

Modest Means

11 November 2009 version
Forthcoming in *Risk Magazine*

Please do not quote or distribute this draft

Jon Frye
Senior Economist
Federal Reserve Bank of Chicago
230 South LaSalle Street
Chicago, IL 60604
Jon.Frye@chi.frb.org
312-322-5035

The author thanks numerous colleagues who have commented on earlier drafts. These include Michael Gordy, Mark Levonian, Jim Nelson, Ed Pelz, two anonymous referees, and participants in the FRB-Chicago Quantitative Congress and the Federal Interagency Quantification Forum. If errors remain, they are the author's. Any views expressed are the author's and do not necessarily represent the views of the management of the Federal Reserve Bank of Chicago or the Federal Reserve System.

Modest Means

This study highlights a mistake that is long overdue for correction: credit loss models are not protected against Type I error. As such, they can mislead their users. It is shown here that Type I error can be controlled, even in a credit loss model.

Scientific researchers protect against Type I error by using statistical hypothesis tests. Each of two models is fit to data. The simpler model is rejected only if it seems inadequate. The probability of falsely rejecting the simpler model—that is, the probability of making a Type I error—is usually kept at 5% or less.

To date, credit loss models have not been protected against Type I error because they have not been tested. Instead, researchers fit default distributions to default data and LGD distributions to LGD data. They stop short of producing a credit loss model, let alone testing one. Practitioners assemble their loss models from default models, LGD models, and an assumed copula, as if the resulting combination were automatically protected against Type I error.

This study presents the first effort to control Type I error in a credit loss model. It connects a default model to an LGD model and calibrates the resulting credit loss model to credit loss data. Then it tests against a suitable null hypothesis. This demonstrates that Type I error can be controlled. The analysis is performed twice: first using an extremely simple LGD model, then using a more nuanced one.

Although the main contribution is to introduce hypothesis testing and control of Type I error, the results of the tests shed light on the current science of credit loss modeling. Neither credit loss model is able to reject the null hypothesis. Since this might suggest a defect in the data, an experiment is performed with simulated data from one of the credit models. Most often, hundreds of years of simulated data are required to reject the null hypothesis. Various models are compared graphically. None appear quite different from the null. All three strands of evidence—historical data, simulated data, and graphical comparison—suggest that the null hypothesis may be genuinely difficult to improve upon with the short data sets now available.

Therefore, the null hypothesis developed here is put forward as a standard against which other combinations of theory and data might be tested. Outperforming such a null hypothesis is the minimum performance needed to protect against Type I error.

Credit loss with fixed LGD

The first credit loss model assumes that LGD is fixed. Such an assumption, fortunately, is not often considered these days. It is now generally acknowledged that LGD varies with the default rate. The idea of fixed LGD has been repudiated by work with LGD data and LGD models.

The current study has a different subject, credit loss models. Since a credit loss model is different from an LGD model (and different from a default model), it is worthwhile to consider the hypothesis of fixed LGD in this new context. A more nuanced LGD model is considered later.

To complete the credit loss model, we choose the most common model of default, the Vasicek distribution.¹ This applies only to a large homogeneous portfolio, which is assumed.

In most cases a choice must be made of the copula that connects the distribution of default and the distribution of LGD. Those three components imply the joint distribution of the default rate and the LGD rate. Assuming all exposure amounts are equal, credit loss is the product of default and LGD. The distribution of credit loss can then be inferred by conventional techniques.

In the current model, there is a shortcut. The credit loss rate equals the product of the LGD rate (which is fixed at the value of expected LGD, ELGD) and the default rate. The default rate depends on the probability of default (PD), correlation (ρ), and an unobserved standard normal systematic risk factor, Z . Specifically,

$$(1) \quad Loss3(Z; PD, \rho, ELGD) = ELGD \Phi \left[\frac{\Phi^{-1}(PD) - \sqrt{\rho} Z}{\sqrt{1-\rho}} \right],$$

where $\Phi[\cdot]$ symbolizes the standard normal cumulative distribution function. This loss function is converted to a probability density function (PDF) using the change-of-variables technique:

$$(2) \quad f_{Loss3}(loss; PD, \rho, ELGD) = \frac{\sqrt{1-\rho}}{\sqrt{\rho}} \frac{\phi \left[\frac{\Phi^{-1}(PD) - \sqrt{1-\rho} \Phi^{-1}(loss / ELGD)}{\sqrt{\rho}} \right]}{ELGD \phi \left[\Phi^{-1}(loss / ELGD) \right]},$$

where $\phi[\cdot]$ symbolizes the standard normal PDF. We refer to Expression (2) as PDF3, to signify that it contains three parameters.

