

# Credit risk contagion

*In a recession, company defaults increase due to both the worsening economic environment and the specific links between customers and suppliers. Banks intuitively know that customer default can cause supplier default. Duncan Martin and Chris Marrison provide analytical support for this intuition with a simple extension of the Merton portfolio model*

## Current approaches

to credit risk modelling typically explain correlation between companies by their exposure to common macroeconomic and financial risk factors. However, this explanation must be incomplete as common sense tells us that a credit event at one company affects the solvency of related companies directly. This effect is known intuitively by risk managers and regulators, which is why they track the interconnectedness of the companies they oversee.<sup>1</sup> In this article, we suggest a way of adding the idiosyncratic interconnectedness of companies to existing portfolio models. In these models, the default of individual companies is considered to be independent, conditional on the overall state of the market. In the model presented in this article, the default of individual companies also depends on the default state of all related companies in the previous time step. By recognising that credit events can be transmitted by business relationships between companies, clusters of defaults become more likely. This in turn increases the concentration of loss events and the capital requirements.

This article presents an approach to modelling this 'credit contagion', the spread of credit events between related companies. Previous approaches have considered adding contagion effects into credit risk models, but have generally required a complex implementation.<sup>2</sup> In our approach, we combine credit contagion with the general factor correlation of a standard Merton portfolio model to create a model that is relatively straightforward to implement and parameterise. In this model, if one firm is 'infected' by a credit event, there will be a knock-on effect on related firms, and the effect is proportional to the strength of the relationship. Specifi-

cally, in the time period following the credit event, the asset value of the infected firm's suppliers will decrease by an amount proportional to the strength of the sales relationship between them.

The contagion risk part of the model is inspired by modern epidemic models (see, for example, Halloran *et al*, 2002). These models relate the risk of contagion between two individuals in a population to the closeness of their relationship. In general, infected individuals are most likely to infect their family members, somewhat less likely to infect their fellow workers and students, and relatively unlikely to infect random members of the public. Each new infection sets off another chain of contagion. Readers who have school-age children, or who have co-workers who do, will be aware of this effect.<sup>3</sup>

This framework can be implemented as an addition to a Merton portfolio simulation model. Probably the best known implementation of the Merton model is the JP Morgan/CreditMetrics model. The Merton model simulates changes in the net asset value of all the companies in a portfolio by linking changes in their asset value to changes in market indexes. If a company's gross asset value falls below the value of its debts<sup>4</sup>, the company is said to default. By mapping companies to a common set of equity return indexes, correlation is created between the companies' asset values and hence their defaults.<sup>5</sup> Our approach then adds contagion – changes to asset values of related companies as a result of defaults. We measure the strength of the relationship based on a variable representing sales. A weakness of this approach is that the data required to parameterise the model is only occasionally available publicly and requires detailed information on companies that is typically only known by their accountants and bankers. For those banks that do know the company operations, our model turns this data into a source of competitive advantage.

### Construction of the model

Starting from company relationship data and probabilities of default (PDs), our combined model has seven steps:

■ 1. Use a Merton model to translate each company's PD into a default threshold for the asset value (this is also known as the distance to default (D2D))<sup>6</sup>:

$$D2D = N^{-1}(PD)$$

■ 2. For the given time step, simulate changes in market and macroeconomic factors, and thus in the asset value of each firm.

■ 3. Add idiosyncratic shocks to the company asset values to account for company-specific events.

■ 4. Flag default events for companies whose asset value falls below their default threshold.

