \\ 0 \end{bmatrix}, \quad A = \begin{bmatrix} \alpha+\beta & \alpha\beta \\ -1 & 0 \end{bmatrix}, \quad \Sigma = \begin{bmatrix} \sigma \\ 0 \end{bmatrix} \tag{3}$$

and the initial condition reads $y_0 = (r_0, 0)^T$. Thus the SRMR model is a two-dimensional (but one-factor) OU process (Gardiner 2009) with a very particular choice of parameters that give rise to certain desirable properties. Consequently, the process has a tractable closed-form expression for the rate at any time, which we derive below.

3. Closed-Form Expressions

If $\alpha, \beta > 0$ and $\alpha \neq \beta$ then we have

$$y_t = e^{-At}y_0 + (I - \mathrm{e}^{-At})A^{-1}\Gamma + \int_0^t e^{A(\tau-t)}\Sigma dW_\tau$$

(an analogous formula can be derived for the case where $\alpha = \beta$). More explicitly, the evolution of the return r_t and the change in the log-spread $\log(s_t/s_0) = \int_0^t r_\tau d\tau$

are given by

$$r_t = \gamma(\mathrm{e}^{-\alpha t} - \mathrm{e}^{-\beta t})/(\beta - \alpha) + r_0(\beta\mathrm{e}^{-\beta t} - \alpha\mathrm{e}^{-\alpha t})/(\beta - \alpha)$$
$$+ \sigma \int_0^t ((\beta\mathrm{e}^{\beta(\tau-t)} - \alpha\mathrm{e}^{\alpha(\tau-t)})/(\beta - \alpha))dW_\tau$$

and

$$\int_0^t r_\tau d\tau = (\gamma/\alpha\beta)(1 - (\beta\mathrm{e}^{-\alpha t} - \alpha\mathrm{e}^{-\beta t})/(\beta - \alpha)) + r_0(\mathrm{e}^{-\alpha t} - \mathrm{e}^{-\beta t})/(\beta - \alpha)$$
$$+ \sigma \int_0^t ((\mathrm{e}^{\alpha(\tau-t)} - \mathrm{e}^{\beta(\tau-t)})/(\beta - \alpha))dW_\tau. \tag{4}$$

The asymptotic behavior of the mean is given by

$$\mathbf{E}(y_t) \to A^{-1}\Gamma = \begin{bmatrix} 0 \\ \gamma/\alpha\beta \end{bmatrix},$$

confirming that spreads are mean reverting. The variance is given by

$$\mathbf{var}(y_t) = \sigma^2/2 \begin{bmatrix} \dfrac{1 - f(t)}{\alpha + \beta} & \dfrac{g(t)}{(\beta - \alpha)^2} \\ \dfrac{g(t)}{(\beta - \alpha)^2} & \dfrac{1 - p(t)}{\alpha\beta(\alpha + \beta)} \end{bmatrix},$$

where

$$f(t) = \frac{(\beta\mathrm{e}^{-\beta t} - \alpha\mathrm{e}^{-\alpha t})^2 + \alpha\beta(\mathrm{e}^{-\alpha t} - \mathrm{e}^{-\beta t})^2}{(\beta - \alpha)^2},$$

$$p(t) = \frac{(\alpha\mathrm{e}^{-\beta t} - \beta\mathrm{e}^{-\alpha t})^2 + \alpha\beta(\mathrm{e}^{-\alpha t} - \mathrm{e}^{-\beta t})^2}{(\beta - \alpha)^2}$$

and

$$g(t) = (\mathrm{e}^{-\alpha t} - \mathrm{e}^{-\beta t})^2.$$

Since $f(t), p(t), g(t) \to 0$ the variances of both the spread and return are bounded for $\alpha, \beta > 0$. However, this is lacking in the OU model for returns (1): if we take $\beta = 0$ we recover the process on returns described by Eq. (1), and if we take the limit of the variance as $\beta \to 0$, we obtain

$$\mathbf{var}\left(\int_0^t r_s ds\right) = \frac{\sigma^2}{\alpha^2}\left(t - \frac{3}{2\alpha} + \frac{4\mathrm{e}^{-\alpha t} - \mathrm{e}^{-2\alpha t}}{2\alpha}\right),$$

which grows with time, as illustrated in Fig. 1.

