Default Correlation in Reduced-Form Models

Fan Yu*

University of California, Irvine

This Version: September 6, 2003

*Address correspondence to Fan Yu, UCI-GSM, Irvine, CA 92697-3125, E-mail: fanyu@uci.edu.

Default Correlation in Reduced-Form Models

Abstract

Reduced-form models have proven to be a useful tool for analyzing the dynamics of credit spreads. However, some have recently questioned their ability to match the level of empirical default correlation. The key concern appears to be the assumption that defaults are independent conditional on the state variables driving the default intensity. In this paper, I use numerical examples calibrated to recent studies to show that the model-implied default correlation can be quite sensitive to the common factor structure imposed on the default intensity. Therefore, the "inability" of reduced-form models to generate sufficient default correlation has more to do with a restrictive common factor structure than the assumption of conditional independence.

1 Introduction

Reduced-form models have recently been used to study the behavior of credit spreads. In contrast to the structural models pioneered by Merton (1974), this approach treats default as a jump process with an exogenous intensity.[1] As long as the intensity is assumed to be a linear function of affine diffusion state variables, the methodology of Duffie and Singleton (1999) can be used to econometrically identify the intensity from observed prices and spreads, much like the estimation of affine term structure models of default-free bonds. Examples of this approach include Duffee (1999) on corporate spreads and Duffie, Pedersen and Singleton (2003) on sovereign spreads.

Rarely mentioned, however, is the fact that reduced-form models can also be used to study default correlation. At the heart of reduced-form models lies the assumption that multiple defaults are independent conditional on the sample paths of the default intensities. This is the assumption that facilitates the construction of the doubly stochastic Poisson processes of default [also called the "Cox process" in Lando (1998)]. It implies that in such models, default correlation is synonymous with the correlation of the default intensities. Naturally, one could ask whether this type of models can reproduce empirically estimated default correlations.

Despite the apparent scope for further research into this issue, recent studies of default correlation appear to have written off the reduced-form approach. For example, Hull and White (2001) suggest that "...the range of default correlations that can be achieved is limited. Even when there is a perfect correlation between two hazard rates, the corresponding correlation between defaults in any chosen period of time is usually very low. This is liable to be a problem in some circumstances." Schonbucher and Schubert (2001) comment that "...the default correlations that can be reached with this approach are typically too low when compared with empirical default correlations, and furthermore it is very hard to derive and analyze the resulting default dependency structure." Casual conversations with practitioners indicate that this belief is widely held in the industry as well.

In this paper, I argue that the default correlation in reduced-form models can be quite sensitive to the common factor structure imposed on individual default intensities. This is illustrated using numerical examples calibrated to two recent studies—Duffee (1999), where there are two common

[1] Earlier examples of reduced-form credit risk models include Jarrow and Turnbull (1995), Jarrow, Lando and Turnbull (1997), Lando (1998), and Duffie and Singleton (1999).

factors, both of which extracted from Treasury yields, and Driessen (2002), where two additional common factors capture the co-movement of corporate credit spreads. I show that the first case implies a default correlation much lower than empirical observations, while the second case implies comparable or even higher values. Therefore, the alleged low default correlation in reduced-form models may have more to do with an inadequate common factor structure than the assumption of conditional independence.

The rest of the paper is organized as follows. In Section 2, I outline the procedure for imputing default correlation from existing studies of reduced-form models. In Section 3, I apply this procedure to the estimated default intensities from Duffee (1999) and Driessen (2002). I conclude with Section 4.

2 Calibration Procedure

To start, assume the existence of two default stopping times, τ^1 and τ^2, with physical intensities λ^1 and λ^2. The precise meaning of this statement is that

$$1_{\{t \geq \tau^i\}} - \int_0^t \lambda_s^i 1_{\{s \leq \tau^i\}} ds \tag{1}$$

is a martingale under the physical measure.[2]

The intensities are assumed to be $\mathcal{F}_t^X$-adapted where X_t is a vector-valued process representing state variables driving changes in default rates. These can be common macroeconomic factors such as the Treasury term structure level and slope, or firm-specific characteristics such as book to market and leverage ratios. When reduced-form models are used to fit credit spreads, one can often infer the value of X_t from bond prices rather than assuming X_t to be observable. In theory, this latent variables approach can be used with bankruptcy data to estimate the physical intensity, although it has not been applied in this way to my knowledge.

