The Journal of Credit Risk (109–140) Volume 8/Number 1, Spring 2012

Credit loss and systematic loss given default

Jon Frye
Federal Reserve Bank of Chicago, 230 South LaSalle Street, Chicago,
IL 60604, USA; email: jon.frye@chi.frb.org

Michael Jacobs Jr.
Office of the Comptroller of the Currency, 1 Independence Square, Washington,
DC 20219, USA; email: michael.jacobs@occ.treas.gov

This paper presents a simple and effective model of systematic loss given default. It is simple in that it uses only parameters appearing in standard models. It is effective in that it survives statistical testing against more complicated models.

1 INTRODUCTION

Credit loss varies from period to period, both because the default rate varies and because the loss given default (LGD) rate varies. The default rate has been tied to a firm's probability of default (PD) and to factors that cause default. The LGD rate has proved more difficult to model because continuous LGD is more subtle than binary default and because LGD data is scarcer and lower in quality.

Studies show that the two rates vary together systematically (see Altman and Karlin (2010) and Frye (2000)). Systematic variation works against the lender, who finds that an increase in the number of defaults coincides with an increase in the percentage that is lost in a default. Lenders should therefore anticipate systematic LGD within their credit portfolio loss models, which are required to account for all material risks.

This paper presents a model of systematic LGD that is simple and effective. It is simple in that it uses only parameters that are already part of standard models. It is effective in that it survives statistical testing against more complicated models. It may,

The authors thank Irina Barakova, Terry Benzschawel, Andy Feltovich, Brian Gordon, Paul Huck, J. Austin Murphy, Ed Pelz, Michael Pykhtin and May Tang for comments on previous versions of this paper, as well as participants at the 2011 Federal Interagency Risk Quantification Forum, the 2011 International Risk Management Conference and the First International Conference on Credit Analysis and Risk Management. Any views expressed are the authors' own and do not necessarily represent the views of the management of the Federal Reserve Bank of Chicago, the Federal Reserve System, the Office of the Comptroller of the Currency or the US Department of the Treasury.

therefore, serve for comparison in tests of other models of credit risk as well as for the purposes of advancing models of credit spreads that include premiums for systematic LGD risk.

The LGD model is derived in Section 2. Section 3 discusses both the style of the statistical tests to and their direct focus on credit loss modeling. Sections 4–6 develop the model of credit loss for the finite portfolio, discuss the data to be used for calibration and testing, and introduce alternative hypotheses. Sections 7 and 8 perform the statistical tests. Section 7 separately tests each exposure "cell", which is a particular combination of rating grade and seniority. Section 8 brings together sets of cells: all loans, all bonds or all instruments. Having survived statistical testing, the LGD model is applied in Section 9, which precedes the conclusion in Section 10.

2 THE LOSS GIVEN DEFAULT MODEL

This section derives the LGD model. It begins with the simplest portfolio of credit exposures and assumes that loss and default vary together. This assumption by itself produces a general formula for the relationship of LGD to default. The formula depends on the distributions of loss and default. We note that different distributions of loss and default produce similar relationships, so we specify a distribution based on convenience and ease of application. The result is the specific LGD function that appears as (2.3).

The "asymptotically fine-grained homogeneous" portfolio of credit exposures has enough same-sized exposures that there is no need to keep track of individual defaults (see Gordy (2003)). Only the rates of loss, default and LGD matter. These rates are random variables, and each one has a probability distribution.

Commonly, researchers develop distributions of LGD and default; from these and a connection between them, a distribution of credit loss can be simulated. The loss model might or might not be inferred explicitly, and it would not be tested for statistical significance. This is one way to generate credit loss models, but it does not guarantee that a model is a good one that has been properly controlled for type I error. This is unfortunate, because credit loss is the variable that can cause the failure of a financial institution.

Because loss is the most important random variable, it is the loss distribution that we wish to calibrate carefully, and it is the loss model that we wish to control for error. We denote the cumulative distribution functions (CDFs) of the rates of loss and default by $\mathrm{CDF}_{\mathrm{loss}}$ and $\mathrm{CDF}_{\mathrm{DR}}$.