4. Connection to other Processes

It is clear that setting $\beta = 0$ in the SRMR model (2) yields the OU model on returns (1). Performing an eigenvalue decomposition on the matrix A in Eq. (3) allows us to decompose the above system into two one-dimensional stochastic processes corresponding to the eigenstates of the system. This implies that the process r_t can be written (for $\beta \neq \alpha$)

$$r_t = (\beta b_t - \alpha a_t)/(\beta - \alpha),$$

where

$$\begin{aligned} da_t &= (\gamma - \alpha a_t)dt + \sigma dW_t, \\ db_t &= (\gamma - \beta b_t)dt + \sigma dW_t \end{aligned}$$

and where $r_0 = b_0 = a_0$, i.e. the process (2) can be interpreted as the weighted difference of two standard OU processes.

We can also recover the BK model (Black & Karasinki 1991) as a special case of the SRMR model. By fixing $\alpha > 0$, letting $\beta \to \infty$ and varying γ, σ so that $\gamma/\alpha\beta \equiv \hat{\mu}$ and $\sigma/(\beta - \alpha) \equiv \hat{\sigma}$, where $\hat{\mu}$ and $\hat{\sigma}$ are some constants, Eq. (4) gives

$$\int_0^t r_s ds = \hat{\mu}(1 - \mathrm{e}^{-\alpha t}) + \hat{\sigma} \int_0^t \mathrm{e}^{\alpha(\tau - t)} dW_\tau,$$

and therefore $\int_0^t r_s ds$ must obey the process

$$d\left(\int_0^t r_s ds\right) = \alpha\left(\hat{\mu} - \int_0^t r_s ds\right) dt + \hat{\sigma} dW_t.$$

If we define $X_t = \log s_t = \int_0^t r_s ds + \log s_0$, then the following process describes the evolution of X_t:

$$\begin{aligned} dX_t &= \alpha(\hat{\mu} - X_t + X_0)dt + \hat{\sigma} dW_t \\ &= \alpha(\mu - X_t)dt + \hat{\sigma} dW_t \end{aligned}$$

with $\mu = \hat{\mu} + X_0$, which is precisely the BK model.

In summary, the BK model for spreads and the OU model on returns are in fact limiting cases of the SRMR model (2). Therefore, we may expect that for intermediate values of α and β we would obtain a hybrid between the BK model and OU model for returns and spreads, maintaining the desirable analytic tractability of both and eliminating some of the weaknesses.

5. Parameter Estimation

In this section, we detail how to fit parameters to the SRMR model given the observation data. The variant we consider is the SRMR model with jumps:

$$dr_t = \left(\gamma - (\alpha + \beta) r_t - \alpha\beta \int_0^t r_\tau d\tau\right) dt + \sigma dW_t + dZ_t, \tag{5}$$

where Z_t is an independent compound Poisson process with rate λ and some jump-size distribution that we make no assumptions on *a priori*. For more discussion on the use of jumps to model CDS spreads see, e.g., Madan & Schoutens (2008) and Martin (2009). Due to the assumed independence of the jump and diffusive component the moment generating function of the process is still available in closed form, thus making the implementation of the model in Monte Carlo or multinomial trees straightforward. The incorporation of jumps allows the model to better fit the heavy-tails of the data, as observed in Cont & Kan (2011).

The task of fitting the parameters of the process (5) to observed data can be cast as a convex optimization problem; for alternative approaches for fitting Lévy processes to empirical data see, e.g., Barndorff-Nielsen (1997), Barndorff-Nielsen & Shephard (2003), and Schoutens (2003). Given N samples of the returns at times t_i, $i = 1, \ldots, N$, with time difference between successive samples given by Δt, the model (5) in discrete time is written

$$\Delta r_i = \left(\gamma - (\alpha + \beta) r_i - \alpha\beta \sum_{j=0}^{i} r_j \Delta t\right) \Delta t + \sigma w_i + z_i, \tag{6}$$

where r_i is the return at time t_i, $w_i \sim \mathcal{N}(0, \Delta t)$ is the diffusive component of the innovation at time t_i and z_i is the jump sample at time t_i.