<sup>1</sup> For example, the SEC mandates disclosure of material customer sales concentrations, and many banks keep track of these relationships as part of their credit review processes

<sup>2</sup> There are several previous approaches using the concept of infection or contagion in credit risk. Davis & Lo (2001) use the concept of infection to estimate the diversity score that should be used in a binomial model to represent the default correlation between companies whose assets are in an asset-backed security. This diversity score does not differentiate between the effects of market conditions and the effect of specific company defaults. Giesecke & Weber (2004) use the continuous-time Markov process to analytically calculate the joint evolution of firms' liquidity state. This model does not easily lend itself to parameterisation using the data that is normally available to credit managers. Egloff, Leippold & Vanini (2004) use a graph of business interdependencies and then take a highly mathematical approach to estimate the effect on credit transition matrices. Compared with these papers, we give a direct simulation approach that credit managers can quickly implement with the information already available to them


<sup>3</sup> There are many other aspects of epidemic models that bear examination in a credit context, particularly the framework for the spread of different infections through a population based on their incubation time, infective period and infection channels, and the notion of immunity, whether acquired from a vaccine or from recovery from infection

<sup>4</sup> The likelihood of assets being less than liabilities is the firm's PD. In the Merton model, the PD is converted to a 'distance to default'. The distance to default is used as the threshold change in asset value for default occurring in the simulation

<sup>5</sup> For a longer explanation of the Merton portfolio model, see chapter 21 of Marrison (2002)

<sup>6</sup> By using the inverse of the standard normal distribution ( $N^{-1}$ ), we are expressing the distance to default in terms of the number of standard deviations of the annual uncertainty in the company's net asset value

1 Asset values for three companies in a normal Merton model compared with a contagion model



- 5. If a default occurs, identify companies that will be ‘infected’ by the default. The channels for contagion are the sales relationships between companies; the strength of contagion is a function of the percentage of sales from each company to each of the other companies.
- 6. Subtract the impact of contagion from the asset values of the infected companies in the next time step.
- 7. Repeat.

This is illustrated in figure 1. The solid lines show the paths of asset values for three companies in a normal Merton model. Companies 2 and 3 depend on company 1. The dashed lines show the effect of explicitly adding the effect of this dependency to the asset value paths of companies 2 and 3. The result is that when company 1 defaults, the asset values of both companies 2 and 3 decrease, and the incremental decrease in value for company 3 is sufficient to push it into default.

In terms of equations, the two models are given in matrix form as follows:

- Merton:

$$\mathbf{v}_{t+1} = \mathbf{v}_t + \mathbf{W}\mathbf{i}_t + \mathbf{e}_t$$

- Combined Merton plus contagion:

$$\mathbf{v}_{t+1} = \mathbf{v}_t + \mathbf{W}\mathbf{i}_t - \mathbf{v}_0 \cdot \mathbf{I} \cdot \mathbf{S}\mathbf{d}_t + \mathbf{e}_t$$

Here  $t$  is time.  $\mathbf{v}_t$  is the vector of company net asset values at time  $t$ . As we are using the usual D2D concept, the net asset values are normalised by dividing the actual net asset value by the annual standard deviation of the company’s net asset value. The initial value of  $\mathbf{v}$  is the company’s D2D. Default occurs if the value becomes less than zero, that is, the company’s liabilities are worth more than its assets.

$\mathbf{W}$  is the matrix of weights relating the change in each company’s asset value to the change in its industry index (that is, how much the value of the company changes for a given change in the index).  $\mathbf{i}_t$  is the vector of changes in the index during time period  $t$ .

The difference between the pure Merton model and this contagion model is contained in the term  $\mathbf{v}_0 \cdot \mathbf{I} \cdot \mathbf{S}\mathbf{d}_t$ . Here  $\mathbf{v}_0$  is the company value at the initial time. In the Merton model, the company value is equated to the initial distance to default.<sup>7</sup>  $\mathbf{I}$  is a vector describing the extent to which the sudden loss of all current customers would affect the asset value of each company.  $\mathbf{S}$  is a square matrix describing the dependency of each company on each other company. For example, if 30% of the sales of company  $x$  went to company  $y$ , then 30% would be the entry in row  $y$  of column  $x$ .  $\mathbf{d}_t$  is the vector of default flags at time  $t$ , which equal one for any company that had an initial default at time  $t$ , but otherwise equal zero.