To construct the stopping time with the given intensity, define as in Lando (1998):

$$\tau^i \equiv \inf \left\{ t : \int_0^t \lambda_s^i ds \geq E^i \right\}, \tag{2}$$

where E^i is a unit exponential random variable independent of X_t. This construction satisfies the

[2] Although I focus on modeling correlated defaults under the physical measure, the framework is also applicable to the construction of default times under the risk-neutral measure, in particular, for the valuation of credit derivatives.

martingale property of equation (1), and leads to the following distribution for the stopping time:

$$\Pr\left(\tau^i > t|\mathcal{F}_t^X\right) = \exp\left(-\int_0^t \lambda_s^i ds\right). \tag{3}$$

Furthermore, as in Lando (1994), assume that E^1 is independent of E^2, so that the two default times are conditionally independent given the history of X_t. This defines the class of reduced-form models.

The default correlation between the two stopping times is commonly defined as

$$\rho(t) \equiv \text{Corr}\left(1_{\{\tau^1 \leq t\}}, 1_{\{\tau^2 \leq t\}}\right). \tag{4}$$

Note that the default correlation is a function of the horizon under consideration. Utilizing the above construction, this can be rewritten as

$$\rho(t) = \frac{E\left(y_t^1 y_t^2\right) - E\left(y_t^1\right) E\left(y_t^2\right)}{\sqrt{E\left(y_t^1\right) - \left(E\left(y_t^1\right)\right)^2}\sqrt{E\left(y_t^2\right) - \left(E\left(y_t^2\right)\right)^2}}, \tag{5}$$

where

$$y_t^i = \exp\left(-\int_0^t \lambda_s^i ds\right). \tag{6}$$

This equation shows that in standard reduced-form models, default correlation is completely determined by the individual default intensities.

Various reduced-form models can be distinguished for their choice of the state variables and the processes that they follow. In the models considered below, the intensity is a linear function of X_t and X_t are diffusions in the affine class as defined in Duffie and Kan (1996). Therefore, the expectations above are exponentially linear in X_t, which facilitates computation.

An important issue is that the above framework requires the physical default intensity, while existing studies invariably estimate the risk-neutral intensity implicit in bond prices. The connection between risk-neutral and physical default intensities is studied in Jarrow, Lando and Yu (2001), who show that with a large number of conditionally independent default times, the functional forms of the two intensities are identical in an asymptotic sense. The empirical validation of this conjecture, however, is clouded by the fact that what one calls "credit spread" may contain tax and liquidity components. The empirical analyses in Jarrow, Lando and Yu (2001) and Driessen (2002) provide no definitive answers, but do suggest that the equivalence might hold when non-default components of the credit spread can be more accurately accounted for.

The overall procedure for imputing default correlation from risk-neutral default intensities is as follows:

1. Obtain the risk-neutral intensity function $\widetilde{\lambda}$ estimated for various credit ratings as well as the estimated physical dynamics of the state variables X_t from the empirical literature. Assume the physical intensity function λ to be identical to $\widetilde{\lambda}$.

2. Define the adjusted physical intensity function λ^{adj} as

$$\lambda_t^{\text{adj}} \equiv \lambda_t - \frac{a}{t+b}, \tag{7}$$

   where constant coefficients a and b are determined by minimizing the sum of squared differences between the model-implied conditional default rates and those inferred from historical default experience.[3]

3. Substitute the λ^{adj} function for two different ratings and the physical dynamics of X_t into equation (5) to compute the default correlation between two generic rated issuers.

