Reconcilable differences

H Ugur Koyluoglu and Andrew Hickman explore the common ground between the new credit risk models and the implications for risk management and regulatory capital reform

In the past few years, major advances in credit risk analytics have led to the proliferation of a new breed of sophisticated credit portfolio risk models. Several models have been developed, including proprietary applications developed for internal use by leading-edge financial institutions, and third-party applications intended for sale or distribution as software. Several have received a great deal of public attention, including JP Morgan's CreditMetrics/CreditManager, Credit Suisse Financial Products' CreditRisk+, McKinsey & Company's CreditPortfolioView and KMV's PortfolioManager. These new models allow the user to measure and quantify credit risk comprehensively at both the portfolio and contributory level. As such, they have the potential to cause profound changes to the lending business, accelerating the shift to active credit portfolio management[1], and eventually leading to an "internal models" reform of regulatory credit risk capital guidelines[2].

But before these models can deliver on their promise, they must earn the acceptance of credit portfolio managers and regulators. To these practitioners, this seemingly disparate collection of new approaches may be confusing, or may appear as a warning sign of an early developmental stage in the technology. While these misgivings are understandable, this paper will demonstrate that these new models in fact represent a remarkable consensus in the underlying framework, differing primarily in calculation procedures and parameters rather than financial intuition.

This paper explores both the similarities and the differences among the new credit risk portfolio models, focusing on three representative models:

- "Merton-based", eg, CreditMetrics and PortfolioManager[3];
- "econometric", eg, CreditPortfolioView; and
- "actuarial", eg, CreditRisk+.

Note that this paper examines only the default component of portfolio credit risk. Some models incorporate credit spread (or ratings migration) risk, while others advocate a separate model. In this aspect of credit risk there is less consensus in modelling techniques, and the differences need to be explored and resolved in future research. The reader should strictly interpret "credit risk" to mean "default risk" throughout.

Additionally, for comparability, the models have been restricted to a single-period horizon, a fixed recovery rate and fixed exposures.

Underlying framework

At first, the models appear to be quite dissimilar – CreditMetrics is based on a microeconomic causal model of default; CreditPortfolioView is a macroeconomic causal model; and CreditRisk+ is a top-down model, making no assumptions about causality. Despite these apparent differences, the models fit within a single generalised underlying framework, consisting of three components:

- Joint-default behaviour. Default rates vary over time, intuitively as a result of varying economic conditions. Each borrower's default rate is conditioned on the "state of the world" for the relevant economic conditions. The degree of "correlation" in the portfolio is reflected by borrowers' conditional default rates varying together in different states.
- Conditional distribution of portfolio default rate. For each state, the conditional distribution of a homogeneous sub-portfolio's default rate can be calculated as if borrowers are independent, as the joint-default behaviour is accounted for in generating conditional default rates.
- Convolution/aggregation. The unconditional distribution of portfolio defaults is obtained by combining homogeneous sub-portfolios' conditional default rate distributions in each state, and then simply averaging across states.

This generalised framework allows a structured comparison of the models, as follows.

- **Conditional default rates and probability distribution of default rate.** All three models explicitly or implicitly relate default rates to variables describing the relevant economic conditions ("systemic factors"). This relationship can be expressed as a "conditional default rate" transformation function (see figure 1). The systemic factors are random and are usually assumed to be normally distributed. Since the conditional default rate is a function of these random systemic factors, the default rate will also be random.

The Merton-based model relies on Merton's model of a firm's capital structure[4]: a firm defaults when its asset value falls below its liabilities. Default probability then depends on the amount by which assets exceed liabilities, and the volatility of those assets. If standardised changes in asset value ΔA_i are normally distributed, the default probability can be expressed as the probability of a standard normal variable falling below some critical value c. Joint-default events among borrowers in the portfolio are related to the extent that the borrowers' changes in asset value are correlated.

