Structural credit risk using time-changed Brownian motions: a tale of two models

T. R. Hurd* and Zhuowei Zhou
Dept. of Mathematics and Statistics
McMaster University
Hamilton ON L8S 4K1
Canada

September 13, 2011

Abstract

We consider a structural credit risk framework introduced in [5] where the log-leverage ratio of the firm is a Lévy process in the form of a time-changed Brownian motion (TCBM) where the time-change process has identically distributed independent increments. In models of this type, "vanilla" credit derivative pricing formulas are in closed form in terms of an explicit one-dimensional Fourier transform. Our primary aim is to investigate whether two very simple specifications of the time change process, namely the variance gamma (VG) model and the exponential jump model (EXP), can lead to good fits to CDS data for a representative firm with an interesting credit history, Ford Motor Co. Statistical inference in this class of intrinsically non-Gaussian hidden state models presents some new challenges which are a second main focus of the paper. We consider several variations of nonlinear filtering and statistical inference for TCBM models applied to a 4.5 year time series of credit default swap (CDS) prices on Ford, with the goal of finding a fast, accurate, robust scheme. The main conclusion is that the two TCBM models significantly outperform the classic Black-Cox model. Secondly, we show that a new inference method called the "linearized measurement scheme" is much faster than a standard numerical integration scheme (up to 100 times faster), and yields equivalent performance. The statistical methodology proposed in this paper can be effectively implemented for many other variations of TCBMs and applied to a wide range of firms, and opens the door to far-ranging explorations of a new class of structural credit risk models.

Key words: Credit risk, structural model, first passage, Lévy process, jump process, fast Fourier transform, credit default spread, nonlinear filtering, maximum likelihood estimation.

---

*This research is supported by the Natural Sciences and Engineering Research Council of Canada and MITACS, Mathematics of Information Technology and Complex Systems Canada.

AMS Subject Classification: 91G40, 91G70, 91G20, 60G35, 60G51

1 Introduction

Next to the Merton credit model of 1974 [13], the Black-Cox (BC) model [2] is perhaps the best known structural credit model. It models the time of a firm's default as the first passage time for the firm's log-leverage process, treated as an arithmetic Brownian motion, to cross zero. The BC model is conceptually appealing, but its shortcomings, such as the rigidity of credit spread curves, the counterfactual behaviour of the short end of the credit spread curve and the difficulty of computing correlated multifirm defaults, have been amply discussed elsewhere, see e.g. [9]. Indeed remediation of these different flaws has been the impetus for many of the subsequent developments in credit risk.

One core mathematical difficulty that has hampered widespread implementation of Black-Cox style first passage models has been the computation of first passage distributions for a richer class of processes one might want to use in modeling the log-leverage process. This difficulty was circumvented in [5], enabling us to explore the consequences of using processes that lead to a variety of desirable features: more realistic credit spreads, the possibility of strong contagion effects, and "volatility clustering" effects. [5] proposed a structural credit modeling framework where the log-leverage ratio $X_t := \log(V_t/K(t))$, where V_t denotes the firm asset value process and $K(t)$ is a deterministic default threshold, is a time-changed Brownian motion (TCBM). The time of default is the first passage time of the log-leverage ratio across zero. In that paper, the time change was quite general: our goal in the present paper is to make a thorough investigation of two simple specifications in which the time change is of Lévy type that lead to models that incorporate specific desirable characteristics. We focus here on a single company, Ford Motor Co., and show that with careful parameter estimation, TCBM models can do a very good job of explaining the observed dynamics of credit spreads. TCBMs have been used in other credit risk models, for example [14], [4], [1] and [12].

One new model we study is an adaptation of the variance gamma (VG) model introduced by [11] in the study of equity derivatives, and remaining very popular since then. We will see that this infinite activity pure jump Lévy model of the log-leverage ratio adapts easily to the structural credit context, and that the extra degrees of freedom it allows over and above the rigid structure of geometric Brownian motion correspond to desirable features of observed credit spread curves. The other model, the exponential (EXP) model, is a variation of the Kou-Wang double exponential jump model [8]. Like the VG model it is a Lévy model, but now with a finite activity exponential Lévy distribution. We find that the EXP model perfoRMSE remarkably similarly to the VG model when fit to our dataset.

We apply these two prototypical structural credit models to a dataset, divided into 3 successive 18 month periods, that consists of weekly quotes of credit default swap spreads (CDS) on Ford Motor Company. On each date, seven maturities are quoted: 1, 2, 3, 4, 5, 7, and 10 years. The main advantages of CDS data over more traditional debt instruments such as coupon bonds are their greater price transparency, greater liquidity, their standardized structure, and the fact that they are usually quoted for more maturities.

Our paper presents and compares three alternative statistical inference schemes applied to this time series of credit data. One scheme, the benchmark method, is rather slow (with typical run times of one or two hours) while the other two schemes take full advantage of the fast Fourier transform to speed up the large number of pricing formula evaluations. Note that in our inference method, the model parameters are taken as constants to be estimated for each 18 month time period: in contrast to "daily calibration" methods, only the natural dynamical variables, not the parameters, are allowed to be time varying.

Section 2 of this paper summarizes the financial case history of Ford Motor Co. over the global credit crisis period, and describes the CDS and Treasury data we used. Section 3 reviews the TCBM credit modeling framework introduced in [5]. There we include the main formulas for default probability distributions, defaultable bond prices and CDS spreads. Each such formula is in the form of an explicit Fourier transform representation: This is important for achieving a fast statistical algorithm. Section 4 gives the detailed specification of the two TCBM models under study. Section 5 outlines three possible statistical inference schemes appropriate for this type of model. Section 6 reports the results achieved when these schemes were implemented for the TCBM and Black-Cox credit models on the Ford CDS dataset. Some unresolved issues concerning these methods are discussed in the concluding Section 7.

2 Ford: The Test Dataset

We chose to study the credit history of Ford Motor Co. over the 4.5 year period from January 2006 to June 2010 and to determine the performance of several new structural credit models on this dataset. The case history of Ford over this period spanning the global credit crisis represents the story of a major firm and its near default, and is thus full of financial interest. Although not reported in this paper, we note that we have also applied these models to the credit data for several other types of firm over this period, and achieved similarly robust parameter estimation results. Thus our study of Ford truly exemplifies the capabilities of our modeling and estimation framework.

We divided the period of interest into three nonoverlapping successive 78 week intervals, one immediately prior to the 2007-2008 credit crisis, another starting at the outset of the crisis, the third connecting the crisis and the early recovery period. We used Ford CDS and US Treasury yield data, taking only Wednesday quotes in order to remove weekday effects.

1. Dataset 1 consisted of Wednesday midquote CDS swap spreads $\widehat{\mathrm{CDS}}_m^T$ and their bid-ask spreads w_m^T on weeks $m = 0, 1, \ldots, M-1$ for maturities $T \in \mathcal{T} := \{1, 2, 3, 4, 5, 7, 10\}$ years for Ford Motor Co., for the $M = 78$ consecutive Wednesdays from January 4th, 2006 to June 27, 2007, made available from Bloomberg.

2. Dataset 2 consisted of Wednesday midquote CDS swap spreads $\widehat{\mathrm{CDS}}_m^T$ and their bid-ask spreads w_m^T on weeks $m = M, \ldots, 2M-1$ for maturities $T \in \mathcal{T} := \{1, 2, 3, 4, 5, 7, 10\}$ years for Ford Motor Co., for the $M = 78$ consecutive Wednesdays from July 11, 2007 to December 31, 2008, made available from Bloomberg.

3. Dataset 3 consisted of Wednesday midquote CDS swap spreads $\widehat{\mathrm{CDS}}_m^T$ and their bid-ask spreads w_m^T on dates $m = 2M, \ldots, 3M-1$ for maturities $T \in \mathcal{T} := \{1, 2, 3, 4, 5, 7, 10\}$ years for Ford Motor Co., for the $M = 78$ consecutive Wednesdays from January 7th, 2009 to June 30, 2010, made available from Bloomberg.

4. The US treasury dataset[1] consisted of Wednesday yield curves (the "zero curve") on dates $m = 0, 1, \ldots, 3M-1$, for maturities

$$T \in \tilde{\mathcal{T}} := \{1m, 3m, 6m, 1y, 2y, 3y, 5y, 7y, 10y, 20y, 30y\}$$

for the period January 4th, 2006 to June 30, 2010.