The static adjustment to the intensity in Step 2 is an attempt to eliminate the effect of liquidity and taxes from credit spreads and implied physical default rates. Currently there is no consensus on how to accurately estimate these components. What we do know, however, is that the tax spread is roughly constant [see Elton, Gruber, Agrawal and Mann (2001)] and the liquidity spread is perhaps decreasing with maturity [see Ericsson and Renault (2001) and Perraudin and Taylor (2002)], both of which are incorporated into (7). The approach here assumes the equivalence of the risk-neutral and physical intensities, and ignores any dynamics of the liquidity spread. It implies that default rates and credit spreads would fluctuate in exactly the same manner. A casual comparison between spreads and default rates in Yu (2002, Figure 1) shows that this may not be a poor assumption. In any case, since the focus is on the role of the assumed common factor structure on the determination of default correlation, the accuracy of the outlined procedure should be of secondary concern.

3 Examples

In this section, I apply the calibration procedure given above to the models of Duffee (1999) and Driessen (2002). Several observations are noted. First, as these models are estimated across credit

[3] The historical conditional default rates are computed from the average cumulative default rates published in Hamilton (2001, Exhibit 41).

ratings using a large cross-section of issuers, the average intensity for each rating is representative of what one would encounter in a well-diversified portfolio. Second, since the intensities are estimated firm-by-firm, these models are consistent with an interpretation based on conditional independence—that default correlation is built into the common variation of the individual default intensities. Third, since the models are estimated using credit spreads, the estimated risk-neutral intensity needs to be transformed into the physical intensity required for computing default correlation.

3.1 Duffee (1999)

Duffee (1999) assumes the risk-neutral intensity to be

$$\widetilde{\lambda}_t = \alpha + \lambda_t^* + \beta_1 \left(s_{1t} - \overline{s_{1t}} \right) + \beta_2 \left(s_{2t} - \overline{s_{2t}} \right), \tag{8}$$

where s_{1t} and s_{2t} are default-free factors inferred from Treasury yields through the short rate model $r_t = \alpha_r + s_{1t} + s_{2t}$, and $\overline{s_{1t}}$ and $\overline{s_{2t}}$ are the respective sample means. The firm-specific factor λ_t^* and coefficients α, β_1 and β_2 are in turn inferred from the prices of corporate bonds issued by a given firm, taking the Treasury term structure dynamics as given.

All factors are assumed to be square-root diffusions. The independent Treasury factors are common to every firm and the firm-specific factors are independent across firms. This specification is consistent with Duffee's firm-by-firm estimation approach. In a setting with conditionally independent default times, default correlation is generated by the dependence on the common factors through the β coefficients. On the other hand, the firm-specific factor λ_t^* should reduce the default correlation. Intuitively, this is because the firm-specific component contributes little to the covariance while increasing the variance of default rates.

The physical dynamics of each square-root diffusion can be designated by the triple (κ, θ, σ), where κ is the speed of mean-reversion, θ is the long-run mean and σ is the volatility of the process. To calibrate the liquidity adjustments, one also needs the initial values of the state variables. For the firm-specific factors, these are taken to be their sample means. For the Treasury factors, since their sample means are not provided by Duffee, I use their long-run means instead.[4] Table 1 summarizes the estimates from Duffee (1999) and the liquidity adjustments.

[4] The precise values of $\overline{s_{1t}}$ and $\overline{s_{2t}}$ are not essential to the calculation as they are constants that can be compounded into the liquidity adjustment term.

Using the necessary inputs from Table 1, I compute the default correlation between generic rated issuers. The results, presented in Table 2, bear two distinctive patterns. First, adjusting for liquidity and tax effects appears to be very important when inferring default correlations from credit spreads. Moving from Panel B to Panel A, default correlation can increase by more than ten-fold in some instances. This is because a major part of the short-term credit spread is due to liquidity and state taxes. The adjustments to the intensity, for instance, can be more than 100 basis points at zero maturity. A deterministic reduction of the magnitude of the default intensity increases the proportion of the intensity that is stochastic, thereby increasing default correlation. This effect is the strongest for short-term default correlation between high-quality issuers. A second pattern is that default correlation increases monotonically with maturity. This is not surprising given that default over a very long horizon is almost a certainty.

To better gauge the results, I compare them with those of Zhou (2001) and Lucas (1995). Zhou's calculation is based on a first passage time approach in the Black and Cox (1976) structural framework, while Lucas' results are obtained from historical default data. A comparison between Table 2 and these two studies shows that the procedure used here slightly overestimates the default correlations for short horizons, while significantly underestimate the default correlations for longer horizons. For references, the results from Zhou and Lucas are reproduced here as Table 5.