Our first assumption is that greater default rates and greater loss rates go together. This assumption is much less restrictive than the common assumption that greater default rates and greater LGD rates go together. The technical assumption is that the asymptotic distributions of default and loss are comonotonic. This implies that the loss rate and the default rate take the same quantile, q, within their respective distributions:

$$\mathrm{CDF}_{\mathrm{loss}}[\mathrm{loss}] = \mathrm{CDF}_{\mathrm{DR}}[\mathrm{DR}] = q \tag{2.1}$$

The product of the default rate and the LGD rate equals the loss rate. Therefore, for any value of q, the LGD rate equals the ratio of loss to default, which in turn depends on q and on inverse CDFs:

$$\mathrm{LGD} = \frac{\mathrm{CDF}_{\mathrm{loss}}^{-1}[q]}{\mathrm{CDF}_{\mathrm{DR}}^{-1}[q]} = \frac{\mathrm{CDF}_{\mathrm{loss}}^{-1}[\mathrm{CDF}_{\mathrm{DR}}[\mathrm{DR}]]}{\mathrm{DR}} \tag{2.2}$$

This expresses the asymptotic LGD rate as a function of the asymptotic default rate and it holds true whenever the distributions of loss and default are comonotonic. This function might take many forms depending on the forms of the distributions. Since LGD is a function of default, one could use (2.2) to infer a distribution of LGD and study it in isolation; however, we keep the focus on the distribution of loss and on the nature of an LGD function that is consistent with the distribution of loss.

In particular, we begin with a loss model having only two parameters. If a two-parameter loss model were not rich enough to describe credit loss data, a more complicated model could readily show this in a straightforward statistical test. The same is true of the default model. Therefore, our provisional second assumption is that both credit loss and default have two-parameter distributions in the asymptotic portfolio.

Testing this assumption constitutes the greater part of this study. Significant loss models with more than two parameters are not found; a two-parameter loss model appears to be an adequate description of the loss data used here. Therefore, this section carries forward the assumption of two-parameter distributions of loss and default.

In principle, any two-parameter distributions could be used for the CDFs in (2.2). In practice, we compare three distributions: Vasicek, beta and lognormal. We ensure that each has the same mean and the same standard deviation. To obtain values that are economically meaningful, we turn to the freely available credit loss data published by Altman and Karlin (2010) for high-yield bonds for the period 1989–2007. The means and standard deviations appear in the first column of Table 1 on the next page. The other three columns describe distributions that share these statistics. Figure 1 on page 113 compares the variants of (2.2) that result.

As Figure 1 on page 113 illustrates, the three distributions produce approximately the same LGD–default relationship. They differ principally when the default rate is low. This is the range in which they would be the most difficult to distinguish empirically, because a low default rate generates few defaults and substantial random variation in annual average LGD. The lognormal distribution produces the relationship with the lowest overall slope; however, of the three distributions, the lognormal has the fattest tail.

TABLE 1 Calibration of three possible distributions to loss and default: data from Altman and Karlin (2010).

	Vasicek distribution $\mathbf{0 < x < 1}$	**Beta distribution** $\mathbf{0 < x < 1}$	**Lognormal distribution** $\mathbf{0 < x < \infty}$
(a) Functional forms			
PDF $[x]$	$\dfrac{\sqrt{1-\rho}}{\sqrt{\rho}}\phi\dfrac{[(\Phi^{-1}[\mathrm{EL}]-\sqrt{1-\rho}\Phi^{-1}[x])/\sqrt{\rho}]}{\phi[\Phi^{-1}[x]]}$	$\dfrac{x^{a-1}(1-x)^{b-1}}{\mathrm{Beta}[a,b]}$	$\dfrac{\exp[-(\log[x]-\mu)^2/2\sigma^2]}{x\sqrt{2\pi}\sigma}$
CDF $[x]$	$\Phi\left[\dfrac{\sqrt{1-\rho}\Phi^{-1}[x]-\Phi^{-1}[\mathrm{EL}]}{\sqrt{\rho}}\right]$	$\displaystyle\int_0^x \frac{y^{a-1}(1-y)^{b-1}}{\mathrm{Beta}[a,b]}\,\mathrm{d}y$	$1-\Phi\left[\dfrac{\mu-\log[x]}{\sigma}\right]$
CDF$^{-1}[q]$	$\Phi\left[\dfrac{\Phi^{-1}[\mathrm{EL}]+\sqrt{\rho}\Phi^{-1}[q]}{\sqrt{1-\rho}}\right]$	x such that $q = \displaystyle\int_0^x \frac{y^{a-1}(1-y)^{b-1}}{\mathrm{Beta}[a,b]}\,\mathrm{d}y$	$\exp[\mu-\sigma\Phi^{-1}[1-q]]$
(b) Parameter settings to match mean and standard deviation of loss data			
Mean $= 2.99\%$	EL $= 0.0299$	$a = 0.9024$	$\mu = -3.867$
SD $= 3.05\%$	$\rho = 0.1553$	$b = 29.28$	$\sigma = 0.8445$
(c) Parameter settings to match mean and standard deviation of default data			
Mean $= 4.59\%$	PD $= 0.0459$	$a = 1.180$	$\mu = -3.369$
SD $= 4.05\%$	$\rho = 0.1451$	$b = 24.52$	$\sigma = 0.7588$