We note that Ford Motor Company experienced a large number of credit rating changes during this four-and-a-half year period. The recent history of Standard & Poors (S & P) ratings is as follows: BB+ to BB- on January 5, 2006; BB- to B+ on June 28, 2006; B+ to B on September 19, 2006; B to B- on July 31, 2008; B- to CCC+ on November 20, 2008. The downgrades continued into 2009, with a move from CCC+ to CC on March 4, 2009 and to SD ("structural default") on April 6, 2009. The latest news was good: on April 13, 2009, S & P raised Ford's rating back to CCC, on November 3, 2009 to B-, and on August 2, 2010 to B+, the highest since the onset of the credit crisis.

In hindsight we see that Ford never actually defaulted, although it came close. At the climax of its financial distress, Ford's short term CDS spreads exceeded 100%. In the following estimation methodology, we consider the non-observation of default as an additional piece of information about the firm.

3 The TCBM Credit Setup

The time-changed Brownian motion credit framework of [5] starts with a filtered probability space $(\Omega, \mathcal{F}, \mathcal{F}_t, \mathbb{P})$, which is assumed to support a Brownian motion W and an independent increasing process G where the natural filtration $\mathcal{F}_t$ contains $\sigma\{G_u, W_v : u \leq t, v \leq G_t\}$ and satisfies the "usual conditions". $\mathbb{P}$ is taken to be the physical probability measure.

Assumptions 1.
1. The log-leverage ratio of the firm is a TCBM of the form $X_t := \log(V_t/K(t)) := x + \sigma W_{G_t} + \beta_P \sigma^2 G_t$ with parameters $x > 0, \sigma > 0$ and β_P. The time change G_t is characterized by its Laplace exponent $\psi(u,t) := -\log \mathbb{E}[e^{-uG_t}]$ which is assumed to be known *explicitly* and has average speed normalized to 1 by the condition
$$\lim_{t\to\infty} t^{-1}\partial\psi(0,t)/\partial u = 1.$$
2. The time of default of the firm is the *first passage time of the second kind* for the log-leverage ratio to hit zero (see the definition that follows). The recovery at default is modelled by the "recovery of treasury" mechanism[2] with constant recovery fraction $R \in [0, 1)$.

[1]Obtained from US Federal Reserve Bank, `http://www.federalreserve.gov/datadownload/`
[2]See [9].

3. The family of default-free zero-coupon bond price processes $\{B_t(T), 0 \le t \le T < \infty\}$ is free of arbitrage and independent of the processes W and G.

4. There is a probability measure $\mathbb{Q}$, equivalent to $\mathbb{P}$ and called the risk-neutral measure, under which all discounted asset price processes are assumed to be martingales. Under $\mathbb{Q}$, the distribution of the time change G is unchanged while the Brownian motion W has constant drift.[3] We may write $X_t = x + \sigma W^Q_{G_t} + \beta_Q \sigma^2 G_t$ for some constant β_Q where $W^Q_u = W_u + \sigma(\beta_P - \beta_Q)u$ is driftless Brownian motion under $\mathbb{Q}$.

We recall the definitions from [5] of first passage times for a TCBM X_t starting at a point $X_0 = x \ge 0$ to hit zero.

Definition 2.

- The standard definition of first passage time is the $\mathcal{F}$ stopping time
$$t^{(1)} = \inf\{t | X_t \le 0\} \ . \tag{1}$$
The corresponding stopped TCBM is $X^{(1)}_t = X_{t \wedge t^{(1)}}$. Note that in general $X^{(1)}_{t^{(1)}} \le 0$.

- The first passage time of the second kind is the $\mathcal{F}$ stopping time
$$t^{(2)} = \inf\{t | G_t \ge t^*\} \tag{2}$$
where $t^* = \inf\{t | x + \sigma W_t + \beta_P \sigma^2 t \le 0\}$. The corresponding stopped TCBM is
$$X^{(2)}_t = x + \sigma W_{G_t \wedge t^*} + \beta_P \sigma^2 (G_t \wedge t^*) \tag{3}$$
and we note that $X^{(2)}_{t^{(2)}} = 0$.

The general relation between $t^{(1)}$ and $t^{(2)}$ is studied in detail in [6] where it is shown how the probability distribution of $t^{(2)}$ can approximate that of $t^{(1)}$. For the remainder of this paper, however, we consider $t^{(2)}$ to be the definition of the time of default.

The following proposition, proved in [5], is the basis for computing credit derivatives in the TCBM modeling framework.

Proposition 3. *Suppose the firm's log-leverage ratio X_t is a TCBM with $\sigma > 0$ and that Assumptions 1 hold.*

1. *For any $t > 0, x \ge 0$ the survival probability $P^{(2)}(t, x) := \mathbb{E}_x[\mathbf{1}_{\{t^{(2)} > t\}}]$ is given by*
$$\frac{e^{-\beta x}}{\pi} \int_{-\infty}^{\infty} \frac{u \sin(ux)}{u^2 + \beta^2} e^{-\psi(\sigma^2(u^2+\beta^2)/2, t)} du + (1 - e^{-2\beta x}) \mathbf{1}_{\{\beta > 0\}}, \tag{4}$$
The density for X_t conditioned on no default is
$$\begin{aligned} \rho(y; t, x) &:= \frac{d}{dy} \mathbb{E}_x[\mathbf{1}_{\{X_t \le y\}} | t^{(2)} > t] \qquad (5) \\ &= P^{(2)}(t, x)^{-1} \mathbf{1}_{\{y > 0\}} \frac{e^{\beta(y-x)}}{2\pi} \int_{\mathbb{R}} \left[e^{iu(y-x)} - e^{-iu(y+x)} \right] e^{-\psi(\sigma^2(u^2+\beta^2)/2, t)} du \end{aligned}$$

[3]This assumption can be justified by a particular version of the Girsanov theorem. A more general assumption would be to allow the distribution of G to be different under $\mathbb{Q}$, but to avoid this added complexity we do not consider this possibility further here.

The characteristic function for X_t conditioned on no default is

$$\begin{aligned}\mathbb{E}_x[e^{ikX_t}|t^{(2)} > t] &= P^{(2)}(t,x)^{-1}\mathbb{E}_x[e^{ikX_t} \cdot \mathbf{1}_{\{t^{(2)}>t\}}] \qquad (6)\\ &= P^{(2)}(t,x)^{-1}\frac{e^{-\beta x}}{\pi}\int_{\mathbb{R}} \frac{u\sin(ux)}{(\beta+ik)^2+u^2}e^{-\psi(\sigma^2(u^2+\beta^2)/2,t)}du\\ &\quad + \left(e^{ikx} - e^{-ikx-2\beta x}\right) e^{-\psi(\sigma^2(k^2-2i\beta k)/2,t)} \left(\frac{1}{2}\mathbf{1}_{\{\beta=0\}} + \mathbf{1}_{\{\beta>0\}}\right)\end{aligned}$$

2. *The time 0 price $\bar{B}^{RT}(T)$ of a defaultable zero coupon bond with maturity T and recovery of treasury with a fixed fraction R is*

$$\bar{B}^{RT}(T) = B(T)[P^{(2)}(T,x)) + R(1 - P^{(2)}(T,x))] \qquad (7)$$

3. *The fair swap rate for a CDS contract with maturity $T = N\Delta t$, with premiums paid in arrears on dates $t_k = k\Delta t, k = 1,\ldots,N$, and the default payment of $(1-R)$ paid at the end of the period when default occurs, is given by*

$$CDS(x,T) = \frac{(1-R)\left[\sum_{k=1}^{N-1}[1-P^{(2)}(t_k,x)][B(t_k)-B(t_{k+1})] + B(T)[1-P^{(2)}(T,x)]\right]}{\Delta t\sum_{k=1}^{N} P^{(2)}(t_k,x)B(t_k)} \qquad (8)$$

Remarks 4.

- We shall be using the above formulas in both probability measures $\mathbb{P}$ and $\mathbb{Q}$, as appropriate.
- We observe in (4) that the survival and default probabilities are invariant under the following joint rescaling of parameters

$$(x,\sigma,\beta) \to (\lambda x, \lambda\sigma, \lambda^{-1}\beta), \text{ for any } \lambda > 0. \qquad (9)$$

It follows that all pure credit derivative prices are invariant under this rescaling.