The overestimation at short horizons is unlikely to be a serious concern here. For Lucas' historical estimation, there are simply not enough observations to pin down the default correlations precisely for, say, investment-grade issuers at a one-year horizon. The quality of the historical estimates for lower-rated issuers at longer horizons, however, is much higher due to the larger sample size. Hence the underestimation at longer horizons is a major concern that demands our full attention.

One promising explanation of the underestimation seems to be the insufficient specification of the common factor structure in Duffee's model. In the standard reduced-form framework, default correlation is attributed wholly to the common factors in the intensity function. It is critical, then, that the specification of the intensity function adequately capture the sources of common variation in yield spreads. In Duffee, however, the only common factors are related to the Treasury term structure. To the extent that other important common factors are omitted, using Duffee's model is likely to underestimate the default correlation. Indirectly supporting this conjecture, Duffee shows that the market prices of risk associated with the firm-specific factors are statistically significant, and that the firm-specific factors are also correlated across firms, suggesting an active role for "missing" common factors.

3.2 Driessen (2002)

The inadequacies of Duffee's specification are dealt with by Driessen (2002), who assumes the risk-neutral intensity as

$$\widetilde{\lambda}_t = \alpha + \beta_r r_t + \beta_v v_t + \gamma_1 F_{1t} + \gamma_2 F_{2t} + G_t. \tag{9}$$

In this setup, r_t and ν_t are short-rate factors in an affine Treasury term structure model, F_{1t} and F_{2t} are factors common to every firm, and G_t is a firm-specific factor. Both the common factors and the firm-specific factors are assumed to follow independent square-root diffusions, and hence can be specified by the triple (κ, θ, σ). As evidence of improvement over Duffee (1999), Driessen shows that the estimated market prices of risk for the firm-specific factors are close to zero, and that the firm-specific factors have negligible cross-firm correlations.

For the purpose of calibrating default correlation, I ignore the short-rate factors. The subsequent results show that the inclusion of the two common factors elevates the default correlation dramatically relative to the results obtained using Duffee's model. The portion of the default correlation generated by short-rate dependence is likely to be relatively small.

Table 3 presents Driessen's estimates and the liquidity adjustments. Since Driessen does not aggregate his estimates for the firm-specific factor by credit rating, his median estimates are used for all three ratings. I also divide the parameter estimates of the instantaneous spread by one minus the recovery rate, assumed to be 44 percent for all three ratings, to obtain estimates of the default intensity function. Furthermore, I set the initial values of all factors to be their long-run means.

The default correlations, presented in Table 4, are much higher than those in Table 2 with or without intensity adjustments. In fact, the default correlation at long horizons now appears to be too high relative to the estimates of Lucas (1995).[5] The key point is clear—by incorporating more of the common variation in credit spreads, the default correlation implied by a reduced-form model with conditionally independent defaults can easily be made higher. At a minimum, the evidence

[5] A more recent study by De Servigny and Renault (2002) find speculative-grade default correlation at the five-year horizon that could easily exceed 10 percent for some industries.

here suggests that more work is needed before dismissing such models as unsuitable for describing generic default correlations.

4 Conclusion

This paper examines the popular notion that reduced-form models based on conditional independence cannot generate empirically observed levels of default correlation. A simple procedure is undertaken to compute default correlations implied from existing empirical studies of intensity-based credit risk models. This exercise suggests that the root of the problem is an insufficient specification of the common factor structure of the default intensity, and not the reduced-form framework per se. In fact, with just two common factors driving the intensity, a reduced-form model estimated from individual credit spreads implies too much, rather than too little, default correlation.

References

Black, F., and J. C. Cox, 1976, Valuing corporate securities: Some effects of bond indenture provisions, *Journal of Finance* 31, 351-367.

De Servigny, A., and O. Renault, 2002, Default correlation: Empirical evidence, Working paper, Standard and Poor's.

Driessen, J., 2002, Is default event risk priced in corporate bonds? Working paper, University of Amsterdam.