$\phi[\cdot]$ symbolizes the standard normal probability density function. $\Phi[\cdot]$ symbolizes the standard normal CDF.

FIGURE 1 LGD–default relationships for three distributions.

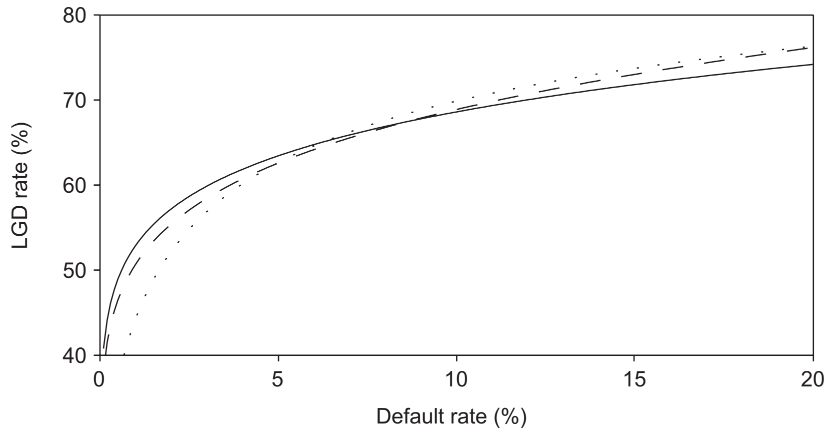

Solid line: lognormal distribution. Dashed line: Vasicek distribution. Dotted line: beta distribution.

Our choice between the distributions is guided by practical considerations. Unlike the beta distribution, the Vasicek distribution has explicit formulas for its CDF and its inverse CDF. Unlike the lognormal distribution, the Vasicek distribution constrains all rates to be less than 100%. Importantly, estimates of the Vasicek correlation parameter already exist within current credit loss models. This makes the Vasicek distribution by far the easiest for a practitioner to apply. Therefore, our third assumption is that loss and default obey the Vasicek distribution.

Our fourth assumption is that the value of ρ in $\mathrm{CDF}_{\mathrm{loss}}$ equals the value of ρ in $\mathrm{CDF}_{\mathrm{DR}}$. This assumption is testable. Alternative E, introduced later, tests by allowing the values to differ, but it does not find that the values are significantly different. Therefore, we carry forward the assumption that the values of ρ are the same.

Substituting the expressions for the Vasicek CDF and the inverse CDF into (2.2) produces the LGD function:

$$\left.\begin{aligned} \mathrm{LGD} &= \Phi\left[\Phi^{-1}[\mathrm{DR}] - \frac{\Phi^{-1}[\mathrm{PD}] - \Phi^{-1}[\mathrm{EL}]}{\sqrt{1-\rho}}\right]\Big/\mathrm{DR} \\ &= \frac{\Phi[\Phi^{-1}[\mathrm{DR}] - k]}{\mathrm{DR}} \\ k &= \frac{\Phi^{-1}[\mathrm{PD}] - \Phi^{-1}[\mathrm{EL}]}{\sqrt{1-\rho}} \end{aligned}\right\} \tag{2.3}$$

This expresses the asymptotic LGD rate as a function of the asymptotic default rate. These rates equal the conditionally expected rates for a single exposure. Equation (2.3) underlies the null hypothesis in the tests that follow.

