Joint Default and Recovery Risk Estimation: An Application to CDS Data

Jens Henrik Eggert Christensen†

Federal Reserve Bank of San Francisco
101 Market Street
San Francisco, CA 94105
(415) 974-3115
jens.christensen@sf.frb.org

Abstract

Motivated by the results in Christensen (2006b) this paper tests the separation of default and recovery risk in a reduced-form affine credit risk model using CDS data for Ford Motor Co. The results indicate that default and recovery risk can in fact be jointly identified from cross-sections of CDS rates in this case. Furthermore, it is investigated whether changes in the risk premium specification has any impact on the separation. To that end, the essentially affine and the extended affine risk premium specifications are analyzed. The results show that the Feller conditions required in the extended affine risk premium specification are binding restrictions which significantly change the outcome of the CDS rate decomposition. For that reason the essentially affine specification is preferred for credit risk modeling purposes.

†I would like to thank David Lando for many helpful discussions and suggestions. I would also like to thank Peter Feldhütter for useful conversations about the Kalman filter. Finally, a special thank goes to Autria Christensen for her patience in reading earlier drafts of this paper.

Date: May 2005. This version: January 22, 2007.

1 Introduction

Motivated by the results in Christensen (2006b) the primary purpose of this paper is to test the separation of default and recovery risk in a reduced-form credit risk setting using actual CDS data for Ford Motor Co. A secondary purpose is to investigate whether the choice of risk premium specification has any impact on the outcome of such separations.

A number of studies have treated the recovery rate as a fixed constant and used corporate bond yields or CDS rates to derive the dynamics of the default intensity process. This is equivalent to modeling the dynamics of the credit spread itself. Excellent examples are Duffee (1999), Driessen (2005), and Pan and Singleton (2005). However, there is a growing literature on realized recovery rates documenting significant and even business-cycle related variation in observed recovery rates on defaulted debt claims, for examples see Covitz and Han (2004), Altman et. al. (2005), and Acharya et al. (2006). Faced with these empirical results it would be natural for investors to build similar correlation structures for the expected future recovery rate into current CDS rates. The question is whether the dynamics of such market-implied recovery rates can be estimated along with the dynamics of the market-implied default intensity process based on cross-sections of corporate bond yields or CDS rates.

Bakshi, Madan, and Zhang (2004) are the first to attempt a joint estimation of default and recovery risk in a reduced-form credit risk setting. Unfortunately, their model only has one factor, the interest rate, to explain both interest rate, default, and recovery risk. Therefore, default and recovery risk are only represented by constants in their model. More recently, Das and Hanouna (2006) apply a reduced-form calibration method to obtain a joint derivation of the market-implied term structure of forward hazard rates and forward recovery rates from cross-sections of CDS rates. However, their approach does not contain any dynamic description of the default and recovery risk components. Christensen (2006b) is the first to look into the feasibility of performing a full dynamic joint estimation of stochastic default and recovery risk. In his simulation study where interest rate, default, and recovery risk are all taken into consideration, evidence is found that a joint estimation is possible, if only measurement noise is not too high and observed spread levels not too low. However, the result in Christensen (2006b) is only tentative in the sense that it is based solely on simulated data and the parameters, though picked with great care from a number of studies, may not combined constitute a realistic choice. Thus, the general validity of the result cannot be asserted. The purpose of this paper is to provide evidence that a dynamic joint estimation of default and recovery risk is possible based on actual panel data of CDS rates.

The first contribution of this paper is to set up a reduced-form model that allows for stochastic variation in the risk-free rate, the default intensity, and the recovery rate processes. To that end, a four-factor affine model is introduced. The two interest rate factors are assumed to follow a so-called $A_1(2)$-model similar to Driessen (2005).[1] The default intensity risk factor is assumed to be a CIR-process, while the recovery risk factor is assumed to be Gaussian. The interest rate, the default intensity, and the recovery rate processes are then assumed to be affine functions of these four risk factors.

The second and most important contribution of this paper is to estimate this model which features stochastic default and recovery risk. In addition, the impact of the choice of risk premium

[1] This is the notation introduced by Dai and Singleton (2000) in their canonical categorization and characterization of affine term structure models.

specification on the separation of default and recovery risk is analyzed. In the affine term structure literature two risk premium specifications are dominant, the '*essentially affine specification*' introduced by Duffee (2002) and the '*extended affine specification*' introduced and proved to be arbitrage-free by Cheridito, Filipović, and Kimmel (2005). In the essentially affine specification the square root-type processes, of which there are two in the models presented in this paper, cannot have freely varying mean and mean-reversion rates when going from the real world probability measure P to the pricing measure Q. This is because the risk premium for such processes is restricted to being proportional to the volatility of the process itself. However, in this specification the lower boundary for the square root-type processes is only required to be reflecting. In the extended affine specification, the mean and the mean-reversion rate of such processes can vary freely when moving from one probability measure to the other. However, that flexibility comes at the cost of requiring that the Feller condition[2] be satisfied under both the P- and the Q-measure. In this paper the focus is on whether the decomposition of CDS rates into default and recovery risk can take advantage of the added flexibility in the extended affine specification or is impeded by having to satisfy the Feller conditions, and if so how it impacts the separation of default and recovery risk.

The data set used in the estimations consists of weekly observations of the yield of the second most-recently issued treasury bonds and quotes of CDS contracts for Ford Motor Co. from March 2, 2001 to June 24, 2005. The choice of using treasury data as a proxy for the risk-free rate deserves some comments. Ideally, general collateral rates should be used since this contract type is closest in spirit to the concept of a pure risk-free rate. Unfortunately, these rates are only observed for very short maturities. Conversely, yields on treasury bonds are easily observable for all relevant maturities, but they are depressed relative to the pure risk-free rate for three primary reasons

- There is a tax effect because interest earned on treasury securities is only taxable at the federal level.
- There are liquidity effects related to the size and efficiency of the treasury bond market.
- There is the effect of 'specialness' in the repo market, in particular with respect to on-the-run treasury instruments as documented in Duffie (1996b).

These factors have made researchers turn to interest rate swap data as a substitute for the risk-free rate.[3] There is evidence in the empirical literature on the relationship between corporate bond spreads and CDS rates that interest rate swap data may be more appropriate to serve as the risk-free benchmark than treasury yields (for examples see Longstaff, Mithal, and Neis (2005) and Blanco, Brennan, and Marsh (2005)). However, because the estimations performed in this paper are latent factor Kalman filter estimations, the results will not change in any significant way by using swap data in stead of treasury bond yields. Furthermore, CDS rates are already spreads. Thus, the choice of risk-free benchmark is without consequences for the distribution of the observed CDS rates unlike the case of corporate bond yields, where a risk-free yield must be chosen in order to convert the observed yields into credit spreads. Finally, in order to mitigate

[2]This condition implies that the lower boundary is an entrance barrier for the square root-type processes, and without this restriction the model is not arbitrage-free as proved by Cheridito, Filipović, and Kimmel (2005).

[3]Note that interest rate swap data are not free from default risk since they are tied to LIBOR rates, which are essentially the rates at which AA-rated financial institutions can borrow in the interbank market.

the effect of liquidity and repo 'specialness' only yields of the second-most-recently issued treasury bonds are used.

The estimation method used is a latent factor quasi maximum likelihood Kalman filter estimation along the lines of Duffee (1999) and Driessen (2005). First, the two-factor risk-free interest rate model is estimated using the treasury yield data. Second, given the result of the first estimation, a similar Kalman filter estimation of the paths and parameters of the default intensity and recovery risk factors is performed using the CDS data. In theory, it would be possible to make a joint estimation of all four factors. However, given that one objective of this paper is to compare performance across models, the two-step estimation strategy is the only consistent way to proceed. In addition, little would be gained by performing a joint estimation provided treasury yields vary close to independently of the corporate performance of individual firms like Ford Motor Co.

The primary findings of the analysis are: 1) It is feasible to make a joint estimation of default and recovery risk in the case of the Ford Motor Co. CDS data. 2) The results indicate that the risk premia associated with recovery risk make a significant contribution to the observed CDS rates. 3) With respect to risk premium specifications the Feller conditions are binding for the extended affine specification in the estimations of the CDS models. This causes the essentially affine specification to be preferred for credit risk pricing purposes.

The remainder of the paper is organized as follows. Section 2 describes the general pricing of CDS contracts in a reduced-form setting. Section 3 sets up the specific CDS models analyzed in this paper and describes how CDS rates are calculated in the models. Section 4 describes the data, while Section 5 explains the estimation procedure. Section 6 analyzes the results of the estimations. Finally, Section 7 concludes the paper. Appendices contain additional technical details.

2 Credit default swap pricing

In this section reduced-form pricing of credit default swaps is described. It is assumed that there is a model for the instantaneous risk-free interest rate r_t, a model for the default intensity under the Q-measure of the firm considered λ_t^Q,[4] and a model for the recovery rate given default under the pricing measure π_t^Q. Finally, it is assumed that any bonds issued by the firm are priced by the market using the assumption of Recovery of Face Value.[5] Thus, it is assumed that at default bond holders will receive a fraction π_τ^Q of the face value of the bond.[6] The default time τ is modeled in a doubly-stochastic point process framework where the timing is given by the first jump of a point process with stochastic intensity λ_t^Q. This implies that conditional on the entire path of λ_t^Q the default time is an independent process.[7]

[4] There is evidence in Driessen (2005) and Berndt et al. (2005) that a so-called default event risk premium exists and is significant. If so, the default intensity under the Q-measure will vary from the default intensity under the P-measure by a factor. As this paper only uses market data, all modeling refers to the default intensity under the Q-measure as indicated by the notation.

[5] A recent empirical study by Guha (2002) documents that, at and in a period of time after a default announcement, bonds of the same seniority with different coupon sizes and different remaining time to maturity are trading at practically identical prices. He finds that the assumption of Recovery of Face Value is the recovery assumption best in line with this observation.

[6] Any accrued interest is neglected here but can easily be incorporated into the pricing formulas to follow. One interpretation is that it is only the 'total' market-implied expected recovery rate that is being modeled. How that recovery rate may actually be distributed in a future default situation between recovered principal and recovered accrued interest is not modeled.

[7] See Lando (1998) for details.

This setup may appear somewhat simplistic. In a typical CDS contract there are several events that can trigger the default event clauses of the contract. The most frequently used are:[8]

- Bankruptcy,
- Failure to Pay,
- Obligation Default,
- Obligation Acceleration,
- Restructuring.

The occurrence of each of these events can be modeled with their individual default intensity $\lambda_t^{Q,i}$ and their individual recovery rates $\pi_t^{Q,i}$ for the $i = 1, \ldots, N$ events. The first jump of N independent point processes is itself a jump process with a stopping time $\tau = \min(\tau_1, \ldots, \tau_N)$ and an intensity given by the sum of the individual intensities[9]

$$\lambda_t^Q = \lambda_t^{Q,1} + \ldots + \lambda_t^{Q,N}.$$

Given the first jump has occurred, the probability that it was the ith event that triggered the jump is simply the ratio $\frac{\lambda_\tau^{Q,i}}{\lambda_\tau^Q}$. The expected instantaneous recovery rate given default is the weighted sum

$$\pi_\tau^Q = \frac{\lambda_\tau^{Q,1}}{\lambda_\tau^Q}\pi_\tau^{Q,1} + \ldots + \frac{\lambda_\tau^{Q,N}}{\lambda_\tau^Q}\pi_\tau^{Q,N}.$$

In the pricing formulas below only λ_t^Q and π_t^Q will appear. Thus, it is possible to imagine a complex underlying structure with different recovery distributions depending on the event that triggered the default.[10] However, by solely observing the term structure of CDS rates across time it is not possible to identify these individual components. As a consequence the default intensity in this analysis is interpreted as the sum of the underlying default event intensities, while the recovery rate is interpreted as the weighted average of the expected recovery rate from each of those events where the weights are the relative intensity with which each event occurs.

A final technical detail is that the recovery rate must be a predictable process, i.e. $\pi_t^Q = E^Q[\pi_\tau^Q|\mathcal{F}_{t-}]$, such that the recovery rate is known an instant before the default jump occurs. However, since the state variables considered in this paper are continuous diffusion processes, this technical requirement is trivially satisfied.

2.1 The swap premium of a plain vanilla CDS

Now, consider a plain vanilla credit default swap with the above issuer as reference credit. Let T denote the time to maturity of the CDS contract and let $t_1, \ldots, t_N$ denote the swap premium payment dates. In case of default before T the payment on the default leg is assumed to be 1 minus the recovery rate times the size of the notional, consistent with an assumption of recovery of face value at default for any bonds issued by the company. Given these assumptions the value

[8] For more details, see Bomfin (2005), chapter 24.

[9] See Duffie (1998) for details.

[10] This would be a theoretical setup that could match the bi-modal empirical distribution of recovery rates, see Schuermann (2005) for details.

of the default leg of a plain vanilla CDS contract can be calculated as

$$\begin{aligned} V_{def}^{CDS}(0,T) &= E^Q\Big[\int_0^T (1-\pi_s^Q)\mathbf{1}_{\{s<\tau\leq s+ds\}}e^{-\int_0^s r_u du}ds\Big|\mathcal{F}_0\Big] \\ &= E^Q\Big[\int_0^T (1-\pi_s^Q)\lambda_s^Q e^{-\int_0^s (r_u+\lambda_u^Q)du}ds\Big|\mathcal{F}_0\Big]. \end{aligned}$$

In return for the loss protection, the protection buyer has to pay a premium, here denoted by $S^{CDS}(0,T)$ and quoted at an annual rate. Define by $\delta_i = t_i - t_{i-1}$ the time between the ith and the (i-1)th payment dates, then the contractual payment on the ith payment date is equal to $\delta_i S^{CDS}(0,T)$. In addition, the market convention is that immediately upon default the accrued swap premium since the last payment date is exchanged in return for the default leg payment. Given this convention the value of the premium leg can be calculated as

$$\begin{aligned} V_{prem}^{CDS}(0,T) &= E^Q\Big[S^{CDS}(0,T)\sum_{i=1}^N \delta_i \mathbf{1}_{\{\tau>t_i\}}e^{-\int_0^{t_i} r_u du}|\mathcal{F}_0\Big] \\ &+ E^Q\Big[\sum_{i=1}^N \int_{t_{i-1}}^{t_i} S^{CDS}(0,T)(s-t_{i-1})\mathbf{1}_{\{s<\tau\leq s+ds\}}e^{-\int_0^s r_u du}ds|\mathcal{F}_0\Big] \\ &= S^{CDS}(0,T)E^Q\Big[\sum_{i=1}^N \delta_i e^{-\int_0^{t_i}(r_u+\lambda_u^Q)du} + \sum_{i=1}^N \int_{t_{i-1}}^{t_i}(s-t_{i-1})\lambda_s^Q e^{-\int_0^s (r_u+\lambda_u^Q)du}ds|\mathcal{F}_0\Big]. \end{aligned}$$

At inception the swap premium is set to give the contract a value of zero:

$$S^{CDS}(0,T) = \frac{E^Q\Big[\int_0^T (1-\pi_s^Q)\lambda_s^Q e^{-\int_0^s (r_u+\lambda_u^Q)du}ds|\mathcal{F}_0\Big]}{E^Q\Big[\sum_{i=1}^N \delta_i e^{-\int_0^{t_i}(r_u+\lambda_u^Q)du} + \sum_{i=1}^N \int_{t_{i-1}}^{t_i}(s-t_{i-1})\lambda_s^Q e^{-\int_0^s (r_u+\lambda_u^Q)du}ds|\mathcal{F}_0\Big]}.$$

3 Model description

This section describes the modeling frameworks used in the subsequent empirical estimations. In Section 3.1 the model for the risk-free rate will be described. As for the CDS data two types of models are estimated in order to get a good understanding of the impact of changing the risk premium specification and of making different recovery rate assumptions. Section 3.2 has the details on what will be denoted here as a standard CDS model, i.e. a model where only the interest rate and the default intensity is modeled, while the recovery rate is left as a constant. Section 3.3 describes the four-factor CDS model that incorporates interest rate risk, default intensity risk as well as recovery risk. Finally, Section 3.4 describes how CDS contracts are priced in this model framework.

3.1 The model for the risk-free rate

The primary purpose of including a model for the risk-free rate is to be able to include the level and the slope factor of the yield term structure as explanatory factors in the subsequent modeling of the default intensity and the recovery rate processes. For that reason a two-factor model is adequate. In the specific choice of a two-factor model, this paper follows Driessen (2005) and

models the default-free instantaneous short rate as an affine function

$$r_t = \delta_0 + \delta_1 X_t^1 + \delta_2 X_t^2,$$

where (X_t^1, X_t^2) are the solutions of a so-called $A_1(2)$ (see Dai and Singleton(2000)) affine system of stochastic differential equations with the following structure under the P-measure

$$\begin{pmatrix} dX_t^1 \\ dX_t^2 \end{pmatrix} = \begin{pmatrix} \kappa_{11} & 0 \\ \kappa_{21} & \kappa_{22} \end{pmatrix} \begin{pmatrix} \theta_1 - X_t^1 \\ -X_t^2 \end{pmatrix} dt + \begin{pmatrix} \sqrt{X_t^1} & 0 \\ 0 & \sqrt{1+\beta_{21} X_t^1} \end{pmatrix} \begin{pmatrix} dW_t^1 \\ dW_t^2 \end{pmatrix}.$$

Here, W_t^1 and W_t^2 are independent standard Brownian motions. This model allows for stochastic volatility in both factors and for two ways of interaction between the two factors: 1) Directly, through the presence of X_t^1 in the drift and volatility of X_t^2. 2) Indirectly through their presence in the affine expression for r_t.

As for the risk premium structure leading to the measure change

$$d\widehat{W}_t = \Lambda_t dt + dW_t,$$

two specifications will be estimated. The first is the essentially affine specification where

$$\Lambda_t^{ess} = \begin{pmatrix} \sqrt{X_t^1} & 0 \\ 0 & \sqrt{1+\beta_{21} X_t^1} \end{pmatrix} \begin{pmatrix} \lambda_1^1 \\ \lambda_2^1 \end{pmatrix} + \begin{pmatrix} 0 & 0 \\ 0 & \frac{1}{\sqrt{1+\beta_{21} X_t^1}} \end{pmatrix} \begin{pmatrix} 0 & 0 \\ \lambda_{21}^{(2)} & \lambda_{22}^{(2)} \end{pmatrix} \begin{pmatrix} X_t^1 \\ X_t^2 \end{pmatrix}.$$

This specification was introduced by Duffee (2002) and applied by Driessen (2005) in his estimation based on US treasury data. For this model to be well-defined, it is only necessary to impose the following parameter restriction

$$\kappa_{11}\theta_1 > 0, \tag{1}$$

which ensures that the zero-boundary for the CIR-process X_t^1 is a reflecting boundary. The intuition behind this constraint is simply to prevent the X_t^1-process from becoming negative in which case its volatility would not be well-defined.

The second risk premium structure is the extended affine risk premium specification. In this specification, Λ_t is given by

$$\begin{aligned} \Lambda_t^{Ext} &= \begin{pmatrix} \sqrt{X_t^1} & 0 \\ 0 & \sqrt{1+\beta_{21} X_t^1} \end{pmatrix} \begin{pmatrix} 0 \\ \lambda_2^1 \end{pmatrix} + \begin{pmatrix} 0 & 0 \\ 0 & \frac{1}{\sqrt{1+\beta_{21} X_t^1}} \end{pmatrix} \begin{pmatrix} 0 & 0 \\ \lambda_{21}^{(2)} & \lambda_{22}^{(2)} \end{pmatrix} \begin{pmatrix} X_t^1 \\ X_t^2 \end{pmatrix} \\ &+ \begin{pmatrix} \frac{1}{\sqrt{X_t^1}} & 0 \\ 0 & 0 \end{pmatrix} \left[\begin{pmatrix} \lambda_1^3 \\ 0 \end{pmatrix} + \begin{pmatrix} \lambda_{11}^4 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} X_t^1 \\ X_t^2 \end{pmatrix} \right]. \end{aligned}$$

This specification was introduced by Cheridito, Filipović, and Kimmel (2005), and they prove that this structure is arbitrage-free provided the following parameter restrictions are imposed

$$\kappa_{11}\theta_1 > \frac{1}{2} \quad \text{and} \quad \kappa_{11}\theta_1 - \lambda_1^3 > \frac{1}{2}. \tag{2}$$

These two restrictions are identical to the Feller condition for the X_t^1-process under the P- and the Q-measure. The Feller condition ensures that the zero-boundary is an entrance boundary for the CIR-process X_t^1. Thus, with these restrictions imposed the process will remain strictly positive

and never hit zero. The intuition for this requirement is that the risk premium term $\frac{1}{\sqrt{X_t^1}}\lambda_1^3 dt$ cannot be allowed to explode P-a.s. With only the restriction in Equation (1), this risk premium will go to infinity whenever X_t^1 hits the zero-boundary. This is equivalent to an infinite premium for carrying zero risk which constitutes an arbitrage as explained by Cheridito, Filipović, and Kimmel (2005). The model becomes arbitrage-free by imposing the constraints in Equation (2).[11] The practical implication of the extended specification is that it provides more flexibility in the sense that the mean level and the mean reversion rate of the X_t^1-process under the Q-measure can vary independently from their respective values under the P-measure. The trade-off is that the increased flexibility comes at the cost of imposing more strict parameter constraints. Estimation of both specifications can reveal which is the more appropriate for the treasury yields considered here.[12] Finally, a salient feature of both risk premium specifications is that the processes remain affine under both measures which is convenient from a technical point of view for both pricing and estimation purposes.