Since the Merton model neither assigns the transformation function, nor assumes a probability distribution for default rates explicitly, these relationships must be derived. The change in asset value can be decomposed into a set of normally distributed orthogonal systemic factors, x_k, and a normally distributed idiosyncratic component ε_i:

$$\Delta A_i = b_{i,1}x_1 + b_{i,2}x_2 + \ldots + \sqrt{1-\sum_k b_{i,k}^2}\,\varepsilon_i$$

where $b_{i,k}$ are the factor-loadings, and $x_k, \varepsilon_i \sim \text{iid } N[0,1]$.

Given the values of the systemic factors, the change in asset value will be normally distributed with a mean given by the factor loadings and factor values, and a standard deviation given by the weight of the idiosyncratic factor. The default rate, conditioned on the systemic factors' values, can then be expressed as[5]:

$$p_i\big|_x = \Phi\left[\frac{c-\sum_k b_{i,k}x_k}{\sqrt{1-\sum_k b_{i,k}^2}}\right]$$

For the single borrower or homogeneous portfolio case, the systemic factors can be summarised by a single variable, m, reducing the transformation function to:

$$p\big|_m = \Phi\left[\frac{c-\sqrt{\rho}\,m}{\sqrt{1-\rho}}\right]$$

[1] For example, see Kuritzkes (1998)
[2] See International Swaps and Derivatives Association (1998)
[3] The discussion that follows will focus on CreditMetrics as the example, but will also apply reasonably well to PortfolioManager
[4] See Merton (1974), Kealhoffer (1995) and Gupton, Finger & Bhatia (1997)
[5] Vasicek (1987) develops this representation of the Merton model for a single factor

where $m \sim N[0,1]$ and

$$\rho = \sum_k b_k{}^2$$

is the asset correlation.

Since the cumulative normal function is bounded $[0,1]$ and concave in the relevant region, the resulting default rate distribution is bounded $[0,1]$ and skewed right, as in figure 1.

The probability density function for the default rate, $f(p)$, can be derived explicitly, as follows:

$$f(p) = \varphi\big(m(p)\big)\left|\frac{dm}{dp}\right| = \frac{\sqrt{1-\rho}\,\varphi\left(\frac{c-\sqrt{1-\rho}\,\Phi^{-1}(p)}{\sqrt{\rho}}\right)}{\sqrt{\rho}\,\varphi\big(\Phi^{-1}(p)\big)}$$

where $\varphi(z)$ is the standardised normal density function.

The econometric model[6] drives the default rate, $p_{i,t}$, according to an "index", $y_{i,t}$, of macroeconomic factors. The index is expressed as a weighted sum of macroeconomic variables, $x_{k,t}$, each of which is normally distributed and has lagged dependency.

$$x_{k,t} = a_{k,0} + a_{k,1}x_{k,t-1} + a_{k,2}x_{k,t-2} + \ldots + \varepsilon_{k,t}$$

and

$$y_{i,t} = b_{i,0} + b_{i,1}x_{1,t} + b_{i,2}x_{2,t} + \ldots + \upsilon_{i,t}$$

where $\varepsilon_{k,t}$ and $\upsilon_{i,t}$ are normally distributed random innovations.

The index is transformed to a default probability by the Logit function:

$$p_{i,t} = \frac{1}{1+e^{y_{i,t}}}$$

The index and macroeconomic variables can be combined to a single equation:

$$y_{i,t} = \left(b_{i,0} + \sum_k b_{i,k}\left(a_{k,0} + \sum_j a_{k,j}x_{k,t-j}\right)\right) + \sum_k b_{i,k}\varepsilon_{k,t} + \upsilon_{i,t}$$

consisting of a constant term and random terms representing systemic and index-specific innovations. For the single borrower or homogeneous portfolio case, these random terms can be summarised by a single normally distributed variable, m, so that the conditional default rate can then be expressed as:

$$p|_m = \frac{1}{1+e^{U+Vm}}$$

where $m \sim N[0,1]$, and U and V represent the summarised constant term and coefficient to the random term, respectively.

Since the Logit function is bounded $[0,1]$ and concave, the resulting distribution is bounded $[0,1]$ and skewed, as in figure 1.