PDF3 is called the alternative hypothesis, or H1. Under the rules of hypothesis testing, it must be compared to a simpler null hypothesis. Quite often, the null hypothesis is that a parameter equals zero. However, this approach is not suitable here. If either PD or ELGD is zero, no credit loss is possible. If correlation is zero, credit loss is always observed at exactly its expected level, and no credit risk is possible.

The null hypothesis in this case is that credit loss, itself, has a Vasicek distribution. The mean of this distribution is expected loss, symbolized as EL:


¹ Vasicek.

$$(3) \quad f_{Loss2}(loss; EL, \rho) = \frac{\sqrt{1-\rho}}{\sqrt{\rho}} \frac{\phi \left[\frac{\Phi^{-1}(EL) - \sqrt{1-\rho} \Phi^{-1}(loss)}{\sqrt{\rho}} \right]}{\phi[\Phi^{-1}(loss)]}$$

We refer to this expression as H0 (the null hypothesis that the credit loss distribution has this two-parameter form) or PDF2 (to signify that Expression (3) has two parameters). Statistically speaking, PDF2 represents a restriction on PDF3. When ELGD is set equal to 1.0 it drops out of PDF3. The Vasicek distribution then becomes a model of loss rather than a model of default, and its mean becomes EL rather than PD. PDF2 has no explicit role for the default rate, the LGD rate, or a copula.

The hypothesis test is conducted using data provided by Dr. Edward Altman.² Dr. Altman has computed sub-investment grade default rates and LGD rates and has published these data in his papers. Using Altman's published data makes the statistical test completely transparent. The data appear in the Sidebar and are graphed in Figure 1.

Figure 1



Under the alternative hypothesis, the log likelihood of the sample equals LnL3

$$(4) \quad \text{LnL3}(\{loss_t | 1982 \leq t \leq 2005\}; PD, \rho, ELGD) = \text{Log} \left[\prod_{t=1982}^{2005} f_{Loss3}(loss_t; PD, \rho, ELGD) \right]$$

² Altman, Brady, Resti, and Sironi; Altman and Pasternak.

Figure 2: Maximized values of LnL3




Figure 2 shows the result of maximizing this expression. For each value of ELGD, LnL3 is maximized in PD and ρ . H1 places no restriction on ELGD. Maximum LnL3 is about 67.3, if ELGD is allowed to take the unrealistic and impractical value of 1,000,000. Under H0, ELGD equals 1. Then, LnL3 equals 66.575.

A likelihood ratio test relies on the fact that, under the null hypothesis with many years of data, twice the difference of log likelihoods has a χ^2_1 distribution. With probability 95% this distribution produces values less than 3.84. Therefore, if the difference of log likelihoods exceeds 1.92 there is a rejection that a χ^2_1 variable has been observed. This implies a rejection of the null hypothesis. The decision rule is that H0 is rejected if H1 can produce a pickup of LnL3 greater than 1.92.

Figure 2 shows that LnL3 is never much more than 66.575. In the modest parlance of statistical hypothesis testing, there is no rejection of the null hypothesis. In more natural language, H0 is accepted. Of course, the acceptance of any statistical hypothesis is provisional on the development of theories and the testing of new data sets.

The development and testing of credit loss models is at an early stage. At a later stage, a picture might emerge. If there prove to be many cases where PDF3 rejects PDF2 for various sets of data, a single failure to reject would not be an important result. On the other hand, if it is found that PDF2 is not rejected by PDF3 using several data sets, a single rejection would not overturn the idea that PDF2 is sufficient. Under the rules of statistical hypothesis testing, null hypotheses can be rejected in 5% of cases for purely random reasons.

A credit loss distribution has been tested for the first time, using the Altman data. The result of the test is conclusive. There is no rejection of the null hypothesis. The null hypothesis states that a two-parameter distribution is adequate to describe the data. The two parameters are expected loss and correlation. Acceptance of the alternative would fail to protect against Type I error.