3.2 A standard CDS model with constant recovery rate

The first reduced-form credit risk model considered in this paper is similar to the model applied by Duffee (1999). Thus, the default intensity is an affine function of the interest rate risk factors and an idiosyncratic risk factor denoted X_t^λ

$$\lambda_t^Q = \lambda_0 + \lambda_r^1(X_t^1 - \theta_1) + \lambda_r^2 X_t^2 + X_t^\lambda,$$

where X_t^λ is assumed to follow a CIR-process

$$dX_t^\lambda = \kappa_\lambda(\theta_\lambda - X_t^\lambda)dt + \sigma_\lambda\sqrt{X_t^\lambda}dW_t^\lambda,$$

while the market-implied recovery rate is kept as a constant π_0.

The interest rate risk factors are included because empirical studies such as Duffee (1998), Duffee (1999), and Driessen (2005) have all found evidence of statistically significant negative correlation between treasury yield factors and US corporate bond yield spreads. However, the specific nature of that correlation varies across those three studies. In Duffee (1999) it is the slope factor that is significant, while in Driessen (2005) it is the level factor, which is similar to the findings of Duffee (1998) in his regression analyses.[13] At the other extreme, Pan and Singleton (2005) just use a constant interest rate $\overline{r}$. They base their choice on the observation that the fixed spread on a floating rate corporate bond only has little sensitivity to changes in the interest rate, and as the CDS rate theoretically is identical to the par floating rate spread as demonstrated by Duffie (1999), the CDS rate should only be mildly correlated with the interest rate factors, thus

[11] As two measures are only equivalent if they assign zero probability to the same events, it is also necessary to require that X_t^1 never hits zero under the Q-measure.

[12] Cheridito, Filipović, and Kimmel (2005) do make a comparison between the two specifications for all affine subclasses $A_i(n)$, $i = 0, 1, \ldots, n$ for $n \in \{1, 2, 3\}$, i.e. all one-, two-, and three-factor models. However, they impose the Feller condition even on the essentially affine risk premium specifications where it is not needed. This implies that the essentially affine specification becomes nested by the extended affine specification, and thus allows for direct inference on the validity of the added parameters in the extended specification, but it is not a fair comparison if the true maximum of the likelihood function for the essentially affine specification is violating the Feller condition. Unfortunately, they do not discuss this issue.

[13] The results in Duffee (1998) even indicate that the lower the rating, the larger are the negative changes in credit spreads from positive changes in the level of the treasury yield curve.

not much would be lost by fixing the interest rate. Furthermore, their CDS data concern three foreign sovereigns (Russia, Mexico, and Turkey), so it is not evident that, say, the US treasury or swap yields should play any significant role for the CDS spread development for these three countries. In conclusion, there is no firm prior expectation about the relationship between the risk-free rate factors and the CDS rate dynamics.

Turning to the risk premium specification, Driessen (2005) observes that provided sufficiently many common risk factors (beyond the risk-free rate factors) are included in the estimation, the median company does not carry any significant risk premium for its idiosyncratic risk factor. However, in the current study, only the two risk-free rate factors are included as opposed to Driessen (2005) who includes two additional common risk factors. Therefore the idiosyncratic risk factor will most likely contain some non-diversifiable risk that will carry a risk premium. For that reason risk premia for X_t^λ are included in the estimation.

In the empirical credit spread literature, Duffee (1999), Driessen (2005), Zhang (2003), Longstaff, Mithal, and Neis (2005) along with many others before them only estimate the essentially affine risk premium specification for the default intensity factor

$$d\widetilde{W}_t^\lambda = dW_t^\lambda + \gamma_\lambda^1 \sqrt{X_t^\lambda} dt.$$

This structure implies that the dynamics under the equivalent measure is given by

$$dX_t^\lambda = (\kappa_\lambda \theta_\lambda - (\kappa_\lambda + \gamma_\lambda^1 \sigma_\lambda) X_t^\lambda) dt + \sigma_\lambda \sqrt{X_t^\lambda} d\widetilde{W}_t^\lambda.$$

Thus, when changing probability measure, it is not possible to change the mean without also changing the mean-reversion rate. Another potential drawback is that the risk premium is unable to change sign across time since its sign is given by the sign of γ_λ^1. The advantage of this specification is that the only restriction is to ensure that zero is a reflecting boundary for the X_t^λ-process, which is obtained by imposing the inequality

$$\kappa_\lambda \theta_\lambda \geq 0. \tag{3}$$

The second risk premium structure implemented in this paper is the extended affine specification proposed by Cheridito, Filipović, and Kimmel (2005) and applied to sovereign CDS spreads by Pan and Singleton (2005). Here, the measure change is described by

$$d\widetilde{W}_t^\lambda = dW_t^\lambda + \Big(\frac{\gamma_\lambda^0}{\sqrt{X_t^\lambda}} + \gamma_\lambda^1 \sqrt{X_t^\lambda} \Big) dt.$$

This structure implies the following dynamics under the Q-measure

$$dX_t^\lambda = (\kappa_\lambda \theta_\lambda - \gamma_\lambda^0 \sigma_\lambda - (\kappa_\lambda + \gamma_\lambda^1 \sigma_\lambda) X_t^\lambda) dt + \sigma_\lambda \sqrt{X_t^\lambda} d\widetilde{W}_t^\lambda.$$

For this measure change not to allow any arbitrages, zero has to be an entrance boundary for the X_t^λ-process under both measures, which is obtained by the two Feller conditions

$$\kappa_\lambda \theta_\lambda > \frac{1}{2}\sigma_\lambda^2 \quad \text{and} \quad \kappa_\lambda \theta_\lambda - \gamma_\lambda^0 \sigma_\lambda > \frac{1}{2}\sigma_\lambda^2.$$

This parameter restriction is more strict than the constraint in Equation (3). Driessen (2005) finds in his sample of 104 firms that the Feller condition is binding for only 7 of them. On the other hand, Duffee (1999) who, unlike Driessen (2005), does not include common risk beyond the risk-free factors finds that the median constellation of parameters in the Aa-, the A-, and the Baa-rating categories, respectively, does not satisfy the Feller condition. These two results combined indicate that the Feller condition is not likely to be satisfied for the Ford Motor CDS data analyzed here. This implies that there may be a real trade-off to imposing the Feller condition in the extended affine specification. Since Pan and Singleton (2005) in their estimation of the affine model for X_t^λ only apply the extended affine specification,[14] there is no real empirical evidence on the relative performance of these two risk premium specifications. By estimating both specifications the nature of the potential trade-off between the two can be studied closer.

Turning to the recovery rate π_0 it is treated as a constant parameter to be estimated which is similar to the estimations in Zhang (2003) and Pan and Singleton (2005). Pan and Singleton (2005) demonstrate convincingly with analytical results for the first partial derivatives as well as with a simulation study that all of the above parameters are well identified when the data consist of time series of CDS contracts with several different maturities observed at each point in time.

Finally, it should be emphasized that these intermediate estimations are performed partly to shed some light on the impact of varying the risk premium specification and partly to have a more familiar background against which the results for the full four-factor model with stochastic default and recovery risk can be compared.

3.3 A CDS model with stochastic recovery

In this section a joint model for the default intensity and the market-implied expected recovery rate given default is proposed. As for the default intensity it is still assumed to be an affine function of the interest rate factors and an idiosyncratic default risk factor

$$\lambda_t^Q = \lambda_0 + \lambda_r^1(X_t^1 - \theta_1) + \lambda_r^2 X_t^2 + X_t^\lambda,$$

where X_t^λ is a CIR-process

$$dX_t^\lambda = \kappa_\lambda(\theta_\lambda - X_t^\lambda)dt + \sigma_\lambda\sqrt{X_t^\lambda}dW_t^\lambda.$$

In the empirical literature on recovery rates, there is evidence that the recovery rates observed immediately at default are positively correlated with short term interest rates.[15] To take this evidence into account and implicitly assume that the participants in the CDS market price a similar relationship into the market-implied expected recovery rates, the stochastic recovery rate is modeled as an affine function of the risk-free rate factors plus an additional idiosyncratic recovery

[14] Besides an affine formulation for X_t^λ, they also estimate a log-normal model and a three-halves model. They find that the log-normal is the best performing of the three models. However, the log-normal is also the least mathematically tractable model of the three, so it does not immediately lend itself to an analysis of the type performed in this paper.

[15] Covitz and Han (2004) use 3-month Treasury Bill rates as a proxy for r_t and regress them on observed recovery rates across more than 1200 defaulted debt instruments. They find a highly significant positive factor loading on the interest rate in their regressions.

risk factor[16]

$$\pi_t^Q = \pi_0 + \pi_r^1(X_t^1 - \theta_1) + \pi_r^2 X_t^2 + X_t^\pi,$$

where the idiosyncratic recovery risk factor X_t^π is assumed to follow a standard Ornstein-Uhlenbeck process

$$dX_t^\pi = \kappa_\pi(\theta_\pi - X_t^\pi)dt + \sigma_\pi dW_t^\pi.$$

For reasons of identification, the mean of the idiosyncratic factor is fixed at zero, i.e. $\theta_\pi = 0$.

As explained in Section 2 on the pricing of CDS contracts, the market-implied recovery rate process π_t^Q that will be estimated in this paper can be thought of as a weighted average recovery rate across several potential future default scenarios, each with their own stochastic recovery rate process. For that reason, the Gaussian distribution of the recovery risk factor X_t^π seems like a good first approximation. And for the same reason there is no theoretical inconsistency between the modeling approach taken here and the very noisy data on observed recovery rates reported in studies like Altman and Kishore (1996) that easily can be interpreted as indicating that observed recovery rates follow a bi- or even multi-modal distribution. These observed draws from different default conditions wash out in the pricing formulas even if they are included in the model.

In the framework above, the default risk factor and the recovery risk factor are assumed to be independent. However, in a structural model of the Merton (1974)-type there is a negative relationship between default probabilities and the expected recovery rate given default driven by changes in the underlying asset value process of the firm. This relationship is not modeled here, but if present in the data it should show up as a negative correlation between the estimated paths of X_t^λ and X_t^π.

As for risk premia, it has already been argued in the last section that the idiosyncratic risk factors will likely contain some systematic risk not caught by the two risk-free interest rate factors. However, ex ante there is no reason to believe that the default intensity risk factor would be the sole idiosyncratic risk factor to carry a risk premium. In fact, Bakshi, Madan, and Zhang (2004) provide an example in a utility-based setting with stochastic loss rates and risk-averse investors where both the expected default probability and the expected loss rate are higher under the pricing measure than the corresponding expected values under the P-measure. For this reason risk premia are attached to both idiosyncratic factors. As is the case with the fixed-recovery CDS model discussed in the last section, this model will be estimated with two risk premium specifications. The essentially affine specification given by the measure change

$$\begin{aligned} d\widetilde{W}_t^\lambda &= dW_t^\lambda + \gamma_\lambda^1\sqrt{X_t^\lambda}dt, \\ d\widetilde{W}_t^\pi &= dW_t^\pi + (\gamma_\pi^0 + \gamma_\pi^1 X_t^\pi)dt, \end{aligned}$$

with the parameter restriction

$$\kappa_\lambda\theta_\lambda \geq 0.$$

[16]Based on the results in Acharya, Bharath, and Srinivasan (2006) it could be argued that an industry factor should be included in addition to the interest rate factors. However, with data from only one firm it is not possible to identify such factors. If present in the data, they will be an implicit part of the idiosyncratic risk factors X_t^λ and/or X_t^π.

And the extended affine risk premium specification given by[17]

$$\begin{aligned} d\widetilde{W}_t^{\lambda} &= dW_t^{\lambda} + \Big(\frac{\gamma_{\lambda}^0}{\sqrt{X_t^{\lambda}}} + \gamma_{\lambda}^1\sqrt{X_t^{\lambda}}\Big)dt, \\ d\widetilde{W}_t^{\pi} &= dW_t^{\pi} + (\gamma_{\pi}^0 + \gamma_{\pi}^1 X_t^{\pi})dt, \end{aligned}$$

with the more demanding parameter restrictions

$$\kappa_{\lambda}\theta_{\lambda} > \frac{1}{2}\sigma_{\lambda}^2 \quad \text{and} \quad \kappa_{\lambda}\theta_{\lambda} - \gamma_{\lambda}^0\sigma_{\lambda} > \frac{1}{2}\sigma_{\lambda}^2.$$

The primary purpose of estimating this model is to provide evidence that the dynamics of the default risk factor X_t^{λ} and the recovery risk factor X_t^{π} can be jointly identified from cross-sections of observed CDS rates. In addition, the analysis can shed some light on the advantages and disadvantages of applying the extended affine specification in a credit risk setting. Finally, this set of estimations can provide insights about the relative proportion of risk premia attached to the default risk and the recovery risk factors and their respective impact on the level of observed CDS rates.

3.4 CDS pricing

This section provides the details of pricing CDS contracts in the modeling frameworks introduced in the last two sections.

The premium of a default swap is given by the formula

$$S^{CDS}(0,T) = \frac{E^Q\Big[\int_0^T (1-\pi_s^Q)\lambda_s^Q e^{-\int_0^s (r_u+\lambda_u^Q)du} ds|\mathcal{F}_0\Big]}{E^Q\Big[\sum_{i=1}^N \delta_i e^{-\int_0^{t_i}(r_u+\lambda_u^Q)du} + \sum_{i=1}^N \int_{t_{i-1}}^{t_i}(s-t_{i-1})\lambda_s^Q e^{-\int_0^s (r_u+\lambda_u^Q)du} ds|\mathcal{F}_0\Big]},$$

where T is the time to maturity, $t_1, \ldots, t_N$ are the swap premium payment dates, and $\delta_i = t_t - t_{i-1}$ denotes the time between the ith and the (i-1)th payment dates. It follows from the formula above that in order to price CDS contracts in the proposed reduced-form framework expectations of the following types have to be calculated

$$E^Q\big[\lambda_s^Q e^{-\int_0^s (r_u+\lambda_u^Q)du}|\mathcal{F}_0\big] \quad \text{and} \quad E^Q\big[(1-\pi_s^Q)\lambda_s^Q e^{-\int_0^s (r_u+\lambda_u^Q)du}|\mathcal{F}_0\big].$$

To that end Christensen (2006a) has the following result.

Proposition 1:

Let the state variables X_t be described by an affine diffusion process

$$dX_t = [\mu^0 + \mu^1 X_t]dt + \Sigma D(X_t)dW_t \tag{4}$$

defined on a set $M \subset \mathbf{R}^n$. Here, $D : M \to \mathbf{R}^{n\times n}$ is assumed to have the following diagonal

[17] Note that because X_t^{π} is Gaussian the essentially affine and the extended affine specifications are identical for this factor.

structure

$$D(X_t) = \begin{pmatrix} \sqrt{\gamma^1 + \delta_1^1 X_t^1 + \ldots + \delta_n^1 X_t^n} & \ldots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \ldots & \sqrt{\gamma^n + \delta_1^n X_t^1 + \ldots + \delta_n^n X_t^n} \end{pmatrix}.$$

To simplify notation below, γ and δ are defined as

$$\gamma = \begin{pmatrix} \gamma^1 \\ \vdots \\ \gamma^n \end{pmatrix} \quad \text{and} \quad \delta = \begin{pmatrix} \delta_1^1 & \ldots & \delta_n^1 \\ \vdots & \ddots & \vdots \\ \delta_1^n & \ldots & \delta_n^n \end{pmatrix}.$$

In addition, assume the discount process to be affine in the state variables

$$R(X_t) = \rho_0 + \rho_1' X_t.$$

Then the expectation

$$G(X_t, t, T) = E\Big[e^{-\int_t^T R(X_s)ds} e^{\overline{B}' X_T + \overline{C}} \big[X_T' \overline{D} X_T + \overline{E}' X_T + \overline{F}\big] \Big| \mathcal{F}_t\Big],$$

where $\overline{D} \in \mathbf{R}^{n\times n}$, $\overline{B}, \overline{E} \in \mathbf{R}^n$, and $\overline{C}, \overline{F} \in \mathbf{R}$, is given by

$$G(X_t, t, T) = \exp\big(B(t,T)' X_t + C(t,T)\big)\big[X_t' D(t,T) X_t + E(t,T)' X_t + F(t,T)\big],$$

if the following conditions are satisfied.

(i). There exists a unique solution X_t for the stochastic differential equation (4) for all $0 \le t \le T$.

(ii). There exist functions $B(t,T)$, $C(t,T)$, $D(t,T)$, $E(t,T)$, and $F(t,T)$ which are the unique solutions for the following system of ordinary first order differential equations[18]

$$\begin{aligned}
\frac{dB(t,T)}{dt} &= \rho_1 - (\mu^1)' B(t,T) - \frac{1}{2}\sum_{j=1}^n (\Sigma' B(t,T) B(t,T)' \Sigma)_{j,j} (\delta^j)', \\
\frac{dC(t,T)}{dt} &= \rho_0 - B(t,T)' \mu^0 - \frac{1}{2}\sum_{j=1}^n (\Sigma' B(t,T) B(t,T)' \Sigma)_{j,j} \gamma^j, \\
\frac{dD(t,T)}{dt} &= -(\mu^1)' D(t,T) - D(t,T)\mu^1 - \sum_{j=1}^n (\Sigma' B(t,T))_j (D(t,T)\Sigma)_{\cdot,j} \delta^j \\
&\quad - \sum_{j=1}^n (\Sigma' B(t,T))_j (D(t,T)'\Sigma)_{\cdot,j} \delta^j, \\
\frac{dE(t,T)}{dt} &= -(\mu^1)' E(t,T) - D(t,T)' \mu^0 - D(t,T)\mu^0 - \sum_{j=1}^n (\Sigma' B(t,T) E(t,T)' \Sigma)_{j,j} (\delta^j)' \\
&\quad - \sum_{j=1}^n (\Sigma' D(t,T)\Sigma)_{j,j} (\delta^j)' - \sum_{j=1}^n (\Sigma' B(t,T))_j (D(t,T)\Sigma)_{\cdot,j} \gamma^j \\
&\quad - \sum_{j=1}^n (\Sigma' B(t,T))_j (D(t,T)'\Sigma)_{\cdot,j} \gamma^j, \\
\frac{dF(t,T)}{dt} &= -E(t,T)' \mu^0 - \sum_{j=1}^n (\Sigma' D(t,T)\Sigma)_{j,j} \gamma^j - \sum_{j=1}^n (\Sigma' B(t,T) E(t,T)' \Sigma)_{j,j} \gamma^j,
\end{aligned}$$

[18] Here, δ^j denote the jth row of the δ-matrix.

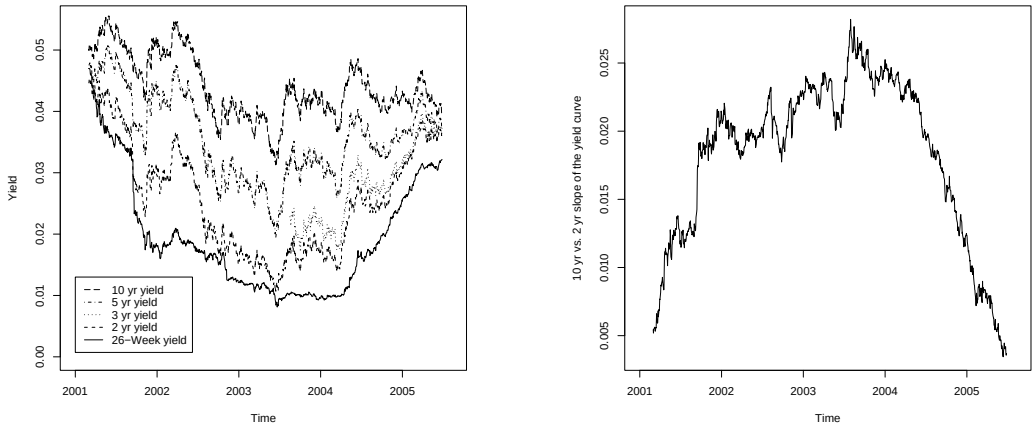

(a) The observed yields on treasury bills.

(b) The 10-year vs. 2-year slope of the treasury yield curve.

Figure 1: The observed treasury yield data for the period March 2, 2001 through June 24, 2005. Note that the 3-year Treasury Notes are not issued before May 15, 2003, which means that the second most recent 3-year yield is not observed until August 15, 2003 where the second series of 3-year Treasury Notes is issued.

(*iii*). The following technical conditions are met

(a) $E[|\Phi_T|] < \infty$,

(b) $E\Big[\big(\int_0^T \eta_t \eta_t' dt\big)^{1/2}\Big] < \infty$ for $\eta_t = \big(\Phi_t B(t,T)' + \Psi_t[E(t,T)' + X_t'(D(t,T) + D(t,T)')]\big)\Sigma D$,

where $\Phi_t = e^{-\int_0^t r(X_s,s)ds} e^{B(t,T)'X_t + C(t,T)}[X_t' D(t,T) X_t + E(t,T)' X_t + F(t,T)]$ and

$\Psi_t = e^{-\int_0^t r(X_s,s)ds} e^{B(t,T)'X_t + C(t,T)}$ for all $0 \le t \le T$.

This proposition implies that it only requires the solution of a system of ODEs to calculate the expectations relevant for pricing CDS contracts. In this paper this is done by standard fourth order Runge Kutta methods.