4 Two TCBM Credit Models

The two credit models we introduce here generalize the standard Black-Cox model that takes $X_t = X_0 + \sigma W_t + \beta_P\sigma^2 t$. They are chosen to illustrate the flexibility inherent in our modeling approach. Many other specifications of the time change are certainly possible and remain to be studied in more detail. The following models are specified under the measure $\mathbb{P}$: by Assumption 1 they have the same form under the risk-neutral measure $\mathbb{Q}$, but with β_P replaced by β_Q.

4.1 The Variance Gamma Model

The VG credit model with its parameters $\Theta = (X_0, \sigma, \beta_P, \beta_Q, b, c, R)$ arises by taking G to be a gamma process with drift defined by the characteristic triple $(b, 0, \nu)_0$ with $b \in (0,1)$ and Lévy measure $\nu(z) = ce^{-z/a}/z, a > 0$ on $(0,\infty)$. The Laplace exponent of G_t is

$$\psi^{VG}(u,t) := -\log E[e^{-uG_t}] = t[bu + c\log(1+au)]. \qquad (10)$$

and by choosing $a = \frac{1-b}{c}$ the average speed of the time change is $t^{-1}\partial\psi^{VG}(0,t)/\partial u = 1$. This model and the next both lead to a log-leverage process of Lévy type, that is, a process with identical independent increments that are infinitely divisible.

4.2 The Exponential Model

The EXP credit model with its parameters $\Theta = (X_0, \sigma, \beta_P, \beta_Q, b, c, R)$ arises taking by G to be a Lévy process with a characteristic triple $(b, 0, \nu)_0$ with $b \in (0,1)$ and Lévy measure $\nu(z) = ce^{-z/a}/a, a > 0$ on $(0, \infty)$. The Laplace exponent of G_t is

$$\psi^{Exp}(u,t) := -\log E[e^{-uG_t}] = t\left[bu + \frac{acu}{1+au}\right].$$

and by choosing $a = \frac{1-b}{c}$ the average speed of the time change is $t^{-1}\partial\psi^{VG}(0,t)/\partial u = 1$.

5 The Statistical Method

The primary aim of this exercise is to develop a robust, efficient and theoretically sound statistical method that will determine whether our two TCBM credit models fit market CDS data on a single firm, in this case Ford Motor Company, better than the original Black-Cox structural model. The three important challenging aspects of the problem are: (i) there is a latent (unobserved) process X_t; (ii) the pricing formula is nonlinear in the state variable; (iii) the transition densities are far from Gaussian. We summarize the modeling ingredients and dataset and make some important remarks:

- model parameters $\Theta \in D \subset \mathbb{R}^n$, where D is a choice of hyperrectangle in parameter space. We augment the vector $\Theta \to (\Theta, \eta)$ to include an additional measurement error parameter η (which could be a matrix or scalar);
- an unobserved Markov process $X_t \in \mathbb{R}$ with initial condition $X_0 = x$ taken to be one component of Θ;
- the model formula (8) $F^k(X, \Theta)$ giving theoretical CDS spreads for different tenors k;
- a dataset consisting of spreads $Y := \{Y_m\}$ observed on weeks $m = 0, \ldots, M-1$ where $Y_m = \{Y_m^T\}_{F\in\mathcal{T}}$ for a term structure $\mathcal{T}$ with 7 tenors, plus their associated quoted bid/ask spreads w_m^T. As well, we observed the US Treasury dataset, assumed to give exact information about the discount factors entering into (8);
- We use a shorthand notation for collections of variables $Y_{\le m} := \{Y_0, \ldots, Y_m\}$ and $X_{<m} := \{X_1, \ldots, X_{m-1}\}$ etc.

The primary output of the inference methodology will be maximum likelihood estimators for Θ and the corresponding observed Fisher information matrix:

$$\widehat{\Theta} = \mathrm{argmax}_{\Theta\in D} \log \mathcal{L}(\Theta|Y_{<M}), \quad \widehat{\mathcal{I}} := -\left[\partial^2_\Theta \log \mathcal{L}(\widehat{\Theta}|Y_{<M})\right]. \tag{11}$$

In addition we will compute filtered estimates $\widehat{X}_{<M} = (\widehat{X}_1, \ldots, \widehat{X}_{M-1})$ of the hidden Markov process X_t. The likelihood function is computed via intermediate conditioning on the values of the hidden Markov state process:

$$\mathcal{L}(\Theta|Y_{<M}) := \rho(Y_{<M}|\Theta) = \int_{\mathbb{R}^{M-1}} \mathcal{F}(Y_{<M}|X_{<M}, \Theta)\ \mathcal{P}(X_{<M}|\Theta)\ dX_1 \ldots dX_{M-1} \qquad (12)$$

Here the "full measurement density" of the observations Y conditioned on X, Θ has the form

$$\mathcal{F}(Y_{<M}|X_{<M}, \Theta) := \prod_{m=0}^{M-1} f_m(Y_m|X_m, \Theta) \qquad (13)$$

where the measurement density f_m for a given date will depend on our assumptions on the distributional nature of the measurement errors. The multiperiod transition density conditioned on nondefault is

$$\mathcal{P}(X_{<M}|\Theta) = \prod_{m=1}^{M-1} p(X_m|X_{m-1}, \Theta) \qquad (14)$$

where $p(X_m|X_{m-1}, \Theta)$ is the one week Markov transition density given by (5) with $t = 1/52$.

We now make some remarks on the nature of the inference problem (11).

Remark 5. This is some variant of nonlinear filtering. Existing nonlinear filtering methods include the extended Kalman filter, the particle filter, the unscented Kalman filter and the Markov Chain Monte Carlo method, but each of these seems likely to be either inaccurate or difficult to implement for our models. [4] The X integration in (12) can be done as iterated one-dimensional integrals as follows:

$$\mathcal{L}(\Theta|Y_{<M}) = \int_0^\infty f(Y_{M-1}|X_{M-1}, \Theta)\ \rho_{M-1}(X_{M-1}, Y_{<M-1}|\Theta)\ dX_{M-1} \qquad (15)$$

where for $m = 0, \ldots, M-2$

$$\rho_{m+1}(X_{m+1}, Y_{\le m}|\Theta) = \begin{cases} \int_0^\infty p(X_{m+1}|X_m, \Theta)\ f_m(Y_m|X_m, \Theta)\ \rho_m(X_m, Y_{<m}, \Theta)\ dX_m, & m > 0 \\ p(X_1|X_0, \Theta) f_0(Y_0|X_0, \Theta) & m = 0 \end{cases} \qquad (16)$$

Being analytically intractable, each integral must be dealt with numerically, and each integrand requires evaluation of the model formula (8) for all values of X.

[4] We have also implemented a nonfiltering method similar to one given in [16] by adopting a hypothesis that the hidden state variable X can be observed exactly as implied by the 5 year CDS spread (which is in practice the most liquid tenor). Although such an approach is easy to implement and yields similar but less accurate results, we elected to treat the CDS data symmetrically with respect to tenor and are thus lead to filtering methods.

Remark 6. The time series of filtered estimates of the state variables $X_1, \ldots, X_{M-1}$ are the solutions $\hat{X}_1, \ldots, \hat{X}_{M-1}$ of

$$\hat{X}_m = \text{argmax}_{x \in \mathbb{R}_+} \log \left(f_m(Y_m | x, \widehat{\Theta}) \rho_m(x, Y_{<m} | \widehat{\Theta}) \right) \tag{17}$$

Following common practice, we also adopt the filtered estimate for X_0 instead of the maximum likelihood estimate coming from (11):

$$\hat{X}_0 = \text{argmax}_{x \in \mathbb{R}_+} \log \left(f_0(Y_0 | x, \widehat{\Theta}) \right) \tag{18}$$

Remark 7. An essential characteristic of TCBM models is that the transition density (5) can be far from Gaussian: Fig 1a plots the one-week log-transition density that arises from the VG model calibrated to Ford on a typical day, and shows the default barrier at $x = 0$.

Remark 8. The Fast Fourier transform (FFT) is a natural choice for evaluating the integral in the model CDS formula (8) for a linear grid of X values, that can then be used for numerical evaluation of the integrals in (15). But there will be challenges to control errors uniformly over the range of parameters Θ and X values: The basic FFT error analysis for this formula is given in the Appendix.