The model is therefore sensitive to the parameter choices for  $\mathbf{I}$  and  $\mathbf{S}$ .  $\mathbf{I}$  is the ratio of the lost income that would result before all current customers could be replaced to the value of the company.  $\mathbf{S}$  is the percentage of sales to specific companies in the portfolio. In the section below on sensitivity testing, we show how different assumptions for these parameters affect the results. In the appendix, we give a simple method of calculating the elements of  $\mathbf{I}$ . The values of the elements of  $\mathbf{S}$  are typically available to banks, and in some cases are publicly available. As an example, in Case 4 (see below) we parameterise  $\mathbf{S}$  using publicly available information for companies who depend on US retailer Home Depot.

$\mathbf{e}_t$  is the vector of idiosyncratic random shocks at time  $t$  that are not explained by the other terms in the pure Merton model.  $\mathbf{e}_t$  is the vector of shocks in the contagion model. If  $\mathbf{e}_t$  equals  $\mathbf{\epsilon}_t$ , the addition of the contagion terms not only increases the correlation between defaults, but also increases the PD for each individual asset. This is valid if the original estimate of the company’s PD did not consider the effects of a major customer defaulting. However, if the effect of customer defaults is already included in the PD, then adding the contagion effect to the portfolio model should not be allowed to increase the PD of the individual companies. If the user believes that this is the case, the PD can be returned to its original value by multiplying  $\mathbf{e}_t$  by a factor less than one to obtain  $\mathbf{\epsilon}_t$ . This is equivalent to saying that the contagion is explaining some of the risk that was previously considered to be unexplained.

If needed, this factor can be found either by trial and error, by estimating the correlation between the Merton effects and contagion effects<sup>8</sup>, or by a mildly complex calculation based on the dif-

<sup>7</sup> From the probability of default for each company, we use the inverse normal function to determine how many standard deviations of a fall in the equity price are required to cause a default. This number of standard deviations is called the distance to default

<sup>8</sup> Because the distance-to-default equation is normalised to have a standard deviation of one, the random shock  $\mathbf{e}_t$  in the usual Merton model should have a standard deviation that accounts for the residual randomness after taking into account the variation caused by random changes in the market index:

$$\sigma_{e_t} = \sqrt{1 - a^2}$$

Here  $a$  is the correlation between the company’s net asset value and the market index. Similarly, in the Merton plus contagion model the random shock  $\mathbf{\epsilon}_t$  could be scaled to account for the residual uncertainty after the market index terms: (a) and the contagion terms (c):

$$\sigma_{\epsilon_t} = \sqrt{1 - a^2 - c^2 + 2ac\rho_{ac}}$$


This then would give us a scaling factor from  $\mathbf{e}_t$  to  $\mathbf{\epsilon}_t$ :

$$f = \frac{\sqrt{1 - a^2 - c^2 + 2ac\rho_{ac}}}{\sqrt{1 - a^2}}$$

However, there is some difficulty in estimating the correlation between the market effects and the contagion effects,  $\rho_{ac}$ . This requires us to estimate the correlation between the change in the market at time  $t$  and the change in asset values caused by defaults at  $t - 1$ . This could be estimated empirically by running the simulation multiple times and iterating towards a result. However, even if this factor can be estimated, the PD from the contagion model will not quite match the PD from the Merton model as the Merton model assumes a purely normal distribution of asset values, but the contagion terms introduce significant non-normality

A. Results from the pure Merton model							
Case	1	2	3	4	5	6	7
Mean (EL)	1.0	1.0	1.0	1.0	0.5	1.0	1.0
Std dev	2.1	2.1	2.1	2.1	1.3	2.1	2.1
Skew	4.5	4.5	4.5	4.5	5.7	4.5	4.5
Kurtosis	36.6	36.6	36.6	36.6	62.3	36.6	36.6
99%ile	10.0	10.0	10.0	10.0	6.0	10.0	10.0
99.9%ile	19.4	19.4	19.4	19.4	13.0	19.4	19.4
99.97%ile	24.8	24.8	24.8	24.8	16.0	24.8	24.8