The three parameters PD, expected loss (EL) and ρ combine to form a single quantity that we refer to as the LGD risk index and symbolize by k. If $\text{EL} = \text{PD}$ (that is, if expected LGD (ELGD) equals 1.0), then $k = 0$ and $\text{LGD} = 1$, irrespective of DR. Except when the LGD risk index equals 0, LGD is a strictly monotonic function of DR as shown in Appendix A. For commonly encountered values of PD, EL and ρ, k is between 0 and 2.

To recap, we derive the LGD function by making four assumptions. The first assumption is that a greater rate of credit loss accompanies a greater rate of default. This plausible starting place immediately produces a general expression for LGD, (2.2). The second assumption is that the distributions of loss and default each have two parameters. Later sections of this paper attempt, unsuccessfully, to find a statistically significant loss model with more parameters. The third assumption is that the distributions are specifically Vasicek. This assumption is a matter of convenience; distributions such as beta and lognormal produce similar relationships but they would be more difficult to implement. The fourth assumption is that the value of ρ estimated from default data also applies to the loss distribution. This assumption is testable, and it survives testing in later sections. The four assumptions jointly imply (2.3), which expresses the LGD rate as a function of the default rate. This LGD function is consistent with the assumption that credit loss has a two-parameter Vasicek distribution.

3 RESEARCH METHODS

This section discusses two research methods employed by this paper. We proceed in a slightly unusual way. Rather than showing the statistical significance of (2.3), we show the lack of significance of more complicated models that allow the LGD–default function to be steeper or flatter than (2.3). Also, we calibrate credit loss models to credit loss data. Rather than assume that the parameters of a credit loss model have been properly established by the study of LGD, we investigate credit loss directly.

This study places its preferred model in the role of the null hypothesis. The alternatives explore the space of differing sensitivity by allowing the LGD function to be equal to, steeper than or flatter than (2.3). The tests show that none of the alternatives have statistical significance compared with the null hypothesis. This does not mean that the degree of systematic LGD risk in (2.3) can never be rejected, but a workmanlike attempt has not yet met with success. Acceptance of a more complicated model that had not demonstrated significance would accept an uncontrolled probability of type I error.

A specific hypothesis test has already been alluded to. Equation (2.3) assumes that the parameter ρ appearing in CDF_{loss} takes the same value as the parameter ρ

appearing in CDF_{DR}. An alternative allows the two values of correlation to differ. This alternative is not found to be statistically significant in tests on several different data sets; the null hypothesis survives statistical testing.

We do not try every possible alternative model, nor do we test using every possible data set. It is impossible to exhaust all the possibilities. Still, these explorations and statistical tests have content. The function for systematic LGD variation is simple, and it survives testing. A risk manager could use the function as it is. If they prefer, they could test the function as we do. A test might show that (2.3) can be improved. Given, however, that several alternative LGD models do not demonstrate significance on a relatively long, extensive and well-observed data set, an attitude of heightened skepticism is appropriate. In any case, the burden of proof is always on the model that claims to impart a more detailed understanding of the evidence.

The second method used in this paper is to rely on credit loss data and credit loss models to gain insight into credit risk. By contrast, the models developed in the last century, such as CreditMetrics (Gupton *et al* (1997)), treat the distribution of credit loss as something that can be simulated but not analyzed directly. This, perhaps, can be traced back to the fact that twentieth century computers ran at less than 1% of the speed of current ones, and some shortcuts were needed. But the reason for modeling LGD and default is to obtain a model of credit loss. The model of credit loss should be the focus of credit loss research, and these days it can be.

We make this difference vivid using a comparison. Suppose a risk manager wants to quantify the credit risk for a specific type of credit exposure. Having only a few years of data, they may find it quite possible that the pattern of LGD rates has arisen by chance. It is concluded that the rates of LGD and default are independent, and a credit loss model is run accordingly. This two-stage approach never tests whether independence is valid using credit loss data and a credit loss model, and it provides no warrant for this elision.

Single-stage methods are to be preferred because each stage of statistical estimation introduces uncertainty. A multistage analysis can allow the uncertainty to grow uncontrolled. Single stage methods can control uncertainty. One can model the target variable, credit loss, directly, and quantify the control of type I error.