Intuition and Type I error

A somewhat dispiriting aspect of the test result is best to confront directly. PDF3 contains a poor model of LGD, but at least it contains one. It also contains Vasicek's model of the default rate, which can be rigorously derived from assumptions. PDF2, by contrast, has no concept of default or LGD. It is not derived from a theory. It arises by imposing a restriction on PDF3. One might expect that the theory-based PDF3 must outperform the impoverished PDF2 using any possible data set. Such is not the case.

A second aspect produces a degree of cognitive dissonance that is also best to confront directly. Researchers are familiar with calibrating distributions of default and LGD. They can feel that something is left out, or that information is wasted, when a credit loss distribution is calibrated to credit loss data. After all, the intuition goes, don't we know something about LGD rates from LGD data, and can't we use that to refine the estimates in a credit loss model?

There are two ways that this intuition can mislead. The first way is direct and easy to see. Any value of ELGD, other than its MLE, reduces the value of the likelihood function. For example, if ELGD is set equal to the average LGD rate (which is 57.3%), the maximum value of LnL is considerably less than under the null hypothesis, as can be seen in Figure 2. Although intuition may say that the parameter named ELGD should take the value 57.3% in any model, that value produces a credit loss distribution that is not as good as it can be.

Different likelihood functions have maxima at different parameter values; for example, the same explanatory variable can have different coefficients in different regression equations. In the current case, perhaps some model produces a maximum likelihood estimate of ELGD equaling 57.3%. This does not mean that each and every credit loss model will achieve its maximum when its ELGD parameter is set equal to 57.3%. Nothing is gained—and something is lost—if one overrides the data and the model with a preconception of the parameter values.

Besides compromising the fit, there is a second way that setting $ELGD = 57.3\%$ misleads: it asserts something that has no grounding in fact. Specifically, it asserts that three parameters are needed to describe the data. This is not supported by facts; it is an assumption. Generalizing this assumption can distort future decisions. Instead, decisions should be based on evidence. It has been shown that the Altman data do not reflect fixed LGD. To assert otherwise is too likely to commit a Type I error.

Hypothesis testing proceeds very deliberately. Theories are chosen. The theories are reduced to a PDF. The PDF contains everything that is theorized. Data are then chosen. Finally, the likelihood is maximized. Nothing is left out if these steps are followed. If they are not followed, something can be left out—namely, the control of Type I error.

The only real choices are the theory and the data. The fact that H_0 was accepted, when challenged by one theory and one set of data, does not mean that it will survive challenge by any theory and any data. However, rejecting a null hypothesis is a minimum criterion to assert that an alternative hypothesis has controlled for Type I error.

A good statistician would not assert that PDF3 is valid unless it proves useful. In general, statistical decision making refrains from accepting extra parameters unless they help make sense of the world. Statistical procedures manifest the modest principle of Occam's razor.

Statistical studies assume that what has been learned from the past can be applied today. This is an extrapolation, which is always a dangerous undertaking. The danger is greatest when the extrapolation does not even reflect learning, but instead reflects assumptions that are not grounded in fact. Control of Type I error is a most important tenet of statistical research.

Simulation experiment

The data of Altman and Pasternak (2005) reflect sub-investment grade borrowers. It is natural to enquire whether data on investment grade borrowers might not produce a different result. Perhaps there would be rejection of H_0 using data that are collected by other means, or that reflect other products, other times, or other places. These possibilities might all be investigated, even though any set of real-world data has shortcomings that can be criticized such as inconsistent definitions, noisy observation, limited quantity, and so forth.

It is more efficient in this case to simulate data that are free of real-world problems. This section performs the simulation experiment. To summarize the results, over 100 years of data are usually required to reject H_0 , even though PDF3 generates the data. This suggests that the result using Altman's data is not solely attributable to defects that might exist in that particular data set. On the contrary, the null hypothesis might be truly difficult to reject using the data sets available at present.

The simulation experiment assumes that PDF3 generates annual credit loss data. By making random draws from this distribution we obtain 20 years of simulated loss data. These data are employed in a hypothesis test with the customary 5% size. Both PDF2 and PDF3 are fit to the data using MLE. If the null hypothesis is not rejected, the data sample is augmented by 20 more years and the test is performed again. If there are continued failures to reject, the simulated sample expands in steps up to a maximum of 200 years. If a rejection does occur, a record is made of the number of years that give rise to the rejection. The whole procedure is repeated 100 times to obtain a distribution of the sample sizes that are sufficient to reject the null hypothesis.