The implied probability density function for the default rate, $f(p)$, is

$$f(p) = \varphi\big(m(p)\big)\left|\frac{dm}{dp}\right| = \frac{1}{Vp(1-p)}\,\varphi\left(\frac{1}{V}\ln\left(\frac{1-p}{p}\right) - \frac{U}{V}\right)$$

The actuarial model[7] assumes explicitly that the default rate distribution follows the gamma distribution. Joint-default behaviour is incorporated by treating the default rate as a random variable common to multiple borrowers. Borrowers are allocated among "sectors", each of which has a gamma-distributed default rate with specified mean and volatility. A borrower's conditional default rate is a scaled weighted average of sector default rates:

$$p|_x = \bar{p}\sum_k \omega_k \frac{x_k}{\mu_k}$$

where $\bar{p}$ is the borrower's unconditional default rate, ω_k represents the weight in sector k,

1. Conditional default rate transformation

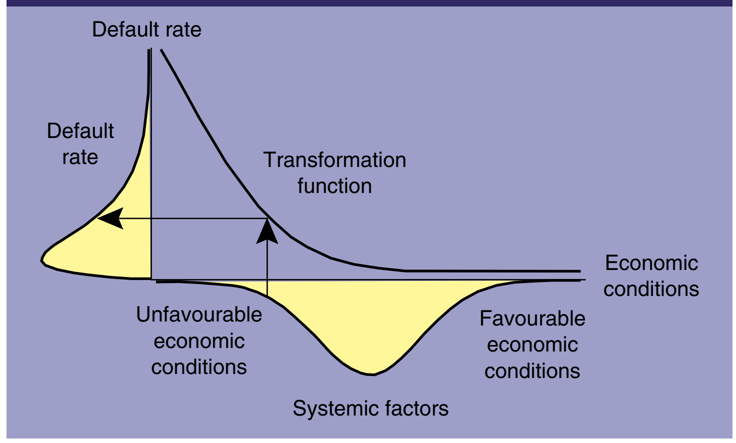

$$\sum_k \omega_k = 1$$

and:

$$x_k \sim \Gamma\big[\alpha_k, \beta_k\big] \text{ with } \alpha_k = \frac{\mu_k{}^2}{\sigma_k{}^2} \text{ and } \beta_k = \frac{\sigma_k{}^2}{\mu_k}$$

The gamma is skewed right as in figure 1, but has unbounded positive support.

It is possible to derive the actuarial model's implied transformation function such that when applied to a normally distributed systemic factor, m, it results in a gamma-distributed default rate. The transformation function consists of all points (χ,ξ) that satisfy:

$$\int_0^{\xi} \Gamma(p;\alpha,\beta)dp = \int_{\chi}^{\infty} \varphi(m)dm$$

Hence, the transformation function is given by:

$$p|_m = \Psi^{-1}\big(1-\Phi(m);\alpha,\beta\big)$$

where

$$\alpha = \frac{\bar{p}^2}{\sigma^2}, \beta = \frac{\sigma^2}{\bar{p}}$$

$m \sim N[0,1]$ and $\Psi(z;\alpha,\beta)$ is the cumulative density function of the gamma distribution.

□ **Conditional distribution of portfolio default rate.** Given fixed or conditional default rates, a homogeneous sub-portfolio's distribution of defaults follows the binomial distribution $B(k;n,p)$, which provides the probability that k defaults will occur in a portfolio of n borrowers if each has default probability p. CreditMetrics implicitly uses the binomial distribution by calculating the change in asset value for each borrower and testing for default – exactly equivalent to the binomial case of two states with a given probability. CreditPortfolioView explicitly uses the binomial distribution by iteratively convoluting the individual obligor distributions, each of which is binomial.

CreditRisk+ uses the Poisson distribution $P(k;pN)$, which provides the probability that k defaults will occur in a portfolio of n borrowers given a rate of intensity per unit time p. The binomial and Poisson distributions are quite similar; indeed, the Poisson distribution is the limiting distribution for the binomial distribution.[8]

□ **Aggregation.** The unconditional probability distribution of portfolio defaults is obtained by combining the conditional distributions of homogeneous sub-portfolio defaults across all "states of the world". Mathematically, this is expressed as a convolution integral.