The first known study to do this is Frye (2010), which tests whether credit loss has a two-parameter Vasicek distribution. One alternative is that the portfolio LGD rate is independent of the portfolio default rate.[1] This produces an asymptotic distribution

[1] The LGD of an individual exposure, since it is already conditioned on the default of the exposure, is independent of it. Nonetheless, an LGD can depend on the defaults of other firms or on their LGDs. This dependence between exposures can produce correlation between the portfolio LGD rate and the portfolio default rate, thereby affecting the systematic risk, systematic risk premium and total required credit spread on the individual loan.

of loss that has three parameters: ELGD, PD and ρ. The tests show that, far from being statistically significant, the third parameter adds nothing to the explanation of loss data used.

This illustrates the important difference that we touched upon earlier. If LGD and default are modeled separately, the implied credit loss distribution tends to contain all the parameters stemming from either model. By contrast, this paper begins with a parsimonious credit loss model and finds the LGD function consistent with it. If a more complicated credit loss model were to add something important, it should demonstrate statistical significance in a test.

We hypothesize that credit loss data cannot support extensive theorizing. This hypothesis is testable, and it might be found wanting. Nevertheless, the current research represents a challenge to portfolio credit loss models running at financial institutions and elsewhere. If those models have not demonstrated statistical significance against this approach, they can be seriously misleading their users.

The current paper extends Frye (2010) in three principal ways. First, it derives and uses distributions that apply to finite-size portfolios. Second, it controls for differences of rating and differences of seniority by using Moody's exposure-level data. Third, it develops alternative models that focus specifically on the steepness of the relationship between LGD and default. These are the topics of the next three sections.

4 THE DISTRIBUTION OF CREDIT LOSS IN A FINITE PORTFOLIO

This section derives the distribution of loss for a portfolio with few exposures, taking a different approach from the pioneering work by Pykhtin and Dev (2002). Later sections use this distribution to test the LGD function against alternatives.

As usual, we must keep separate the concepts related to the population and the concepts related to the sample. Economic and financial conditions give rise to the population variables. These are the conditionally expected default rate, denoted by DR, and the conditionally expected LGD rate, denoted by LGD. The conditionally expected rates are connected by (2.3) or by one of the alternatives developed later in this paper. In a sample of credit loss data, the quantities of interest are the number of defaults, D, and the portfolio average LGD rate, $\overline{\mathrm{LGD}}$.

The derivation begins with DR. Conditioned on DR there is a distribution of D. Conditioned on D, there is a distribution of $\overline{\mathrm{LGD}}$. These distributions are independent. Their product is the joint distribution of D and $\overline{\mathrm{LGD}}$ conditioned on DR. The joint distribution of D and $\overline{\mathrm{LGD}}$ is transformed to the joint distribution of D and loss in the usual way. The marginal distribution of loss is found by summing over the number of defaults and removing the conditioning on DR. This produces the distribution of credit loss when the portfolio is finite.

At the outset we recognize two cases. In the first case, $D = 0$. In this case, credit loss equals zero. This case has probability equal to $(1 - \mathrm{DR})^N$. In the second case, when $D = d > 0$, a distribution of the portfolio average LGD rate is produced. Average LGD would approach normality for large D, according to the central limit theorem. We assume normality for all D for two reasons: convenience and the practical benefit that normality allows average LGD outside of the range [0,1]. This is important because the credit loss data includes portfolio years where $\overline{\mathrm{LGD}}$ is negative. The variance of the distribution is assumed equal to σ^2/d:

$$f_{\overline{\mathrm{LGD}}|D=d}[\overline{\mathrm{LGD}}] = \frac{1}{\sigma/\sqrt{d}}\phi\left[\frac{\overline{\mathrm{LGD}} - \mathrm{LGD}}{\sigma/\sqrt{d}}\right] \tag{4.1}$$

The conditional distribution of D and $\overline{\mathrm{LGD}}$ is then the product of the binomial distribution of D and the normal distribution of $\overline{\mathrm{LGD}}$:

$$f_{D,\overline{\mathrm{LGD}}|\mathrm{DR}}[d, \overline{\mathrm{LGD}}] = \mathrm{DR}^d(1 - \mathrm{DR})^{N-d}\binom{N}{d}\frac{1}{\sigma/\sqrt{d}}\phi\left[\frac{\overline{\mathrm{LGD}} - \mathrm{LGD}}{\sigma/\sqrt{d}}\right] \tag{4.2}$$

