CreditRisk + Crisis = (solutions)^2

Yimin Liu Noel McWilliam

In 1997 Credit Suisse Financial Products released CreditRisk$^+$ its seminal offering for the determination of losses due to default within a portfolio of obligors. The simplicity and elegance of the mathematical foundations of CreditRisk$^+$ afford it a particular ease of use and degree of transparency not shared by other more elaborate credit loss models. This aspect of its formulation ensures that data requirements and user inputs are minimal so that it can easily be deployed as a stand alone component or interfaced with other products. Analyticity is a key feature, giving rise to closed-form solutions for loss distributions that may be computed rapidly and unambiguously, allowing for a level of dissection, drill-down and analysis that would be too computationally intensive for brute-force Monte-Carlo methods. For these and other reasons CreditRisk$^+$ has become increasingly popular, especially in environments where speed and performance are critical.

The State Of Modelling

Of course no model is perfect. In the intervening decades since its release much work has focused on addressing some of the issues presented by the original model framework. One area that has received particular attention regards the manner in which default correlations are generated: by allowing obligor default probabilities to depend on a shared set of "sector risk factors" one can impose dependencies between defaults. Unfortunately the range of attainable default correlations is restricted by the assumption of independence between sector values, which is clearly at odds with empirical observation and would otherwise lead to an an underestimation of risk.

Subsequent additions to the model have relaxed this assumption, and although many solutions have been proposed — such as the "Hidden Gamma" model we later reference — none permit a good range of sector (and hence default) correlations in an analytical setting.

An important aspect of the model framework which has not received deserved attention is the issue of extreme value risk. Recent events have highlighted the failure of many models to account for large and impactful movements in market and credit factors and have underscored the clear need to identify and quantify these risks.

In this article we outline a state-based approach that not only has the ability to represent extreme risks but also presents a flexible means to addresses the problem of sector dependence, which for the first time permits the specification of negative sector correlations. In doing so we touch upon many issues that have come to the fore since the credit crisis. We demonstrate the ability of our construct to: mitigate the pro-cyclical nature of myopic model calibration practices, address the "through-the-cycle" nature of default probabilities, readily obtain conditional VaR estimates (not to be confused with CVaR), identify previously undisclosed concentration risks via tail-dependence analysis and to inform realistic stress-testing methodologies. All of our results are formulated in a unified model consistent with the analytic framework of CreditRisk$^+$.

The Modelling Of States

Any casual reader of financial literature will be aware of the pervasive notion of 'market state'. Words such as bullish and bearish embed within them common perceptions of distinct operational modes of the market each characterised by different risk-return profiles and underpinned by associated economic factors and psychological/behavioural regimes.

It might seem surprising therefore that despite general recognition of these concepts relatively few financial models attempt to explicitly represent them. Fundamental ignorance of market dynamics in risk management, of its cyclical nature and the feedback mechanisms existing between market stresses has played an important role in the current economic downturn.

To see this one need only consider the manner in which market models are typically calibrated. By using only the most recent data to determine model parameters institutions ensure that pricing formulae and risk statistics reflect the current market situation. Inevitably in doing so past states are "forgotten" and as such cannot be reflected in model predictions. Short-sighted calibration techniques positively reinforce perverse market mechanisms since in peaceful times projections are unduly optimistic, facilitating inappropriate and unmeasured risk taking and thereby laying the foundations of the next bubble. In times of crisis however such techniques lead to unnecessarily pessimistic projections that may prevent investment and stifle growth precisely when it is needed most. Model myopia and pro-cyclicity are two sides of the same coin.

In part the dissemination of state-based models has been impeded by associated computational difficulties. Suppose for example that we wish to model a set risk factors (equities, interest rates, commodities etc) and assume for simplicity that each risk factor can occupy one of two states: "peaceful" (e.g. bullish) or "stressed" (e.g. bearish). When measured over a time interval of sufficient length returns will sometimes be peaceful and sometimes stressed, so that the overall statistical distribution can be thought of as a mixture of component distributions each having possibly distinct means, volatilities etc reflecting the different state characteristics.

Calibration of a mixture is relatively simple when considering risk factors in isolation, however the problem quickly becomes intractable in multiple-dimensions since the addition of each new factor doubles the number of stressed/peaceful combinations and hence parameters within the model. Even for a small set of 20 factors there are over $2^{20} = 1,048,576$ values to estimate, which would otherwise render calibration via established methods all but impossible.

Fortunately one can address this problem using pattern analysis techniques (see McWilliam 2009, and McWilliam and Loh 2008) which are able to identify key stress patterns within and between markets irrespective of the size of the system under consideration. Thus we are able to determine otherwise undisclosed concentration risks [1] in a manner consistent with observed extreme values within each risk factor distribution and pair-wise extreme value risks as implied, for example, from empirical correlations.