Remark 9. The measurement density $\mathcal{F}$ is part of the model specification, and is to some extent controlled by the modeler. Two possible specifications of f_m on a typical day are plotted as a function of X in Fig 1b.

Remark 10. A quasi-Newton method applied to the optimization problem (11) will involve a large number of evaluations of the likelihood function $\mathcal{L}(\Theta | Y_{<M})$ which may potentially have multiple local maxima.

Here we outline three alternative inference schemes that will be implemented and compared in Section 6.

5.1 Standard Measurement (SM) Scheme

The most obvious choice for the measurement densities f_m that relate the model to the observations arises by what we call the "standard measurement hypothesis" that CDS spreads are observed with a Gaussian measurement error:

$$f_m^{SM}(Y_m | X_m, \Theta) = \prod_{T \in \mathcal{T}} \left[\frac{1}{\sqrt{2\pi} \eta w_m^T} \exp \left(- \frac{(Y_m^T - F^T(X_m, \Theta))^2}{2\eta^2 (w_m^T)^2} \right) \right] \tag{19}$$

Here we have adopted a parsimonious specification that the residuals $Y_m^T - F^T(X_m, \Theta)$ are independent across date and tenor, and have standard deviations that are a scalar parameter η times the observed bid/ask spreads w_m^T.

The SM scheme requires numerical evaluation of the integrals (15) for $m = 1, \ldots M-2$. After experimentation with different grid sizes for Simpson's Rule and different FFT parameters, it was found that an inference scheme with adequate stability and robustness requires about $n = 1000$ Simpson grid points and an FFT grid with about 2^{10} points. So specified, we used this method as a benchmark against which to compare to two faster algorithms we describe next.

5.2 Linearized Measurement (LM) Scheme

The computational complexity of the SM scheme led us to consider the following nonlinear mapping from $Y_{<M}, \Theta$ to random variables $Z_{<M} = (Z_0, \ldots, Z_{M-1})$:

$$Z_m^T = G^T(Y_m^T, \Theta) \tag{20}$$

where G^T is the inverse function of F^T (we note that F^T, being monotonic in X, has an inverse). The measurement densities f_m^{SM} are then transformed into

$$\tilde{f}_m^{SM}(Z_m|\Theta) := f_m^{SM}\left(F(Z_m, \Theta)|\Theta\right) J(Z_m, \Theta) \tag{21}$$

where J is the Jacobian determinant:

$$J(Z_m, \Theta) = \prod_{T \in \mathcal{T}} \left[\frac{\partial G^T}{\partial Y} \Big|_{Y = F^T(Z_m^T)} \right] \tag{22}$$

The transformed measurements Z_m^T for different T can be all interpreted as direct measurements of X_m, that is as "market implied log-leverage", and it is very natural to consider the residuals $Z_m^T - X_m$ rather than $F^T(Z_m^T) - F^T(X_m)$. The "linearized measurement hypothesis" that $Z_m^T - X_m$ are independent Gaussians is thus in some sense as natural as the standard measurement hypothesis, and leads to the following alternative measurement densities:

$$\begin{aligned}
\tilde{f}_m^{LM}(Z_m|X_m, \Theta) &= \prod_{T \in \mathcal{T}} \left[\frac{1}{\sqrt{2\pi}\eta \tilde{w}_m^T} \exp\left(-\frac{(Z_m^T - X_m)^2}{2\eta^2 (\tilde{w}_m^T)^2} \right) \right] \\
f_m^{LM}(Y_m|X_m, \Theta) &= \prod_{T \in \mathcal{T}} \left[\frac{1}{\sqrt{2\pi}\eta w_m^T} \exp\left(-\frac{(G^T(Y_m^T, \Theta) - X_m)^2}{2\eta^2 (\tilde{w}_m^T)^2} \right) \right]
\end{aligned} \tag{23}$$

Here the standard deviations of the new residuals are naturally $\eta \tilde{w}_m^T$ where $\tilde{w}_m^T = \frac{\partial G^T}{\partial Y_m^T} w_m^T$ and the inverse Jacobian is $J^{-1} = \prod_{T \in \mathcal{T}} (\tilde{w}_m^T / w_m^T)$.

Note that the linearized measurement scheme is applicable whenever the measurement equation is monotonic in the latent state variable: thus it can be applied in some other important problems in finance, most notably, filtering in stochastic volatility models. The main advantage of this type of scheme is that f_m^{LM} is a Gaussian in X with computable moments, whereas f_m^{SM} is a complicated function. In the next section we capitalize on the explicit Gaussian form of the measurement density f_m^{LM} to implement an effective approximation scheme for the integrals (16). The idea is to inductively approximate the intermediate density ρ_m by a truncated Gaussian density that vanishes for $x < 0$, and to combine this with exact moment formulas for p to compute the first and second moments of the density ρ_{m+1}. Then moment matching can be used to determine the approximation for the next density ρ_{m+1}.

5.3 Linear Kalman Filter (KF) Scheme

Having accepted the linearized measurement hypothesis, it becomes possible to reduce the inference problem to a variation of the linear Kalman filter, by additionally approximating the transition densities $p(X_m|X_{m-1}, \Theta)$ as Gaussian distributions. The resulting algorithm can be expressed as a typical Kalman filter that recursively computes the filtered estimates of X_m and its variance [7]. However, Fig 1a shows that $p(X_m|X_{m-1}, \Theta)$ is typically far from Gaussian in our models, so we need to investigate whether or not this inconsistency turns out to make a big impact on the overall accuracy of the method.

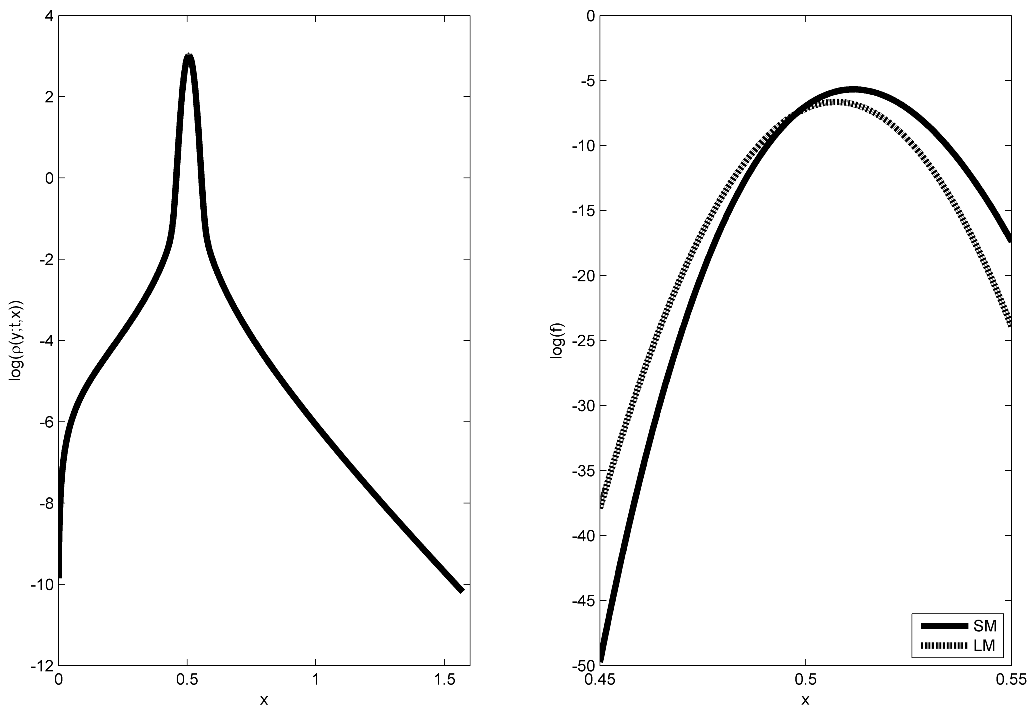

Figure 1: The left hand plot shows the true one-week VG log-transition density $\log p$ computed for 01/01/2006, a typical day in the time series. The right hand plot shows the standard measurement log-density $\log f_m^{SM}$ and linearized measurement log-density $\log f_m^{LM}$ on the same date, as functions of the Markov state X. Model parameters are taken from Table 1.