[6] See Wilson (1997)
[7] See Credit Suisse Financial Products (1997)
[8] See Freund (1992)

2. The models in relation to the generalised framework

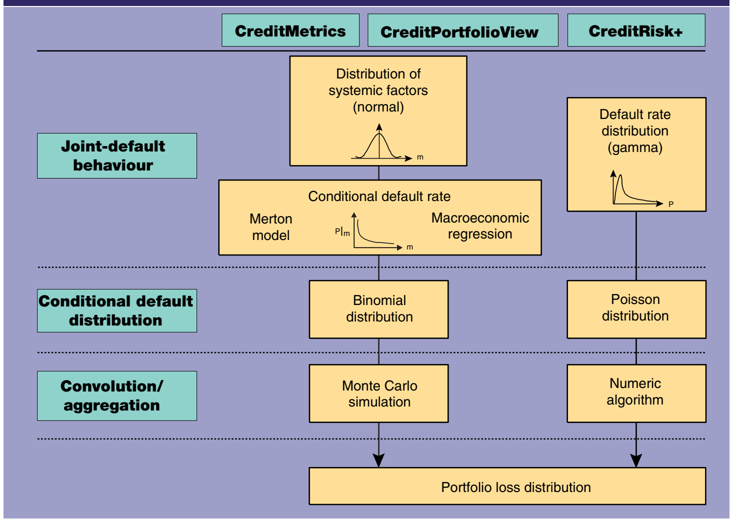

3. Default rate volatility/asset correlation

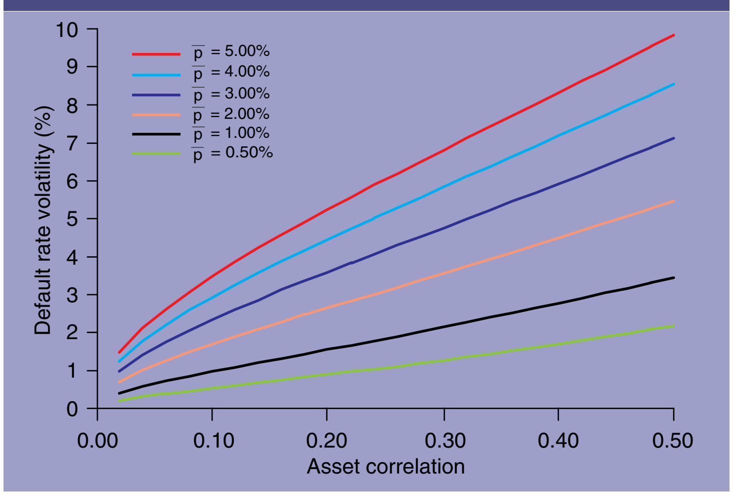

The Merton-based and econometric models are conditioned on normally distributed systemic factors, and the independent loans' defaults are binomially distributed. Hence, the convolution integral for a homogeneous sub-portfolio with a single systemic factor is expressed as:

$$\int_{-\infty}^{\infty} B\left(k;n,p|_m\right)\varphi(m)dm$$

The actuarial model's homogeneous sub-portfolio convolution integral, with gamma-distributed default rate and Poisson-distributed conditional defaults, is:

$$\int_{0}^{\infty} P(k;np)\Gamma(p;\alpha,\beta)dp$$

These integrals are easily evaluated; in particular, the convolution of the Poisson distribution and gamma distribution yields a closed-form distribution, the negative binomial distribution. It is the differences between sub-portfolios – differing exposure size or default probabilities, or multiple systemic factors, complex correlation structure, etc – that create difficulty in aggregation. In practice, then, the convolutions are evaluated by Monte Carlo simulation in CreditMetrics and CreditPortfolioView, while CreditRisk+ uses a numeric algorithm based on "banding" exposures. In all three cases, the procedures are exact in the limit.

Figure 2 depicts the models as they are redefined in relation to the generalised framework.

Harmonisation of parameters

The preceding discussion shows that all three models critically depend on the unconditional default probability and joint-default behaviour. While unconditional default probability is relatively straightforward, joint-default behaviour appears in a different form in each model. The Merton-based model uses pairwise asset correlations; the actuarial model uses sector weightings and default rate volatilities; and the econometric model uses coefficients to common macroeconomic factors. Although these parameters are very different in nature, they contain equivalent information to characterise joint-default behaviour.