The ability to represent multiple tail-dependency/concentration risks is important — the credit impact of a steady stream of independent and heterogeneous defaults or downgrades is likely to have very different operational consequences than that of rare but synchronized large losses occurring in times of systemic crisis. For this reason our model framework has implications within many areas of finance and has particular relevance to pricing tail-sensitive instruments (see McWilliam and Loh, 2008).

With specific reference to CreditRisk^{+} we make use of our techniques and extend the hidden Gamma model of Giese (2003). In keeping with this approach we assume the aforementioned sector variables are a direct sum of a random Gamma variable specific to the sector and an unobserved "hidden" Gamma variable common across the sectors. However we assume this property to hold *conditional the state of the market* so that now the overall (i.e. unconditional) sector distribution will be a mixture of Gammas in which each component Gamma characterises the statistical properties of the sector variables when the market occupies a particular state.

As we later demonstrate the flexibility of this construct permits a much closer fit to empirical sector data than previously possible. Our approach also has significant implications for the derivation of conditional risk statistics and the generation of scenarios for use within stress testing.

[1] That is, risks from simultaneous, grouped stresses arising within asset classes, geographical regions or indeed any subset of risk factors

Stress Testing

A common criticism sounded in the aftermath of the credit crunch is that stress testing prior to the crisis failed to prevent its occurrence and therefore could not have been fit for purpose. Of course one may counter with the adage "history never repeats itself" and note that few people could have had the insight to foresee the magnitude of the looming crisis.

Given the clear difficulties in predicting the specific nature of rare market events, it is tempting to question the usefulness of such techniques especially if the warnings they provide are not properly understood or heeded. Despite these misgivings the irrational argument favouring complete ignorance over partial information belies the truth that stress testing is more of a subjective art-form than it is a science. Nonetheless it is apparent that the methods employed in these analyses could be improved upon.

A particular difficulty in this regard is in specifying a realistic set of worst-case scenarios to test. While this inevitably relies on the intuition of the risk manager and his or her understanding market dynamics both in peaceful times and in times of crisis, one must also concede that the finite bounds of human processing capabilities and comprehension are dwarfed by the sheer volume of information available in the market. Thus one would do well to augment this process with systematic methods that provide analysis and insight into the workings of the market where human minds cannot.

In this respect the output of the pattern analysis routine provides a natural view of market stresses which are not influenced by expectation or personal bias nor by the limits of human ability to assimilate and condense high dimensional data into a comprehensible format.

Our analysis generates a set of stress scenarios together with an indication of the probability, magnitude and direction of the stress. Each stress scenario takes the form of a vector of integer values indicating the state of each of the risk factors. If "0" indicates a peaceful state and "1" a stressed state then a scenario ψ_i involving three risk factors $X_j; j = 1, \ldots, 3$ might have the following form:

$$\psi_i = \left[\begin{array}{c} 1 \\ 1 \\ 0 \end{array} \right] \left. \begin{array}{c} \\ \\ \\ \end{array} \right\} \begin{array}{c} X_1 = stress \\ X_2 = stress \\ X_3 = peace \end{array}$$

To each scenario is accorded a corresponding probability $0 < w_i \leq 1; i = 1, \ldots, m$ (where m is the number of identified stress state combinations) which is also available as an output of the pattern analysis together with a definition of the characteristics of stress and peace for each risk factor:

$$X_j \sim \left\{ \begin{array}{ccc} \textbf{mean} & \textbf{variance} & \\ \mu_0 & \sigma_0^2 & peace \\ \mu_1 & \sigma_1^2 & stress. \end{array} \right.$$

In the context of the moment matching method discussed herein, the above risk factor stress characteristics are chosen to match the moments of each sector distribution, while the set of ψ_i and w_i are chosen to match the empirical correlation structure.

One may then select desired stress pattern(s) feeding them back as inputs into the mixture model to produce a credit loss distribution that would be apparent if the scenario(s) were to dominate in the market. Since one also has access to the individual stress characteristics, one may modify these values to test the corresponding sensitivity of the loss distribution.

Through-the-Cycle Defaults

A subject of further post crunch criticism regards the manner in which default probabilities are calculated. Since they are determined on a periodic basis they are in some sense averages through the business cycle within the period. As such it is claimed they principally describe an historic set of states of the market; having limited predictive power.

A point of interest in the state-based view of the market is that within this framework it is possible to use the current realisation of the sector risk values/market returns and determine the probability of the market occupying any of the stress-state combinations identified in the pattern analysis. Hence we may obtain a conditional set of scenario probabilities that can be used as inputs to the mixture model.