6 Numerical Implementations

We now summarize the results of the SM, LM and KF inference schemes, each applied to the three time periods of the dataset. Before starting, we reduce the complexity of our models with negligible loss in accuracy by removing two "nuisance parameters". First, in view of the rescaling invariance (9), and the interpretation of σ as the volatility of X, without loss of generality we set $\sigma = 0.3$ in all models. Second, since estimating β_P is equivalent to estimating the mean of X_t in the physical measure, and well known to be a poorly determined statistic, we arbitrarily set $\beta_P = -0.5$. So specified, the two TCBM models have five free parameters $\Theta = (b, c, \beta_Q, R, \eta)$ as well as two frozen parameters $\sigma = 0.3, \beta_P = -0.5$. The Black-Cox model with its free parameters $\Theta = (\beta_Q, R, \eta)$ and frozen parameters $\sigma = 0.3, \beta_P = -0.5$ then nests as the $c = 0$ limit inside both the VG and EXP models, for all b. We determined the maximal likelihood estimators by the quasi-Newton optimization algorithm *fmincon* implemented in MATLAB, constrained to the following domain D:

$$b \in [0.2, 1];\ c \in [0.1, 10];\ \beta_Q \in [-3, -0.1];\ R \in [0, 1];\ \eta \in [0.5, 10]\ .$$

Note the error analysis outlined in the Appendix shows that we need to bound b away from 0 if we are to use the FFT to compute (8) sufficiently accurately. Early experimentation showed that the likelihood function is slowly varying in b, and that the b estimator usually drifted slowly to the boundary value $b = 0.2$ with little increase in likelihood, demonstrating that this parameter is also in some sense a "nuisance" parameter. On the given domain D we were able to choose the FFT truncation parameter $\bar{u} = 300$ and depending on Θ, we allowed the size of the FFT lattice to vary from 2^8 to 2^{10}. With these choices, we were able to keep the combined truncation-discretization error within 10^{-10}.

Table 1 summarizes the estimation results for each of the three models, for the three datasets in 2006-2010, using the LM and SM inference methods. Estimated parameter values are given with standard errors, as well as summary statistics for the resulting filtered time series of X_t. We also present the root mean square error (RMSE) defined as the average error of the CDS spreads quoted in units of the bid/ask spread:

$$\text{RMSE} = \sqrt{\frac{1}{7M}\sum_{m=0}^{M-1}\sum_{t\in\mathcal{T}}\frac{\left(F^T(\hat{X}_m, \Theta) - Y_m^T\right)^2}{(w_m^T)^2}}$$

We do not exhibit the results from the KF scheme because in all cases they were very similar to LM scheme results.

For each of the three schemes, the *fmincon* algorithm involved approximately 200 evaluations of the function $\mathcal{L}$ before settling near a well defined maximum. The total computation times for the rather slow benchmark SM method were typically several hours for each dataset on a standard laptop. The total times for both the LM and KF implementations were similar and much faster, about 2 minutes. While a priori we worried that the output from the *fmincon* iteration scheme would be highly dependent on the initial value of Θ, that is that there might be multiple local maxima of the likelihood function, we found the results to be robust over a wide range of initial values. Overall, we see that the SM and LM inference methods give similar results, verifying that for our dataset it is reasonable to apply quite brutal approximations to the partial likelihood functions to facilitate an efficient estimation.

As a check on the consistency of the three inference methods, we tested them on a simulated dataset with VG parameters taken from Table 1 and found that parameter values and the filtered state variable X_t were in all cases reliably estimated on longer datasets (300 weeks gave very accurate estimates) while parameter values estimated over 78 weeks were close to the known values but showed appreciable scatter.

We find that the VG and EXP models give very good, perhaps surprisingly good, qualitative fit of the observed CDS spreads over the 78 week periods considered. Figure 2 shows the in-sample fit of the three models to the observed market data on three typical dates in the time series. One can note that the two TCBM models are better able than the BC model to fit the varying term structure shapes. The finite activity EXP model shares similarities with the infinite activity VG model, both in behavior and performance. For these two TCBM models, model parameters are quite similar between dataset 1 and dataset 3. It is consistent with Ford's history of credit ratings that dataset 3 has lower, more volatile log-leverage ratios and lower recovery rate than dataset 1. We can also see that during the peak of the credit crisis in dataset 2, the estimated parameters show noticeable signs of stress. The mean time change jump size is up by approximately 50%, driven mainly by the increased short term default probability. The recovery rate is also significantly lower for dataset 2. In the very stressed financial environment at that time, a firm's value would be greatly discounted and its capacity to liquidate assets would be limited. On the other hand the risk neutral drift β_Q is significantly higher, reflecting a certain positive expectation on the firm. At the peak of the credit crisis, Ford's annualized credit spreads exceeded 100%. The log-leverage ratios are much suppressed to a level of about 65% of that of dataset 1.

By definition, RMSE measures the deviation of the observed CDS spreads from the model CDS spreads while η measures the deviation of the "observed" log-leverage ratios $\tilde{X}_t$ from the "true" log-leverage ratios X_t. We can see that RMSE and η are very close in all cases, which implies that the objective functions based on the standard CDS measurement density (19) and the linearized measurement density (23) are fundamentally very similar.

In terms of RMSE and η, both TCBM models performed consistently better than the Black-Cox model. The TCBM fit is typically within two times the bid/ask spread across 3 datasets, while the errors of the Black-Cox model are about 30% higher on average. Figure 3 displays histograms of $(w_m^T)^{-1}\big(F^T(X_m,\Theta) - Y_m^T\big)$, the signed error in units of the bid/ask spread, for the three models, for the short and long end of the term structure. For both TCBM models we can see that most errors are bounded by ± 2 and are without obvious bias. By comparison, the errors of the Black-Cox model are highly biased downward in both the short and long terms. For 1-year spreads the majority of errors stay near -2 and for 10-year spreads there is a concentration of errors near -4. Surprisingly, all the three models perform better and more closely to one another during the crisis period of dataset 2. For the TCBM models, the great majority of errors are near 0 and without obvious bias. The Black-Cox model does not have obvious bias either, but there are more errors beyond the ± 2 range. The performance of all three models is better for intermediate tenors between 1 and 10 years, with the mid-range 5-year and 7-year tenors having the best fit. The histograms for these tenors (not shown) do still indicate that the TCBM models perform better than the Black-Cox model, in regard to both bias and absolute error.

The estimation results using the Kalman filter method indicate that for our dataset, the transition density can be safely approximated by a Gaussian density. The Kalman filter is convenient for calculating the weekly likelihood function, which is needed in the Vuong test [17], a test to compare the relative performance of nested models. If $\bar{X}_m$ and $\bar{P}_m$ denote the ex-ante forecast and variance of time m values of the measurement series obtained from Kalman filtering, and $\hat{X}_m$ denotes the filtered estimate, then the weekly log-likelihood function can be written as

$$l_m = -\frac{1}{2}\log|\bar{P}_m| - \frac{1}{2}(\hat{X}_m - \bar{X}_m)^\top(\bar{P}_m)^{-1}(\hat{X}_m - \bar{X}_m) - f_m^{LM}(Y_m|\hat{X}_m, \Theta). \qquad (24)$$

The log-likelihood ratio between two models i and j is

$$\lambda_{ij} = \sum_{m=1}^{M-1}(l_{im} - l_{jm})$$

and the Vuong test statistic is

$$\mathcal{T}_{ij} = \frac{\lambda_{ij}}{\hat{s}_{ij}\sqrt{M-1}},$$

where $\hat{s}_{ij}^2$ is the sample variance of $\{l_{im} - l_{jm}\}_{m=1,\dots,M-1}$. Vuong proved that $\mathcal{T}_{ij}$ is asymptotic to a standard normal under the null hypothesis that models i and j are equivalent in terms of likelihood function. Due to the serial correlation within the log-likelihood functions, Newey and West's estimator [15] is used for $\hat{s}$. The Vuong test results are shown in Table 2 and confirm that the Black-Cox model is consistently outperformed by the two TCBM models. Moreover, by this test, the EXP model shows an appreciable improvement over the VG model that could not be easily observed in the previous comparison. Similar conclusions about the relative performance of the three models are also supported by the Akaike Information Criterion (AIC), also given in Table 2, that compares the values of the likelihood function achieved.

It is interesting to compare the time series of Ford stock prices S_m to the filtered log-leverage ratios $\hat{X}_m$. Fig 4 shows there is a strong correlation between these two quantities, indicating that the equity market and credit market are strongly interconnected. The empirical observations supporting this connection as well as financial modeling that interprets this connection can be found in [12], [3] and their references.