☐ **Coefficients and correlations.** The Merton-based model represents joint-default behaviour with a set of asset factor-loadings or, equivalently, a pairwise asset correlation matrix:

$$\Delta A_i = b_{i,1}x_1 + b_{i,2}x_2 + \ldots + \sqrt{1-\sum_k b_{i,k}^2}\,\varepsilon_i$$

The systemic factors are defined to be orthonormal, so that:

$$\text{correlation}\left[\Delta A_i, \Delta A_j\right] = \frac{E\left[\Delta A_i \Delta A_j\right] - E\left[\Delta A_i\right]E\left[\Delta A_j\right]}{\sqrt{\left(E\left[\Delta A_i^2\right] - E\left[\Delta A_i\right]^2\right)\left(E\left[\Delta A_j^2\right] - E\left[\Delta A_j\right]^2\right)}}$$
$$= b_{i,1}b_{j,1} + b_{i,2}b_{j,2} + \ldots$$

The econometric model's "index" regression coefficients closely resemble the asset factor-loadings of the Merton-based model. An "index correlation" is easily defined in a similar fashion to an asset correlation, and will be treated as equivalent, though they may provide slightly different results to the extent of differences in their respective conditional default rate functions.

☐ **Unconditional default rate and default rate volatility.** The unconditional default rate and default rate volatility are specified directly in the actuarial model. For the Merton-based and econometric models, they are calculated by:

$$\bar{p} = \int_{-\infty}^{\infty} p|_m\, \varphi(m)dm$$

and:

$$\sigma^2 = \int_{-\infty}^{\infty} \left(p|_m - \bar{p}\right)^2 \varphi(m)dm$$

The parameters for the Merton-based (c and ρ) and econometric models (U and V) can then be solved to yield a specified unconditional default rate and default rate volatility. This defines the relationship between default rate volatility and asset correlation (see figure 3).

☐ **Default correlation.** Some models take a Markowitz variance-covariance view of credit risk portfolio modelling. Each borrower has a variance of default given by the variance for a Bernoulli variable:

$$VAR(\text{default}_i) = \bar{p}_i\left(1 - \bar{p}_i\right)$$

For a large homogeneous portfolio, the portfolio variance approaches:

$$\sigma^2 = \bar{p}(1-\bar{p})\rho_{\text{default}}$$

This provides the relationships between default correlation and default rate volatility and, therefore, asset correlation.

Mappings such as these allow parameter estimates to be "triangulated" by multiple methods, to the extent that model differences are not significant. For example, default rate volatilities can be used to estimate implied asset correlations in the absence of asset value data.

Differences in default rate distribution

The discussion above ("Underlying framework") demonstrates that substantial model differences could arise only from the differing treatment of joint-default behaviour – the conditional default distributions are effectively the same and the aggregation techniques are all exact in the limit. The sec-

[9] *These parameters were selected to match Moody's Investors Service's "All Corporates" default experience for 1970–1995, as reported in Carty & Lieberman (1996)*

tion "Harmonisation of parameters" provides the means to compare the joint-default behaviour on an apples-to-apples basis.

This comparison will be illustrated for a homogeneous portfolio with an unconditional default rate, $\bar{p}$, of 116 basis points and a standard deviation of default rate, σ, equal to 90bp[9]. Since each model produces a two-parameter default rate distribution, the mean and standard deviation are sufficient statistics to define the relevant parameters for any of the models, as above. To yield $\bar{p} = 116\text{bp}$ and $\sigma = 90\text{bp}$, the parameters for each model are as follows:

- Merton-based: $c = -2.27$, $\rho = 0.073$
- econometric: $U = 4.684$, $V = 0.699$
- actuarial: $\alpha = 1.661$, $\beta = 0.0070$.

In this example, the models' conditional default rate functions are virtually indistinguishable when the systemic factor is greater than negative two standard deviations. For extremely unfavourable economic conditions, the econometric model predicts a somewhat higher default rate, and the actuarial model predicts a somewhat lower default rate. The default rate distributions (see figure 4) are also very similar, with only minor discrepancies in the tails.