Equation (6) can be written in a simpler form if we make the following substitutions; let $x_i = \Delta r_i$,

$$y_i = \left(1, r_i, \sum_{j=0}^{i} r_j \Delta t\right)^T$$

and

$$\theta = (\gamma \Delta t, -(\alpha + \beta)\Delta t, -\alpha\beta\Delta t)^T.$$

Now we can rewrite (6) as

$$x_i = \theta^T y_i + \sigma w_i + z_i. \tag{7}$$

We assume that $\lambda \Delta t$ is small, and thus only one jump is likely to occur in each Δt interval, and the probability of observing a jump in an interval is approximately $\lambda \Delta t$. Since we make no assumptions on the jump-size distribution we take the following uninformative prior on z_i:

$$p(z_i) = \begin{cases} \lambda \Delta t & \text{if } z_i \neq 0, \\ 1 - \lambda \Delta t & \text{if } z_i = 0, \end{cases}$$

or equivalently

$$p(z_i) \propto \exp(-c \mathcal{I}_{z_i \neq 0}(z_i)),$$

where $c = -\log((\lambda \Delta t)/(1 - \lambda \Delta t))$ and

$$\mathcal{I}_{z_i \neq 0}(z_i) = \begin{cases} 1 & \text{if } z_i \neq 0, \\ 0 & \text{if } z_i = 0, \end{cases}$$

is the indicator function for nonzero z_i. From (7) we have that $w_i = (x_i - \theta^T y_i - z_i)/\sigma \sim \mathcal{N}(0, \Delta t)$, and therefore the likelihood of observing data sample i is given by

$$p(x_i \,|\, \theta, y_i, z_i) \propto \exp(-(x_i - \theta^T y_i - z_i)^2/2\sigma^2 \Delta t).$$

Each sample, x_i, can be decomposed into the deterministic drift component $\theta^T y_i$, the diffusion component w_i and the jump component z_i, in a maximum *a posteriori* sense given our prior on z_i. The problem of fitting the parameters to the data can be written in this notation as

$$\text{maximize} \prod_{i=1}^{N} p(x_i \,|\, \theta, y_i, z_i) p(z_i).$$

Taking the negative logarithm yields

$$\text{minimize } \|x - Y^T\theta - z\|_2^2 + \mu \sum_{i=1}^{N} \mathcal{I}_{z_i \neq 0}(z_i), \tag{8}$$

where $Y = (y_1, y_2, \ldots, y_N)$ is the matrix of data samples stacked columnwise, $\mu = 2\sigma^2 c\Delta t$, and $\|\cdot\|_2$ denotes the Euclidean norm. The variables in the above optimization problem are $\theta \in \mathbf{R}^3$ and $z \in \mathbf{R}^N$ ($x \in \mathbf{R}^N$ and $Y \in \mathbf{R}^{3\times N}$ are known data). Problem (8) is hard to solve in general, due to the nonconvexity of the term involving the indicator function. One common approach to approximately solve this problem is to replace this term with the ℓ_1 norm, which is a good approximation in many cases, see, e.g., Tibshirani (1994), Candès *et al.* (2006), Donoho (2006), and Candès & Wakin (2008). The ℓ_1 norm is often referred to as a *sparsity promoting* penalty function, as the resulting vector will typically be sparse (which is what we expect for a jump process). This simplification yields

$$\text{minimize} \|x - Y^T\theta - z\|_2^2 + \mu \|z\|_1. \tag{9}$$

Problem (9) is jointly convex in θ and z and thus can be solved efficiently using modern methods, (Nesterov & Nemirovsky 1994, Wright 1997, Sturm 1999, Toh *et al.* 1999, Boyd & Vandenberghe 2004, Grant & Boyd 2011, and O'Donoghue *et al.* 2013). The parameter $\mu \in \mathbf{R}_+$ must be chosen in advance, however it can be selected using a back-testing or cross-validation procedure.

Let $\theta^\star$ and $z^\star$ be the minimizers of the optimization problem (9). From $\theta^\star$ we can extract estimates of γ, α and β, the parameters that specify the SRMR model (5). The nonzero entries of $z^\star$ are the jump locations, which can be used to estimate the jump rate λ and the jump-size distribution. Finally, since $(x_i - \theta^T y_i - z_i) \sim \mathcal{N}(0, \sigma^2 \Delta t)$, the standard deviation of $(x - Y^T\theta - z)/\sqrt{\Delta t}$ yields an estimate for the diffusion noise level σ.