7 Conclusions

In this paper, we demonstrated that the Black-Cox first passage model can be efficiently extended to a very broad class of firm value processes that includes exponential Lévy processes. We tested the fit of two realizations of Lévy subordinated Brownian motion models to observed CDS spreads for Ford Motor Co., a representative firm with an interesting credit history in recent years. We found that the two Lévy process models can be implemented easily, and give similarly good performance in spite of the very different characteristics of their Lévy measures. Both models outperform the Black-Cox model in fitting the time series of CDS term structures over all three 78 week periods, as measured by both the Vuong statistic and the Akaike Information Criterion. However, they still have limitations in fitting all tenors of the CDS term structure, suggesting that further study is needed into models with more flexible time changes. In this preliminary study we focused on just one firm, Ford, and an important next step should be an extended survey of a wide variety of firms from different countries and different economic sectors.

We tested three methods for filtered statistical inference, and found that both the linearized measurement (LM) scheme and a Kalman Filter (KF) scheme gave similar results whereas the parameters determined by the standard measurement scheme while different, were similar in accuracy. By their strategic use of the fast Fourier transform, both the LM and KF approximation methods can be made rather efficient: a complete parameter estimation and filtering for a time series of term structures for 78 weeks can be computed on a laptop in about two minutes using these schemes, compared to over two hours for our realization of the Standard Measurement (SM) scheme. One should note that the LM and KF schemes can be implemented for filtering problems in finance whenever the measurement equation is monotonic in the latent variable, for example, stochastic volatility models.

Finally, we observe a strong correlation between Ford's stock price and the filtered values of its unobserved log-leverage ratios. This final observation provides the motivation for our future research that will extend these TCBM credit models to TCBM models for the joint dynamics of credit and equity.

Appendix: Numerical Integration

Statistical inference in these models requires a large number of evaluations of the integral formula (4) that must be done carefully to avoid dangerous errors and excessive costs. To this end, we approximate the integral by a discrete Fourier transform over the lattice

$$\Gamma = \{u(k) = -\bar{u} + k\delta | k = 0, 1, \ldots, N-1\}$$

for appropriate choices of $N, \delta, \bar{u} := N\delta/2$. It is convenient to take N to be a power of 2 and lattice spacing δ such that truncation of the u-integrals to $[-\bar{u}, \bar{u}]$ and discretization leads to an acceptable error. If we choose initial values x_0 to lie on the reciprocal lattice with spacing $\delta^* = 2\pi/N\delta = \pi/\bar{u}$

$$\Gamma^* = \{x(\ell) = \ell\delta^* | \ell = 0, 1, \ldots, N-1\}$$

then the approximation is implementable as a fast Fourier transform (FFT):

$$P^{(2)}(t,x(\ell)) \sim \frac{-i\delta e^{-\beta x(\ell)}}{\pi}\sum_{k=0}^{N-1}\frac{u(k)e^{iu(k)x(\ell)}}{u(k)^2+\beta^2}\exp[-\psi(\sigma^2(u(k)^2+\beta^2)/2,t)] \quad (25)$$

$$= -i(-1)^n\delta e^{-\beta x(\ell)}\sum_{k=0}^{N-1}\frac{u(k)e^{2\pi ik\ell/N}}{u(k)^2+\beta^2}\exp[-\psi(\sigma^2(u(k)^2+\beta^2)/2,t)] \quad (26)$$

Note that we have used the fact that $e^{-iN\delta x(\ell)/2}=(-1)^n$ for all $\ell\in\mathbb{Z}$.

The selection of suitable values for N and δ in the above FFT approximation of (8) is determined via general error bounds proved in [10]. In rough terms, the pure truncation error, defined by taking $\delta\to 0, N\to\infty$ keeping $\bar{u}=N\delta/2$ fixed, can be made small if the integrand of (4) is small and decaying outside the square $[-\bar{u},\bar{u}]$. Similarly, the pure discretization error, defined by taking $\bar{u}\to\infty, N\to\infty$ while keeping δ fixed, can be made small if $e^{-|\beta|\bar{x}}P^{(2)}(\bar{x},t)$, or more simply $e^{-|\beta|\bar{x}}$, is small, where $\bar{x}:=\pi/\delta$. One expects that the combined truncation and discretization error will be small if $\bar{u}$ and $\delta=\pi/\bar{x}$ are each chosen as above. These error bounds for the FFT are more powerful than bounds one finds for generic integration by the trapezoid rule, and constitute one big advantage of the FFT. A second important advantage to the FFT is its $O(N\log N)$ computational efficiency that yields $P^{(2)}$ on a lattice of x values with spacing $\delta^*=2\pi/N\delta=\pi/\bar{u}$: this aspect will be very useful in estimation. These two advantages are offset by the problem that the FFT computes values for x only on a grid.

We now discuss choices for N and δ in our two TCBM models. For $\beta<0$, the survival function of the VG model is

$$P^{(2)}(0,t,x,\beta)=\frac{e^{-\beta x}}{\pi}\int_{-\infty}^{\infty}\exp[-tb\sigma^2(u^2+\beta^2)/2]\left(1+\frac{a\sigma^2(u^2+\beta^2)}{2}\right)^{-ct}\frac{u\sin ux}{u^2+\beta^2}du$$

while for the EXP model

$$P^{(2)}(0,t,x,\beta)=\frac{e^{-\beta x}}{\pi}\int_{-\infty}^{\infty}\exp\left[-t\left(b\sigma^2(u^2+\beta^2)/2+\frac{ac\sigma^2(u^2+\beta^2)}{2+a\sigma^2(u^2+\beta^2)}\right)\right]\frac{u\sin ux}{u^2+\beta^2}du$$

In both models, the truncation error has an upper bound ϵ when $\bar{u}>C|\Phi^{-1}(\epsilon C')|$, where Φ^{-1} is the inverse normal CDF and C,C' are constants depending on t. On the other hand, provided $\beta<0$, the discretization error will be small (of order ϵ or smaller) if $N>\frac{\bar{u}}{2\pi|\beta|}\log\left(\epsilon^{-1}(1+\exp(-2\beta x))\right)$. Errors for (6) can be controlled similarly.

References

[1] M. Baxter. Dynamic modelling of single-name credits and CDO tranches. working paper downloadable at `www.defaultrisk.com/pp_crdrv_07.htm`, 2006.

[2] F. Black and J. C. Cox. Valuing corporate securities. *J. Finance*, 31:351–367, 1976.

[3] P. Carr and L. Wu. Stock options and credit default swaps: A joint framework for valuation and estimation. *J. Fin. Economet.*, doi:10.1093/jjfinec/nbp010:1–41, 2009.

[4] Fang Fang, Henrik Jönsson, Cornelis W. Oosterlee, and Wim Schoutens. Fast Valuation and Calibration of Credit Default Swaps Under Levy Dynamics. *Journal of Computational Finance*, 14, 2010.

[5] T. R. Hurd. Credit risk modeling using time-changed Brownian motion. *Int. J. Theor. App. Fin.*, 12:1213–1230, 2009.

[6] T. R. Hurd and A. Kuznetsov. On the first passage time for Brownian motion subordinated by a Lévy process. *J. Appl. Probab.*, 46:181–198, 2009.

[7] Kazufumi Ito and Kaiqi Xiong. Gaussian filters for nonlinear filtering problems. *IEEE Transactions on Automatic Control*, 45(5):910–927, 2000.

[8] S. G. Kou and H. Wang. First passage times of a jump diffusion process. *Adv. in Appl. Probab.*, 35(2):504–531, 2003.

[9] D. Lando. *Credit Risk Modeling: Theory and Applications.* Princeton University Press, 2004.

[10] R. Lee. Option pricing by transform methods: Extensions, unification, and error control. *Journal of Computational Finance*, 7:51–86, 2004.

[11] D. Madan and E. Seneta. The VG model for share market returns. *Journal of Business*, 63:511–524, 1990.

[12] Rafael Mendoza-Arriaga, Peter Carr, and Vadim Linetsky. Time-changed markov processes in unified credit-equity modeling. *Mathematical Finance*, 20(4):527–569, 2010.

[13] R. C. Merton. On the pricing of corporate debt: the risk structure of interest rates. *J. Finance*, 29:449–470, 1974.

[14] T. Moosbrucker. Pricing CDOs with correlated Variance Gamma distributions. Working paper `http://www.defaultrisk.com/pp_crdrv103.htm`, 2006.