The degree of agreement in the tails of these distributions can be assessed with the following statistic:

$$\Xi_z(f,g) = 1 - \frac{\int_z^\infty \left|f(x) - g(x)\right| dx}{\int_z^\infty f(x)dx + \int_z^\infty g(x)dx}$$

where $f(x)$ and $g(x)$ are probability density functions and z defines the lower bounds of the "tail", which will be defined arbitrarily as the area more than two standard deviations above the mean, ie, $z = \bar{p} + 2\sigma$. This statistic measures the amount of the probability distributions' mass that overlaps in the tail, normalised to the total probability mass of the two distributions in the tail. The statistic will be bounded $[0,1]$, where zero indicates distributions with no overlapping probability mass, and one indicates exact agreement. Table A provides the tail-agreement statistics for the example distributions.

Without a credible alternative distribution, this tail-agreement statistic provides a relative rather than an absolute measure. However, it can be used to test the robustness of the similarity to the parameters (see table B).

The results in table B demonstrate that the similarity of the models holds for a reasonably wide range of parameters. The models begin to diverge at a very high ratio of default rate volatility to default probability, particularly for very low or very high default probabilities. Accordingly, in very high quality (AA or better) or very low quality (B or worse) portfolios, model selection can make a difference, though there are scant data on which to base such a selection. In a portfolio with only moderate weight in very high or very low quality sub-portfolios, these differences should not be significant in aggregation.

Impact of parameter inconsistency. This finding of similarity should be taken with caution, as it hinges on harmonising parameter values. In practice, the parameters will vary by estimation technique. The different estimation techniques appropriate to different joint-default parameters may result in inconsistent default rate volatility. Even mean default probabilities may vary considerably depending on the estimation technique, sample, etc. Unsurprisingly, when the parameters do not imply consistent mean and standard deviation of default rate distribution, the result is that the models are significantly different. This case is illustrated by an example of three parameter sets that are not consistent, though plausibly obtainable for the same portfolio (see table C). Within any one of these three parameter sets, a comparison of the models yields results similar to figure 4 and table A – tail-agreement statistics average 91% and range from 82% to 95%. Large differences arise when the models are compared across the inconsistent parameter sets (see figure 5) – tail-agreement statistics average only 76% and range from 65% to 85%, even when comparing the same model applied to each of the inconsistent parameter sets. The differences in parameters, well within the typical range of estimation error, have much greater impact than model differences in this example.

4. Default rate distributions compared

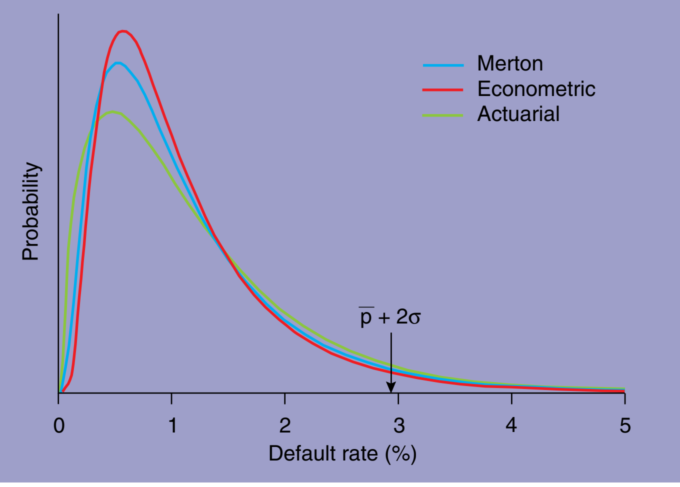

A. Tail-agreement statistics for the example distributions

Merton versus econometric	94.90%
Merton versus actuarial	93.38%
Econometric versus actuarial	88.65%