**Figure 3. Years of data to reject PDF2
when PDF3 is true; PD = 10%**

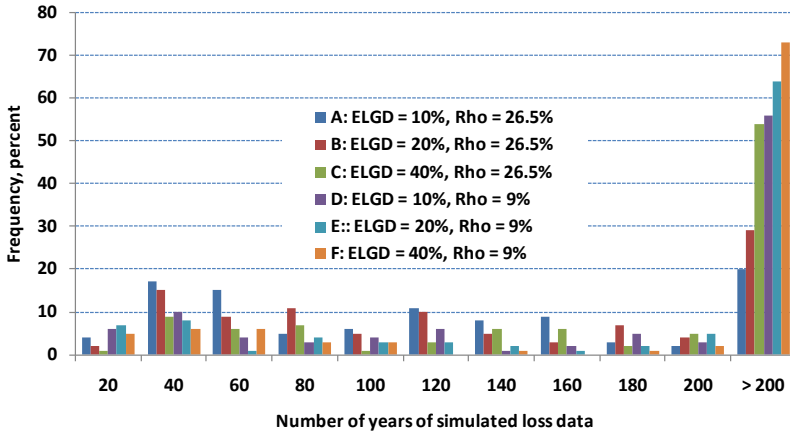


Figure 3 shows the results with PD = 10%. The six sets of bars reflect six combinations of ELGD and ρ . For example, combination “F” reflects ELGD = 40% and $\rho = 9\%$. If this combination of parameters generates the data, over 200 years of data are required to reject H0 in 73% of cases. At the other extreme, combination “A” has the least value of ELGD (10%) and the greatest value of correlation (26.5%). Data from this variant of PDF3 are able to reject H0 more often. Still, the median time to rejection exceeds 100 years.

Similar sets of experiments were conducted assuming PD = 2% or PD = 0.40%. For each combination of ELGD and ρ , lower values of PD require more years of data to reject H0; rejections of the null are even less likely when borrowers are investment grade.

These simulations show that even if PDF3 were true, a short data set is unlikely to lead to a rejection of PDF2. Therefore, a test result accepting the null is not solely due to defects in real-world data, other than the short duration of the data set. With the short data sets now available, the null hypothesis might be truly difficult to improve.

Credit loss with variable LGD

An early criticism of the fixed LGD hypothesis was that LGD rates are apt to vary in concert with default rates.³ Several LGD models have been published to address this. Perhaps the most rigorously developed is that of Michael Pykhtin.⁴ Pykhtin’s model provides a distribution of the LGD rate that Pykhtin and others have calibrated to LGD data. To say the least that can be said, this is a marked improvement over the idea that LGD is fixed. But it is another question whether it brings forth an improvement in a credit loss distribution.

³ Frye 2000

⁴ Pykhtin.

To derive the distribution we choose the comonotonic copula: greater values of the LGD rate accompany greater values of the default rate. This requires no additional parameters and can be modeled using a single standard normal risk factor symbolized as Z . The default rate is assumed to have the Vasicek distribution.

Pykhtin allows three parameters in the relation between Z and the LGD rate. We symbolize the parameters as μ , σ , and ρ_{LGD} , to distinguish the last from ρ . Roughly stated, these parameters measure the log of the initial value of collateral, the dispersion of its ending value, and the correlation between its return and the risk factor Z . Pykhtin states the relation between Z and the conditionally expected LGD rate:

(5)

$$LGD(Z; \mu, \sigma, \rho_{LGD}) = \Phi \left[\frac{-\mu / \sigma - \rho_{LGD} Z}{\sqrt{1 - \rho_{LGD}^2}} \right] - \text{Exp}[\mu + \sigma^2(1 - \rho_{LGD}^2) / 2 + \sigma \rho_{LGD} Z] \Phi \left[\frac{-\mu / \sigma - \sigma(1 - \rho_{LGD}^2) - \rho_{LGD} Z}{\sqrt{1 - \rho_{LGD}^2}} \right]$$

Expression (5) is bounded on $[0, 1]$ and is strictly decreasing in the risk factor Z . Its graph would be described as an “S” curve; Z ranges over the real line but LGD is bounded. If Z takes a large value, both the LGD rate and the default are low. Credit loss is the product. We refer to the loss function as $g(Z)$:

$$(6) \quad g(Z; PD, \rho, \mu, \sigma, \rho_{LGD}) = \Phi \left[\frac{\Phi^{-1}(PD) - \sqrt{\rho} Z}{\sqrt{1 - \rho}} \right] LGD(Z; \mu, \sigma, \rho_{LGD})$$

