Journal of Economic Dynamics & Control 61 (2015) 334–349

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

Estimation of correlations in portfolio credit risk models based on noisy security prices

Mathieu Boudreault [a,*], Geneviève Gauthier [b], Tommy Thomassin [c]

[a] Department of Mathematics, Université du Québec à Montréal, Box 8888, Montreal, Quebec, Canada H3C 3P8
[b] Department of Decision Sciences, HEC Montréal, 3000 Côte-Sainte-Catherine Road, Montreal, Quebec, Canada H3T 2A7 and GERAD
[c] Corporate Actuarial Department, Standard Life Canada, 1245 Sherbrooke Street West, Montreal, Quebec, Canada H3G 1G2

ARTICLE INFO

Article history:
Received 11 February 2015
Received in revised form
21 September 2015
Accepted 7 October 2015
Available online 22 October 2015

JEL classification:
C13
C18
C51
C58

Keywords:
Credit risk
Correlation
Estimation
Noise
Maximum likelihood
Unscented Kalman filter (UKF)

ABSTRACT

Portfolio credit risk models are very often constructed with correlation matrices serving as proxies for interrelations in the creditworthiness of each company. In addition to the size of the matrix, estimation of correlation is also complicated by the fact that defaults are rare and credit-sensitive securities such as stocks, bonds and credit default swaps (CDS) are noisy. Therefore, we present in this paper an estimation approach based on credit-sensitive instruments that accounts for noise and is highly parallelizable, the latter being a very important feature for large portfolios in finance. A simulation study shows that the method is reliable and has better statistical properties when benchmarked against other correlation estimators. In an empirical study based on the CDS premiums and stock prices of 225 firms listed on the CDX North American indices, we analyze the correlations computed using numerous approaches. Overall, we find that ignoring noise severely underestimates correlations, whereas equity correlation is poorly related to the best correlation estimates inferred from the CDS market.

1. Introduction

The solvency of many financial institutions relies on an appropriate assessment of credit risk. For example, banks and insurance companies need to set aside sufficient funds to protect against potential credit losses on the asset side of their balance sheet. To compute portfolio risk measures and the capital requirement, risk managers turn to portfolio credit risk models. A crucial input in these models is indubitably the dependence between the assets that make up this portfolio, which creates clusters of defaults (Lucas, 1995; Das et al., 2004, 2007). This dependence, or what is commonly known as correlation in a Gaussian framework, is not straightforward to estimate, and small changes in this quantity can have tremendous consequences for portfolio risk management.

There is a wide range of portfolio frameworks available to the modeller. Professional models such as Moody's-KMV (by Moody's), CreditMetrics (by RiskMetrics, formerly owned by J.P Morgan) and CreditRisk+ (Crédit Suisse) tend to focus on loss estimation for capital requirements (for a discussion on portfolio credit risk, see Crouhy et al., 2000 and McNeil et al.,

* Corresponding author. Tel.: +1 514 987 3000x8560.
E-mail address: boudreault.mathieu@uqam.ca (M. Boudreault).

http://dx.doi.org/10.1016/j.jedc.2015.10.001
0165-1889/

2005). Many multi-name extensions of structural, reduced-form and hybrid credit risk models also include interconnections between firms. Popular approaches involve using copulas or common factors[1] (see Li, 2000; Schönbucher and Schubert, 2001; Frey and McNeil, 2003; Andersen and Sidenius, 2004; Laurent and Gregory, 2005; Hull et al., 2010), common jumps in the default process (as in Duffie and Singleton, 1999 and Duffie and Gârleanu, 2001) and other default contamination/contagion mechanisms (e.g. Davis and Lo, 2001; Jarrow and Yu, 2001; Giesecke and Weber, 2004 and Giesecke, 2004).

In the large majority of these portfolio models, the joint asset or default intensity dynamics is specified as a function of multivariate Gaussian random vectors for which the estimation of a correlation matrix is required. Because (market) asset values and default intensity are not available to market participants, the direct estimation of underlying correlations is not feasible. Various estimation techniques are found in the literature and the main goal of this paper is to compare them, document some caveats associated with some of them and propose an estimation approach that solves these issues.