4 Data description

The treasury data set consists of yields for US Treasury Bills and Notes downloaded from Bloomberg. First, observed yields on all 2-year, 3-year, 5-year, and 10-year Notes issued during the period from March 2, 2001 until June 24, 2005 were obtained. Second, yields on all 26-Week Treasury Bills issued over that same period were downloaded.[19] Here, it should be noted that the first series of 3-year notes were issued on May 15, 2003, therefore for this specific maturity no data is available before this date. All the observed yields are mid quotes of the yield-to-maturity, i.e. the average of the bid and the ask yield. In order to reduce any effects of specialness that particularly impact the on-the-run treasury bonds as documented by Duffie (1996), only yields of the second most recently issued bond for each of the above maturities are used in the estimation. Finally, to reduce noise in the estimation, only weekly observations were used, i.e. data for every Friday is used.

[19] Treasury Bills with maturities less than 26-Weeks were left out of consideration, as the very short end of the yield curve is known to be subject to special liquidity effects, see Duffee (1996) for details.

Treasury Bill/Note	Mean	St. deviation	No. obs.
26-Weeks	194.2	92.5	226
2-year	258.9	89.8	226
3-year	277.6	62.1	98
5-year	358.6	69.1	226
10-year	440.5	52.1	226

Table 1: Summary statistics for the observed treasury yield data for the period March 2, 2001 through June 24, 2005. Note that the 3-year Treasury Notes are not issued before May 15, 2003, which means that the second most recent 3-year yield is not observed until August 15, 2003.

The Treasury Notes are all bullet bonds with semi-annual coupon payments. This means that for a given bond price $P_t(C, N, \tau)$, where

- C is the annual rate of coupon payments,
- N is the number of remaining coupon payments,
- $\tau = \frac{\text{number of days until the next coupon payment}}{\text{number of days in coupon period}}$ is the fraction of the up-coming coupon payment yet to be accrued,[20]

the annual yield-to-maturity can be calculated as the solution to the following equation

$$\sum_{i=1}^{N} \frac{C}{2}\Big(1+\frac{y_t}{2}\Big)^{-(\tau+0.5(i-1))} + \Big(1+\frac{y_t}{2}\Big)^{-(\tau+0.5(N-1))} = P_t(C, N, \tau). \tag{5}$$

Let $P_{zero,t}^{model}(T)$ denote the model-implied zero-coupon bond price with maturity in T years. The model-implied price for a coupon bond with the above characteristics is then calculated as

$$P_t^{model}(C, N, \tau) = \sum_{i=1}^{N} \frac{C}{2} P_{zero,t}^{model}(\tau + 0.5(i-1)) + P_{zero,t}^{model}(\tau + 0.5(N-1)),$$

and the model-implied yield-to-maturity is obtained by replacing $P_t(C, N, \tau)$ in Equation (5) with $P_t^{model}(C, N, \tau)$ and solve for y_t.

The 26-Week Treasury Bills, on the other hand, are discount bonds where the remaining time to maturity is measured as

$$\tau = \frac{\text{number of days until maturity}}{365}.$$

The quoting mechanism for Treasury Bills is based on the so-called bankers' discount yield

$$b = \frac{\text{face value} - \text{price}}{\text{face value}} \frac{360}{\text{number of days until maturity}}.$$

Normalizing the face value to 1, the model-implied yield based on the bankers' discount convention can be calculated as

$$b_t^{model} = [1 - P_{zero,t}^{model}(\tau)]\frac{360}{n},$$

[20] In principle, the numerator of τ should be given by the number of days from settlement until the next coupon payment (see Jarrow (1996) p.12 for details). However, this minor deviation is not believed to influence the results in any significant way.

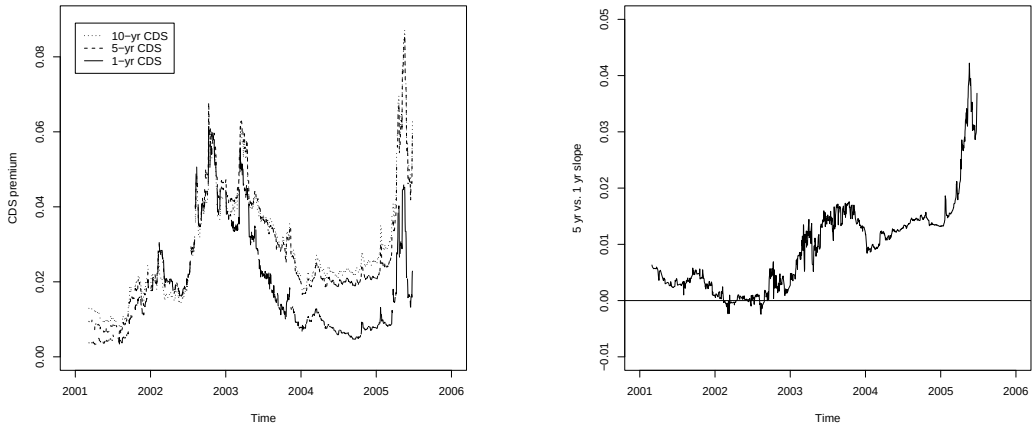

(a) The observed CDS spread for the 1-year, 5-year and 10-year contract.

(b) The 5-year vs. 1-year slope of the CDS spread term structure.

Figure 2: Illustration of the CDS spread data for Ford Motor Company over the period March 2, 2001 through June 24, 2005.

where

$$n = \text{number of days to maturity} = 365 \cdot \tau.$$

Summary statistics for the observed yields over the sample period are provided in Table 1. The time series of the observed yields are illustrated in Figure 1(a), while Figure 1(b) depicts the slope of the yield over the sample period, here measured as the difference between the 10-year and the 2-year treasury yields.

4.1 The CDS data

Daily quotes on CDS contracts for Ford Motor Co. with maturities ranging from 6 months to 30 years were downloaded from the Markit database. Markit, the CDS data provider, gathers daily closing quotes from a wide variety of dealers, but before new data is added to the database, the observations are tested for staleness, flat curves, and outliers. Data failing any one of these tests is discarded entirely. This is done to ensure that the quotes reflect as much real price information as possible. The data analyzed here is the daily average of the quotes passing all the above-mentioned tests.

Given the purpose of the estimations in this paper there is a trade-off to be made. It would be preferable to have as many observations of the CDS term structure as possible in order to better be able to distinguish movements in the default intensity from movements in the expected recovery rate. However, the evidence in Christensen (2006b) indicates that if there is too much noise in the observed CDS rates it could potentially ruin the ability of the estimation to make this distinction. For that reason, only the most traded maturities should be included. These are believed to be the 1-year, 3-year, 5-year, 7-year, and 10-year CDS contracts.[21] Figure 2(a) depicts the time series of CDS rates for the 1-year, 5-year, and 10-year CDS contracts covering the period from March 2, 2001, to June 24, 2005, while Figure 2(b) depicts the slope of the Ford Motor Co. CDS curve over that same period as measured by the difference between the 5-year and the 1-year CDS rates.

[21] Unfortunately, the dataset does not contain any information about the liquidity of the different contracts.

CDS contract	Mean	St. deviation	No. obs.
1-year	187.2	138.6	217
3-year	256.5	157.4	216
5-year	283.3	150.4	218
7-year	286.4	146.6	219
10-year	289.9	142.8	216

Table 2: Summary statistics for the weekly observations of CDS contracts with maturities in 1, 3, 5, 7, and 10 years, respectively, covering the period from March 2, 2001 to June 24, 2005.

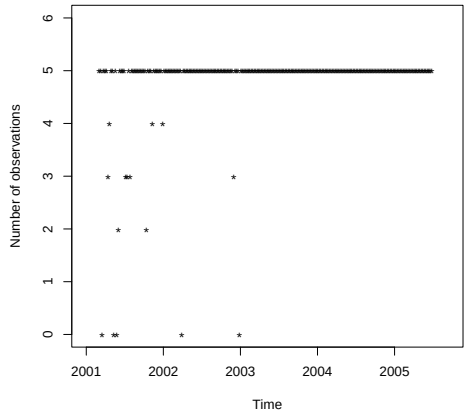

Figure 3: The number of CDS observations across time in the data set.

In order to reduce the noise in the estimation, only weekly observations were used. The weekly data set was constructed as follows. All quotes on Fridays are used, and if no quotes were available for a specific Friday, the quotes on the Thursday before were used. However, if neither the Friday nor the Thursday before it had any quotes, no quote would be observed for that particular contract in that week. The distribution of the total number of observed quotes across maturities in the final data set is provided in Table 2, while the total number of observed quotes across time is illustrated in Figure 3. It is noted that 5 weeks have no observations, and another 10 weeks have between 2 and 4 observations. However, since the beginning of 2003 until the end of the sample period there are 5 CDS quotes for every week.

From Table 2 it follows that on average the CDS term structure has been upward sloping over the 4-year period. However, the time series of the slope depicted in Figure 2(b) reveals that there has been large changes and even inversions in the slope over the period, but since mid-2002 it has had a significant upward trend leaving a steep CDS curve by June 2005.

5 The estimation method

All the models in this paper are estimated with the extended Kalman filter. In the standard Kalman filter with Gaussian affine state equation and normally distributed measurement errors, the Kalman filter is a consistent and efficient estimator. However, in the models estimated here both of the interest rate factors and the default risk factor X_t^λ have stochastic volatilities which

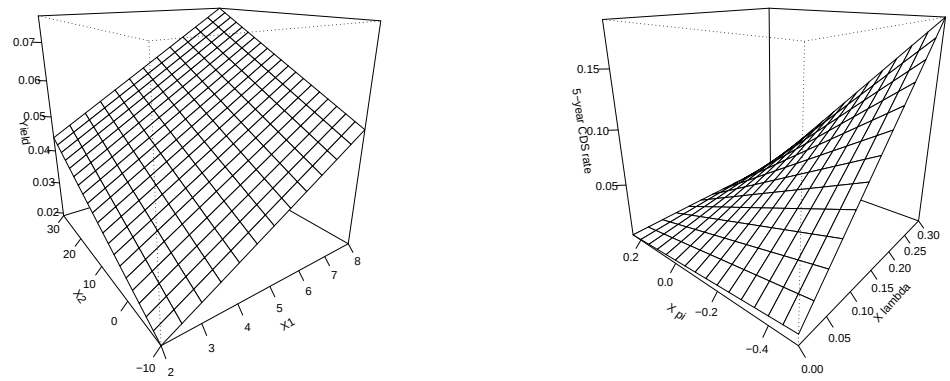

(a) The 5-year Treasury coupon bond yield as a function of (X_t^1, X_t^2).

(b) The 5-year CDS rate as a function of (X_t^λ, X_t^π).

Figure 4: Illustration of the factor sensitivity of the yield of the 5-year 4% Treasury Note to changes in (X_t^1, X_t^2) and of the 5-year CDS rate to changes in (X_t^λ, X_t^π). The parameters of the interest rate factors are $(\kappa_{11}, \kappa_{21}, \kappa_{22}, \theta_1, \beta_{21}, \lambda_2^1, \lambda_{21}^{(2)}, \lambda_{22}^{(2)}, \lambda_1^3, \lambda_{11}^4) = (0.621, -2.99, 0.265, 4.48, 5.06, -0.925, 2.58, 0.119, 1.15, -0.292)$ with factor loadings given by $(\delta_0, \delta_1, \delta_2) = (0.0199, -0.000550, 0.00134)$, while the parameters of the default risk factor are $(\kappa_\lambda, \theta_\lambda, \sigma_\lambda, \gamma_\lambda^1) = (0.178, 0.0433, 0.247, -1.69)$ and the parameters of the recovery risk factor are $(\kappa_\pi, \theta_\pi, \sigma_\pi, \gamma_\pi^0, \gamma_\lambda^1) = (0.608, 0, 0.146, 1.24, 1.18)$. For the CDS rates the values of the interest rate factors are fixed at $(\overline{X}_t^1, \overline{X}_t^2) = (4.260, -0.619)$.

implies that they are non-Gaussian. This has consequences for the properties of the Kalman filter as an estimator. Duan and Simonato (1995) traces the problem back to the discrete-time approximation of the state variables in the prediction step, where the errors in the prediction step are no longer Gaussian and, more importantly, their conditional covariance matrix now depends on the filtered values of the state variables. This implies that the Kalman filter algorithm does not provide a true likelihood function, instead it is referred to as a *quasi maximum likelihood estimation*. The practical implication of this is that the Kalman filter is no longer a consistent estimator. However, based on Monte Carlo simulations, both Duan and Simonato (1995) and Lund (1997) find the bias caused by the non-Gaussian state equations to be small. Finally, the Kalman filter requires the observed data to be affine functions of the state variables, but both the yield-to-maturity of treasury coupon bonds and the formula for the CDS rates are non-linear functions of the state variables. This problem is solved by linearizing the measurement equation through a first-order Taylor approximation around the predicted values of the state variables in each step. Such approximations have not been reported to deteriorate the performance of the Kalman filter. To provide some evidence of the validity of this approximation in the models estimated in this paper, Figure 4(a) illustrates the yield of the 5-year 4% Treasury Note as a function of (X_t^1, X_t^2) while Figure 4(b) illustrates the 5-year CDS rate as a function of (X_t^λ, X_t^π). It is noted that the treasury bond yield is practically linear in the two interest rate factors, while the 5-year CDS rate exhibits some curvature, but locally a linear first-order approximation seem to be satisfactory even in this case.

In terms of estimation strategy, it would technically be feasible to perform a joint estimation

of the interest rate model along with any of the CDS models presented in Section 3. However, given that one purpose of this paper is to compare performance across different CDS models and across different risk premium specifications, the only consistent way to proceed is to first estimate the interest rate model and then estimate the various CDS models using the estimation result of the preferred interest rate model as a fixed input.

In all the estimations the measurement equations take the form

$$y_t = z(X_t; \psi) + \varepsilon_t, \quad E_{t-1}[\varepsilon_t \varepsilon_t'] = H,$$

where ε_t represents white measurement error noise. The measurement errors are assumed to be normally distributed without any serial correlation. However, since the data set in each estimation consists of five time series with fixed maturities it is meaningful to estimate both the variance and the cross-sectional correlation of the errors in each time series. This is implemented by fixing C as a lower triangular matrix and let the measurement error covariance matrix be defined by

$$H = CC'.$$

This structure ensures that H is a positive definite matrix for any choice of parameters in the C-matrix.

For affine state variables, the discrete-time approximation to the true continuous-time formulation is affine in the lagged state variables, and thus the state transition equation is an affine function

$$X_t = \Phi_t^0(\psi) + \Phi_t^1(\psi) X_{t-1} + u_t, \quad E_{t-1}[u_t u_t'] = V_t(X_{t-1}).$$

Given this discrete-time approximation the prediction and update steps in the Kalman filter algorithm do occasionally deliver negative values for the square root-type processes. Whenever this happens the value is truncated at zero and the algorithm is allowed to continue.[22]

In all the estimations the Kalman filter is started at the unconditional mean of the state variables, since the evidence in the empirical literature indicates that both interest rates and corporate bond spreads are stationary processes under the physical measure.

Finally, because the estimations are only quasi maximum likelihood due to the non-Gaussian state variables, the standard deviations of the parameter estimates are calculated with the numerically more robust method based on the outer product of the first derivatives of the likelihood function instead of using the Hessian of the likelihood function

$$\widehat{\Sigma} = \frac{1}{T}\Big[\frac{1}{T}\sum_{i=1}^{T} \frac{\partial \log l_t(\widehat{\psi})}{\partial \psi} \frac{\partial \log l_t(\widehat{\psi})}{\partial \psi}'\Big]^{-1}.$$

Appendix B has more details on the extended Kalman filter estimation of the interest rate model, while Appendix C describes the estimation of the CDS models.

[22] Given that the parameter restriction imposed in the essentially affine risk premium specification only ensures the lower boundary of the square root-type processes to be reflecting, it must be anticipated that the estimated path of such processes will hit the lower boundary during the sample period, even at the optimal parameter set.

Parameter	Driessen (2005)		Essentially affine		Extended affine	
κ_{11}	0.012	(0.007)	0.676	(0.173)	0.621	(0.321)
κ_{21}	-0.182	(0.235)	-3.37	(3.62)	-2.99	(3.65)
κ_{22}	0.619	(0.204)	0.525	(0.189)	0.265	(0.231)
θ_1	2.012	(1.292)	2.84	(2.60)	4.48	(2.28)
β_{21}	0.871	(0.642)	3.45	(1.75)	5.06	(2.34)
δ_0	0.038	(0.002)	0.00572	(0.00631)	0.0199	(0.0199)
δ_1	0.002	(0.001)	-0.000437	(0.000377)	-0.000550	(0.000322)
δ_2	0.004	(0.003)	0.00156	(0.00156)	0.00134	(0.00132)
λ_1^1/λ_1^4	-0.052	(0.041)	-0.334	(0.104)	-0.293	(0.244)
λ_2^1	-3.481	(0.981)	-0.863	(0.184)	-0.925	(0.671)
$\lambda_{21}^{(2)}$	-0.326	(0.682)	2.35	(2.53)	2.58	(3.72)
$\lambda_{22}^{(2)}$	1.444	(1.092)	-0.160	(0.200)	0.119	(0.291)
λ_1^3	n.a.		n.a.		1.15	(1.56)
llk	n.a.		5605.03		5612.51	

Table 3: The estimated parameter values for the two-factor affine model of the default-free instantaneous interest rate process. The second column represents the estimates obtained by Driessen(2005) on weekly data of mid-quotes for six-month Treasury bill rates and Treasury bond yields with maturities 2, 3, 5, 7, 10, and 30 years over the period February 22, 1991 until February 18, 2000. The third and the fourth columns represent the results obtained in this paper for the essentially and extended affine risk premium specifications, respectively, when estimated on weekly data of mid-quotes for six-month Treasury Bills and Treasury Notes with maturities 2, 3, 5, and 10 years over the period March 2, 2001 until June 24, 2005. Note that the 3-year Treasury Notes are only observed from August 15, 2003. The numbers in parentheses are the standard deviations of the parameter estimates.

6 Results

The first step is to perform the estimation of the risk-free rate model based on the treasury yield data. The results of this estimation is described in Section 6.1 and the preferred risk premium specification is chosen. Building on the result of the preferred risk-free rate model, the next step is to estimate the CDS models. In Section 6.2 the results of the estimations of the constant recovery CDS model are discussed, while Section 6.3 contains the results of the estimations of the four-factor CDS model with stochastic default and recovery risk.

6.1 The risk-free rate estimation

Since the choice between the essentially affine and the extended affine risk premium specification is an issue of interest in the subsequent estimations of the CDS models, the interest rate model is also estimated with both specifications. The focus in this result subsection will be on finding the best performing specification for the treasury data to be used as an input in the estimations of the CDS models.

The estimated parameters from the two estimations of the interest rate model are presented in Table 3 along with the result obtained by Driessen (2005) in his estimation of the same $A_1(2)$-model with the essentially affine risk premium specification.

The first thing to note is that the estimates for the essentially affine model obtained here are rather different from the parameter estimates found by Driessen (2005) despite the fact that it is the same model and same quasi-maximum likelihood Kalman filter estimation that has been applied. This goes to show that parameter stability over two so relatively short and non-overlapping time periods is too much to expect from a fairly complex two-factor model like this one.

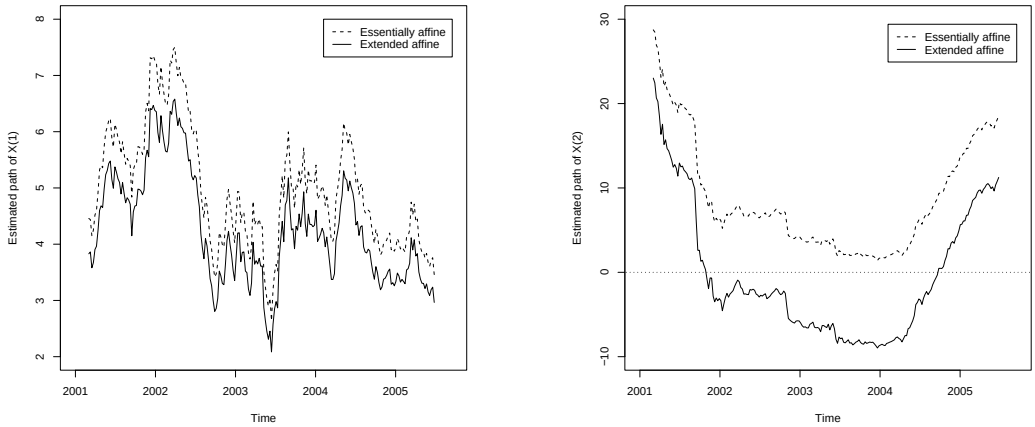

(a) The estimated path of X_t^1.

(b) The estimated path of X_t^2.

Figure 5: The estimated paths of the two risk-free factors, X_t^1 and X_t^2, from the two extended Kalman filter estimations based on the observed US treasury bill yields over the period March 2, 2001 through June 24, 2005.

As for the comparison of the essentially affine and the extended affine risk premium specifications, it is noted that the estimated parameters in both specifications satisfy the Feller condition. Thus, the extended affine risk premium specification is not penalized for having to satisfy the Feller condition. Looking at the estimated parameters across the two specifications, the primary differences relate to the mean of the X_t^1-factor θ_1 and to the volatility β_{21} and the mean-reversion rate κ_{22} of the unrestricted X_t^2-factor. Comparing the two θ_1 estimates with the estimated paths of X_t^1 illustrated in Figure 5(a), it is noted that the mean estimate from the extended model is close to the average model-implied value of X_t^1, whereas the mean estimate in the essentially affine model is below the model-implied path throughout the estimation period. Thus, there appears to be a better internal consistency between parameter estimate and model-implied path for the X_t^1-factor in the extended model. As for the estimated volatility parameter β_{21} it is 40% larger in the extended model. However, the impact of this larger value of β_{21} in the extended affine model is offset by the lower model-implied path for X_t^1, and as a result the average model-implied instantaneous volatility of X_t^2 is not significantly different (mean of 4.73 versus 4.25).[23] Finally, the estimated mean-reversion rate of X_t^2 in the extended model is only half of the estimate in the essentially affine model. The estimate of $\widehat{\kappa}_{22} = 0.265$ in the extended model indicates that the X_t^2-factor should have larger and longer lasting deviations from its mean than the X_t^1-process, whereas the estimate in the essentially affine model indicates that the two factors should have more similar mean-reversion rates. However, from the estimated paths of X_t^2 illustrated in Figure 5(b) it follows that the model-implied path of X_t^2 in both models is much less mean-reverting than the estimated paths for the X_t^1-process. Thus, again, the estimate from the extended model is more in line with the implied paths from the estimation than what is obtained with the essentially affine specification.