As before, the loss function is transformed to a PDF, referred to as PDF5, which contains five parameters and constitutes the alternative hypothesis:

$$(7) \quad f_{Loss5}(loss) = \left| \frac{\partial g^{-1}(loss)}{\partial loss} \right| \phi(g^{-1}(loss))$$

The required derivative can be stated in closed form (it is messy), but the inverse function must be computed numerically. The expected value of the distribution—expected loss—is provided by Pykhtin:

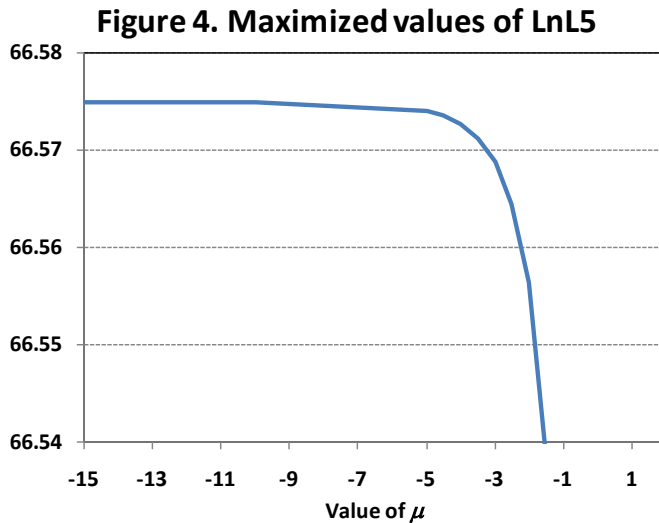
$$(8) \quad \begin{aligned} EL[PD, \rho, \mu, \sigma, \rho_{LGD}] = & \Phi_2[\Phi^{-1}(PD), -\mu / \sigma, \rho, \rho_{LGD}] \\ & - \text{Exp}[\mu + \sigma^2 / 2] \Phi_2[\Phi^{-1}(PD) - \sigma \rho_{LGD}, -\mu / \sigma - \sigma, \rho, \rho_{LGD}], \end{aligned}$$

where $\Phi_2[x, y, \rho]$ refers to the bivariate standard normal CDF with correlation equal to ρ .

The null hypothesis is the same as before. We estimate its parameters as if ELGD equals 1.0. To produce this in Pykhtin’s model requires a low value of μ . Lower μ implies lower collateral and

greater LGD. This happens quickly because it is an exponential relationship, but limiting behavior is involved. As a practical matter, $\mu = -10$ implies $LGD \cong 1.0$. The alternative hypothesis, H1, is that all five parameters of PDF5 are required to give nuance to the loss distribution.

Again the Altman data is employed to perform the statistical test. Figure 4 shows the maximized value of LnL5 for a range of μ .



H1 permits any value of μ , so the maximum is 66.575. H0 restricts μ to something less than -10. Obviously, this produces the same value of LnL5. The pickup is zero, which is as low as it can be. The null hypothesis is accepted.

To summarize this section, a second credit loss distribution has been tested. H1 allows LGD to vary in concert with the default rate. H0 is the same as in the previous test. Using the Altman data, the null hypothesis is again accepted.

This result, perhaps even more than the first, has a dispiriting aspect. Although the LGD model is clearly superior to what has gone before, it is apparently the case that a good LGD model does not guarantee a good credit loss model. Improvements to credit loss models, LGD models, and the default/LGD copula are to be encouraged, but the associated credit loss models must be tested if they are to be trusted.

Graphical comparison

Well-regarded theories of default and LGD have been unable to reject a hypothesis that has no theoretic content. To make this result more understandable, this section compares the PDFs graphically. Of course, a graphical comparison is only suggestive and its suggestions would be overturned by a definitive rejection of H0 using data.