[15] Whitney K Newey and Kenneth D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–08, 1987.

[16] Jun Pan and Kenneth Singleton. Default and recovery implicit in the term structure of sovereign cds spreads. *Journal of Finance*, 63:2345–2384, 2008.

[17] Q. H. Vuong. Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica*, 57:307–333, 1989.

	04/01/06-27/06/07		11/07/07-31/12/08		07/01/09-30/06/10	
VG Model						
	LM scheme	SM scheme	LM scheme	SM scheme	LM scheme	SM scheme
$\hat{\sigma}$	0.3	0.3	0.3	0.3	0.3	0.3
$\hat{b}$	0.2	0.2	0.2	0.2	0.2	0.2
$\hat{c}$	1.039(0.060)	1.093(0.063)	0.451(0.034)	0.519(0.040)	1.08(0.11)	1.21(0.14)
$\hat{\beta}_Q$	-1.50(0.12)	-1.34(0.11)	-0.879(0.061)	-0.888(0.063)	-1.368(0.066)	-1.303(0.053)
$\hat{R}$	0.626(0.026)	0.590(0.029)	0.450(0.029)	0.457(0.029)	0.611(0.018)	0.576(0.016)
$\hat{\eta}$	1.53	1.54	0.897	0.902	1.80	1.93
$\hat{x}_{av}$	0.693	0.699	0.457	0.462	0.480	0.503
$\hat{x}_{std}$	0.200	0.181	0.239	0.234	0.267	0.264
RMSE	1.43	1.43	0.84	0.84	1.79	1.80
EXP Model						
	LM scheme	SM scheme	LM scheme	SM scheme	LM scheme	SM scheme
$\hat{\sigma}$	0.3	0.3	0.3	0.3	0.3	0.3
$\hat{b}$	0.2	0.229(0.038)	0.2	0.240(0.038)	0.2	0.2
$\hat{c}$	2.23(0.12)	2.153(0.11)	1.17(0.07)	1.150(0.06)	2.33(0.20)	2.526(0.26)
$\hat{\beta}_Q$	-1.44(0.12)	-1.29(0.10)	-0.780(0.060)	-0.841(0.073)	-1.286(0.027)	-1.207(0.059)
$\hat{R}$	0.609(0.028)	0.573(0.030)	0.395(0.033)	0.433(0.038)	0.588(0.022)	0.548(0.019)
$\hat{\eta}$	1.503	1.492	0.882	0.835	1.755	1.902
$\hat{x}_{av}$	0.702	0.706	0.479	0.475	0.486	0.506
$\hat{x}_{std}$	0.199	0.180	0.242	0.234	0.266	0.260
RMSE	1.41	1.41	0.821	0.833	1.76	1.78
BC Model						
	LM scheme	SM scheme	LM scheme	SM scheme	LM scheme	SM scheme
$\hat{\sigma}$	0.3	0.3	0.3	0.3	0.3	0.3
$\hat{\beta}_Q$	-2.02(0.10)	-2.02(0.10)	-1.793(0.067)	-1.789(0.064)	-1.78(0.12)	-1.822(0.082)
$\hat{R}$	0.773(0.011)	0.773(0.011)	0.757(0.009)	0.759(0.009)	0.760(0.013)	0.713(0.011)
$\hat{\eta}$	2.38	2.39	1.29	1.24	2.18	2.35
$\hat{x}_{av}$	0.624	0.642	0.406	0.421	0.422	0.493
$\hat{x}_{std}$	0.187	0.189	0.214	0.221	0.237	0.256
RMSE	2.19	2.13	1.19	1.17	2.14	2.20

Table 1: Parameter estimates and filtered X statistics for the VG, EXP and Black-Cox models, using two different filtering methods. The numbers in brackets are standard errors. The estimation uses weekly (Wednesday) CDS data from January 4th 2006 to June 30 2010. $\hat{x}_{std}$ is the square root of the annualized quadratic variation of $\hat{X}_t$. Columns 2, 4, 6 refer to the Linearized Measurement scheme evaluated using the truncated Gaussian approximation; columns 3, 5, 7 refer to the Standard Measurement scheme. Results for the Kalman Filter scheme are very similar to those from the LM scheme and are not recorded here.

		VG	EXP	B-C
VG	Vuong	0	-2.21/-1.41/-2.33	5.42/5.10/2.03
	AIC, LM	0	16/18/12	-424/-348/-184
	AIC, SM	0	12/6/8	-394/-286/-188
EXP	Vuong	2.21/1.41/2.33	0	5.46/5.22/2.19
	AIC, LM	-16/-18/-12	0	-440/-366/-196
	AIC, SM	-12/-6/-8	0	-406/-292/-196
B-C	Vuong	-5.42/-5.10/-2.03	-5.46/-5.22/-2.19	0
	AIC, LM	424/348/184	440/366/196	0
	AIC, SM	394/286/188	406/292/196	0

Table 2: Results of the Vuong and AIC tests for the three models, for dataset 1, dataset 2 and dataset 3. Positive values for the Vuong statistic and negative values for the AIC statistic indicate that the row model is more accurate than the column model.

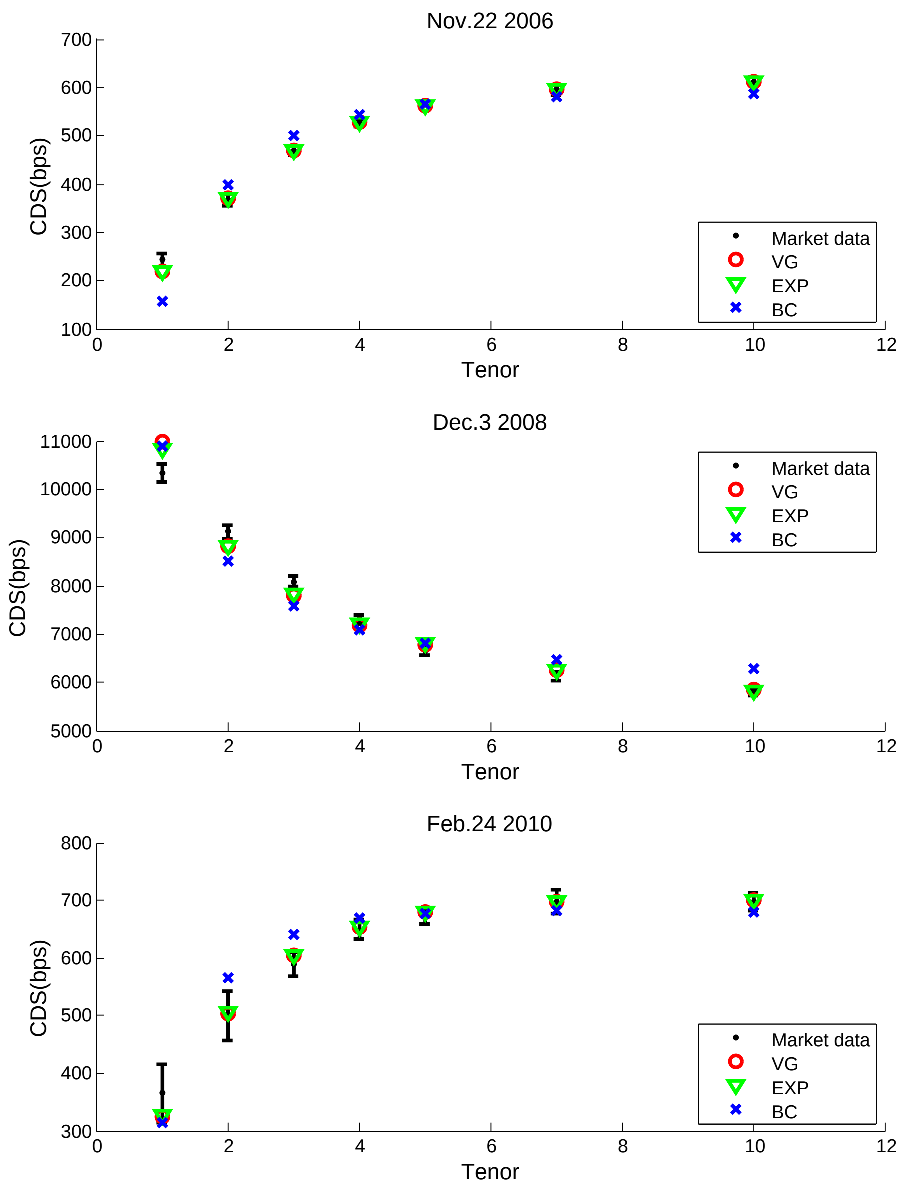

Figure 2: The in-sample fit of the two TCBM models and Black-Cox model to the observed Ford CDS term structure for November 22, 2006 (top), December 3, 2008 (middle) and February 24, 2010 (bottom). The error bars are centered at the mid-quote and indicate the size of the bid-ask spread.