Another way of comparing the two risk premium specifications is to look at their performance in terms of pricing errors. To that end the mean error, the mean absolute error, and the root mean

[23] This is calculated as $\frac{1}{T}\sum_{t=1}^{T}\sqrt{1+\widehat{\beta}_{21}\widehat{X}_t^2}$.

Maturity	Essentially affine			Extended affine		
	Mean error	Mean abs err	RMSE	Mean error	Mean abs err	RMSE
26 Weeks	-1.23	2.13	2.46	-0.02	0.16	0.24
2 years	-12.63	21.39	24.67	-11.56	19.57	22.58
3 years	-1.68	17.81	21.41	-1.62	16.20	19.41
5 years	-5.32	8.89	10.48	-4.52	7.64	9.22
10 years	-0.67	7.97	10.08	-0.11	8.63	11.00

Table 4: The mean and mean absolute pricing error in basis points for the five different maturities in the two extended Kalman filter estimations based on the weekly Treasury data over the period from March 2, 2001 to June 24, 2005.

square error are provided in Table 4 for each of the five maturities in the data set. Judged from the mean pricing errors both models seem to match the observed yields fairly well on average with the exception of the 2-year Treasury Notes that have a mean pricing error close to 12 basis points in both models. However, once the mean absolute pricing error is taken into consideration, the improvement in performance when switching to the extended affine specification becomes evident. Except for the 10-year yields where the performance is slightly in favor of the essentially affine model, the extended affine specification has mean absolute errors that are 1-2 bps lower than the results obtained for the model with the essentially affine risk premium specification. The same impression is obtained when looking at the estimated measurement error covariance matrices from the Kalman filter[24]

$$\widehat{H}^{ess} = \begin{pmatrix} 0.000246 & 0.9997 & 0.9925 & 0.9641 & -0.0257 \\ & 0.002462 & 0.9951 & 0.9654 & -0.0178 \\ & & 0.002400 & 0.9708 & 0.0359 \\ & & & 0.001047 & 0.2398 \\ & & & & 0.001005 \end{pmatrix}$$

and

$$\widehat{H}^{ext} = \begin{pmatrix} 0.000039 & 0.3315 & 0.2731 & 0.3849 & 0.0577 \\ & 0.002255 & 0.9936 & 0.9275 & -0.0917 \\ & & 0.002210 & 0.9409 & -0.0263 \\ & & & 0.0009233 & 0.2833 \\ & & & & 0.001101 \end{pmatrix}.$$

First, looking at the standard deviations of the individual measurement errors in the diagonal, the extended affine estimation has lower estimated standard deviations for every maturity except for the 10-year yield where it is slightly higher. And turning to the cross-sectional correlations, the estimation of the essentially affine model seems to split the bonds into two groups, the 10-year yield in one group and all the other yields in the other group with internal pricing error correlation coefficients in the range from 0.96 to 0.999, whereas the extended affine estimation split the yields into 3 groups, the 26-week yield is one group, the 10-year yield is another, and the 2-year, 3-year, and 5-year yields form a third group with internal correlations in the 0.92-0.99 range.

To round off the comparison of the two risk premium specifications in the risk-free rate model, a final test can be performed. As mentioned in the beginning, for both risk premium specifications the Feller condition is satisfied under both the P- and the Q-measure. Thus, the essentially affine

[24] These two matrices are organized in the following way. The diagonal contains the estimated standard deviations of the measurement errors for the five maturities considered here, 26-week, 2-year, 3-year, 5-year, and 10-year, respectively. The off-diagonal elements are the estimated correlation coefficients between those measurement errors.

Correlation	$\widehat{X}_t^1$	$\widehat{X}_t^2$	$y_{26-Week}$	y_{2-yr}	y_{5-yr}	y_{10-yr}	$slope_{10-yr/2-yr}$
$\widehat{X}_t^1$	1	-0.0031	0.0938	0.3826	0.6790	0.8428	0.1616
$\widehat{X}_t^2$		1	0.9953	0.9001	0.7129	0.4995	-0.9254
$y_{26-Week}$			1	0.9337	0.7757	0.5788	-0.9064
y_{2-yr}				1	0.9319	0.7771	-0.8325
y_{5-yr}					1	0.9440	-0.5822
y_{10-yr}						1	-0.2982
$slope_{10-yr/2-yr}$							1

Table 5: The correlation coefficients between the levels of the following time-series: the estimated paths of X_t^1 and X_t^2 from the extended affine model, the observed yields with maturities 26-Week, 2-year, 5-year, and 10-year, and the slope of the yield curve measured as the difference between the 10-year and the 2-year yields. All variables are observed weekly over the period from March 2, 2001 to June 24, 2005.

Correlation	$\Delta\widehat{X}_t^1$	$\Delta\widehat{X}_t^2$	$\Delta y_{26-Week}$	Δy_{2-yr}	Δy_{5-yr}	Δy_{10-yr}	$\Delta slope_{10-yr/2-yr}$
$\Delta\widehat{X}_t^1$	1	0.0210	0.3083	0.7109	0.8968	0.8665	0.2138
$\Delta\widehat{X}_t^2$		1	0.9568	0.4987	0.3354	0.2340	-0.4834
$\Delta y_{26-Week}$			1	0.6868	0.5820	0.4769	-0.4036
Δy_{2-yr}				1	0.9206	0.8335	-0.3521
Δy_{5-yr}					1	0.9452	-0.0229
Δy_{10-yr}						1	0.2236
$\Delta slope_{10-yr/2-yr}$							1

Table 6: The correlation coefficients between the first-differences of the following time-series: the estimated paths of X_t^1 and X_t^2 from the extended affine model, the observed yields with maturities 26-Week, 2-year, 5-year, and 10-year, and the slope of the yield curve measured as the difference between the 10-year and the 2-year yields. All variables are observed weekly over the period from March 2, 2001 to June 24, 2005.

model is nested within the extended affine model and the validity of imposing the restriction $\lambda_1^3 = 0$ in the essentially affine model can be tested with a likelihood ratio test

$$LR = 2[\text{loglikelihood}^{ext} - \text{loglikelihood}^{ess}] = 2[5612.51 - 5605.03] = 14.96.$$

This ratio is asymptotically χ^2-distributed with one degree of freedom. The probability of observing 14.96 or more in the χ^2-distribution with one degree of freedom is 0.00011, so the H_0-hypothesis that the two models perform equally well is rejected at normal significance levels.

Combining all the results of the above analysis, the extended affine model is judged to be the better model for the risk-free rate for this specific data set. Therefore, it is the result of this specification that will be used as the interest rate input in the subsequent estimations of the CDS models. However, given the similarity in the dynamic structure of the estimated paths across the two estimations,[25] a switch from the extended affine to the essentially affine model will not change in any material way the results of the subsequent estimations of the CDS models.

Based on the result of the preferred extended affine model, the two interest rate factors can be interpreted. When the estimated paths of the two latent factors are compared to the level and the slope of the yields in Figure 1, it is noted that the behavior of X_t^1 seems to be related to the movement of the 5- and 10-year treasury yields, whereas the estimated path of X_t^2 seems to match the slope of the yield curve. To verify this conjecture the correlation coefficients between

[25] The correlation in levels across the two specifications are 0.9991 and 0.9998 for the X_t^1- and X_t^2-factor, respectively, while the correlation in first-differences is 0.9991 and 0.9951, respectively.

Parameter	Essentially affine		Essentially affine Feller cond. imposed		Extended affine	
κ_λ	0.586	(0.365)	1.21		1.11	(1.03)
θ_λ	0.0561	(0.0335)	0.0817	(0.0514)	0.0889	(0.0859)
σ_λ	0.431	(0.0127)	0.445	(0.0192)	0.444	(0.0245)
γ_λ^0	n.a.		n.a.		0.0000539	(0.205)
γ_λ^1	-2.78	(0.813)	-3.04	(1.66)	-2.82	(2.33)
λ_0	0.0502	(0.00579)	-0.00648	(0.00371)	-0.00658	(0.0116)
λ_r^1	0.00671	(0.00321)	0.00963	(0.00297)	0.00988	(0.00820)
λ_r^2	-0.00322	(0.000317)	-0.00569	(0.000458)	-0.00570	(0.000850)
π_0	0.880	(0.00560)	0.855	(0.0131)	0.855	(0.0160)
llk	5392.14		5351.857		5351.863	

Table 7: The estimated parameter values for the three-factor standard CDS model with constant recovery rate. The second column represents the result for the model with the essentially affine risk premium specification, the third column presents the results of the essentially affine specification *with* the Feller condition imposed, while the last column contains the result for the model with the extended affine risk premium specification for the X_t^λ-process. All three estimations are based on the Ford Motor Co. CDS data covering the period from March 2, 2001 until June 24, 2005. The numbers in parentheses are the standard deviations of the parameter estimates.

the estimated paths and the observed yields along with the slope of the yield curve are calculated and presented in Table 5. Similarly, the correlation coefficients between the first-differences of those same time-series are provided in Table 6.

The first thing to notice is that the estimated paths of X_t^1 and X_t^2 are practically uncorrelated in both levels and changes which means that they each catch independent essential characteristics of the yield term structure. Second, the correlation coefficients in Tables 5 and 6 reveal that the estimated path of X_t^1 is highly correlated with the long-term treasury rates. Thus, X_t^1 can be interpreted as a level factor. Third, it is noted that X_t^2 is extremely highly correlated with the dynamics of the 26-Week yield and therefore it is also strongly negatively correlated with the observed slope of the yield curve, in particular in the levels. However, in first-differences X_t^2 is only highly correlated with the 26-Week yield, thus X_t^2 is most appropriately interpreted as a short-term factor, but it has to be said that the time period considered in this study is too short to allow for a distinction between the two.

This concludes the analysis of the results from the estimation of the risk-free interest rate factors.

6.2 Estimation results for the CDS model with constant recovery rate

Building on the result of the extended affine $A_1(2)$-model for the risk-free rate analyzed in the last section, this section will describe the results of estimating the CDS model with constant recovery rate introduced in Section 3.2. Similar to the interest rate model, this CDS model will be estimated with both the essentially affine and the extended affine risk premium specifications. Table 7 contains the results of the estimations, while the estimated paths of the default risk factor X_t^λ and the default intensity process λ_t^Q are illustrated in Figures 6(a) and 6(b), respectively.

The first important observation to take away from Table 7 is that the Feller condition imposed in the extended affine specification is a binding constraint in this model. In the essentially affine specification, this materializes as an outright violation of the Feller condition under the P-measure

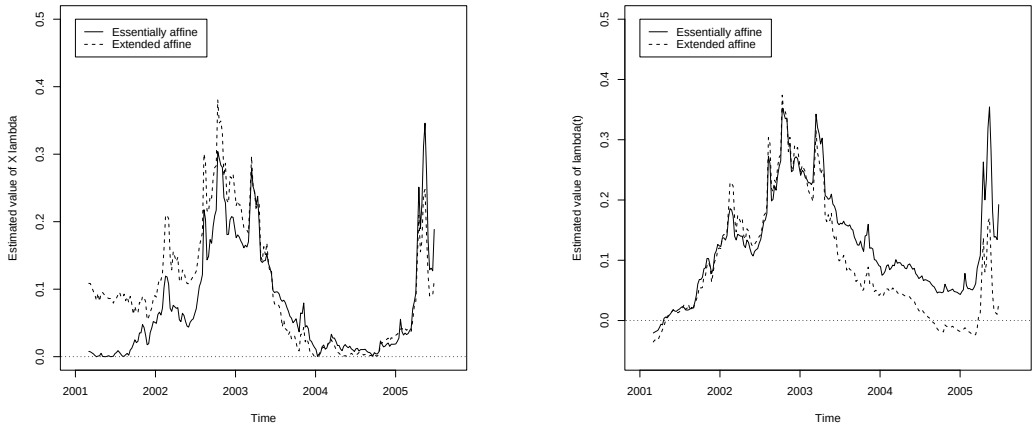

(a) The estimated path of X_t^λ.

(b) The estimated path of λ_t^Q.

Figure 6: Illustration of the estimated paths of X_t^λ and λ_t^Q in the standard CDS model with essentially and extended affine risk premium specifications. Both estimations use the Ford Motor Co. CDS data covering the period from March 2, 2001 until June 24, 2005.

at the optimal parameter values

$$\widehat{\kappa}_\lambda^{ess}\widehat{\theta}_\lambda^{ess} = 0.033 < \frac{1}{2}(\widehat{\sigma}_\lambda^{ess})^2 = 0.093.$$

To investigate the impact of imposing the Feller condition, Table 7 also reports the result of estimating the essentially affine model with the mean-reversion rate defined in such a way that the Feller condition is exactly binding

$$\kappa_\lambda = \frac{1}{2}\frac{\sigma_\lambda^2}{\theta_\lambda}.$$

The result of imposing the Feller condition is that there is a significant loss in likelihood value. In addition, some of the parameters have changed in significant ways. Most notably, the mean-reversion rate κ_λ has more than doubled while the mean parameter θ_λ is increased by more than 40%. On the other hand, the volatility σ_λ is practically the same across all three estimations.

In the extended affine model there is the additional risk premium parameter γ_λ^0 to make up for the loss of having to satisfy the Feller condition. In technical terms the advantage of including γ_λ^0 in the model is that the risk premium

$$\Lambda^{ext} = \frac{\gamma_\lambda^0}{\sqrt{X_t^\lambda}} + \gamma_\lambda^1\sqrt{X_t^\lambda}$$

can change sign.[26] Figure 7 illustrates the estimated paths of the risk premia across the two risk premium specifications. It is noted that the risk premium does change sign across time in the extended affine specification. However, neither the size of the estimate of γ_λ^0 nor the change in the likelihood value indicate that there is any advantage of this added flexibility. Thus, the major differences in the results across the two specifications can be referred to the constraint imposed on the extended affine specification by having to satisfy the Feller condition under both probability measures.

[26] Of course, this requires γ_λ^0 and γ_λ^1 to have opposite signs.

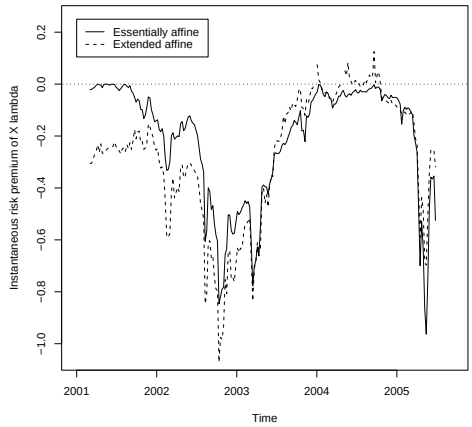

Figure 7: Illustration of the estimated path of the premium in the essentially affine and the extended affine risk premium specifications of the standard CDS model. The two are calculated as $\Lambda_t^{ess} = \widehat{\gamma}_\lambda^1 \sqrt{\widehat{X}_t^\lambda}$ and $\Lambda_t^{ext} = \frac{\widehat{\gamma}_\lambda^0}{\sqrt{\widehat{X}_t^\lambda}} + \widehat{\gamma}_\lambda^1 \sqrt{\widehat{X}_t^\lambda}$. Both estimations use the Ford Motor Co. CDS data covering the period from March 2, 2001 until June 24, 2005.

The higher estimated mean parameter θ_λ in the extended affine specification is also reflected in the estimated path of the default intensity risk factor X_t^λ as illustrated in Figure 6(a), where the estimated path of X_t^λ is notable higher in the extended affine model for the first two years of the sample. However, this does not imply a higher level for the estimated path of the default intensity λ_t^Q because the positive effect from the larger estimated values of X_t^λ is more than offset by a 5.6% difference in the estimate of the constant term λ_0 and by an increase in the absolute size of the factor loadings on the interest rate factors. The combined result is that the estimated path of the default intensity in the extended affine model is actually significantly lower during the second half of the sample period.

Another noticeable feature about the estimated paths of the default intensity is that for both specifications the estimated path is negative for the first few months of the sample period and again in late 2004 and early 2005 for the extended affine model. This is at odds with the notion of an intensity. However, this is not as serious a problem as it may appear because even though X_t^λ is stationary under the P-measure as indicated by the positive estimates of κ_λ, the large estimates of the risk premium parameter γ_λ^1 cause the X_t^λ-process to be non-stationary under the Q-measure. As a result the conditional expected value of λ_t^Q under the Q-measure is actually positive except for very short time horizons even in the early part of the sample. This is illustrated in Figure 8 where the estimated path of λ_t^Q along with the 6-months and 12-months ahead Q-conditional expectations of λ_t^Q are depicted for both estimations. The two figures show that for time horizons longer than 6 months the conditional expected default intensity under the Q-measure is positive throughout the estimation period. Pan and Singleton (2005) also find Q-non-stationary dynamics in their estimations based on CDS data for Russia, Mexico, and Turkey, and their results show that this finding is not confined to the affine CIR-model as they observe the same phenomenon for their log-normal and three-halves models.

The fact that the default intensity in the extended model is lower than the default intensity

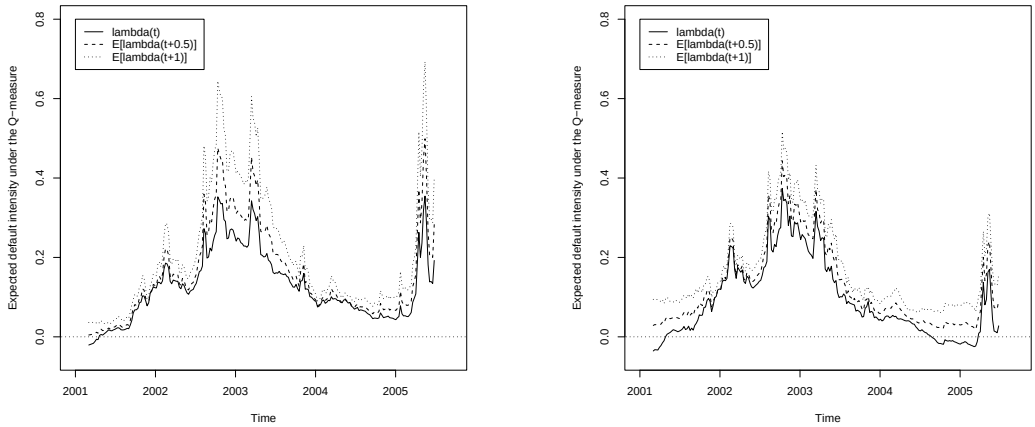

(a) Conditional expected default intensities in the essentially affine model.

(b) Conditional expected default intensities in the extended affine model.

Figure 8: Illustration of the estimated path of the default intensity $\widehat{\lambda}_t^Q$ along with the 6-months and 12-months ahead conditional expectation of λ_t^Q under the Q-measure based on the estimation of the constant-recovery CDS model with essentially and extended affine risk premium specifications. Both estimations use the Ford Motor Co. CDS data covering the period from March 2, 2001 until June 24, 2005.

in the essentially affine model, for an extended period even significantly so, cause the estimated constant recovery rate in the extended affine model to be lower than in the essentially affine model, 85.5% versus 88.0%. These levels of the market-implied recovery rate may seem fairly high relative to the recovery rates reported in the empirical literature. However, Pan and Singleton (2005) find similarly high market-implied recovery rates[27] of 83.9% and 56.1% for sovereign CDS data on Mexico and Russia, respectively, even though Russia actually defaulted on its foreign debt as recently as August 1998 and May 1999 with recovery rates ranging from 20% to 35% of their values just weeks before the default event.[28] Zhang (2003) estimates a constant market-implied recovery rate based on Argentinean CDS data from February 1999 to November 2000 (i.e. until 6 weeks before the default by Argentina on its foreign debt) and obtains an estimate of a mere 27.5%, not far from the actually observed recovery rate of 28% reported by Moody's (2003).[29] Thus, Zhang (2003)'s study provides evidence that the market may get the implied recovery rates right. A more detailed discussion of the model-implied recovery rate is deferred to the subsequent section on the results of the stochastic recovery CDS model.

Finally, based on the results in Table 7 the impact of the interest factors can be determined. With the recovery rate fixed the CDS model is equivalent to a model of the credit spread dynamics. In all three estimations, the factor loading on the first interest rate factor X_t^1 is positive. Since X_t^1 represents the level of the treasury yields, this means that there is a positive correlation between the level of treasury yields and the credit spread of Ford Motor Co. over this period. As for the factor loading on the second interest rate factor, λ_r^2, it is negative in all three estimations. From

[27] See their Table 2 under the unconstrained Affine Model.

[28] See Duffie, Pedersen, and Singleton (2003) for more details on the Russian defaults in this period and a related credit risk analysis on some dollar-denominated Russian bonds outstanding at that time.

[29] Moody's (2003) observed the prices of 12 outstanding bond series issued by the Argentinean government. The average price across the 12 bonds one month after the default event was 28% of face value.