There are basically two sources of data that can be exploited to estimate correlations: default counts (or transitions) and credit-sensitive security prices. To estimate correlations from counts, the portfolio model requires a time series of annual default counts (or transitions) in various categories, which can be ratings, industry sectors, etc. Lucas (1995), CreditMetrics (see Gupton et al., 1997), Gordy (2000), De Servigny and Renault (2002), McNeil and Wendin (2007) and Duffie et al. (2009) used counts (or transitions) to estimate their respective models.

Another approach is to infer the correlation parameters from credit-sensitive security prices. When a portfolio structural credit risk model is used, one needs to estimate the correlation in the evolution of each company's assets, which is also known as asset correlation. However, the true market value of the assets, which is the desired input in these models, is not observed by investors (Duffie and Lando, 2001 and Jarrow and Protter, 2004). To overcome this issue, many authors use a correlation computed directly from stocks (equity correlation) (Gupton et al., 1997 and De Servigny and Renault, 2002). However, this approach is not theoretically well-founded as in many frameworks, equity value can be interpreted as a type of call option (plain vanilla in Merton, 1974, down-and-in in Brockman and Turtle, 2003) on the company's assets and it is not clear how equity correlation preserves asset correlation. In a multivariate version of a reduced-form model, one needs to estimate the correlation between each firm's default intensity. Since it is not directly observed, a similar problem arises with these models.[2]

Duellmann et al. (2010) reviewed the literature for the various estimation approaches and reported correlations in the range of 0.5–26% depending on the specific model, industry and method. Given the importance of correlation estimates for risk-management purposes, a rigorous investigation of estimation techniques and their effectiveness is necessary. In a Merton-type (1974) portfolio model, they studied the performance of default counts and stock prices to estimate the asset correlation. They found that: (1) correlations estimated with default rates are severely biased downward, and (2) asset correlations estimated with a Duan (1994, 2000)-type approach[3] performs better (lower bias and RMSE) than directly inferring asset correlation from stocks (taking equity correlation as equal to asset correlation).

Although variations in the firm's fundamentals (such as creditworthiness) are a major driver of the variability of security prices, other components may affect prices as well. In addition to model error, it is documented that transient shocks (Aït-Sahalia, 2004) and trading noise (Duan and Fulop, 2009; Huang and Yu, 2010) may also affect the observed price and thus influence the parameter estimates of a credit risk model. Duan and Fulop (2009) found that asset volatility in Merton's (1974) model can be overstated when noise is not filtered out. Ignoring noise in that model is similar to assuming that all variations in stock price were due to changes in the market value of the firm's assets.

The recent literature addresses the discrepancy between the model price and the observed price using filtering techniques. Indeed, given a model, the filtering algorithm is able to disentangle changes in fundamentals from noise, especially when there is a whole cross-section of security prices at each given date. Filtering techniques mainly come from the signal processing literature and its popularity is constantly growing in finance when using stock options (see Carr and Wu, 2006 or Christoffersen et al., 2010) or interest rates (Christoffersen et al., 2014) for example. In the estimation of single-name credit risk models with observational errors, Duan and Fulop (2009) proposed an Auxiliary Particle Filter (APF), Huang and Yu (2010) proposed a Monte Carlo Markov Chain algorithm, whereas Boudreault et al. (2013) used an unscented Kalman filter (UKF).

To the best of our knowledge, no author has investigated estimation of correlations when security prices are noisy. Therefore, following from Duan and Fulop (2009) and Duellmann et al. (2010), the contributions of our paper are twofold. First, we present in this paper an algorithm to estimate correlations from noisy credit-sensitive security prices. We show numerically that the method is statistically consistent, highly parallelizable and has better statistical properties than other correlation estimators. Second, we perform an extensive empirical study of correlation estimates with the 225 firms of the CDX North American portfolio. Broadly speaking, we find that ignoring noise severely underestimates correlations, whereas equity correlation is poorly related to the best correlation estimates inferred from the CDS market.

[1] Some factor models can be rewritten as a type of copula.

[2] An alternative to these approaches would be to find the level of correlation that matches theoretical and observed prices of basket credit derivatives (such as Collateralized Debt Obligations (CDOs)). This is known as implied correlation and is similar in spirit to the concept of implied volatility in the Black–Scholes' model. This approach has numerous caveats and lost much of its popularity in the 2008 CDO market collapse.

[3] The derivative price is assumed to be a function of asset value. Inverting that function, one can construct the likelihood function associated with the observed sample to retrieve all parameters, including those associated with the latent variable.