B. Combined Merton plus contagion results relative to pure Merton results							
Case	1	2	3	4	5	6	7
Mean (EL)	1.14	1.11	1.11	1.02	1.02	1.18	1.00
Std dev	1.24	1.37	1.19	1.03	1.03	1.28	1.17
Skew	1.23	1.78	1.16	1.05	1.07	1.13	1.34
Kurtosis	1.58	3.07	1.29	1.12	1.21	1.29	1.87
99%ile	1.20	1.20	1.20	1.00	1.00	1.20	1.20
99.9%ile	1.30	1.85	1.33	1.05	1.08	1.32	1.25
99.97%ile	1.48	2.00	1.32	1.04	1.12	1.30	1.40



ference in PD<sup>9</sup> that results between the two models if the same idiosyncratic factor is used for both models.

One of the challenges is to find values for the entries in **I**, the vector describing the effects of losing sales as a result of a default at a related company. The elements of **I** will vary depending on factors such as the company's cost/revenue ratio, the debt ratio and the amount of time to find new customers. In the appendix, we give a simple method of calculating the elements of **I**. We conclude that a typical company will have a value of around 0.43, that is, all other things being equal, if all the company's customers go bankrupt, the asset value will drop by around 43%.

$$f \approx \sqrt{\frac{2 - a^2 - \left( \frac{N^{-1}(PD_0)}{N^{-1}(PD_1)} \right)}{1 - a^2}}$$

<sup>9</sup> Here *f* is the factor by which the idiosyncratic risk should be multiplied, *a* is the correlation to the market index, *PD*<sub>0</sub> is the target PD and *PD*<sub>1</sub> is the PD that was obtained without reducing the idiosyncratic factor

### Procedure for running the model

The procedure for building the combined model is as follows:

- 1. Construct a standard Merton portfolio simulation model but with monthly time steps. This requires specifying values for each company's PD, the corresponding distance to default, the index dependency (**W**) and the idiosyncratic risk. The idiosyncratic risk has the variance set equal to one minus the square of the company's correlation to the index, so that the overall asset value has a variance of one. This construction produces a correlated random walk in which the probability of the asset value of a company falling below its distance to default equals the company's PD. The result is a set of defaults whose correlation is due to the common dependence on the market index.
- 2. Construct an identical Merton model and add the contagion effect. This requires specifying values for **I** and **S**, the inter-company sales relationships.
- 3. Run both the standard Merton and combined Merton plus contagion models in parallel and calculate the statistics of the loss distribution from each model.
- 4. If required, in the contagion model, reduce the idiosyncratic risk until the mean (expected loss) is equal for each model.
- 5. Observe differences in the loss distributions.
- 6. Test for different parameter assumptions.

Following this procedure, our base case assumptions were:

- A portfolio of 100 companies, each with a 1% PD.
- 100% loss given default.
- A single market index. The random walks for the market index and idiosyncratic company-specific shocks have a standard deviation of one at the 12-month point.
- Each company has a 50% correlation to the market index and therefore a weight of 87% on the uncorrelated idiosyncratic company specific factor.
- In cases 1 to 3 (see below), each company depends on the default of only one other company with a weight of 0.43. This means that the elements of **S** are all either one or zero, and the elements of **I** are all 0.43.