Figure 2 shows the 5Y CDS spread and return time series for Pfizer Inc. from February 2008 to September 2010, in this case Δt was one day. Figure 3 shows the results of the maximum *a posteriori* decomposition; the top trace is the deterministic drift, the middle trace is the diffusion and the bottom trace are the jumps.

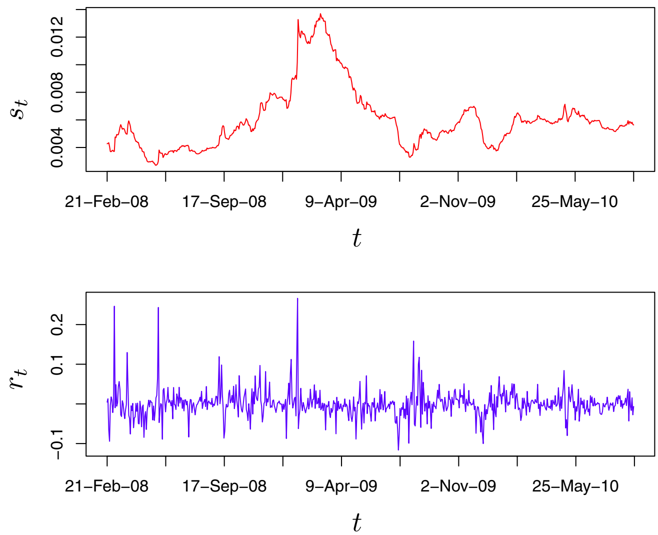

Fig. 2. 5Y CDS spread and return time series for Pfizer Inc. for the period February 2008 to June 2010.

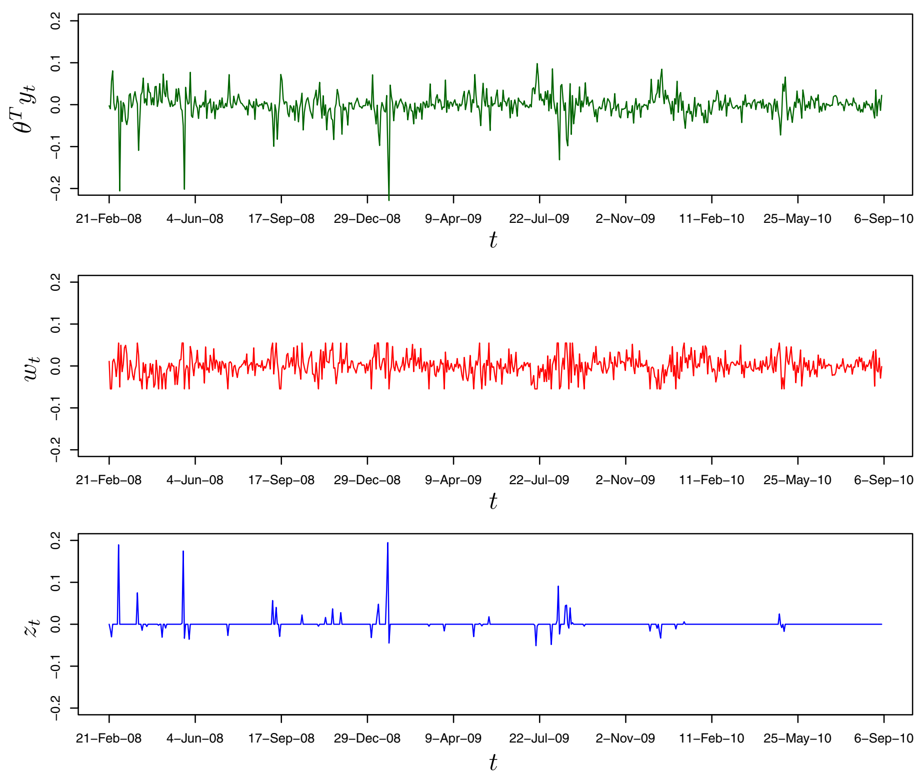

Fig. 3. Return time series for Pfizer Inc. of Fig. 2 decomposed from top to bottom into drift, diffusion and jump processes.