In a portfolio with uniform exposure amounts, the loss rate equals default rate times the LGD rate. We pass from the portfolio's LGD rate to its loss rate with the monotonic transformation:

$$\overline{\mathrm{LGD}} = N\,\mathrm{loss}/D, \qquad D = D \tag{4.3}$$

The Jacobian determinant is N/D. The transformed joint distribution is then:

$$f_{D,\mathrm{loss}|\mathrm{DR}}[d, \mathrm{loss}] = \mathrm{DR}^d(1 - \mathrm{DR})^{N-d}\binom{N}{d}\frac{N}{\sigma\sqrt{d}}\phi\left[\frac{N\,\mathrm{loss}/d - \mathrm{LGD}}{\sigma/\sqrt{d}}\right] \tag{4.4}$$

Summing over d, combining the two cases, and removing the conditioning on DR produces the distribution of credit loss in the finite portfolio:

$$\begin{aligned} f_{\mathrm{loss}}[\mathrm{loss}] = {} & I_{[\mathrm{loss}=0]}[\mathrm{loss}]\int_0^1 f_{\mathrm{DR}}[\mathrm{DR}](1 - \mathrm{DR})^N \,\mathrm{d}\,\mathrm{DR} \\ & + I_{[\mathrm{loss}>0]}[\mathrm{loss}]\int_0^1 f_{\mathrm{DR}}[\mathrm{DR}]\sum_{d=1}^{N} f_{D,\mathrm{loss}|\mathrm{DR}}[d, \mathrm{loss}]\,\mathrm{d}\,\mathrm{DR} \end{aligned} \tag{4.5}$$

where $f_{\mathrm{DR}}[\cdot]$ is the PDF of the Vasicek density with parameters PD and ρ. Distribution (4.5) depends on the parameters of the default distribution, PD and ρ. It also depends on any additional parameters of the LGD function. These consist solely of EL in the null hypothesis of (2.3) but include an additional parameter in the alternatives introduced later. Finally, the distribution depends on N, the number of exposures. As N increases without limit, (4.5) becomes the Vasicek distribution with mean equal to

FIGURE 2 Distributions of loss for asymptotic and finite portfolios.

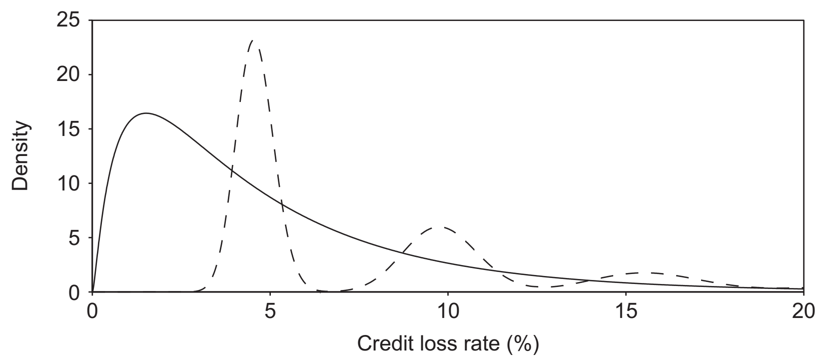

Solid line: asymptotic portfolio, PD = 10%, ELGD = 50%, ρ = 15%. Dashed line: portfolio with ten loans, PD = 10%, ELGD = 50%, ρ = 15%, σ = 1%.

EL. For small N, however, the decomposition of EL into PD and ELGD has an effect on the distribution of loss.

Figure 2 compares the distribution of loss for the asymptotic portfolio with the distribution for a portfolio containing ten exposures. Each distribution has EL $=5\%$ and $\rho=15\%$. Those values completely describe the distribution of credit loss in the asymptotic portfolio. For the sake of comparison, EL is assumed to decompose to PD $=10\%$ and ELGD $=50\%$. Credit loss in the finite portfolio has the distribution in (4.5). The point mass at zero loss has probability 43%; therefore, the area under the curve illustrated in Figure 2 is 57%. Assuming $\sigma=1\%$ produces distinct humps for one, two and three defaults. The hump for one default is centered at less than 5% loss, while the hump for three defaults is centered at greater than 15% loss. In other words, LGD tends be greater when there are more defaults.