B. Tail-agreement statistics v. parameter values

$\bar{p}$	σ/p̄:	0.50	1.00	2.00	3.00
		98.10%	95.94%	91.16%	89.12%
	0.05%	93.53%	88.50%	81.17%	78.99%
		91.71%	84.70%	73.13%	69.04%
		97.92%	95.57%	91.15%	88.40%
	0.10%	93.75%	88.94%	82.16%	80.40%
		91.73%	84.78%	73.95%	69.73%
		97.61%	94.93%	90.35%	87.86%
	0.25%	94.11%	89.73%	83.87%	82.92%
		91.79%	84.97%	74.83%	71.60%
		97.33%	94.41%	89.91%	88.09%
$\bar{p}$	0.50%	94.42%	90.56%	85.71%	85.61%
		91.88%	85.30%	76.15%	74.28%
		97.06%	93.97%	89.89%	88.97%
	1.00%	94.93%	91.69%	88.29%	89.38%
		92.04%	85.93%	78.57%	78.72%
		96.62%	93.62%	90.77%	91.33%
	2.50%	95.82%	93.94%	93.65%	94.68%
		92.62%	87.79%	84.77%	87.53%
		96.33%	93.85%	92.79%	94.59%
	5.00%	97.02%	96.70%	95.24%	81.79%
		93.55%	90.87%	91.38%	79.96%
		96.21%	94.92%	95.79%	na
	10.00%	98.55%	95.57%	72.95%	na
		95.45%	95.22%	72.41%	na

na = not applicable because it is an unreasonable combination of parameters – model results become unstable. Each cell contains tail-agreement statistics for Merton v. econometric, Merton v. actuarial and econometric v. actuarial

Conclusions

On the surface, the credit risk portfolio models studied here seem to be quite different. Deeper examination reveals that the models belong to a single general framework, which identifies three critical points of com-

C. Hypothetical inconsistent parameter values

	$\bar{p}$	σ	c	ρ	α	β	U	V	Model for comparison[1]
1	**2.26%**	**1.70%**	–2.00	8.5%	1.767	0.0128	**4.00**	**0.70**	Econometric
2	**1.52%**	**1.71%**	–2.16	14.4%	**0.790**	**0.0192**	4.60	0.95	Actuarial
3	**1.54%**	**2.63%**	**–2.16**	**26.2%**	0.343	0.0449	4.95	1.30	Merton

[1] In the inconsistent parameter case, the parameter sets' "models for comparison" were selected arbitrarily. Figures in bold indicate parameters appropriate to selected model

parison – the default rate distribution, the conditional default distribution, and the convolution/aggregation technique. Differences were found to be immaterial in the last two of these, so that any significant differences between the models must arise from differences in modelling joint-default behaviour which manifest in the default rate distribution. Further, when the joint-default parameter values are harmonised to a consistent expression of default rate and default rate volatility, the default rate distributions are sufficiently similar as to cause little meaningful difference across a broad range of reasonable parameter values. Any significant model differences can then be attributed to parameter value estimates that have inconsistent implications for the observable default rate behaviour.

Parameter inconsistency is not a trivial issue. A "naïve" comparison of the models, with parameters estimated from different data using different techniques, is quite likely to produce significantly different results for the same portfolio. The conclusions of empirical comparisons of the models will vary according to the degree of difference in parameters.[10] In such comparisons, it is important to understand the proportions of "parameter variance" and "model variance" if different results are produced for the same portfolio. The findings in this paper suggest that "parameter variance" is likely to dominate. Future studies should focus on the magnitude of parameter differences and the sensitivity of results to these differences.

Parameter inconsistency can arise from two sources: estimation error, which could arise from small sample size or other sampling issues; or model mis-specification. While default rate volatility may be immediately observable, even long periods of observation provide small sample size and risk non-stationarity. At the other extreme, asset correlations can be measured with reasonable sample size in much shorter periods, albeit with the risk of mis-specification in the return distributions and default causality assumptions in the translation to default rate volatility. Rather than conclude that parameter inconsistency potentially constitutes irreconcilable differences between the results of these models, this paper concludes that because the models are so closely related, the estimates are complementary and should provide improved accuracy in parameter estimation within the generalised framework as a whole.