### Sensitivity testing

Given these common factors, we ran seven cases:

- **Case 1:** a cascade. The default of company *n* affects company *n* + 1. **S** in this case is a 100 × 100 matrix filled with zeros, apart from the entries just below the diagonal, which equal one, as illustrated by:

$$S_{base} = \begin{bmatrix} 0 & 0 & 0 & 0 & \dots \\ 1 & 0 & 0 & 0 & \\ 0 & 1 & 0 & 0 & \\ 0 & 0 & 1 & 0 & \\ \vdots & & & & \ddots \end{bmatrix}$$

- **Case 2:** all companies depend on company 1. If that company defaults, all others are affected immediately. **S** takes the form of all zeros other than the first column, which is all ones.
- **Case 3:** 10 large companies completely dominating 10 different sectors, with the other nine companies in each sector depending on them 100%. In this case, **S** is block diagonal, with 10 10 × 10 blocks, each with ones in the first column of the block and zeros elsewhere.
- **Case 4:** 10 companies tending to dominate 10 different sectors based on real data. For this example, we took US Securities and Exchange Commission (SEC) data for the dependence on Home Depot of nine companies, including Stanley Works, Toro and US

Home Improvements. The form of the  $S$  matrix is the same as case 3, but the entries are less than one; on average they are 0.34, as shown below:

$$S_{HD} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.37 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.10 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.70 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.18 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.21 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.12 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.27 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.65 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.45 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

■ **Case 5:** a lower risk portfolio. All the other factors are the same as case 4 but the PD is reduced from 1% to 0.5%.

■ **Case 6:** a portfolio of more leveraged companies. This is the same as case 4, but the average debt of the companies is assumed to double, from 42% to 84%, increasing the entries in  $\mathbf{l}$  from 0.43 to 1.56. This means that if a company lost all its sales, it would almost certainly default, unless it had previously had more than a 50% rise in its value due to good economic circumstances.

■ **Case 7:** the same as case 1 except the idiosyncratic company specific random factor is multiplied by 0.975 to bring the expected loss (EL) of the model with contagion equal to the EL in the pure Merton model.

We used 50,000 evaluations for each case. In table A, we show the detailed results for the pure Merton model. In table B, we show the relative results of the combined Merton plus contagion model. Figure 2 shows the difference in the tails of the cumulative distribution for the pure Merton model and for case 1. Figure 3 shows the same information on a log scale. Figure 4 shows the tail of the distribution for all cases (note that case 5 is significantly different because it is the case where the underlying PD is lower).

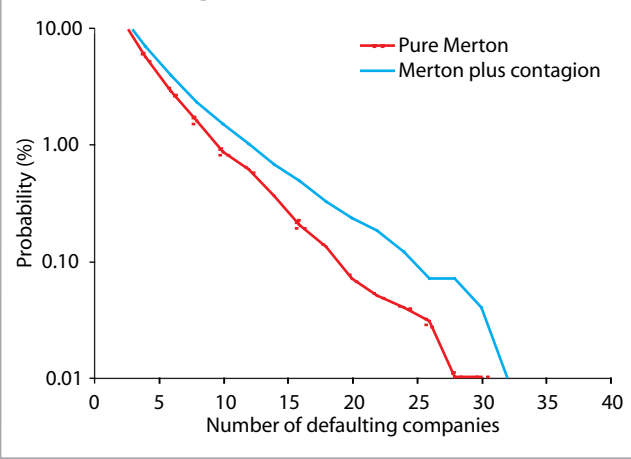
### Description of results

The first thing to notice in table A is that the results from the pure Merton model are identical in all cases other than case 5, in which the PD is reduced. In these cases, the mean loss is 1% and the 99.9 percentile is 19.2%. For comparison with typical results for bank credit portfolios, if we had assumed a loss given default of 50%, this result would imply 99.9% economic capital of 9.6% of portfolio exposure.