Duffee, G., 1999, Estimating the price of default risk, *Review of Financial Studies* 12, 197-226.

Duffie, D., and R. Kan, 1996, A yield-factor model of interest rates, *Mathematical Finance* 6, 379-406.

Duffie, J. D., L. Pedersen, and K. J. Singleton, 2003, Modeling sovereign yield spreads: A case study of Russian debt, *Journal of Finance* 58, 119-160.

Duffie, J. D., and K. J. Singleton, 1999, Modeling term structures of defaultable bonds, *Review of Financial Studies* 12, 687-720.

Elton, E., M. Gruber, D. Agrawal, and C. Mann, 2001, Explaining the rate spreads on corporate bonds, *Journal of Finance* 56, 247-277.

Ericsson, J., and O. Renault, 2001, Credit and liquidity risk, Working paper, McGill University.

Hamilton, D. T., 2001, Default and recovery rates of corporate bond issuers: 2000, Moody's Special Comment.

Hull, J., and A. White, 2001, Valuing credit default swaps II: Modeling default correlations, *Journal of Derivatives* 8(3), 12-22.

Jarrow, R. A., D. Lando, and S. M. Turnbull, 1997, A Markov model for the term structure of credit risk spread, *Review of Financial Studies* 10, 481-523.

Jarrow, R., D. Lando, and F. Yu, 2001, Default risk and diversification: Theory and applications, Working paper, University of California, Irvine.

Jarrow, R. A., and S. M. Turnbull, 1995, Pricing derivatives on financial securities subject to credit risk, *Journal of Finance* 50, 53-85.

Lando, D., 1994, Three essays on contingent claims pricing, Ph.D. dissertation, Cornell University.

Lando, D., 1998, On Cox processes and credit risky securities, *Review of Derivatives Research* 2, 99-120.

Lucas, D., 1995, Default correlation and credit analysis, *Journal of Fixed Income* 4(4), 76-87.

Merton, R. C., 1974, On the pricing of corporate debt: The risk structure of interest rates, *Journal of Finance* 29, 449-470.

Perraudin, W., and A. Taylor, 2002, Liquidity and bond market spreads, Working paper, Birbeck College.

Schonbucher, P., and D. Schubert, 2001, Copula-dependent default risk in intensity models, Working paper, Bonn University.

Yu, F., 2002, Modeling expected return on defaultable bonds, *Journal of Fixed Income* 12(2), 69-81.

Zhou, C., 2001, An analysis of default correlation and multiple defaults, *Review of Financial Studies* 14, 555-576.

	s_1	s_2	Aa	A	Baa
κ	.474	.032	.186	.242	.212
θ	1.003	.060	.00499	.00559	.00628
σ	.0134	.0449	.074	.078	.059
α			.00637	.00739	.00961
β_1			-.077	-.090	-.171
β_2			.001	-.033	-.006
$\overline{\lambda_t^*}$			.00440	.00537	.00864
a			1.601	.519	.204
b			154.253	41.865	13.105
adj_0			.0104	.0124	.0156
adj_5			.0101	.0111	.0113
adj_{10}			.0097	.0100	.0088

Table 1: **Estimates from Duffee (1999)**. The Treasury factors s_1 and s_2 and the firm-specific factors are specified through the triple (κ, θ, σ). The intensity parameters are α, β_1 and β_2. The liquidity adjustment parameters are a and b. $\overline{\lambda_t^*}$ denotes the sample mean of the firm-specific factor. adj_n denotes the liquidity adjustment to the conditional default rate (CDR) at the n-year maturity. The firm-specific factor and the intensity function are specified for three rating categories, Aa, A and Baa.