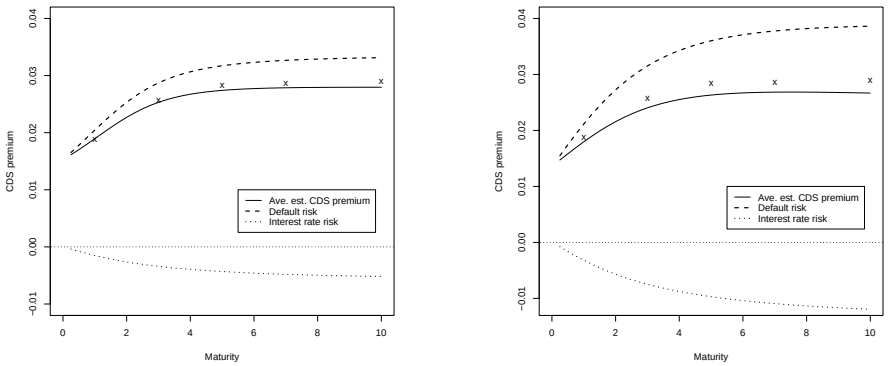

(a) The essentially affine CDS model.

(b) The extended affine CDS model.

Figure 9: Decomposition of the average model-implied CDS term structure into its two components: Interest rate and default risk. The model estimated is the three-factor standard CDS model with fixed recovery rate combined with the essentially and the extended affine risk premium specification, respectively. The data used is the CDS data for Ford Motor Co. over the period March 2, 2001 through June 24, 2005.

the estimation of the interest rate model it is recalled that X_t^2 is a close proxy for the negative of the slope of the treasury yield curve. Thus, a negative loading on this factor implies that the credit spreads of Ford Motor are larger, ceteris paribus, when the treasury curve is steep. Duffee (1999) observes a similar pattern for the factor loading on the slope factor. Across his 161 estimations, 75% of the firms have a negative loading on the slope factor. For the factor loading on the level factor his result is more mixed with little more than half of the companies having a negative loading on the level factor. Driessen (2005) observes negative factor loadings on both interest rate factors across the three investment grade rating categories in his study, with the only exception being a very small statistically insignificant positive slope factor loading for the A-category. Thus, while the evidence for the connection between credit spreads and the interest rate level factor is mixed, the negative factor loading on the interest rate slope factor seems to be a more robust finding.

To determine the relative impact of the default risk factor and the interest rate factors on the Ford Motor Co. CDS rates, the following decomposition is performed.

(i). For each estimated vector of the three factors $(\widehat{X}_t^1, \widehat{X}_t^2, \widehat{X}_t^\lambda)$ the entire term structure of model-implied CDS rates is calculated. For each maturity the average across time is calculated. This gives the solid lines in Figure 9.

(ii). In order to isolate the impact of the default risk factor X_t^λ the above exercise is repeated, but this time the effect of the risk-free factors on the default intensity is eliminated by fixing their factor loadings to zero, i.e. $\lambda_1^r = 0$ and $\lambda_2^r = 0$. This gives the dashed lines in Figure 9.

(iii). As a final step, the net contribution from the interest rate factors are calculated as the difference between the results from step 1 and step 2. This gives the dotted lines in Figure 9.

The decomposition in Figure 9 reveals that the interest rate factors have a negative impact on CDS rates, on average, increasing with the time to maturity, the figure also shows that the overall

Maturity	Essentially affine			Extended affine		
	Mean error	Mean abs error	RMSE	Mean error	Mean abs error	RMSE
1 year	-4.85	26.80	33.06	5.64	24.78	33.73
3 year	1.41	16.38	20.78	15.37	55.38	80.65
5 year	7.04	37.74	44.56	18.94	66.78	94.54
7 year	6.68	44.78	53.21	17.33	69.77	98.00
10 year	9.32	52.05	63.35	22.74	75.45	104.04

Table 8: The mean pricing error, the mean absolute pricing error, and the RMSE, all measured in basis points, for the 5 different maturities in the two extended Kalman filter estimations based on the Ford Motor Co. CDS data over the period from March 2, 2001 to June 24, 2005.

structure of the decomposition gives somewhat the same picture. However, there are noticeable differences across the two specifications in the relative contribution from the interest rate risk factors and the default risk factor. Additionally, it is noted that the average fit of the extended model is worse than the average fit of the essentially affine model.

Based on the results so far it is hard to detect any significant advantage of one risk premium specification over the other beyond the large difference in likelihood value. Therefore the next step is to analyze the performance of the two specifications in terms of fitting the observed CDS rates. In Table 8, the mean error, the mean absolute error, and the root mean square error are provided for both models across the five different maturities in the data set.

First, it is noted that the mean pricing errors are generally smaller for the essentially affine specification with the exception of the 1-year contract where the extended affine model performs almost as well. Second, once measures of dispersion are included in the analysis, much the same picture is observed. The extended affine specification can only match the essentially affine specification for the one-year contract while it has significantly larger standard deviations of its pricing errors for the four other contracts. Third, similar results are observed in the estimated measurement error covariance matrices from the Kalman filter which are given by[30]

$$\widehat{H}^{ess} = \begin{pmatrix} 0.003218 & -0.8915 & -0.8949 & -0.8458 & -0.8260 \\ & 0.001993 & 0.9191 & 0.9168 & 0.9168 \\ & & 0.004300 & 0.9930 & 0.9820 \\ & & & 0.005145 & 0.9860 \\ & & & & 0.006105 \end{pmatrix}$$

and

$$\widehat{H}^{ext} = \begin{pmatrix} 0.003317 & 0.8482 & 0.8775 & 0.9008 & 0.9228 \\ & 0.007930 & 0.9919 & 0.9891 & 0.9802 \\ & & 0.009334 & 0.9982 & 0.9912 \\ & & & 0.009654 & 0.9932 \\ & & & & 0.01020 \end{pmatrix}.$$

As expected, the estimated pricing error standard deviations in the diagonal are similar to the sample equivalent estimates provided in the RMSE columns of Table 8. The off-diagonal corre-

[30] These two matrices are organized in the following way. The diagonal contains the estimated standard deviations of the measurement errors for the five maturities considered here, 1-year, 3-year, 5-year, 7-year, and 10-year, respectively. The off-diagonal elements are the estimated correlation coefficients between those measurement errors over the sample period.

lation coefficients reveal that in the essentially affine model overpricing of the 1-year contract is highly correlated with simultaneous underpricing of the 3-, 5-, 7-, and 10-year contracts, while in the extended model all the errors are highly positively correlated. This indicates a clear difference in the way the two models fit the data. The essentially affine model try, at any given point in time, to balance mispricing of the 1-year contract with opposite mispricing of the four other contracts. In the extended model the balance of mispricing is across time in the following sense. The time series of the pricing errors across the five contracts indicate that until the beginning of 2003 the extended model systematically underprice all five contracts, while it with few exceptions overprice the five contracts in the remaining part of the sample period.

Summing up the results, as for the choice between the two risk premium specifications, the estimated parameters reveal that the estimation of the extended affine risk premium specification does not take advantage of the added flexibility provided by the extra risk premium parameter γ_λ^0, but it does feel the drawback of having to satisfy the Feller condition relative to the essentially affine specification. This constraint forces the estimated path and parameters of the X_t^λ-factor to deviate from the unrestricted estimates obtained in the essentially affine specification. This in turn causes the estimated default intensity as well as the estimated recovery rate to change in important ways. This ultimately deteriorates significantly the performance of the extended affine model in terms of fitting the observed CDS rates. Based on this evidence, the essentially affine specification is found to be the preferred risk premium specification for this CDS model given this specific data set.

A final word of caution can also be drawn from the analysis above. The widely used practice of assuming one of the contracts (typically the 5-year contract) observed without measurement error and use this to back-out the value of the state variable does not seem to be justified in a setting like the one analyzed here. The result of the Kalman filter estimation in both specifications clearly indicate that not a single of the five CDS contracts is measured with close to zero error. This is one of the advantages of the Kalman filter, it lets the data speak on how well each contract is priced by the model.

6.3 Result of the estimation of the stochastic recovery CDS model

This section will go through and analyze the results of the four-factor CDS model with stochastic default and recovery risk factors. As in the last section, the model has been estimated with both the essentially affine and the extended affine risk premium specifications in order to determine which of the two specifications is the more appropriate in a credit risk setting. The first step will be to describe the parameters and relate them to the estimates obtained in the three-factor standard CDS model with constant recovery parameter analyzed in the last section. Second, the contribution of interest rate, default, and recovery risk on average CDS rates is studied. Third, the pricing performance of the two risk premium specifications is discussed in detail. Finally, based on the result from the preferred model, the ability of the model to decompose CDS rates into a default and a recovery risk component is analyzed in more detail.

Table 9 contains the estimated parameters for the essentially affine and the extended affine specifications along with the result of the essentially affine model with the Feller condition imposed on the default risk factor through the restriction $\kappa_\lambda = \frac{1}{2}\frac{\sigma_\lambda^2}{\theta_\lambda}$.

For the parameters of the default risk factor it is first noticed that across all three estimations

Parameter	Essentially affine		Essentially affine Feller cond. imposed		Extended affine	
κ_λ	0.178	(0.128)	0.197		0.200	(0.443)
θ_λ	0.0432	(0.0325)	0.0270	(0.0146)	0.0281	(0.0365)
σ_λ	0.247	(0.0132)	0.103	(0.0312)	0.104	(0.0519)
κ_π	0.608	(0.470)	0.511	(0.468)	0.508	(0.469)
θ_π	0		0		0	
σ_π	0.146	(0.0231)	0.295	(0.155)	0.296	(0.176)
λ_0	0.0116	(0.00594)	-0.0349	(0.0176)	-0.0355	(0.0345)
λ_r^1	0.00118	(0.00367)	0.00656	(0.00462)	0.00650	(0.00480)
λ_r^2	0.000879	(0.000400)	0.000374	(0.000290)	0.000368	(0.000325)
π_0	0.867	(0.117)	0.813	(0.161)	0.82	(0.180)
π_r^1	0.0208	(0.0104)	0.0832	(0.0460)	0.0773	(0.0484)
π_r^2	0.00587	(0.000782)	0.00134	(0.00865)	0.0132	(0.00946)
γ_λ^0	n.a.		n.a.		0.00176	(0.135)
γ_λ^1	-1.69	(0.599)	-1.57	(0.621)	-1.56	(4.30)
γ_π^0	1.24	(0.625)	2.51	(0.491)	2.46	(0.496)
γ_π^1	1.18	(3.43)	0.885	(1.78)	0.866	(1.81)
llk	5658.70		5643.75		5644.19	

Table 9: The estimated parameter values for the four-factor CDS model with stochastic default and recovery risk. The second column represent the estimates obtained with the essentially affine risk premium specification, the third column represents the result of imposing the Feller condition on this model. Finally, the fourth column represents the estimates obtained with the extended affine risk premium specification. All three estimations use weekly CDS data for Ford Motor Co. over the the period March 2, 2001 until June 24, 2005. The numbers in parentheses are the standard deviations of the parameter estimates.

the parameter estimates are lower than those obtained in the standard CDS model. The estimated mean θ_λ is lower, so is the estimated volatility σ_λ, and the mean-reversion rate κ_λ has been reduced significantly. As for the Feller condition it is still violated in the essentially affine model, while the estimated parameters in the extended affine model are very close to the boundary imposed by the Feller condition. So once more the Feller condition is a binding restriction for the extended affine model which causes its likelihood value to be lower than the result obtained for the essentially affine model.

In Figure 10 the estimated paths of the default intensity risk factor X_t^λ and the resulting estimated paths of λ_t^Q are illustrated. For the essentially affine model the violation of the Feller condition means that zero is merely a reflecting boundary. This is reflected in the estimated path of X_t^λ as it actually hits the lower boundary several times during the sample period as observed in Figure 10(a). In Figure 10(b) the estimated paths of the default intensity λ_t^Q are depicted. Here it is noted that on average the model-implied default intensity is significantly lower than the estimated values in the constant-recovery CDS model, in particular for the 2002-2003 period. As will be evident below, this result is caused by depressed market-implied recovery rate levels during this period. Comparing the two specifications, the higher estimated values of X_t^λ in the extended affine model during most of the sample period is off set by the large negative estimate for the constant term λ_0 in the default intensity, a pattern similar to the one found in the constant-recovery CDS model. Furthermore, the negative estimated values of the default intensity observed in the standard CDS model have been eliminated so that the estimated path for the default intensity process is now strictly positive throughout the estimation for both risk premium

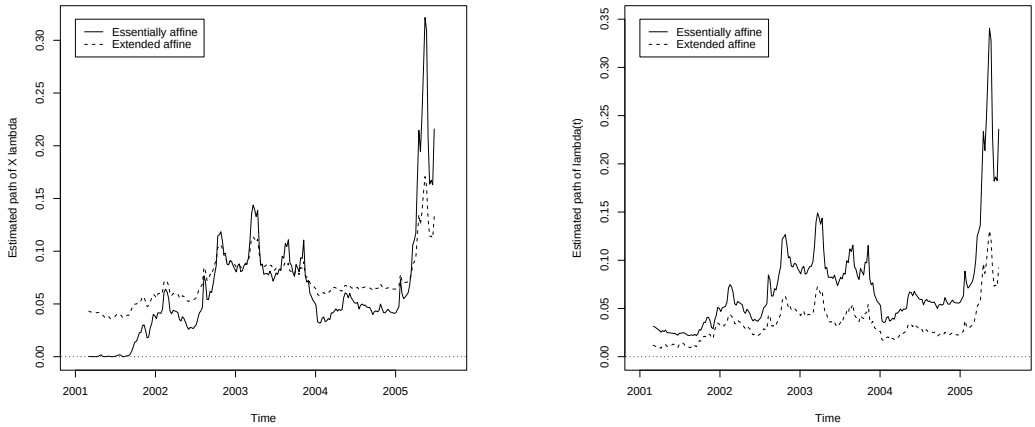

(a) The estimated path of X_t^λ.

(b) The estimated path of λ_t^Q.

Figure 10: Illustration of the estimated paths of X_t^λ and λ_t^Q in the stochastic recovery CDS model with essentially and extended affine risk premium specifications. Both estimations use the Ford Motor Co. CDS data covering the period from March 2, 2001 until June 25, 2005.

specifications.

As for the interest rate factor loadings in the default intensity process, it is observed that the factor loading on the first interest rate factor X_t^1 is positive in both specifications. However, for both specifications the size of the estimated factor loading is lower as compared with the estimates from the fixed-recovery CDS model. For the second interest rate factor, the estimated factor loading λ_r^2 is now positive across all three estimations, but numerically much smaller than the estimates in the constant-recovery CDS model.

As for the risk premium parameters related to the default risk factor X_t^λ there is a significant decrease in the absolute size of the estimates of γ_λ^1 in both specifications relative to the estimates found in the fixed-recovery CDS model. This is only natural as there are now two factors to explain the idiosyncratic risk while there is only one factor in the standard CDS model. In the extended affine specification, the estimate of γ_λ^0 is still small in absolute terms indicating that the estimation is not really benefitting from this additional parameter.

Turning to the recovery risk factor, Figure 11 illustrates the estimated path of the recovery risk factor X_t^π along with the estimated path for the market-implied instantaneous recovery rate π_t^Q. This figure shows the main difference between the constant-recovery three-factor CDS model from the last section and the stochastic recovery CDS model analyzed here. In the standard model, the recovery rate was treated as a fixed constant and the estimated values were in the 85% to 88% range in all three estimations. In the stochastic recovery model the constant term of the recovery rate process is still estimated to be in the 80%-90% range, however, the large factor loadings on the interest rate factors combined with the large variation in the idiosyncratic recovery risk factor cause the recovery rate process to be rather volatile.

One important thing to notice is that the factor loadings on the interest rate factors in the recovery rate process are positive in all three estimations. In order to explain this start with the result from constant-recovery CDS model. There it was observed that across all three estimations the results indicated that the Ford credit spreads in the period considered here are positively

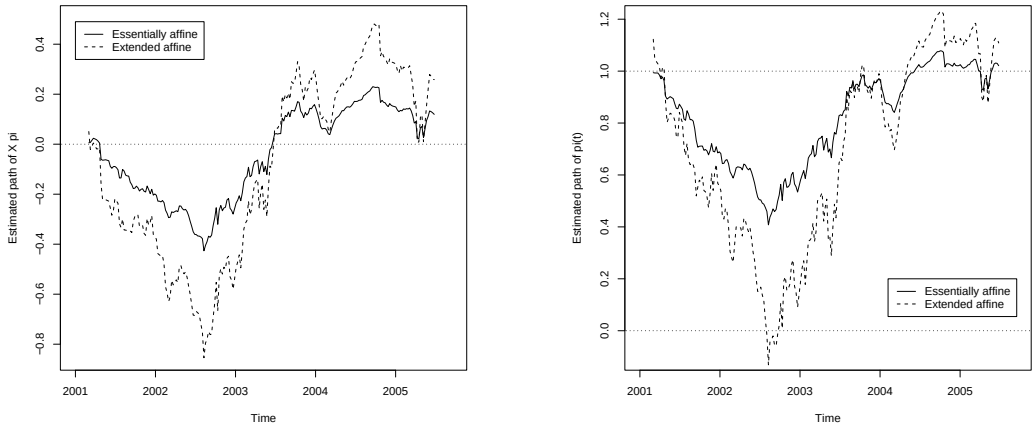

(a) The estimated path of X_t^π.

(b) The estimated path of π_t^Q.

Figure 11: Illustration of the estimated paths of X_t^π and π_t^Q in the stochastic recovery CDS model with essentially and extended affine risk premium specifications. Both estimations use the Ford Motor Co. CDS data covering the period from March 2, 2001 until June 24, 2005.

correlated with the long-term treasury yield level factor X_t^1 and negatively correlated with the short-term yield factor X_t^2. If this is taken as a stylized fact for the data, it is possible based on the estimation results for the stochastic recovery CDS model to derive which factor is the more important in, respectively, the default intensity process and the recovery rate process. Given that the factor loading on the interest rate level factor X_t^1 is positive in both the default intensity and the recovery rate processes across all three estimations of the stochastic recovery CDS model, it can be concluded that the primary channel for the positive correlation between the interest rate level and the credit spread of Ford Motor Co. is through the default intensity. The positive factor loading in the recovery rate process, which by itself indicates a negative relationship between the interest rate level and the credit spread, is not large enough to match the positive effect in the default intensity. For the short-term interest rate factor the opposite result is obtained. The positive factor loading in the default intensity on X_t^2 should indicate a positive correlation between short-term interest rates and the Ford credit spread. Thus, the negative correlation found in the fixed-recovery CDS model can only be explained by the positive factor loading of X_t^2 in the recovery rate process. By implication, the primary channel for the correlation between the short-term interest rates and the credit spread of Ford Motor Co. is through the recovery rate process. A similar, highly significant positive relationship between 3-month T-Bill rates and actually observed recovery rates on defaulted debt is observed by Covitz and Han (2004).

A second important thing to note from Figure 11 is that the instantaneous market-implied recovery rate is above 1 for most of the last year of the sample period in both estimations and briefly below 0 in mid-2002 in the extended affine model. However, once the risk premium parameters are taken into consideration this counterintuitive result is eliminated except for the very short time horizon. This is illustrated in Figure 12. The dash-dotted lines represent the estimated path of the instantaneous recovery rate π_t^Q which is the market-implied expected recovery rate should default happen immediately. The dotted lines represent the estimated path of the Q-conditional expected recovery rate at the 1-year horizon, $E^Q[\pi_{t+1}^Q|\mathcal{F}_t]$. The dashed and solid lines

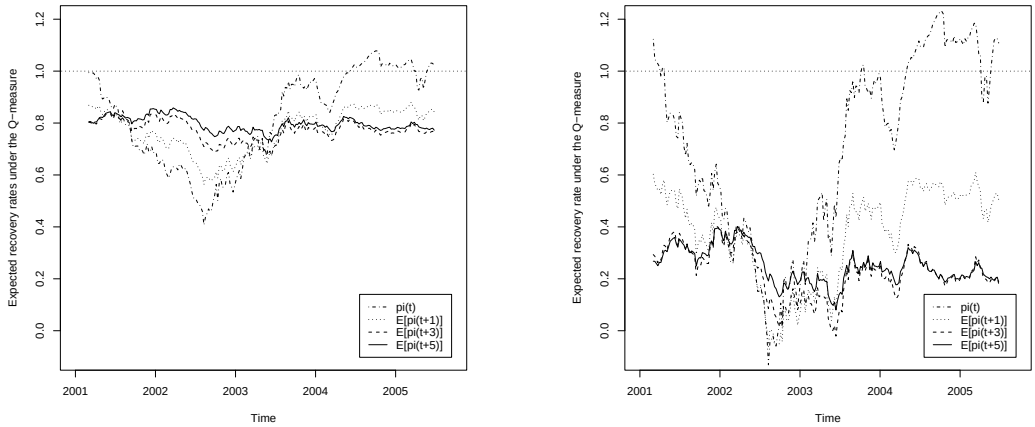

(a) Conditional expected recovery rates from the essentially affine model.

(b) Conditional expected recovery rates from the extended affine model.