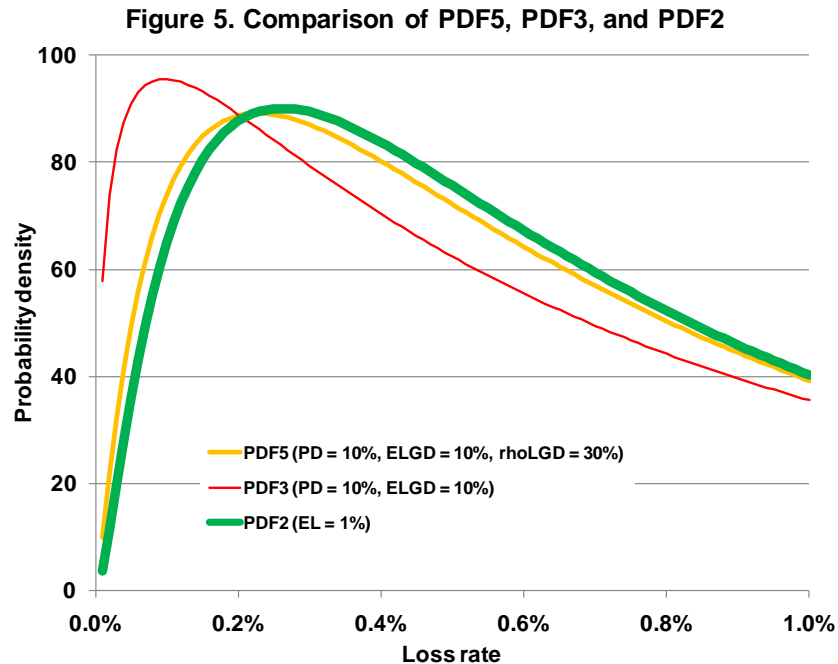


Figure 5 compares variants of PDF5, PDF3, and PDF2. PDF5 has the following parameter values: PD = 10%, $\rho = 9.7\%$, $\mu = .0492$, $\sigma = 30\%$, $\rho_{LGD} = 30\%$. To make PDF3 most closely comparable, it has the same value of PD (10%), the same value of EL (1%, using Expression (8)), and the same variance (0.01%). These constraints imply that the value of ρ in PDF3 is equal to 26.5%. Thus, PDF3 is nothing other than the parameter combination “A” that is simulated in Figure 3. PDF2 shares the same EL and the same variance, which implies ρ equal 10.6%.

The three PDFs have different values of ρ , and this can produce cognitive dissonance that traces back to an assumption. The assumption is that the value of correlation in a credit model need not reflect the distribution of credit data; it can be calibrated to asset returns. This assumption is testable.⁵ The current demonstration compares loss distributions. To achieve a common variance, the values of the respective correlation parameters must differ from each other.

The point of Figure 5 is that PDF2 is closer to PDF5 than to PDF3. This suggests that H0 is even more difficult to reject if PDF5 generates the data than if PDF3 generates the data. But the simulation experiment has already shown that over 100 years of data from PDF3 are usually required to reject H0. This suggests that it might be extremely difficult for a real-world data set to reject H0 if it exhibits some systematic LGD risk.

One can imagine a sequence of variants of PDF5 for a sequence of values of ρ_{LGD} , beginning with $\rho_{LGD} = 0$, that is, beginning with PDF3. As the value of ρ_{LGD} increases from zero, PDF5 moves closer to PDF2. This continues until some moderate value of ρ_{LGD} is reached; for the constraints used here, the value is approximately 30%. Beyond this value of ρ_{LGD} , progress is reversed and PDF5 moves back toward PDF3. At high-enough values of ρ_{LGD} (about 60% for these

⁵ Frye 2008

constraints), PDF5 again resembles PDF3. With extreme values of ρ_{LGD} , PDF5 becomes a highly skewed distribution with mode near 0%. This sequence of PDFs suggests that, if H0 is to be rejected by a statistical test, the data would probably exhibit systematic LGD risk that is either very high or very low.

Together with the simulation experiment, this graphical comparison suggests that H0 might be most difficult to reject when systematic LGD risk is moderate, when PD is less than 10%, when ELGD is greater than 10%, and when correlation is less than 26.5%. This set of characteristics is believed to be common in real-world data sets. Therefore, the null hypothesis employed appears to be a worthy challenge to researchers' model-making abilities. It is put forward as a standard of comparison for future tests of credit loss models.

Conclusion

The last five years of Risk Magazine's Cutting Edge section contains six articles that address modeling the LGD rate and twelve articles that address modeling the default rate. Presumably, each of these is meant to contribute to the understanding of the key variable, credit loss. However, the key variable has remained unanalyzed until now.