A useful metaphor can be drawn from the success of the value-at-risk framework in modelling market risk. VAR has become the industry standard and the basis for regulatory capital requirements. But in practice, VAR encompasses a variety of significantly different modelling and parameter estimation techniques, eg, historical simulation versus variance-covariance, delta-gamma versus exact Monte Carlo simulation, etc. The underlying coherence of the VAR concept – that risk is measured by combining the relationship between the value of trading positions to market variables with the distribution of those underlying market variables – ensures a consistency sufficient for widespread acceptance and regulatory change. Similarly, the underlying coherence of these new sophisticated credit risk portfolio models should allow them to overcome differences in calculation procedures and parameter estimation. Rather than dissimilar competing alternatives, these models represent an emerging industry standard for credit risk management and regulation. ■

5. Conditional default rates for different models compared

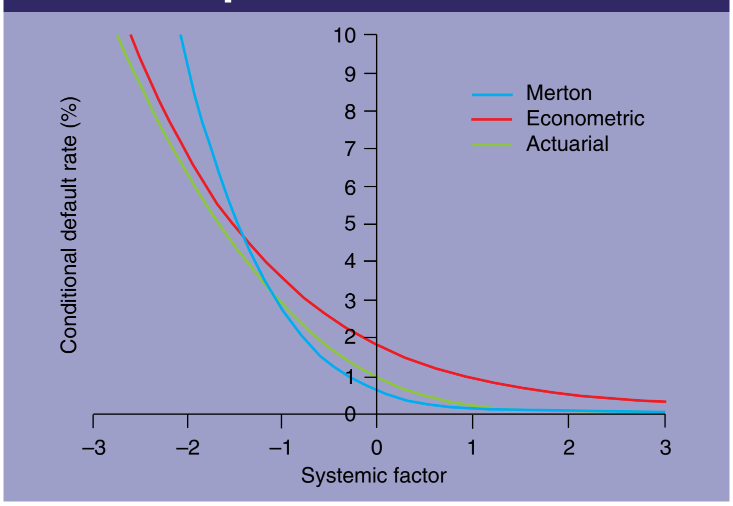

Dr H Ugur Koyluoglu is an associate consultant at Oliver, Wyman & Company. Andrew Hickman is in the risk management and quantitative analysis department at CSFP Capital, an indirect subsidiary of Credit Suisse Financial Products. The authors wish to thank Andrew Cross, Frank Diebold, Tom Garside, Marc Intrater, Andrew Kuritzkes, Hashem Pesaran, Til Schuermann, James Wiener and Tom Wilde for providing helpful comments and discussions. All errors are the responsibility of the authors alone. Opinions expressed herein are those of the authors alone and do

[10] For example, Isda (1998) and Roberts & Wiener (1998) compare the results of several models on test portfolios. The former finds that model results are fairly consistent, while the latter finds that the models may produce quite different results for the same portfolio using parameters independently selected for each model

REFERENCES

Carty L and D Lieberman, 1996
Corporate bond defaults and default rates 1938–1995
Moody's Investors Service Global Credit Research, January

Credit Suisse Financial Products, 1997
CreditRisk+ – a credit risk management framework

Freund J, 1992
Mathematical statistics
Fifth edition, Prentice Hall, New Jersey

Gupton G, C Finger and M Bhatia, 1997
CreditMetrics technical document
Morgan Guaranty Trust, 1997

International Swaps and Derivatives Association, 1998
Credit risk and regulatory capital
March

Kealhoffer S, 1995
Managing default risk in derivative portfolios
In Derivative Credit Risk: Advances in Measurement and Management
Risk Publications, London, 1995

Kuritzkes A, 1998
Transforming portfolio management
Banking Strategies, July–August, pages 57–62

Merton R, 1974
On the pricing of corporate debt: the risk structure of interest rates
Journal of Finance 29, pages 449–470

Roberts J and J Wiener, 1998
Handle with care
Oliver, Wyman & Company working paper, 1998

Vasicek O, 1987
Probability of loss on loan portfolio
KMV Corporation, February 12

Wilson T, 1997
Portfolio credit risk, parts I and II
Risk September 1997, pages 111–117 (part I) and October 1997, pages 56–61 (part II)