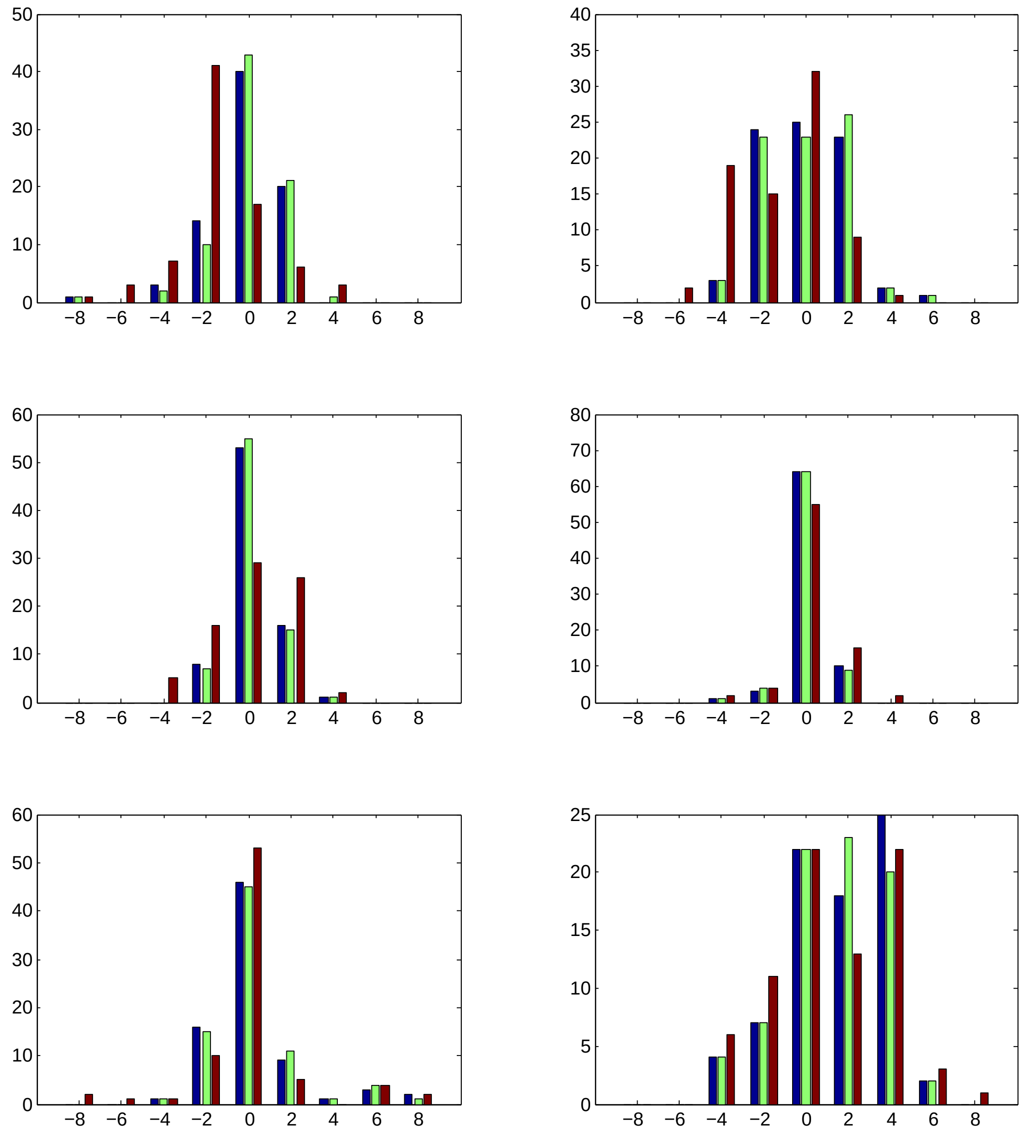

Figure 3: Histograms of the relative errors, in units of bid-ask spread, of the in-sample fit for the VG model (blue bars), EXP model (green bars) and Black-Cox model (red bars) for dataset 1 (top), dataset 2 (middle) and dataset 3 (bottom). The tenor on the left is 1-year and on the right, 10-year.

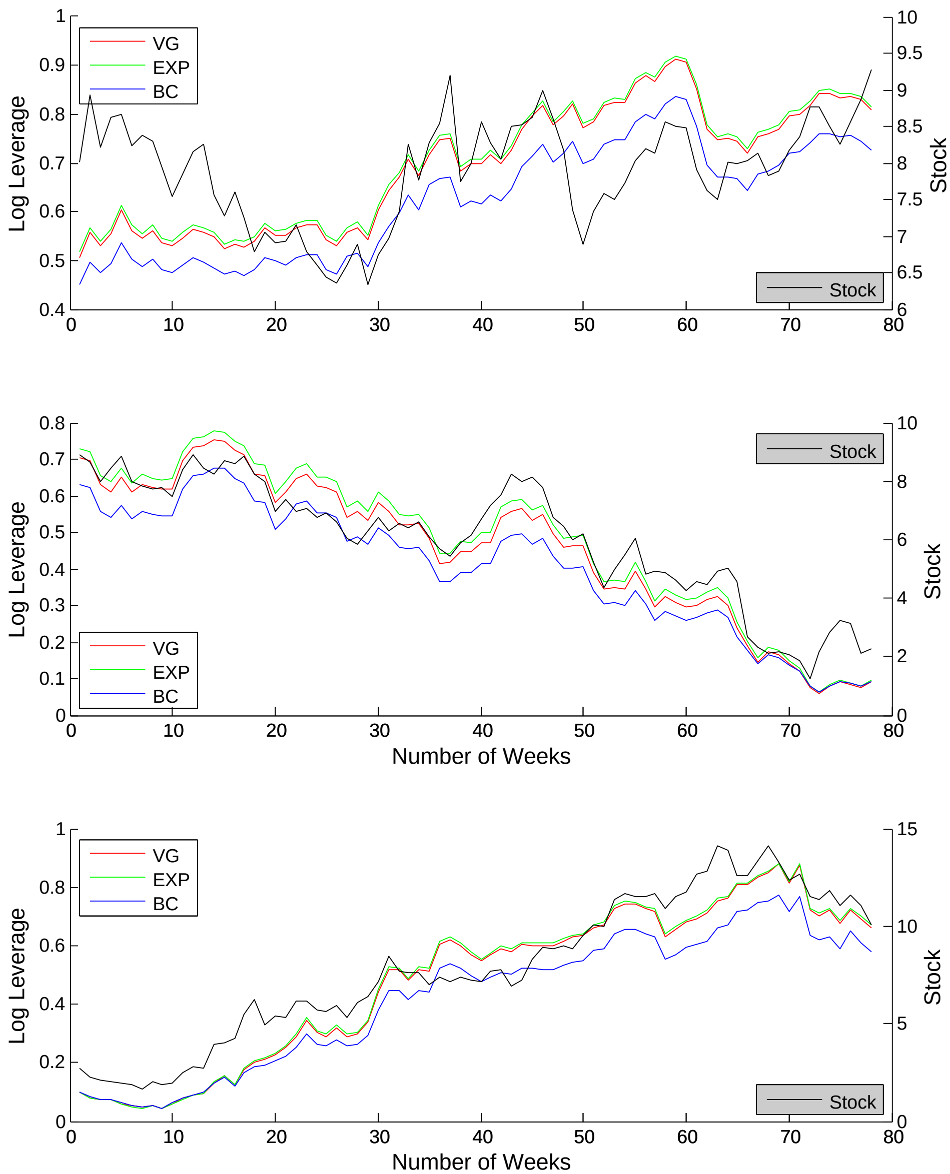

Figure 4: Filtered values of the unobserved log-leverage ratios X_t versus stock price for Ford for dataset 1(top), 2 (middle) and 3 (bottom).