Figure 12: Illustration of the estimated path of the expected recovery rate given default under the Q-measure at the 1-year, 3-year, and 5-year time horizons along with the estimated path of the instantaneous recovery rate. The model estimated is the stochastic default and recovery risk CDS model with essentially and extended affine risk premium specifications using the CDS data for Ford Motor Co. over the period March 2, 2001 through June 24, 2005.

represent similar results at the 3- and 5-year horizons, respectively. Due to the stationarity of both of the interest rate factors and the recovery risk factor under the Q-measure, the process is expected to return to its long-run mean for long horizons. In the essentially affine specification this long-run mean is 84.3%, while it is only 34.8% in the extended affine specification. Given that the estimated path of the default intensity λ_t^Q is systematically higher in the essentially affine specification it is only reasonable also to observe a higher model-implied recovery rate process for this specification. However, this clearly shows that the imposition of the Feller condition on the X_t^λ-factor in the extended affine model forces the dynamics of X_t^λ away from the unconstrained result in the essentially affine specification. And now, as opposed to the situation in the fixed-recovery model where the imposition of the Feller condition in the extended affine model only impacted a few parameters and the path of X_t^λ, it influences in significant ways the estimated dynamics for the recovery risk factor and thereby the dynamics of the recovery rate process itself.

Another important result from the joint estimation of default and recovery risk is the ability to analyze the contribution from each risk factor to the CDS spread term structure. This is illustrated in Figure 13. The decomposition of the average CDS spread curve is performed in the following four steps.

(i). For each estimated vector of the four factors, $(\widehat{X}_t^1, \widehat{X}_t^2, \widehat{X}_t^\lambda, \widehat{X}_t^\pi)$, the whole term structure of model-implied CDS spreads is calculated for maturities from one quarter up to 10 years. For each maturity, the average is calculated. These are the solid lines in Figure 13.

(ii). In order to isolate the net contribution from the default risk factor, X_t^λ, the exercise above is repeated. However, this time the effects from the risk-free factors and the recovery risk factor are eliminated by fixing $\lambda_1^r = 0$, $\lambda_2^r = 0$, $\pi_1^r = 0$, $\pi_2^r = 0$, $\gamma_\pi^0 = 0$, $\gamma_\pi^1 = 0$, and letting X_t^π be fixed uniformly at zero. This gives the dashed lines in Figure 13.

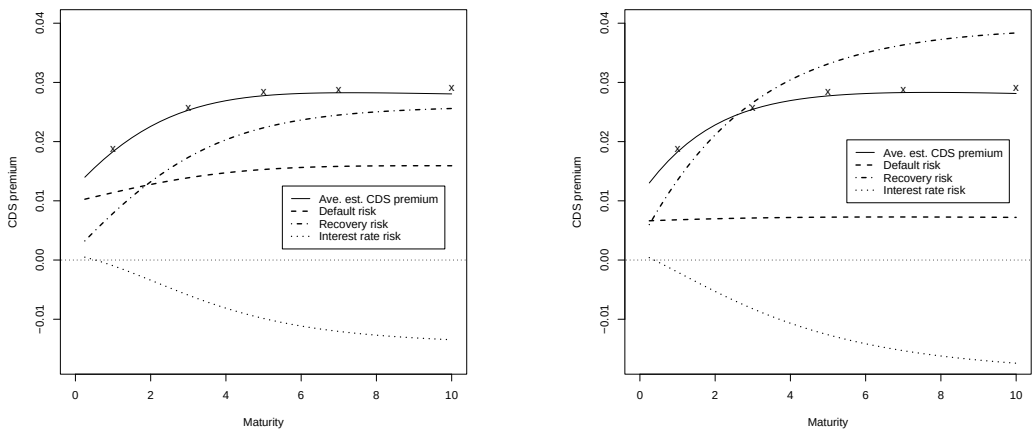

(a) The essentially affine CDS model.

(b) The extended affine CDS model.

Figure 13: Decomposition of the average model-implied CDS term structure into three components: Interest rate, default, and recovery risk. The models estimated are the four-factor CDS model with stochastic default and recovery risk combined with the essentially affine and the extended affine risk premium specifications, respectively. The data used is the CDS data for Ford Motor Co. over the period March 2, 2001 through June 24, 2005.

(*iii*). To measure the impact from the risk-free factors the above exercise is repeated. However, only parameters related to the recovery risk factor are eliminated, i.e. $\gamma_\pi^0 = 0$ and $\gamma_\pi^1 = 0$, and again X_t^π is fixed uniformly at zero. When the result from step 2 is subtracted from the result in step 3, the average net contribution from the risk-free rate factors are found. The result is given by the dotted lines in Figure 13.

(*iv*). Finally, the average net contribution from the recovery risk factor can be calculated as the difference between the result from step 1 and the intermediate result in step 3. The dash-dotted lines in Figure 13 illustrate the net impact of recovery risk on the CDS term structure.

The intuition behind this decomposition is as follows. In step 2 it is assumed that the default intensity is given by $\lambda_t^Q = \lambda_0 + X_t^\lambda$ while the recovery rate is kept constant at the estimated value of π_0. Thus, this step reveals what the CDS rate term structure would be for a hypothetical company where only the estimated dynamics of X_t^λ would enter the CDS model. Here, this is interpreted as the net contribution from the default risk factor to the average CDS rate of Ford Motor Co. over the considered period. In step 3 the interest rate factors are included in the model

$$\begin{aligned} \lambda_t^Q &= \lambda_0 + \lambda_r^1(X_t^1 - \theta_1) + \lambda_r^2 X_t^2 + X_t^\lambda, \\ \pi_t^Q &= \pi_0 + \pi_r^1(X_t^1 - \theta_1) + \pi_r^2 X_t^2. \end{aligned}$$

When the average CDS rate term structure from step 2 is deducted from the average CDS term structure in step 3, the result is the net contribution of the interest rate factors to the average CDS rate term structure. Finally, deducting the result of step 3 from the average CDS term structure for the full model with the stochastic recovery risk factor produces the net contribution from the

recovery risk factor to the average CDS rate term structure over the considered period.

The decomposition of the average CDS rates for Ford Motor reveals significant differences across the two risk premium specifications. In the extended affine model recovery risk constitutes a much larger part of the average CDS rate than in the essentially affine model. The increase in the recovery risk component comes at the expense of the default risk factor, while the negative contribution from the interest rate factors is at about the same level across the two specifications. If we focus solely on the essentially affine model, it is noted that the default risk component contributes a relatively stable amount to the average spread across maturities while recovery risk contributes a larger share of the average CDS rate as maturity increases. For maturities larger than 2 years the recovery risk component actually contributes more to the average CDS rate than the default risk component. This effect is caused by the large risk premia found for the recovery risk factor in all three estimations reported in Table 9. For example in the essentially affine model the estimated values of the recovery risk premia γ_π^0 and γ_π^1 imply that the mean and the mean-reversion rate of the recovery risk factor X_t^π under the Q-measure are -0.266 and 0.772, respectively, while the corresponding values under the P-measure are 0 and 0.608, respectively. Therefore, compared to a situation with a constant recovery rate equal to the estimated value of π_0 (which is the assumption in step 2 and 3 of the CDS rate decomposition), the recovery risk premia produces a steep curve for the net contribution of the recovery risk factor to the average CDS term structure.

As a final comparison of the two risk premium specifications, it is analyzed whether the Feller conditions have any impact on the performance of the extended affine risk premium specification in terms of fitting the observed CDS rates. To that end, the mean error, the mean absolute error, and the root mean square error across the five contracts in the estimations of the essentially affine and extended affine models are provided in Table 10.

Generally speaking, compared to the results obtained with the constant-recovery CDS model, there is a significant improvement in the fit of these models by including the stochastic recovery risk factor. The mean errors are within 2-3 bps in both estimations, except for the 10-year contract. The mean absolute errors have been reduced significantly, in the essentially affine model the absolute errors have been reduced by more than 50% for all five contracts, and for the extended affine model the improvement is even larger due to its very poor performance in the constant-recovery CDS model. This implies that in both specifications the error statistics are not far from the errors found in estimations of regular two-factor yield term structure models such as the one performed in Section 6.1. A similar picture is observed from the estimated measurement error covariance matrix which in the essentially affine specification is estimated as[31]

$$\widehat{H}^{ess} = \begin{pmatrix} 0.001248 & 0.1420 & 0.8523 & 0.9069 & 0.8761 \\ & 0.000952 & 0.4374 & 0.4733 & 0.5655 \\ & & 0.001495 & 0.9494 & 0.8772 \\ & & & 0.001953 & 0.9129 \\ & & & & 0.002908 \end{pmatrix},$$

while the estimated measurement error covariance matrix from the extended affine specification

[31] These two matrices are organized in the following way. The diagonal contains the estimated standard deviations of the measurement errors for the five maturities considered here, 1-year, 3-year, 5-year, 7-year, and 10-year, respectively. The off-diagonal elements are the estimated correlation coefficients between those measurement errors over the sample period.

Maturity	Essentially affine			Extended affine		
	Mean error	Mean abs error	RMSE	Mean error	Mean abs error	RMSE
1 year	1.78	10.09	12.65	2.33	10.93	13.87
3 years	1.19	7.19	9.72	-0.29	9.10	12.29
5 years	3.08	11.87	14.56	3.35	10.10	12.64
7 years	2.45	15.62	19.51	2.10	12.96	16.89
10 years	7.56	23.89	29.89	6.75	20.69	26.92

Table 10: The mean pricing error, the mean absolute pricing error, and the RMSE, all measured in bps, for the 5 different maturities in the two extended Kalman filter estimations based on CDS data for Ford Motor Co. over the period from March 2, 2001 to June 24, 2005.

is given by

$$\widehat{H}^{ext} = \begin{pmatrix} 0.001379 & 0.1723 & 0.8365 & 0.9107 & 0.9031 \\ & 0.001268 & 0.4651 & 0.4511 & 0.5251 \\ & & 0.001368 & 0.9439 & 0.9094 \\ & & & 0.001771 & 0.9070 \\ & & & & 0.002680 \end{pmatrix}.$$

Comparing the two risk premium specifications, it is noted that the extended affine model performs better for the 5-year, the 7-year, and the 10-year contracts, while the essentially affine model is the better at matching the 1-year and 3-year contracts. However, the differences are relatively small, the differences in the mean error is less than 1.5 basis point across the five contracts, and for the mean absolute error the maximum difference is 3.2 bps observed for the 10-year CDS contract.

Despite the slightly better performance of the extended affine model in terms of absolute pricing errors for three of the five contracts, the essentially affine model is considered to be a more true reflection of the underlying dynamics for the following reasons. First, it imposes a minimum constraint on the admissible set of parameters in addition to having one parameter less to be estimated. Second, it has the highest likelihood value, though not directly comparable due to the Feller conditions in the extended affine model. Third, and most importantly, the estimated dynamics for the recovery risk factor is more in line with the estimated recovery rate values obtained in the fix-recovery CDS model. On the other hand, for the extended affine model, the fact that the Feller conditions are binding constraints is found to impact the estimation in important ways and as a consequence the estimated path and parameters for the recovery rate process cannot be considered a true reflection of the market-implied recovery rate process.

For these reasons the remaining discussions in this section will be based on the result from the estimation of the essentially affine specification of the four-factor stochastic recovery CDS model.

6.3.1 The interpretation of X_t^λ and X_t^π

In regular two-factor interest rate term structure estimations one of the factors is usually interpreted as a level factor, while the second factor becomes a slope factor. An example of this is the treasury yield term structure estimation in Section 6.1. Given that the Ford Motor Co. CDS data analyzed in this paper is similar in nature to the treasury data and that a two-factor model has been fitted to both types of data, it would be natural to ask whether the default risk factor and the recovery risk factor can be given similar intuitive interpretations. The discussion here

Correlation	$\widehat{X}_t^\lambda$	$\widehat{X}_t^\pi$	s_{1-yr}	s_{3-yr}	s_{5-yr}	s_{7-yr}	s_{10-yr}	$slope_{5-yr/1-yr}$
$\widehat{X}_t^\lambda$	1	0.144	0.574	0.822	0.896	0.914	0.927	0.692
$\widehat{X}_t^\pi$		1	-0.586	-0.304	-0.136	-0.070	0.006	0.757
s_{1-yr}			1	0.928	0.849	0.811	0.756	-0.142
s_{3-yr}				1	0.982	0.964	0.936	0.237
s_{5-yr}					1	0.996	0.983	0.402
s_{7-yr}						1	0.994	0.465
s_{10-yr}							1	0.536
$slope_{5-yr/1-yr}$								1

Table 11: The correlation coefficients between the levels of the following time-series: the estimated paths of X_t^λ and X_t^π from the essentially affine stochastic recovery CDS model, the observed CDS rates with maturities 1-year, 3-year, 5-year, 7-year and 10-year, and the slope of the CDS curve measured as the difference between the 5-year and the 1-year CDS rates. All variables are observed weekly over the period from March 2, 2001 to June 24, 2005.

will be based on the results from the preferred essentially affine stochastic recovery CDS model, but similar results can be derived from the CDS model with the extended affine risk premium specification.

By mere visual comparison of the estimated path of X_t^λ to the observed CDS rates, it follows that X_t^λ indeed seems to be a level factor. To confirm this conjecture the correlation coefficients between the estimated path of the default intensity risk factor X_t^λ and all the observed CDS rates including the slope of the CDS curve are calculated.[32] The sample estimates are provided in Table 11. The correlation coefficients show that the estimated path of X_t^λ is in fact highly correlated with, in particular, the long-term CDS rates. Performing the same exercise for the estimated path of the recovery risk factor it turns out that X_t^π is not really correlated with any of the observed CDS rates with the exception of the 1-year CDS rate. However, it is strongly correlated with the slope of the CDS curve. Thus, X_t^π can be interpreted as a slope factor.

Another way to study the role of the recovery risk factor X_t^π is to look at its correlation structure with the three other factors in the model. This is done in Table 12. Two things are worth noticing in this table. First, the estimated path of the recovery risk factor X_t^π is uncorrelated with the second factor from the risk-free rate estimation, so the slope factor in the CDS data analyzed here is uncorrelated with the slope factor from the treasury yield estimation. Second, there is negative correlation between the changes of the X_t^λ-process and the X_t^π-process as measured by the correlation coefficient between the first-differences of the two estimated paths of those two factors, despite the assumed independence between the two factors. In a simple structural model of the Merton (1974)-type, however, this would be a natural result as a change in the underlying asset value process would induce a similar change into the expected recovery given default and a change in the opposite direction of the T-year expected default probability.

The ultimate question of course is: Does X_t^π really reflect the market's perception about future expected recovery rates? Or is it just a slope factor that is falsely given the interpretation of being a recovery factor? To see that it makes sense to interpret X_t^π as a recovery factor start by focusing on the mean-reversion rate κ_π of the recovery risk factor. The large estimate of κ_π in the essentially affine specification implies that the recovery rate process mean-reverts fairly quickly. This indicates that recovery risk is only something the market can relate to at the 12- to

[32] Throughout the slope of the CDS curve is measured as the difference between the 5-year and the 1-year CDS rates.

	Essentially affine	Extended affine
$corr(\widehat{X}_t^\pi, \widehat{X}_t^\lambda)$	0.144	0.130
$corr(\Delta\widehat{X}_t^\pi, \Delta\widehat{X}_t^\lambda)$	-0.187	-0.210
$corr(\widehat{X}_t^\pi, \widehat{X}_t^1)$	-0.356	-0.353
$corr(\Delta\widehat{X}_t^\pi, \Delta\widehat{X}_t^1)$	0.141	0.191
$corr(\widehat{X}_t^\pi, \widehat{X}_t^2)$	0.055	0.041
$corr(\Delta\widehat{X}_t^\pi, \Delta\widehat{X}_t^2)$	0.006	-0.007

Table 12: Correlation coefficients between the estimated path of the recovery factor, X_t^π, and the estimated paths of the other three factors in the stochastic recovery CDS model estimated with essentially and extended affine risk premium specifications. Both estimations are based on CDS data for Ford Motor Co. over the period from March 2, 2001 to June 24, 2005.

18-month horizon. For maturities beyond that the market-implied expected recovery is close to the unconditional mean. Accordingly, changes in X_t^π only have a significant impact on that part of the cash flow of the CDS contract that is no more than 12-18 months into the future, and for that reason the recovery risk factor is relatively highly correlated with the 1-year CDS rates, lesser so with the 3-year rates, and practically uncorrelated with the 5- to 10-year rates. In addition, it is worth noting that all the correlations have the right negative sign as an increase in the recovery rate should decrease the CDS rates. Finally, as observed in Table 12, the estimated recovery risk factor is practically uncorrelated with the interest rate slope factor. This can be explained by the fact that the interest rate factors appear as factors in the recovery rate, and seemingly all the material correlation between the interest rate and the recovery rate processes is caught in this way.

6.3.2 The issue of identifying X_t^λ and X_t^π

In this subsection the problem of distinguishing movements in X_t^λ from movements in X_t^π based on cross-sections of CDS rates will be briefly discussed and compared to the problem of identifying the two interest rate factors in the treasury yield model.

Begin by focusing on the estimated parameters and state variables from the preferred interest rate model and the preferred CDS model.[33] Assume that those two models are the true models and fix the date to be September 14, 2001.[34] The first step is to find the combinations of (X_t^λ, X_t^π) that give a perfect fit to each of the five model-implied CDS rates on this date keeping the interest rate factors fixed at the estimated values. The result is illustrated in Figure 14(a). Second, repeat the exercise for the risk-free interest rate factors, i.e. calculate the combinations of (X_t^1, X_t^2) that give a perfect fit to the four model-implied treasury bond yields observed on this date. This is illustrated in Figure 14(b).

The main thing to note when comparing the two figures is the similarity. In both figures it is clear that each of the illustrated contracts react differently to changes in the state factors which is the prerequisite for a successful joint identification. Thus, the ability to perform a joint estimation seem to be about the same in the two types of estimations. Therefore, in a world where the true

[33] For the interest rate model, this means the model with the extended affine risk premium specification. For the CDS model, this means the four-factor stochastic recovery model with the essentially affine risk premium specification.

[34] This is the first day where the estimate of X_t^λ is above 25 bps and a day with a relatively moderate CDS level compared to the rest of the sample. For the point to be made in this section, higher values of X_t^λ will only add to the ability of identifying both factors, ceteris paribus.

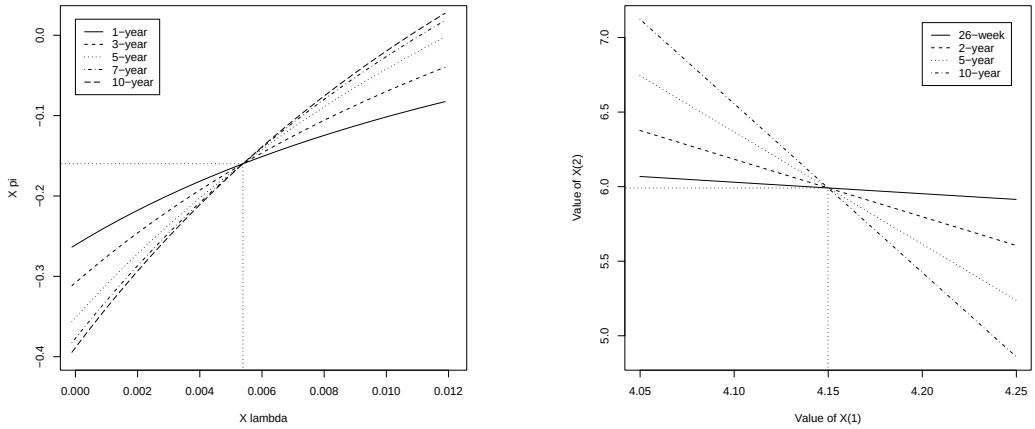

(a) (X_t^λ, X_t^π)-combinations giving a perfect fit for the model-implied CDS rates on September 14, 2001.

(b) (X_t^1, X_t^2)-combinations giving a perfect fit for the model-implied Treasury yields on September 14, 2001.

Figure 14: The figure to the left illustrates the combinations of (X_t^λ, X_t^π) that gives a perfect fit to each of the five model-implied CDS rates on September 14, 2001. The figure to the right illustrates the combinations of the two risk-free factors (X_t^1, X_t^2) that give a perfect fit for the four model-implied treasury yields on September 14, 2001. The models used are the stochastic recovery CDS model with essentially affine risk premium specification and the interest rate model with the extended affine risk premium specification.

model is known and there is a given level of measurement noise, the joint estimation of (X_t^λ, X_t^π) based on observing time series of CDS rates for five different maturities is about as well-identified as the estimation of the risk-free rate factors (X_t^1, X_t^2) based on time series of treasury yields across a similar number of bonds with different maturities. In addition, it is noted that given the set of estimated parameters obtained for the specific data sets used in this paper, the short term contracts, i.e. the 1-year CDS rate and the 26-Week T-bill yield, determine the values of X_t^π and X_t^2, respectively, while the long-term contracts are the primary determinants of X_t^λ and X_t^1, respectively. For the risk-free rate factors this is in line with the interpretation of X_t^1 as being a level factor, while X_t^2 was a slope and/or short-term factor. Similarly, in the last section, X_t^λ was interpreted as a level factor in the CDS estimation, while X_t^π was shown to be a short-term factor in the CDS estimation.

The simulation study performed in Christensen (2006b) indicates that the level of measurement error noise is important for the ability to decompose CDS rates into default and recovery risk components. In order to illustrate the impact of varying the size of the measurement error standard deviation, start by considering the combinations of (X_t^λ, X_t^π) that will cause all of the five CDS contracts to deviate less then a certain number of basis points from their corresponding model-implied CDS rates. For threshold values of 1 bp, 5 bps, and 10 bps, respectively, these combinations are illustrated in Figure 15(a).[35] In the figure, the full grey line encircles the set of (X_t^λ, X_t^π)-combinations that will cause all five CDS contracts simultaneously to be priced with an absolute error smaller than 1 bp relative to the estimated model-implied CDS rates on September 14, 2001. The black dashed lines illustrate the set of (X_t^λ, X_t^π)-combinations that will lead to an absolute

[35] I thank Mike Gibson for suggesting this way of illustrating the impact of differences in the level of measurement noise.