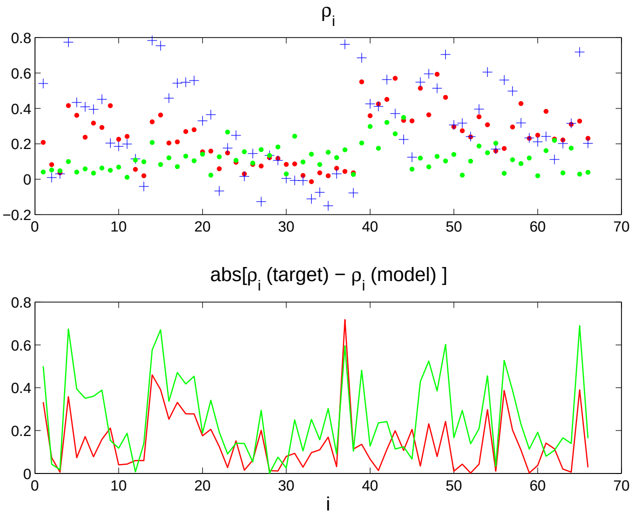

Figure 1: Model (green: hidden Gamma, red: Gamma Mix) Vs. empirical sector correlations (blue) & Errors.

While the unconditional probabilities are chosen to match the moments of each risk factor — so that the mean default probability is equal to the through-the-cycle mean — the conditional probabilities will infer a set of modified mean default probabilities conditioned upon and immediately responsive to market factors.

Test Case

To demonstrate the practical benefits of a state-based model we provide a comparative analysis of model predictions between the "standard" (independent sectors), hidden Gamma and Gamma mixture models. Our portfolio of 5767 obligors having total exposure of \$1bn is constructed using default probabilities, sector weights and rating proportions in keeping with composition characteristics derived from real data[2]. Model calibration is performed using real default and sector data quoted annually between 1970 and 2006. Data Source: Moody's "Corporate Default and Recovery Notes", February 2007, Report No 102071.

	Skew	Kurt	Corr
HG	2.30	9.80	0.26
GMix	0.38	2.59	0.14

Table 1: Comparison of average moment errors for hidden Gamma (HG) and Gamma mixture (Gmix) fits to Moody sector data.

Any meaningful test should also offer a comparison of real-world calibration techniques pertinent to the models under consideration. For this reason we calibrate both standard (Std) and hidden gamma (HG) models to the means, variances and correlations of normalised exponentially weighted (Exp) with exponent coefficient $\lambda = 0.07$ and non-weighted (NoW) data. All calibrations use data from 1970 up until the time at which VaR is measured. The values for the mixture model (Gmix) use non-weighted data, but with stress-scenario probabilities conditioned on the current sector realisation. (In this sense the resulting VaR are actually conditional VaR values.)

It is hoped that these measures will ensure our analysis provides a genuine picture of model performance within as realistic a setting as possible.

Correlation, VaR And Conditional VaR

Figure 1 and Table 1 illustrate the results of our moment matching and pattern analysis techniques in matching the skewness, kurtosis and correlation values for 12 sector time series. Clearly the mixture model substantially improves the parity between model and empirical correlations and higher order moments. Average (absolute) correlation errors of the mixture are roughly half the hidden Gamma model with almost all correlation values in the lat-

[2] In summary: the number and proportion of obligors was chosen in accordance with 2006 Moody rating data. Average default values for each rating (Aa – C) were determined using 37 years of historical default data (1970-2006) provided by the same. Each obligor was assigned a default in accordance with its rating. Sector weights for each class were determined via multi-linear regression ($R^2 = 88\%[Baa] - 44\%[Caa - C]$) of the historical rate against the set of sector specific default rates on the same period. Obligors were assigned weights appropriate to the rating. Loss given defaults were randomly assigned from a uniform distribution.