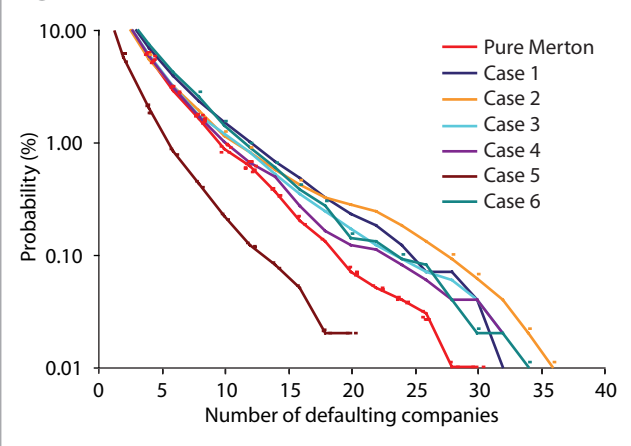
Looking at table B, we see that the mean (EL) is generally higher than the pure Merton results because we have added the contagion effect without reducing the idiosyncratic risk. Case 7 shows that, if required, the ELs can be made to match by reducing the idiosyncratic risk. In this case, the idiosyncratic risk is multiplied by a factor of 0.975 to bring EL in line.

The factors for the standard deviation, skew, kurtosis and percentiles are all increased by adding contagion. For example, in case 1, with the simple diagonal cascade, table B shows that the standard deviation is 24% greater and the 99.9 percentile is 30% greater. The intuition behind this is that if any one company defaults, it tends to take many other companies with it. In the extreme (case 2) where there is a single dominant company, the 99.9 percentile is 85% larger than that from the pure Merton model and the 99.97 percentile doubles.

3 Comparison of the tails of the base case loss distributions with and without contagion effects (shown with a log scale)



4 Comparison of the tails of all the cases (shown with a log scale)



The results are less dramatic in case 3, with 10 dominant companies where the 99.9 percentile only increases by 33%. In the more real-world example of case 4, the effect is still more modest, only increasing the 99.9% capital by 11%.

Case 5 halves the probability of default for each asset, reducing the increase in the 99.9 percentile to 8%. However, the greater sensitivity to lost revenues in case 6 produces a 99.9 percentile increase of 32%. Case 7 shows that reducing the idiosyncratic risk to align EL still leaves a significant difference in the 99.9 percentile.

The divergence between the lines in the charts shows that further out in the tails, the difference in the results becomes more pronounced. Overall, the results show that in extreme cases, adding contagion effects could almost double the required capital. In more realistic cases, the increase is still 10–30%. Since these cases represent chains of quite closely linked companies, this is intuitively sensible: the tails of the distribution represent default ‘cascades’ where related companies suffer from a domino effect.

### Conclusions

Explicitly adding contagion effects to credit portfolio modelling can increase the required capital by 5% to 30% (cases 4 and 6) or by up to 100% in the extreme case 2. This implies that standard

**Appendix: parameterisation of I**

Here, we present a simple method for relating the loss in asset value of one company to the default of a related company. The assumptions are simplistic and stylised, but illustrate a line of reasoning that could be used to link the default in one company to a drop in distance to default of another company. There will also be other valid approaches to calculating these parameters.

The market value of a company can be approximated as the net present value of its future profits. For simplicity, here we approximate it as the sum of the next 10 years of profits, that is,  $10 \times (\text{revenue} - \text{costs})$ . This is equivalent to using a discount rate of 10% and assuming only inflationary growth. If the costs of the company are 80% of revenues, then the company's value will be around twice annual revenue, that is,  $10 \times (20\% \text{ of annual revenue}) = 200\%$  of annual revenue. If we posit that for this company it would take around six months to replace the revenues from a significant customer if it were to default, then the fall in value will be 50% of annual revenues multiplied by the percentage of sales to the defaulted company. For example, if all sales were to one company, then if that company defaulted, the gross value would drop from two times earnings to 1.5 times earnings, that is, a 25% drop in gross value. To get the effect on the net asset value, we need to add the effect of debt. If the company initially had the S&P 500 average of 42% of debt to gross value, the 25% drop in gross value would translate to a 43% drop in net asset value, that is, the element in vector **I** corresponding to this company will be 0.43. To test the effect of varying the parameters, two other cases are shown in the table.