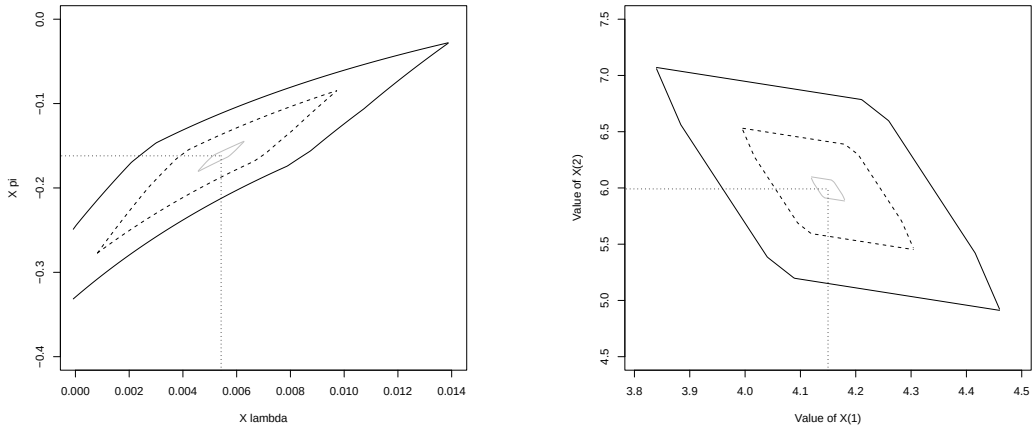

(a) (X_t^λ, X_t^π)-combination matching observed CDS rates on September 14, 2001.

(b) (X_t^1, X_t^2)-combination matching observed Treasury bond yields on September 14, 2001.

Figure 15: The figure to the left illustrates the combinations of (X_t^λ, X_t^π) that cause all five CDS contracts to simultaneously deviate less than 1 bp, 5 bps, and 10 bps, respectively, from the model-implied CDS rates on September 14, 2001. The figure to the right illustrates the combinations of (X_t^1, X_t^2) that cause all four treasury yields to simultaneously deviate less than 1 bp, 5 bps, and 10 bps, respectively, from the model-implied yields observed on September 14, 2001. Grey lines represent the 1 bp error set, while the dashed black lines represent the 5 bps error set, and full black lines represent the 10 bps error set. The models used are the stochastic recovery CDS model with essentially affine risk premium specification and the interest rate model with the extended affine risk premium specification.

pricing error of less than 5 bps across all five CDS contracts. Finally, the full black lines illustrate the combinations that give absolute pricing errors less than 10 bps for all five contracts. In Figure 15(b) the same exercise is repeated for the two-factor model of the treasury bond yields. The impact of measurement error across the two types of estimations seem to be structured in much the same way. Thus, given the fact that researchers have not been worried about identification issues in yield term structure estimations of the above two-factor kind, this analysis provides evidence that a similar degree of joint identification can be obtained in the estimation of the four-factor CDS model with joint modeling of default and recovery risk, at least for the CDS data used in this paper.

In terms of generalizing the result, there are two factors that are of importance. First, if the data is more noisy than the Ford Motor Co. CDS data applied here, clearly Figure 15 shows that such noise can deteriorate any type of estimation. Second, and more importantly in a credit risk setting, if the overall CDS rate level is lower than the level of the CDS data in this paper, the combinations of (X_t^λ, X_t^π) that provide a perfect fit for each CDS contract will be closer to each other, thus deteriorating the ability to distinguish movements in X_t^λ from movements in X_t^π.[36]

Based on the evidence in this paper and the results obtained by Christensen (2006b) it seems fair to conclude that for speculative grade issuers and the lower end of the BBB-rated segment it will make sense to perform a joint estimation of default and recovery risk, and on the other hand for firms with an A-rating or higher any expected default event will be far in the future and any

[36]Compare for example Figure 14(a) with similar figures in Christensen (2006b), where the average spread level is the neighborhood of 50-70 bps.

Time	Moody's rating
March 2001	A2
October 2001	A3
January 2002	Baa1
May 2005	Baa3
August 2005	Ba1
January 2006	Ba3

Time	S&P rating
March 2001	A
October 2001	BBB+
October 2002	BBB
November 2003	BBB-
May 2005	BB+
January 2006	BB-

Table 13: The ratings assigned by Moody's and S&P to Ford Motor Co. over the period from March 2001 to May 2006.

changes in the recovery risk factor will only have a very marginal impact on CDS rates of any maturity. Thus, it does not make much sense to perform this type of estimation for high-quality firms, instead researchers will be better off using the three-factor CDS model with a constant recovery rate where the recovery rate parameter can still be estimated as argued by Pan and Singleton (2005).

6.3.3 Uncertainty about the situation of Ford Motor Co.

Combining the existing available information on Ford Motor Co. leaves a mixed picture that the estimation result of the stochastic recovery CDS model may be able to shed some light on.

In Table 13 the history of the credit ratings assigned by Moody's and S&P to Ford Motor Co. over the last five years are shown. There is a clear downward trend in the ratings assigned by both rating agencies, a natural consequence of the reported negative operating income from Ford Motor Co. over the last few years.

In figure 16 the 1-year EDF for Ford Motor Co. is illustrated. This one-year default probability is estimated by Moody's KMV based on the stock price and the liability structure of Ford Motor Co. This curve has a clear trend in the opposite direction indicating increased confidence by the stock market that Ford Motor Co. will be able to manage its liability problems in an orderly fashion. Actually, the current 1-year EDF estimate is in the 0.1%-0.15% range which is close to the 1-year EDF of the average A-rated company. The primary causes for the good EDF estimates over the most recent two years are: 1) The estimated asset volatility has been cut in half from 0.08 to 0.04 over the period. 2) The liabilities of Ford and therefore the implied default point have followed the estimated, model-implied asset value down as can be seen in Figure 17.

This begs the question: How can the low default probability from the Moody's KMV EDF estimates be reconciled with the relatively poor ratings assigned by all the major rating agencies? Based on the result of the CDS estimations performed in this paper the answer seem to be that there certainly has been a significant increase in the market-implied default intensity since the beginning of 2005 when the speculation of a downgrade to below investment grade started. In that sense the rating agencies are right and the EDF estimates are wrong. However, should a default happen then the market seem to expect to recover between 80% and 90% on the dollar. This high number is also consistent with the low volatility of Ford's asset value process as estimated by Moody's KMV. Thus, both the estimations performed in this paper and the analyses performed by Moody's KMV indicate that the market is confident that Ford Motor has sufficient means to be able to service all of its liabilities in the short term, while the long-term outlook is more uncertain with a higher probability of default.

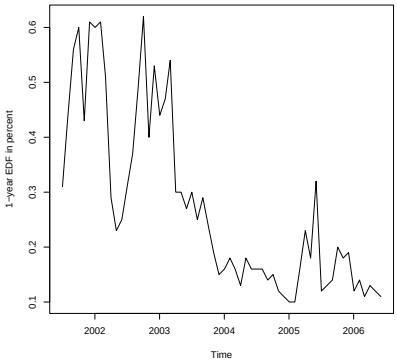

Figure 16: The one-year EDF estimated by Moody's KMV for Ford Motor Co. over the period from July 2001 to May 2006.

7 Conclusion

This paper has attempted to decompose time series of observed CDS rates into a default and recovery risk component based on a reduced-form credit risk modeling framework. For the CDS data for Ford Motor CO. analyzed in this study, evidence is provided that it is possible to perform such joint estimations. The default risk factor acts as a level factor, while the recovery risk factor is a slope factor with a high rate of mean-reversion in line with an interpretation that fluctuations in recovery risk is only of importance for that part of the cash flow that becomes due within the first 12 to 18 months. In addition, decomposition of the average CDS rates indicates that recovery risk is a significant component of average CDS rates, actually surpassing default risk at the 2 year maturity for the Ford Motor Co. CDS data analyzed in this paper.

Furthermore, the issue of the most appropriate affine risk premium specification was investigated. For the modeling of the risk-free rate, the extended affine risk premium specification was the preferred. However, in the modeling of CDS rates, it turns out that the Feller conditions required to make the extended affine specification arbitrage-free are binding constraints which not only change the estimates of important parameters, but also deteriorate the performance of the estimated model.

Finally, the important task of studying systematic patterns in the market-implied default intensities and recovery rates across firms that have identical credit ratings or belong to the same industry is left for future research.

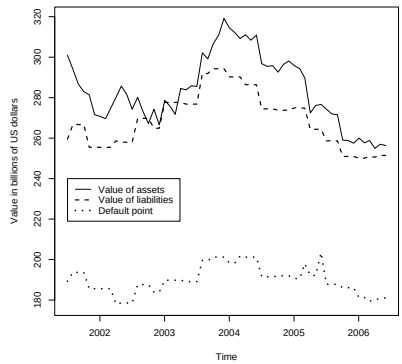

Figure 17: The total value of assets and liabilities of Ford Motor Co. as calculated by Moody's KMV. Included is also their estimate of the default point. All three series cover monthly observations over the period from July 2001 to May 2006.

8 Appendix A: First and second moments in the risk-free rate model

In the $A_1(2)$-model for risk-free rate it is immediately noted that the X_t^1-process is an independent CIR-process

$$dX_t^1 = \kappa_{11}(\theta_1 - X_t^1)dt + \sqrt{X_t^1}dW_t^1.$$

The conditional mean of this process is

$$E[X_T^1|\mathcal{F}_t] = \theta_1[1 - e^{-\kappa_{11}(T-t)}] + X_t^1 e^{-\kappa_{11}(T-t)}.$$

Provided that $\kappa_{11} > 0$ the X_t^1-process is stationary with an unconditional mean that can be derived from the above expression by letting T go to infinity. This limit is equal to θ_1.

The conditional variance of the CIR-process is equal to

$$V[X_T^1|\mathcal{F}_t] = \frac{1}{2\kappa_{11}}\theta_1[1 - e^{-\kappa_{11}(T-t)}]^2 + \frac{1}{\kappa_{11}}X_t^1 e^{-\kappa_{11}(T-t)}[1 - e^{-\kappa_{11}(T-t)}],$$

which implies that the unconditional variance of X_t^1 is $\frac{\theta_1}{2\kappa_{11}}$, again under the assumption that $\kappa_{11} > 0$.

Further, simple calculations lead to the conditional mean of X_t^2

$$E[X_T^2|\mathcal{F}_t] = \frac{\kappa_{21}}{\kappa_{22} - \kappa_{11}}(\theta_1 - X_t^1)e^{-\kappa_{11}(T-t)} + \left[X_t^2 - \frac{\kappa_{21}}{\kappa_{22} - \kappa_{11}}(\theta_1 - X_t^1)\right]e^{-\kappa_{22}(T-t)},$$

and the conditional expectations of the quadratic terms

$$
\begin{aligned}
E[(X_T^1)^2|\mathcal{F}_t] &= \theta_1^2[1-e^{-\kappa_{11}(T-t)}]^2 + \frac{1}{2\kappa_{11}}\theta_1[1-e^{-\kappa_{11}(T-t)}]^2 \\
&+ (X_t^1)^2 e^{-2\kappa_{11}(T-t)} + [\frac{1}{\kappa_{11}} + 2\theta_1]X_t^1 e^{-\kappa_{11}(T-t)}[1-e^{-\kappa_{11}(T-t)}], \\
E[X_T^1 X_T^2|\mathcal{F}_t] &= X_t^1 X_t^2 e^{-(\kappa_{11}+\kappa_{22})(T-t)} \\
&+ \frac{\kappa_{21}\theta_1}{\kappa_{22}}(X_t^1-\theta_1)e^{-\kappa_{11}(T-t)}[1-e^{-\kappa_{22}(T-t)}] \\
&+ \frac{\kappa_{21}}{\kappa_{22}-\kappa_{11}}\frac{\kappa_{11}\theta_1}{\kappa_{22}}(\theta_1-X_t^1)e^{-\kappa_{11}(T-t)}[1-e^{-\kappa_{22}(T-t)}] \\
&+ \Big[X_t^2 - \frac{\kappa_{21}}{\kappa_{22}-\kappa_{11}}(\theta_1-X_t^1)\Big]\theta_1 e^{-\kappa_{22}(T-t)}[1-e^{-\kappa_{11}(T-t)}] \\
&- \frac{\kappa_{21}\theta_1}{2\kappa_{11}}\frac{1}{\kappa_{11}+\kappa_{22}}[1-e^{-(\kappa_{11}+\kappa_{22})(T-t)}] \\
&- \frac{\kappa_{21}}{\kappa_{22}}\Big[-2\theta_1^2 - \frac{\theta_1}{\kappa_{11}} + [\frac{1}{\kappa_{11}}+2\theta_1]X_t^1\Big]e^{-\kappa_{11}(T-t)}[1-e^{-\kappa_{22}(T-t)}] \\
&- \frac{\kappa_{21}}{\kappa_{22}-\kappa_{11}}\Big[\theta_1^2 + \frac{\theta_1}{2\kappa_{11}} + (X_t^1)^2 - [\frac{1}{\kappa_{11}}+2\theta_1]X_t^1\Big]e^{-\kappa_{11}(T-t)}[e^{-\kappa_{11}(T-t)} - e^{-\kappa_{22}(T-t)}],
\end{aligned}
$$

$$
\begin{aligned}
E[(X_T^2)^2|\mathcal{F}_t] &= (X_t^2)^2 e^{-2\kappa_{22}(T-t)} + \frac{1}{2\kappa_{22}}[1-e^{-2\kappa_{22}(T-t)}] \\
&+ \frac{\theta_1\beta_{21}}{2\kappa_{22}}[1-e^{-2\kappa_{22}(T-t)}] - \frac{\beta_{21}}{2\kappa_{22}-\kappa_{11}}(\theta_1-X_t^1)[e^{-\kappa_{11}(T-t)} - e^{-2\kappa_{22}(T-t)}] \\
&+ 2\kappa_{21}\theta_1\frac{\kappa_{21}}{\kappa_{22}-\kappa_{11}}\frac{1}{2\kappa_{22}-\kappa_{11}}(\theta_1-X_t^1)[e^{-\kappa_{11}(T-t)} - e^{-2\kappa_{22}(T-t)}] \\
&+ \frac{2\kappa_{21}\theta_1}{\kappa_{22}}\Big[X_t^2 - \frac{\kappa_{21}}{\kappa_{22}-\kappa_{11}}(\theta_1-X_t^1)\Big]e^{-\kappa_{22}(T-t)}[1-e^{-\kappa_{22}(T-t)}] \\
&+ \frac{\kappa_{21}^2\theta_1}{2\kappa_{11}\kappa_{22}}\frac{1}{\kappa_{11}+\kappa_{22}}[1-e^{-2\kappa_{22}(T-t)}] \\
&- 2\kappa_{21}\Bigg(\frac{\kappa_{21}\theta_1}{\kappa_{22}}(X_t^1-\theta_1) + \frac{\kappa_{21}}{\kappa_{22}-\kappa_{11}}\frac{\kappa_{11}\theta_1}{\kappa_{22}}(\theta_1-X_t^1) \\
&- \frac{\kappa_{21}}{\kappa_{22}}\Big[-2\theta_1^2 - \frac{\theta_1}{\kappa_{11}} + [\frac{1}{\kappa_{11}}+2\theta_1]X_t^1\Big]\Bigg)\frac{1}{2\kappa_{22}-\kappa_{11}}[e^{-\kappa_{11}(T-t)} - e^{-2\kappa_{22}(T-t)}] \\
&- \frac{2\kappa_{21}\theta_1}{\kappa_{22}}\Big[X_t^2 - \frac{\kappa_{21}}{\kappa_{22}-\kappa_{11}}(\theta_1-X_t^1)\Big]e^{-\kappa_{22}(T-t)}[1-e^{-\kappa_{22}(T-t)}] \\
&- \frac{2\kappa_{21}}{\kappa_{22}-\kappa_{11}}\Bigg(X_t^1X_t^2 - \frac{\kappa_{21}\theta_1}{\kappa_{22}}(X_t^1-\theta_1) - \frac{\kappa_{21}}{\kappa_{22}-\kappa_{11}}\frac{\kappa_{11}\theta_1}{\kappa_{22}}(\theta_1-X_t^1) \\
&\quad -\Big[X_t^2 - \frac{\kappa_{21}}{\kappa_{22}-\kappa_{11}}(\theta_1-X_t^1)\Big]\theta_1 + \frac{\kappa_{21}\theta_1}{2\kappa_{11}}\frac{1}{\kappa_{11}+\kappa_{22}} \\
&\quad +\frac{\kappa_{21}}{\kappa_{22}}\Big[-2\theta_1^2 - \frac{\theta_1}{\kappa_{11}} + [\frac{1}{\kappa_{11}}+2\theta_1]X_t^1\Big] \\
&\quad +\frac{\kappa_{21}}{\kappa_{22}-\kappa_{11}}\Big[\theta_1^2 + \frac{\theta_1}{2\kappa_{11}} + (X_t^1)^2 - [\frac{1}{\kappa_{11}}+2\theta_1]X_t^1\Big]\Bigg)e^{-\kappa_{22}(T-t)}[e^{-\kappa_{11}(T-t)} - e^{-\kappa_{22}(T-t)}] \\
&+ \frac{\kappa_{21}^2}{(\kappa_{22}-\kappa_{11})^2}\Big[\theta_1^2 + \frac{\theta_1}{2\kappa_{11}} + (X_t^1)^2 - [\frac{1}{\kappa_{11}}+2\theta_1]X_t^1\Big][e^{-2\kappa_{11}(T-t)} - e^{-2\kappa_{22}(T-t)}].
\end{aligned}
$$

The unconditional moments

The unconditional moments are thus given by

$$
\begin{aligned}
\lim_{T\to\infty} E[X_T^1|\mathcal{F}_t] &= \theta_1, \\
\lim_{T\to\infty} E[X_T^2|\mathcal{F}_t] &= 0, \\
\lim_{T\to\infty} E[(X_T^1)^2|\mathcal{F}_t] &= \theta_1^2 + \frac{\theta_1}{2\kappa_{11}}, \\
\lim_{T\to\infty} E[X_T^1 X_T^2|\mathcal{F}_t] &= -\frac{\kappa_{21}\theta_1}{2\kappa_{11}}\frac{1}{\kappa_{11}+\kappa_{22}}, \\
\lim_{T\to\infty} E[(X_T^2)^2|\mathcal{F}_t] &= \frac{1}{2\kappa_{22}} + \frac{\theta_1\beta_{21}}{2\kappa_{22}} + \frac{\kappa_{21}^2\theta_1}{2\kappa_{11}\kappa_{22}}\frac{1}{\kappa_{11}+\kappa_{22}}.
\end{aligned}
$$

This implies that the unconditional variances and covariances are given by

$$
\begin{aligned}
\lim_{T\to\infty} V[X_T^1|\mathcal{F}_t] &= \frac{\theta_1}{2\kappa_{11}}, \\
\lim_{T\to\infty} V[X_T^2|\mathcal{F}_t] &= \frac{1}{2\kappa_{22}} + \frac{\theta_1\beta_{21}}{2\kappa_{22}} + \frac{\kappa_{21}^2\theta_1}{2\kappa_{11}\kappa_{22}}\frac{1}{\kappa_{11}+\kappa_{22}}, \\
\lim_{T\to\infty} cov[X_T^1, X_T^2|\mathcal{F}_t] &= -\frac{\kappa_{21}\theta_1}{2\kappa_{11}}\frac{1}{\kappa_{11}+\kappa_{22}}.
\end{aligned}
$$

9 Appendix B: The Kalman filter estimation of the risk-free rate

Building on the conditional moments provided in Appendix A, the discrete-time transition equation for the state variables in the risk-free rate process is given by

$$X_t = \Phi_t^0(\psi) + \Phi_t^1(\psi) X_{t-1} + u_t, \quad u_t \sim N(0, V_t(\psi)),$$

where

$$\Phi_t^0(\psi) = \begin{pmatrix} \theta_1[1 - e^{-\kappa_{11}\Delta t}] \\ \frac{\kappa_{21}}{\kappa_{22}-\kappa_{11}}\theta_1[e^{-\kappa_{11}\Delta t} - e^{-\kappa_{22}\Delta t}] \end{pmatrix}$$

and

$$\Phi_t^1(\psi) = \begin{pmatrix} e^{-\kappa_{11}\Delta t} & 0 \\ \frac{\kappa_{21}}{\kappa_{22}-\kappa_{11}}[e^{-\kappa_{22}\Delta t} - e^{-\kappa_{11}\Delta t}] & e^{-\kappa_{22}\Delta t} \end{pmatrix},$$

and Δt is the time between observations.

As for the conditional variance, it is calculated using the updated value $\widehat{X}_{t-1}$ from the last step of the Kalman filter algorithm

$$V_t(\psi) = \begin{pmatrix} V[X_t^1|\widehat{X}_{t-1}^1] & cov[X_t^1 X_t^2|\widehat{X}_{t-1}^1, \widehat{X}_{t-1}^2] \\ cov[X_t^1 X_t^2|\widehat{X}_{t-1}^1, \widehat{X}_{t-1}^2] & V[X_t^2|\widehat{X}_{t-1}^1, \widehat{X}_{t-1}^2] \end{pmatrix}.$$

The model-implied price of a treasury bond with coupon rate C paid semi-annually and τ years until next coupon payment date on which date there will be N remaining coupon payments is

$$\begin{aligned} V_0(X_0^1, X_0^2, \tau, N, C; \psi) &= E_0^Q[e^{-\int_0^{\tau+0.5(N-1)} r_u du}] + \sum_{i=1}^{N} \frac{C}{2} E_0^Q[e^{-\int_0^{\tau+0.5(i-1)} r_u du}] \\ &= e^{B(\tau+0.5(N-1);\psi)' X_0 + C(\tau+0.5(N-1);\psi)} \\ &+ \sum_{i=1}^{N} \frac{C}{2} e^{B(\tau+0.5(i-1);\psi)' X_0 + C(\tau+0.5(i-1);\psi)}, \end{aligned}$$

where $B(t;\psi)$ and $C(t;\psi)$ are the solutions for a system of ODEs that can be derived from the result in Proposition 1 in Section 3.4.