5.1. *Comparison with the empirical properties of CDS spreads*

After fitting the parameters of the SRMR model (2) to the historical time series as described above, we can simulate the process and compare to the empirical returns. Figure 4 shows the return autocorrelations generated by the SRMR model for the 5Y CDS spreads of Southwest Airlines Co. The comparison with the partial

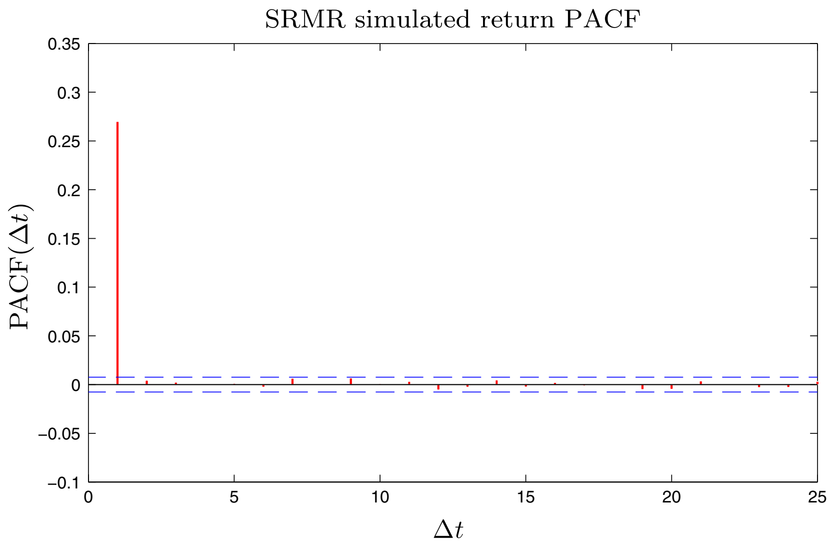

Fig. 4. The partial autocorrelation function estimated from samples generated by the SRMR model fit to Southwest 5Y CDS data for the period February 2008 to August 2011.

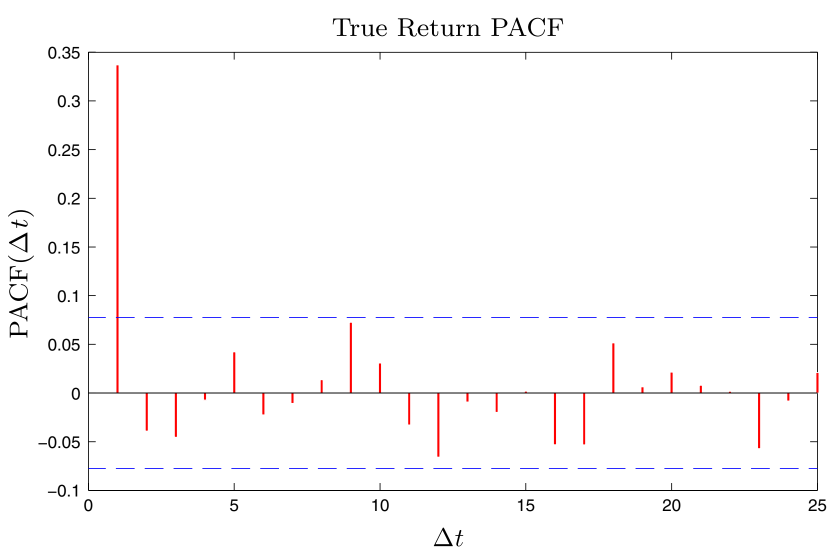

Fig. 5. The true partial autocorrelation function for the 5Y CDS spread returns of Southwest Airlines Co. for the period February 2008 to August 2011.

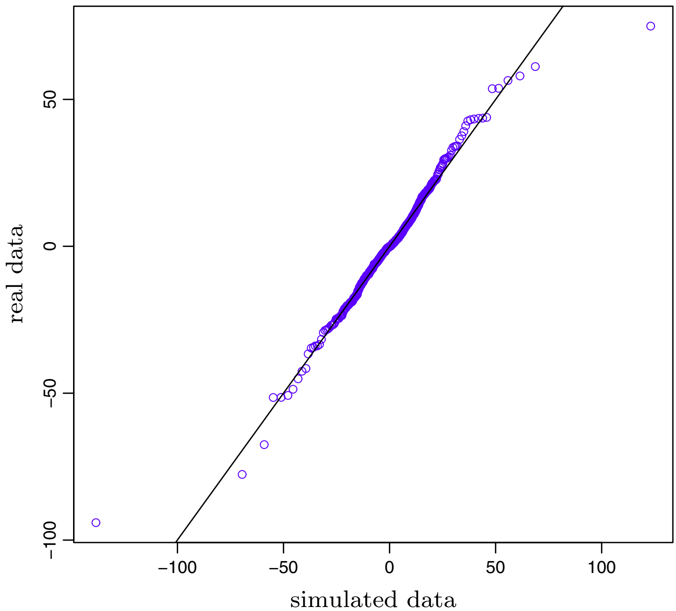

Fig. 6. Q–Q plot of real versus simulated return innovations, for the 5Y CDS spreads of IBM during the period September 2008 to August 2011.