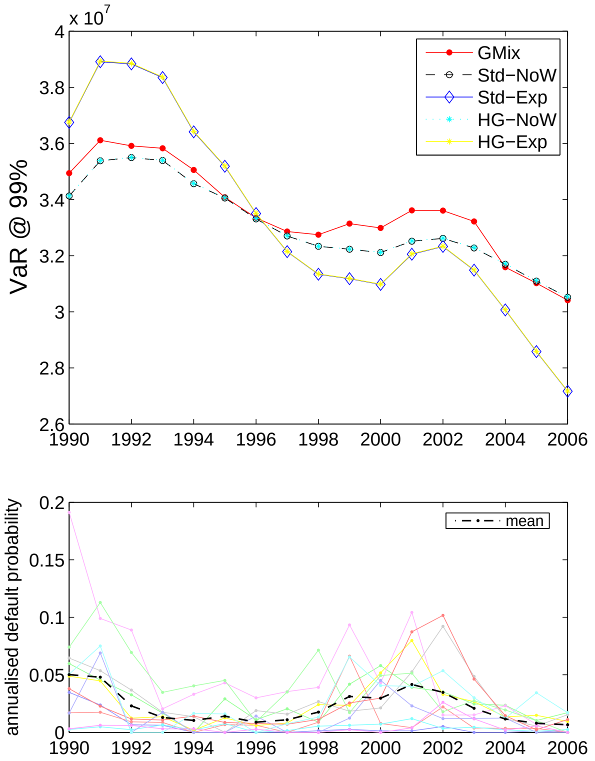

Figure 2: Top: comparison of 99 % VaR 1990-2006. Bottom: Raw sector data. Note: GMix is immediately responsive to sector variation, time-lags in competing methods. Exp VaR overcompensates. Sectors: Banking, Products, Energy, Finance, Hotels, Gaming, Industrial, Media, Misc, Retail, Tech, Transport, Utils.

ter within the range $0 - 20\%$ even though only 23% of empirical correlations fall within these bounds.

Although not immediately apparent by inspection our model allows for the specification of negative correlations between sectors — to the authors' knowledge no other models can accommodate this possibility.

Figure 2 shows a time-line of 1-year VaR values for all three models under each of the calibration methodologies. Immediately below is a graph showing the sector values at these times. We note first of all that the standard and the hidden Gamma models show very little difference under either calibration. We would expect this result when hidden Gamma correlations are near zero, as is the case here. The hidden gamma model performs particularly poorly because the inner product form of the model correlation is too restrictive when default variances are small and correlations high.

A second point of interest is the high degree of variation in VaR under the exponential calibration. One might expect such sensitivity since the weighted data is dominated by the most recent observations. As already stated this may exacerbate pro-cyclicity since projections will tend to be over-optimistic in the good times and overly pessimistic in the bad. Indeed, the greater the degree of data attenuation the more violent VaR oscillations become and since the choice of weighting coefficient is usually nothing more than arbitrary it seems that one has little recourse to a robust analysis of VaR within this calibration framework.

Calibrating to the non-weighted data gives VaR values that show markedly less fluctuation. Here the model parameters apportion equal weighting to all observations such that the resulting VaR reflects a historically balanced view of risk arising from variation in sector and default data. It is important to note however that the marginal impact of new information diminishes through the course of time as the historic data set increases in size, which in the limit would lead to risk measures that are unresponsive to current conditions.

A significant advantage of the state-based method is the ease in which we can condition on current sector values. In this way we can provide estimates for VaR that embody the balanced view of the non-weighted model as a "base-line" VaR, upon which fluctuations are imposed in accordance with stresses currently at-play within the market.

Comparing the top and bottom plots of Figure 2 we see that in the periods $1990-1994$ and $1998-2003$ the mixture model suggests elevated risk with respect to the base-line which corroborate the elevated status of the sector risk factors during these periods. In contrast to the mixture model the exponential calibration considers 1995 a year of relatively high risk and $1997-2001$ as a period of relatively low (and mostly decreasing!) risk despite the overall increase in sector factor values; which here are considered to be annualised sector specific default rates. Thus our results demonstrate that even though this calibration technique is sensitive to changes in sector/default values, variation is nonetheless smoothed over the exponential time window thereby muting the immediacy and intensity of incoming data.

Lastly we note that in peaceful times the gamma mixture VaR values are only mildly lower than baseline. This is principally due to our choice of moment matching using two components: since in this instance the probability of peace is approximately 90% it generally follows that the peaceful component mean and volatility will be close to the unconditional (i.e. base-line) values relative to that of the stress state. In this way our model is more responsive to the occurrence of market stress.

A Sting In The Tail

Perhaps the most striking and publicly embarrassing aspect of the recent crisis was the revelation that banks had failed to retain sufficient capital to operate in adverse conditions. Without adequate provisions many institutions have been forced to rely upon government rescue packages, leaving them ever-more vulnerable to political and reputational risk and raising the spectre of unwelcome and unwieldy intervention in the future.

In a world in which ignorant models facilitate excessive risk taking in times of plenty, shocks come fast and hit hard. The relevance and need for accurate and robust measures of extreme value risk and the dependencies that exist between them has never been more apparent. In response to these concerns we have developed a state-based model together with necessary calibration techniques that allow one to identify, capture, intuit and mitigate these risks. We have implemented a simple application of this approach within the CreditRisk$^+$ framework and have demonstrated the potential to address key problems regarding default dependence structures present within even the most advanced variants of this model.

Contact

For further information regarding Misys Risk solutions including CreditRiskVantage and other developments please contact Ellen Stars (e-mail: Ellen.Stars@misys.com).