Since the yield to maturity $y_0(\tau, N, C)$ on such a bond is the solution to the equation

$$\sum_{i=1}^{N} \frac{C}{2} e^{-y_0(\tau,N,C)(\tau+0.5(i-1))} + e^{-y_0(\tau,N,C)(\tau+0.5(N-1))} = V_0(X_0^1, X_0^2, \tau, N, C; \psi),$$

it is a non-linear function of the state variables (X_t^1, X_t^2)

$$y_0(\tau, N, C) = z(X_t^1, X_t^2, \tau, N, C; \psi).$$

In addition, it is assumed that there is error in the measurement of yields, so the measurement equation becomes

$$y_t(\tau, N, C) = z(X_t^1, X_t^2, \tau, N, C; \psi) + \varepsilon_t, \quad \varepsilon_t \sim N(0, H_t(\varepsilon)).$$

Here, the measurement errors are assumed to be serially uncorrelated, but cross-sectionally correlated.

In order to apply the extended Kalman filter a linear measurement equation is needed. To this end the model-implied yields are approximated by a first order Taylor expansion around the best guess of (X_t^1, X_t^2) in the prediction step of the Kalman filter algorithm (see below for details). The best guess is denoted by $(\widehat{X}_{t|t-1}^1, \widehat{X}_{t|t-1}^2)$ whereby the approximation becomes

$$\begin{aligned} z(X_t^1, X_t^2, \tau, N, C; \psi) &\approx z(\widehat{X}_{t|t-1}^1, \widehat{X}_{t|t-1}^2, \tau, N, C; \psi) \\ &+ \frac{\partial z(X_t^1, X_t^2, \tau, N, C; \psi)}{\partial X_t^1}\Big|_{(X_t^1, X_t^2)=(\widehat{X}_{t|t-1}^1, \widehat{X}_{t|t-1}^2)} (X_t^1 - \widehat{X}_{t|t-1}^1) \\ &+ \frac{\partial z(X_t^1, X_t^2, \tau, N, C; \psi)}{\partial X_t^2}\Big|_{(X_t^1, X_t^2)=(\widehat{X}_{t|t-1}^1, \widehat{X}_{t|t-1}^2)} (X_t^2 - \widehat{X}_{t|t-1}^2). \end{aligned}$$

Defining

$$\begin{aligned} A_t(\psi) &\equiv z(\widehat{X}_{t|t-1}^1, \widehat{X}_{t|t-1}^2, \tau, N, C; \psi) - \frac{\partial z(X_t^1, X_t^2, \tau, N, C; \psi)}{\partial X_t^1}\Big|_{(X_t^1, X_t^2)=(\widehat{X}_{t|t-1}^1, \widehat{X}_{t|t-1}^2)} \widehat{X}_{t|t-1}^1, \\ &- \frac{\partial z(X_t^1, X_t^2, \tau, N, C; \psi)}{\partial X_t^2}\Big|_{(X_t^1, X_t^2)=(\widehat{X}_{t|t-1}^1, \widehat{X}_{t|t-1}^2)} \widehat{X}_{t|t-1}^2, \end{aligned}$$

and

$$B_t(\widehat{X}_{t|t-1}^1, \widehat{X}_{t|t-1}^2, \tau, N, C; \psi) \equiv \begin{pmatrix} \frac{\partial z(X_t^1, X_t^2, \tau, N, C; \psi)}{\partial X_t^1}\Big|_{(X_t^1, X_t^2)=(\widehat{X}_{t|t-1}^1, \widehat{X}_{t|t-1}^2)} \\ \frac{\partial z(X_t^1, X_t^2, \tau, N, C; \psi)}{\partial X_t^1}\Big|_{(X_t^1, X_t^2)=(\widehat{X}_{t|t-1}^1, \widehat{X}_{t|t-1}^2)} \end{pmatrix},$$

the measurement equation can be given on an affine form as

$$y_t(\tau, N, C) = A_t(\psi) + B_t(\psi)' \begin{pmatrix} X_t^1 \\ X_t^2 \end{pmatrix} + \varepsilon_t.$$

In the following, the steps of the standard Kalman filter are described. Define the total information available at time t by

$$Y_t = (y_1, y_2, \ldots, y_t).$$

Now, assume that the best guess of the value of the state variables in the last period is $\widehat{X}_{t-1}$, and the best guess for its mean square error matrix is $\widehat{\Sigma}_{t-1}$. In the prediction step of the Kalman filter the best guess of X_t given the information set Y_{t-1} can be shown to be

$$\begin{pmatrix} \widehat{X}_{t|t-1}^1 \\ \widehat{X}_{t|t-1}^2 \end{pmatrix} = E^P\left[\begin{pmatrix} X_t^1 \\ X_t^2 \end{pmatrix}\Big| Y_{t-1}\right] = \Phi_t^0(\psi) + \Phi_t^1(\psi) \begin{pmatrix} \widehat{X}_{t-1}^1 \\ \widehat{X}_{t-1}^2 \end{pmatrix}.$$

The corresponding best guess for the mean square error matrix is

$$\widehat{\Sigma}_{t|t-1} = \Phi_t^1(\psi)\widehat{\Sigma}_{t-1}\Phi_t^1(\psi)' + V_t(\psi).$$

In the update step the guess is improved upon by using the additional information contained in Y_t. It can be shown that the updated values of the state variables and the mean square error matrix are

$$\begin{aligned}
\begin{pmatrix} \widehat{X}_t^1 \\ \widehat{X}_t^2 \end{pmatrix} &= \left[\begin{pmatrix} X_t^1 \\ X_t^2 \end{pmatrix} \middle| Y_t \right] = \begin{pmatrix} \widehat{X}_{t|t-1}^1 \\ \widehat{X}_{t|t-1}^2 \end{pmatrix} + \widehat{\Sigma}_{t|t-1} B_t(\psi)' F_t^{-1} v_t, \\
\widehat{\Sigma}_t &= \widehat{\Sigma}_{t|t-1} - \widehat{\Sigma}_{t|t-1} B_t(\psi)' F_t^{-1} B_t(\psi) \widehat{\Sigma}_{t|t-1},
\end{aligned}$$

where[37]

$$\begin{aligned}
v_t &= y_t - E[y_t|Y_{t-1}] = y_t - z(\widehat{X}_{t|t-1}^1, \widehat{X}_{t|t-1}^2, \tau, N, C; \psi), \\
F_t &= cov(v_t) = B_t(\psi) \widehat{\Sigma}_{t|t-1} B_t(\psi)' + H_t(\psi).
\end{aligned}$$

The log-likelihood function for a given vector of parameters ψ can be calculated as[38]

$$\log L(y_1, \ldots, y_N; \psi) = \sum_{t=1}^{N} \Big(-\frac{m_t}{2} \log(2\pi) - \frac{1}{2} \log |F_t| - \frac{1}{2} v_t' F_t^{-1} v_t \Big).$$

Finally, since the interest rate process is assumed to be stationary under the P-measure, the algorithm is started at the unconditional mean and covariance matrix of the state variables under the P-measure

$$\begin{pmatrix} \widehat{X}_0^1 \\ \widehat{X}_0^2 \end{pmatrix} = \begin{pmatrix} \theta_1 \\ 0 \end{pmatrix}, \quad \widehat{\Sigma}_0 = \begin{pmatrix} \frac{\theta_1}{2\kappa_{11}} & -\frac{\kappa_{21}\theta_1}{2\kappa_{11}} \frac{1}{\kappa_{11}+\kappa_{22}} \\ -\frac{\kappa_{21}\theta_1}{2\kappa_{11}} \frac{1}{\kappa_{11}+\kappa_{22}} & \frac{1}{2\kappa_{22}} + \frac{\theta_1 \beta_{21}}{2\kappa_{22}} + \frac{\kappa_{21}^2 \theta_1}{2\kappa_{11}\kappa_{22}} \frac{1}{\kappa_{11}+\kappa_{22}} \end{pmatrix}.$$

The optimal choice of the parameters and the corresponding most likely path of r_t is found by maximizing the value of the log-likelihood function.

9.1 Heterogeneous measurement errors

Allowing for different measurement error variances across the different yield time series and also allowing for cross-sectional correlation between those errors is easily handled in this setup. A positive definite covariance matrix is obtained by defining a lower triangular matrix[39]

$$C = \begin{pmatrix} \sigma_{1,1}^{\varepsilon} & 0 & 0 & 0 & 0 \\ \sigma_{2,1}^{\varepsilon} & \sigma_{2,2}^{\varepsilon} & 0 & 0 & 0 \\ \sigma_{3,1}^{\varepsilon} & \sigma_{3,2}^{\varepsilon} & \sigma_{3,3}^{\varepsilon} & 0 & 0 \\ \sigma_{4,1}^{\varepsilon} & \sigma_{4,2}^{\varepsilon} & \sigma_{4,3}^{\varepsilon} & \sigma_{4,4}^{\varepsilon} & 0 \\ \sigma_{5,1}^{\varepsilon} & \sigma_{5,2}^{\varepsilon} & \sigma_{5,3}^{\varepsilon} & \sigma_{5,4}^{\varepsilon} & \sigma_{5,5}^{\varepsilon} \end{pmatrix},$$

and letting the covariance matrix be given by

$$H_t(\psi) = CC'.$$

[37] The dimension of $H_t(\psi)$ is determined by the number of bond yields observed at time t.

[38] m_t is the number of observed yields at time t.

[39] Here, the numbers 1 through 5 refer to the 6-month, 2-year, 3-year, 5-year, and 10-year Treasury yields, respectively.

9.2 Standard errors of the estimated parameters

Following advice from Peter Feldhütter, a colleague at the Copenhagen Business School, the covariance matrix of the estimated parameters are calculated using the numerically more robust method based on the first derivatives of the log likelihood function[40]

$$\widehat{\Sigma} = \frac{1}{T}\Big[\frac{1}{T}\sum_{i=1}^{T}\frac{\partial \log l_t(\widehat{\psi})}{\partial \psi}\frac{\partial \log l_t(\widehat{\psi})}{\partial \psi}'\Big]^{-1}.$$

In calculating this matrix, the first step is to find the step size $\Delta\psi$ that gives the most stable result for the first derivatives. Once this is done, any remaining numerical instability is averaged out by calculating the covariance matrix for 20 different step sizes around the 'optimal' step size using $(0.8 + 0.02i)\Delta\psi$, $i = 1, \ldots, 20$. The final result is obtained by taking the mean $\widehat{\Sigma} = E[\widehat{\Sigma}_i]$.

[40] See Feldhütter and Lando (2005) for another application of this method.

10 Appendix C: The Kalman filter estimation of the CDS models

In the four-factor CDS model with stochastic default and recovery risk described in Section 3.3, the discrete-time version of the transition equation for the state variables is given by

$$\begin{pmatrix} X_t^\lambda \\ X_t^\pi \end{pmatrix} = \Phi_t^0(\psi) + \Phi_t^1(\psi) \begin{pmatrix} X_t^\lambda \\ X_t^\pi \end{pmatrix} + u_t, \quad u_t \sim N(0, V_t(\psi)),$$

where

$$\Phi_t^0(\psi) = \begin{pmatrix} \theta_\lambda[1 - e^{-\kappa_\lambda \Delta t}] \\ 0 \end{pmatrix}$$

and

$$\Phi_t^1(\psi) = \begin{pmatrix} e^{-\kappa_\lambda \Delta t} & 0 \\ 0 & e^{-\kappa_\pi \Delta t} \end{pmatrix},$$

and Δt is the time between observations.

The covariance matrix conditional on $\widehat{X}_{t-1}$ is given by

$$V_t(\psi) = \begin{pmatrix} V[X_t^\lambda | \widehat{X}_{t-1}^\lambda] & 0 \\ 0 & V[X_t^\lambda | \widehat{X}_{t-1}^\lambda] \end{pmatrix},$$

where

$$\begin{aligned} V[X_t^\lambda | \widehat{X}_{t-1}^\lambda] &= \frac{\sigma_\lambda^2}{2\kappa_\lambda} \theta_1 [1 - e^{-\kappa_\lambda \Delta t}]^2 + \frac{\sigma_\lambda^2}{\kappa_\lambda} \widehat{X}_{t-1}^\lambda e^{-\kappa_\lambda \Delta t} [1 - e^{-\kappa_\lambda \Delta t}], \\ V[X_t^\pi | \widehat{X}_{t-1}^\pi] &= \frac{\sigma_\pi^2}{2\kappa_\pi} [1 - e^{-2\kappa_\pi \Delta t}]. \end{aligned}$$

For the three-factor CDS model with constant recovery the above expressions are easily reduced to accommodate the fact that only X_t^λ appears in that model.

Given the estimated values of the interest rate factors (X_t^1, X_t^2), the measurement equation for the CDS rates is a function of the two state variables (X_t^λ, X_t^π)

$$S_t^{CDS}(T) = z(X_t^\lambda, X_t^\pi, T; \psi) + \varepsilon_t, \quad \varepsilon_t \sim N(0, H_t(\varepsilon)).$$

See Appendix B for the details on how this equation is linearized.

As for the measurement noise, the estimation will allow for cross-sectional correlation between the measurement errors while still assuming no serial correlation. This is implemented using the Choleski decomposition

$$H_t(\psi) = CC'$$

where[41]

$$C = \begin{pmatrix} \sigma^{\varepsilon}_{1,1} & 0 & 0 & 0 & 0 \\ \sigma^{\varepsilon}_{2,1} & \sigma^{\varepsilon}_{2,2} & 0 & 0 & 0 \\ \sigma^{\varepsilon}_{3,1} & \sigma^{\varepsilon}_{3,2} & \sigma^{\varepsilon}_{3,3} & 0 & 0 \\ \sigma^{\varepsilon}_{4,1} & \sigma^{\varepsilon}_{4,2} & \sigma^{\varepsilon}_{4,3} & \sigma^{\varepsilon}_{4,4} & 0 \\ \sigma^{\varepsilon}_{5,1} & \sigma^{\varepsilon}_{5,2} & \sigma^{\varepsilon}_{5,3} & \sigma^{\varepsilon}_{5,4} & \sigma^{\varepsilon}_{5,5} \end{pmatrix},$$

Finally, given that both Duffee (1999) and Driessen (2005) find that all the companies in their samples have stationary spread processes under the P-measure, it seems reasonable to assume in the case of Ford Motor Co. that both the default risk factor and the recovery risk factor are stationary processes under the P-measure, and thus the Kalman filter can be started at the unconditional mean and covariance matrix

$$\begin{pmatrix} \widehat{X}^{\lambda}_0 \\ \widehat{X}^{\pi}_0 \end{pmatrix} = \begin{pmatrix} \theta_{\lambda} \\ \theta_{\pi} \end{pmatrix}, \quad \widehat{\Sigma}_0 = \begin{pmatrix} \frac{\sigma^2_{\lambda}\theta_{\lambda}}{2\kappa_{\lambda}} & 0 \\ 0 & \frac{\sigma^2_{\pi}}{2\kappa_{\pi}} \end{pmatrix}.$$

All other details of the Kalman filter remain the same as those described in Appendix B for the estimation of the interest rate model.

[41] Here, the numbers 1 through 5 refer to the 1-year, 3-year, 5-year, 7-year, and 10-year CDS contracts, respectively.

References

[1] Acharya, Viral V., Sreedhar T. Bharath, and Anand Srinivasan, "Does Industry-wide Distress Affect Defaulted Firms? Evidence from Creditor Recoveries", forthcoming Journal of Financial Economics, April 2006.

[2] Altman, Edward I., Brooks Brady, Andrea Resti, and Andrea Sironi, "The Link between Default and Recovery Rates: Theory, Empirical Evidence, and Implications", Journal of Business, Vol. 78, No. 6, 2005.

[3] Bakshi, Gurdip, Dilip Madan, and Frank Zhang, "Understanding the Role of Recovery in Default Risk Models: Empirical Comparisons and Implied Recovery Rates", unpublished working paper, Department of Finance, University of Maryland, 2004.

[4] Berndt, Antje, Rohan Douglas, Darrell Duffie, Mark Ferguson, and David Schranz, "Measuring Default Risk Premia from Default Swap Rates and EDFs", Working paper, Stanford University, November 2005.

[5] Blanco, Roberto, Simon Brennan and Ian W. Marsh, "An Empirical Analysis of the Dynamic Relationship between Investment-Grade Bonds and Credit Default Swaps", Journal of Finace, Vol. 60, No. 5, 2255-2281, October 2005.

[6] Bomfim, Antulio N., "Understanding Credit Derivatives and Related Instruments", Elsevier Academic Press, 2005.

[7] Cheridito, Patrick, Damir Filipović, and Robert L. Kimmel, "Market Price of Risk Specifications for Affine Models: Theory and Evidence", forthcoming Journal of Financial Economics, November 2005.

[8] Christensen, Jens, "The Quadratic and Affine-Quadratic Asset Pricing Models", unpublished working paper, Copenhangen Business School, 2006a.

[9] Christensen, Jens, "Joint Estimation of Default and Recovery Risk: A simulation Study", unpublished working paper, Copenhangen Business School, 2006b.

[10] Covitz, Daniel and Song Han, "An Empirical Analysis of Bond Recovery Rates: Exploring a Structural View of Default", Working Paper, The Federal Reserve Board, Washington DC, December 2004.

[11] Dai, Qiang and Kenneth J. Singleton, "Expectation puzzles, time-varying risk premia, and affine models of the term structure", Journal of Financial Economics, Vol. 63, 415-441, 2002.

[12] Driessen, Joost, "Is Default Event Risk Priced in Corporate Bonds?", Review of Financial Studies, Vol. 18, No. 1, 165-195, January 2005.

[13] Duan, Jin-Chuan and Jean-Guy Simonato, "Estimating and Testing Exponential-Affine Term Structure Models by Kalman Filter", Working paper, CIRANO, October 1995.

[14] Duffee, Gregory R., "Idiosyncratic variation of Treasury bill yields," Journal of Finance, Vol. 51, No. 2, 527-552, June 1996.

[15] Duffee, Gregory R., "The Relation Between Treasury Yields and Corporate Bond Yield Spreads", Journal of Finance, Vol. 53, No. 6, 2225-2241, December 1998.

[16] Duffee, Gregory R., "Estimating the Price of Default Risk", Review of Financial Studies, Vol. 12, No. 1, 197-226, Spring 1999.

[17] Duffee, Gregory R., "Term Premia and Interest Rate Forecasts in Affine Models", Journal of Finance, Vol. 57, No. 1, 405-443, February 2002.

[18] Duffee, Gregory R. and Richard H. Stanton, "Estimation of Dynamic Term Structure Models", working paper Haas School of Business, UC Berkeley, March 2004.

[19] Duffie, Darrell, "Special Repo Rates", Journal of Finance, vol. 51, No. 2, 493-526, June 1996.

[20] Duffie, Darrell, "First-to-Default Valuation", Working paper, Institut de Finance, University of Paris, Dauphine, and Graduate School of Business, Stanford University, May, 1998.

[21] Duffie, Darrell, "Credit Swap Valuation", Financial Analysts Journal, 73-87, January-February 1999.

[22] Duffie, Darrell, Lasse Pedersen, and Ken Singleton, "Modeling Sovereign Yield Spreads: A Case Study of Russian Debt", Journal of Finance , Vol. 58 , No. 1, 119-160, February 2003.

[23] Feldhütter, Peter and David Lando, "Decomposing Swap Spreads", unpublished working paper, Copenhagen Business School, September 2005.

[24] Guha, Rajiv, "Recovery of Face Value at Default: Theory and Empirical Evidence", unpublished working paper, London School of Business, December 2002.

[25] Jarrow, Robert A., "Modelling Fixed Income Securities and Interest Rate Options", McGraw-Hill, 1996.

[26] Lando, David, "On Cox Processes and Credit Risky Securities", Review of Derivatives Research, Vol 2, 99-120, 1998.

[27] Longstaff, Francis A., Sanjay Mithal and Eric Neis, "Corporate Yield Spreads: Default Risk or Liquidity? New Evidence from the Credit-Default Swap Market", Journal of Finace, Vol. 60, No. 5, 2213-2253, October 2005.

[28] Merton, Robert C., "On the pricing of corporate debt: The risk structure of interest rates", Journal of Finance, Vol. 29, 449-470, 1974.

[29] Moody's, "Sovereign Bond Defaults, Rating Transitions, and Recoveries (1985-2002)", Special Comment, February 2003.

[30] Pan, Jun and Kenneth Singleton, "Default and Recovery Implicit in the Term Structure of Sovereign CDS Spreads", unpublished working paper, Stanford University, July 2005.

[31] Schuermann, Til, "What Do we know about Loss Given Default?", In: Altman, Edward I., Andrea Resti, and Andrea Sironi (Eds.), "Recovery Risk - The Next Challenge in Credit Risk Management", Risk Books, London, 3-24, 2005.

[32] Unal, Haluk, Dilip Madan, and Levent Güntay, "Pricing the risk of recovery in default with absolute priority rule violation", Journal of Banking and Finance, Vol. 27, 1001-1025, 2003.

[33] Zhang, Frank X., "What did the Credit Market Expect from Argentina Default? Evidence from Credit Default Swap Data", unpublished working paper, April 2003.