autocorrelation function of Fig. 5 indicate that the SRMR model is able to reproduce the same structure of return autocorrelations observed empirically. The BK model on spreads is unable to capture this autocorrelation on returns.

Figure 6 shows the Q–Q-plot of real return data and the simulated return data from the SRMR model with Laplace-distributed jump sizes for 5YR CDS spreads of IBM. The goodness-of-fit is evidenced by the straight line fit with slope of approximately one. This suggests that the distribution of returns implied by the model is very close to what we observe in the empirical data.

6. An Example Application: VaR

An accurate description of the statistical properties of CDS spreads is important for obtaining meaningful risk quantification, e.g., in VaR measures (Alexander 2009). To test the SRMR model we selected 20 5YR CDS spread series and fit three models to each name using empirical data: the BK model, the BK model with jumps (BK + J) and the SRMR model with jumps (SRMR + J); we used a Laplace distribution to model the jump sizes in all cases. We then created six portfolios of those same 20 names; one portfolio was all long positions (protection buyer), one all short (protection seller) and four were of randomly mixed long-short positions. We simulated the spreads using the calibrated models and we computed the present value of each portfolio implied by the simulated spreads using the ISDA CDS standard model (O'Kane 2008). Using the obtained time series of 1-day P&Ls we estimated

Table 1. Proportion of days each portfolio experienced losses exceeding the 95% VaR threshold during the period February 2008 to August 2011.

	BK (%)	BK + J (%)	SRMR + J (%)
Long 20	3.5	8.1	6.3
Short 20	3.3	6.2	5.4
Random mixed A	1.6	2.4	3.3
Random mixed B	4.0	3.2	4.3
Random mixed C	4.1	5.1	5.2
Random mixed D	1.9	7.1	3.2
Average	3.1	5.4	4.7
Std. error	2.2	2.1	1.2

the 95% 1-day VaR threshold (Alexander 2009) for each portfolio under each model, and compared the results obtained to the historical data. In Table 1, we show the fraction of the time that the portfolios had real daily losses greater than the estimated 95% VaR threshold (a perfect model with infinite data would have 5% in every row). Note that although the portfolios are not real, we are using real data to both fit the models and estimate their performance.

As shown in Table 1, the BK model slightly overestimates the risk, as the variance of the diffusion fit is very large in order to 'explain' the jumps. Adding jumps to BK improves the performance somewhat, resulting in a better average and a slightly smaller standard error. However, the SRMR model outperforms both other models, having an average closer to 5% and a significantly smaller standard error.

7. Conclusions

An accurate description of credit spreads dynamics is essential for the effective risk management of portfolios of credit derivatives. In this paper, we developed a continuous-time, one-factor stochastic model for the dynamics of CDS spreads and their returns, dubbed the SRMR model. The SRMR model can be seen as a hybrid of the BK model for spreads and the OU model for spread returns, and is able to capture the desirable properties of both, while overcoming many of their respective weaknesses. As such the SRMR model is better able to capture important statistical properties that we observe in real data. These include nonnegativity of spreads, spread and return mean-reversion, bounded variance with time, heavy tails of the return distribution and return autocorrelation. Although we presented an application in the context of VaR calculation for CDS portfolios, the proposed model is also suitable for other applications, including portfolio optimization and derivative pricing.

Finally, although we focused on the use of the SRMR model solely in the context of modeling CDS spreads and returns, we note that it has applications in other areas, in particular for interest rate modeling.

Acknowledgments

We are very grateful to Lukasz Wiechecki for a careful reading of the paper and useful feedback. We are also thankful to an anonymous reviewer for his/her insightful comments and suggestions.

The opinions and views expressed in this paper are uniquely those of the authors, and do not necessarily represent those of Credit Suisse Group.