If a credit loss model is assembled from pieces, it contains a default rate, an LGD rate, and a copula to connect them. The default model must contain at least two parameters, usually, PD and ρ . The LGD model must also contain at least two parameters, unless LGD is imagined to be a single fixed number. The copula might contain additional parameters that need to be estimated. At the minimum, the resulting credit loss distribution has four parameters: two from the default side and two from the LGD side. These four parameters must be calibrated with a data set spanning a limited historical period.

A hypothesis test challenges a credit loss distribution to reject a suitable null hypothesis. The null hypothesis employed here has two parameters. It is not rejected using real-world data. Rejection is not frequent using simulated data. Graphical comparisons suggest a reason: credit loss distributions do not look much different from the null hypothesis. It is far from certain that this null hypothesis can be improved with the short data sets that are now available.

To fill the practical need for credit loss models, practitioners have made assumptions. Assumptions cause little harm if the consequence is limited to understanding the past. Unfortunately, assumptions propagate to the present and future. This danger is great enough to warrant the name "Type I error," and control of Type I error is a precept of statistical research. This study shows that Type I error can be controlled in a credit loss model. Until credit loss models routinely protect against Type I error, risk managers are wise to calibrate not only their models, but also the credence they give to the results.

This study is the first to derive, calibrate, and test distributions of the credit loss rate. Neither distribution significantly improves upon an extremely simple null hypothesis. Simulation experiments and graphical evidence suggest that the null hypothesis might be truly difficult to

reject with the short data sets that are now available. The null hypothesis is therefore put forward as a standard of comparison for future research. Though the null hypothesis might be rejected by a more nuanced alternative, it is dangerous to assume that particular nuances are true if they are not supported by facts. A more modest approach, based on standard statistical decision making, is less likely to mislead.

References

Edward I. Altman, Brooks Brady, Andrea Resti, and Andrea Sironi
The Link between Default and Recovery Rates: Theory, Empirical Evidence, and Implications
Journal of Business, 2005, vol. 78, no. 6

Edward I. Altman with Brent Pasternack
High Yield and Distressed Debt Default and Return Report: Third-Quarter 2005 Update
November 2005

Jon Frye
Collateral Damage
Risk Magazine
April 2000

Jon Frye
Correlation and asset correlation in the structural portfolio model
Journal of Credit Risk
Volume 4/Number 2, Summer 2008

Alexander Mood, Franklin Graybill, and Duane Boes
Introduction to the Theory of Statistics

Michael Pykhtin
Unexpected recovery risk
Risk Magazine
August 2003

Oldrich Vasicek
Loan portfolio value
Risk Magazine
December 2002

Sidebar

Default, LGD, and credit loss data

	<u>Annual rates</u>		
	Default	LGD	Loss
1982	3.19%	61.40%	1.96%
1983	1.09%	44.30%	0.48%
1984	0.84%	51.40%	0.43%
1985	1.71%	54.10%	0.93%
1986	3.50%	65.50%	2.29%
1987	1.34%	38.00%	0.51%
1988	2.66%	56.40%	1.50%
1989	4.29%	61.70%	2.65%
1990	10.14%	76.60%	7.77%
1991	10.27%	64.00%	6.57%
1992	3.40%	49.90%	1.70%
1993	1.11%	43.40%	0.48%
1994	1.45%	60.60%	0.88%
1995	1.90%	59.40%	1.13%
1996	1.23%	48.10%	0.59%
1997	1.25%	45.80%	0.57%
1998	1.60%	64.10%	1.03%
1999	4.15%	72.10%	2.99%
2000	5.07%	73.60%	3.73%
2001	9.80%	74.50%	7.30%
2002	12.79%	74.70%	9.55%
2003	4.66%	54.50%	2.54%
2004	1.25%	42.30%	0.53%
2005	2.26%	39.60%	0.89%
Annual averages:	3.79%	57.33%	2.46%

Source: Edward I. Altman with Brent Pasternack

High Yield and Distressed Debt Default and Return Report: Third-Quarter 2005 Update
November 2005

Figure 7. Default rates and Losses. "Default Rate %" and "Weighted Average Price after Default"

Source for data 1987 (correction for strategic default):

The Link between Default and Recovery Rates: Theory, Empirical Evidence, and Implications

Edward I. Altman, Brooks Brady, Andrea Resti, and Andrea Sironi

Journal of Business, 2005, vol. 78, no. 6, Table 1, page 2210