Panel A: Adjusted intensity

	One Year			Two Years			Five Years			Ten Years		
	Aa	A	Baa	Aa	A	Baa	Aa	A	Baa	Aa	A	Baa
Aa	.05	.06	.05	.13	.14	.12	.30	.28	.27	.36	.30	.33
A	.06	.07	.05	.14	.17	.13	.28	.36	.27	.30	.48	.31
Baa	.05	.05	.04	.12	.13	.11	.27	.27	.25	.33	.31	.31

Panel B: Unadjusted intensity

	One Year			Two Years			Five Years			Ten Years		
	Aa	A	Baa	Aa	A	Baa	Aa	A	Baa	Aa	A	Baa
Aa	.00	.00	.00	.01	.01	.01	.02	.02	.03	.03	.03	.05
A	.00	.00	.00	.01	.01	.01	.02	.03	.04	.03	.06	.06
Baa	.00	.00	.01	.01	.01	.02	.03	.04	.06	.05	.06	.09

Table 2: **Default correlation inferred from Duffee (1999)**. Values are in percentages, given for three rating categories, Aa, A and Baa, and horizons of one, two, five and ten years. Panel A includes liquidity and tax adjustments, while Panel B does not.

	F_1	F_2	Aa	A	Baa
κ	.030	.490	.017	.017	.017
θ	.005	.005	.00655	.00655	.00655
σ	.016	.046	.017	.017	.017
α			.00271	.00409	.00400
γ_1			2.991	3.520	5.707
γ_2			.391	.714	.988
a			.5008	.2604	.0450
b			56.590	25.084	5.001
adj_0			.0088	.0104	.0090
adj_5			.0081	.0087	.0045
adj_{10}			.0075	.0074	.0030

Table 3: **Estimates from Driessen (2002)**. The common factors F_1 and F_2 and the firm-specific factors are specified through the triple (κ, θ, σ). The intensity parameters are α, γ_1 and γ_2. The liquidity adjustment parameters are a and b. adj_n denotes the liquidity and tax adjustment to the conditional default rate (CDR) at the n-year maturity. The firm-specific factor and the intensity function are specified for three rating categories, Aa, A and Baa.

Panel A: Adjusted intensity

	One Year			Two Years			Five Years			Ten Years		
	Aa	A	Baa	Aa	A	Baa	Aa	A	Baa	Aa	A	Baa
Aa	.85	1.08	.76	2.83	3.22	2.39	12.01	11.69	9.63	32.13	28.41	25.83
A	1.08	1.41	.98	3.22	3.74	2.75	11.69	11.50	9.44	28.41	25.25	22.92
Baa	.76	.98	.68	2.39	2.75	2.03	9.63	9.44	7.76	25.83	22.92	20.84

Panel B: Unadjusted intensity

	One Year			Two Years			Five Years			Ten Years		
	Aa	A	Baa	Aa	A	Baa	Aa	A	Baa	Aa	A	Baa
Aa	.04	.05	.08	.17	.19	.31	.93	1.04	1.68	3.16	3.48	5.67
A	.05	.06	.10	.19	.22	.35	1.04	1.17	1.89	3.48	3.85	6.27
Baa	.08	.10	.15	.31	.35	.56	1.68	1.89	3.05	5.67	6.27	10.23

Table 4: **Default correlation inferred from Driessen (2002)**. Values are in percentages, given for three rating categories, Aa, A and Baa, and horizons of one, two, five and ten years. Panel A includes the liquidity and tax adjustments, while Panel B does not.

Panel A: Zhou (2001)

	One Year			Two Years			Five Years			Ten Years		
	Aa	A	Baa	Aa	A	Baa	Aa	A	Baa	Aa	A	Baa
Aa	0.00	0.00	0.00	0.00	0.00	0.01	0.59	0.92	1.24	4.66	5.84	6.76
A	0.00	0.00	0.00	0.00	0.02	0.05	0.92	1.65	2.60	5.84	7.75	9.63
Baa	0.00	0.00	0.00	0.01	0.05	0.25	1.24	2.60	5.01	6.76	9.63	13.12

Panel B: Lucas (1995)

	One Year			Two Years			Five Years			Ten Years		
	Aa	A	Baa	Aa	A	Baa	Aa	A	Baa	Aa	A	Baa
Aa	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.0	1.0	1.0	2.0	1.0
A	0.0	0.0	0.0	0.0	0.0	0.0	1.0	1.0	1.0	2.0	2.0	1.0
Baa	0.0	0.0	0.0	0.0	0.0	0.0	1.0	1.0	0.0	1.0	1.0	0.0

Table 5: **Default correlations from Zhou (2001) and Lucas (1995)**.