References

C. Alexander (2009) *Market Risk Analysis*. Chichester: J. Wiley & Sons.

O. Barndorff-Nielsen, (1997) Normal inverse gaussian distributions and stochastic volatility modelling, *Scandinavian Journal of Statistics* **24** (1), 1–14.

O. Barndorff-Nielsen & N. Shephard (2003) Integrated OU processes and non-Gaussian OU-based stochastic volatility models, *Scandinavian Journal of Statistics* **30** (2), 277–295.

F. Black & P. Karasinki (1991) Bond and option pricing when short rates are lognormal, *Financial Analysts Journal* **47** (4), 52–59.

S. Boyd & L. Vandenberghe (2004) *Convex Optimization*. Cambridge, UK: Cambridge University Press.

D. Brigo & A. Alfonsi (2005) Credit default swaps calibration and option pricing with the SSRD stochastic intensity model, *Finance and Stochastics* **9** (1), 29–42.

D. Brigo and F. Mercurio (2006) *Interest Rate Models-Theory and Practice*: *With Smile, Inflation and Credit*. Berlin: Springer.

E. Candès, J. Romberg & T. Tao (2006) Stable signal recovery from incomplete and inaccurate measurements, *Communications on Pure and Applied Mathematics* **59** (8), 1207–1223.

E. Candès & M. Wakin (2008) An introduction to compressive sampling, *IEEE, Signal Processing Magazine* **25** (2), 21–30.

R. Cont & Y. Kan (2011) Statistical modeling of credit default swap portfolios, Available at SSRN: http://ssrn.com/abstract=1771862.

J. Cox, J. Ingersoll & S. Ross (1985) A theory of the term structure of interest rates, *Econometrica* **53** (2), 385–407.

D. Donoho (2006) Compressed sensing, *IEEE Transactions on Information Theory* **52** (4), 1289–1306.

D. Duffie & K. Singleton (1999) Modeling term structures of defaultable bonds, *Review of Financial Studies* **12** (4), 687–720.

C. Gardiner (2009) *Stochastic Methods*. Berlin-Heidelberg: Springer.

M. Grant & S. Boyd (2011) CVX: Matlab software for disciplined convex programming, version 1.21, Available at http://cvxr.com/cvx.

T. Ho & S.-B. Lee (1986) Term structure movements and pricing interest rate contingent claims, *The Journal of Finance* **41** (5), 1011–1029.

J. Hull (2012) *Options*, *Futures*, *and Other Derivatives*. New York: Prentice Hall.

J. Hull & A. White (1990) Pricing interest-rate-derivative securities, *Review of Financial Studies* **3** (4), 573–592.

J. James & N. Webber (2000) *Interest Rate Modelling*. Chichester: Wiley-Blackwell Publishing Ltd.

D. Madan & W. Schoutens (2008) Break on through to the single side, *Journal of Credit Risk* **4** (3), 3–20.

R. Martin (2009) Credit spread shocks: How big and how often? *Risk* **22** (8), 84–89.

Y. Nesterov & A. Nemirovsky (1994) *Interior-Point Polynomial Methods in Convex Programming*. Philadelphia: SIAM.

B. O'Donoghue, E. Chu, N. Parikh & S. Boyd (2013) Operator splitting for conic optimization via homogeneous self-dual embedding, Available at http://www.stanford.edu/~boyd/papers/scs.html.

D. O'Kane (2008) *Modelling Single-name and Multi-name Credit Derivatives*. Chichester: J. Wiley & Sons.

R. Rendleman & B. Bartter (1980) The pricing of options on debt securities, *Journal of Financial and Quantitative Analysis* **15** (1), 11–24.

W. Schoutens (2003) *Lévy Processes in Finance*. Chichester: Wiley.

J. Sturm (1999) Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optimization Methods and Software* **11**, 625–653, Available at http://sedumi.ie.lehigh.edu/.

R. Tibshirani (1994) Regression shrinkage and selection via the lasso, *Journal of the Royal Statistical Society, Series B* **58** (1), 267–288.

K. Toh, M. Todd & R. Tütüncü (1999) SDPT3 — A Matlab software package for semidefinite programming, version 1.3, *Optimization Methods and Software* **11** (1), 545–581.

O. Vasicek (1977) An equilibrium characterization of the term structure, *Journal of Financial Economics* **5** (2), 177–188.

S. Wright (1997) *Primal-Dual Interior-Point Methods*. Philadelphia: SIAM.