Company B is highly leveraged. If it was 100% dependent on another company, the default of that company would drop the asset value of B by 500%, causing a certain default (in reality, the drop in value is capped at 100%). Com-

pany C, on the other hand, would hardly be affected by a customer default.

The calculation of **I** can be expressed as follows:

$$\frac{[\text{Replacement time} \times \text{Revenue}]}{[\text{NPV of profit} \times (1 - \text{Leverage}\%)]}$$

Or, with the simple assumption for net present value:

$$\frac{\text{Replacement time}}{[10 \times (1 - \text{Costratio}) \times (1 - \text{Leverage}\%)]}$$

In our example model for the base case, we assume that the effect of a default of one company is to reduce the distance to default of related companies by 43% times the percentage of the defaulted companies' sales that were going to the defaulted company. We then try different levels of dependency and diversification of sales to see the effect on the portfolio's loss distribution.

Effect of company profiles on I			
Company	A	B	C
Cost/income	80%	90%	50%
Replacement time	6 months	12 months	1 month
Leverage	42%	80%	0%
Result for I	43%	500%	2%

portfolio models substantially underestimate the risk of portfolios that contain many closely related companies. Many institutions hold such portfolios, namely those that are rich in exposure to:

- Integrated supply chains, such as in the automotive and retail industries. This effect will be particularly pronounced where competing supply chains are segregated and/or localised, leading to buyer concentration. In addition, results will be different where supply chains cross borders, as the dynamics of these trans-national relationships will not be picked up by standard portfolio models.
- Families of related businesses within a holding company such as those that dominate emerging markets such as Mexico and Russia.

We recommend that institutions with these types of portfolios enhance their risk modelling by collecting data on dependency and using it in the framework outlined above.

**Directions for additional development**

The most obvious area is to create better estimates for the key parameters, notably asset value reductions (see Appendix) and contagion channels.

In addition, we assumed there was no overlap between Merton model correlations and contagion, and adjusted idiosyncratic risk as a plug. It is probable that a more sophisticated approach would yield increased accuracy.

In some cases, the default of one company could be beneficial to another, competing company. If this data could be sourced, it could easily be included in the approach as negative values in **S**.

In this article, we examined the effect of contagion on the creditworthiness of companies based on the strength of their business relationships. The approach may also be applicable to assessing the contagion effect that occurs between stocks due to the actions of investors. There are probably substantial direct contagion effects between certain high-risk stocks, such as distressed com-

panies, emerging markets and advanced technologies. These effects are driven not by business relationships between companies but by investor clientele effects: commonalities in stock selection across risk-seeking fund managers. The result is that the same stocks end up in many different portfolios. Consequently, buyers and sellers of the equities create channels for transmission of equity price moves across markets. As one high-risk stock falls, others in the same portfolio are sold to meet margin requirements, and contagion occurs. ■

Duncan Martin is a senior credit manager at Royal Bank of Scotland. Chris Marrison is chief executive and founder of Risk Integrated. They would like to thank Til Schuermann, Adam Kleczkowski and the anonymous referees for their help in the preparation of this article. Email: duncan.martin@rbs.co.uk, Chris.Marrison@RiskIntegrated.com

**References**

Davis M and V Lo, 2001  
*Infectious defaults*  
Quantitative Finance 1, pages 382–387

Egloff D, M Leippold and P Vanini, 2004  
*A simple model of credit contagion*  
EFA 2004 Maastricht Meetings Paper 1,142

Giesecke K and S Weber, 2004  
*Cyclical correlations, credit contagion and portfolio losses*  
Journal of Banking & Finance 28, pages 3,009–3,036

Halloran E, I Longini, A Nizam and Y Yang, 2002  
*Containing bioterrorist smallpox*  
Science, November 15, 298(5,597), pages 1,428–1,432

Marrison C, 2002  
*The fundamentals of risk measurement*  
McGraw Hill