Under the usual statistical assumptions (the parameters are stable over time and the variable DR is independent each year) the log of the likelihood function of the data is:

$$\mathrm{LnL}_{\mathrm{loss}}[\mathrm{loss}_1, \mathrm{loss}_2, \dots, \mathrm{loss}_T] = \sum_{t=1}^{T} \log[f_{\mathrm{loss}}[\mathrm{loss}_t]] \tag{4.6}$$

5 DATA

The data is twenty-seven years' worth of data drawn from Moody's Corporate Default Rate Service. An exposure "cell", the intersection of a rating grade and a seniority class, controls for both borrower quality and for exposure quality. A cell is assumed to be a homogenous portfolio of statistically identical exposures, as called for in the loss models.

Distributions of credit loss can say nothing about cases where the loss amount is unknown. Therefore, we restrict the definition of default to cases where Moody's observes a post-default price. By contrast, studies of default in isolation can include defaults that produce unknown loss. We refer to this less restrictive definition as "nominal default" and note that it produces default rates that are generally greater than the ones we present.

We delimit the data set in several ways. To have notched ratings available at the outset, the data sample begins with 1983. To align with the assumption of homogeneity, a firm must be classified as industrial, public utility or transportation and be headquartered in the United States. Ratings are taken to be Moody's "senior" ratings of firms, which usually correspond to the rating of the firm's long-term senior unsecured debt if such exists. To focus on cells that have numerous defaults, we analyze firms rated Baa3 or lower. We group the ratings C, Ca, Caa, Caa1, Caa2 and Caa3 into a single grade we designate "C". This produces five obligor rating grades altogether: Ba3, B1, B2, B3 and C.

To align with the assumption of homogeneity, debt issues must be dollar denominated, intended for the US market, and not guaranteed or otherwise backed. We define five seniority classes:

(1) senior secured loans (senior secured instruments with debt class designated "bank credit facilities");

(2) senior secured bonds (senior secured instruments with debt class designated "equipment trusts", "first mortgage bonds" or "regular bonds/debentures");

(3) senior unsecured bonds ("regular bonds/debentures" or "medium-term notes");

(4) senior subordinated bonds ("regular bonds/debentures");

(5) subordinated bonds ("regular bonds/debentures").

This excludes convertible bonds, preferred stock and certain other instruments.

A firm is defined as exposed in a cell-year if, on January 1, the firm has one of the five obligor ratings, is not currently in default and has a rated issue in the seniority class. A firm is defined to default if there is a record of nominal default and one or more post-default prices are observed. The LGD of the obligor's exposures in the cell equals 1.0 minus the average of such prices expressed as a fraction of par; there is exactly one LGD for each default. The default rate in the cell-year is the number of LGDs divided by the number of firms that are exposed, and the loss rate is the sum of the LGDs divided by the number of firms that are exposed. There is no correction for firms that are exposed to default for only part of the year, perhaps because their debts mature or because their ratings are withdrawn.

To make ideas concrete, consider the most populated cell: senior secured loans made to B2-rated firms. This cell has 1842 cell-years of exposure. However, public agencies began rating loans only in the latter half of the data sample; of the twenty-seven years of the data sample in total, only fourteen years contain loans to B2-rated firms. Of those fourteen years, only six record a default by a B2-rated firm that had a rated loan outstanding. Those six years contain all the information about the LGD–default relationship that is contained within the cell. In all, the cell generates fourteen annual observations on the three variables needed to calibrate the distribution of loss:

(1) N (the number of exposures);

(2) D (the number of defaults);

(3) loss (the sum of the LGDs divided by N, or zero if $D = 0$).

6 ALTERNATIVES FOR TESTING

This section presents alternative LGD functions that have an additional parameter and might provide a better fit to the data. Designed to focus on a particular question, the alternatives necessarily have a functional forms that appear more complicated than (2.3).

In general, a statistical alternative could have any number of functional forms. For example, one might test (2.3) against a linear LGD hypothesis:

$$\mathrm{LGD} = u + v\,\mathrm{DR} \tag{6.1}$$

The linear equation (6.1) can be mentally compared with the curved function for the Vasicek distribution that is illustrated in Figure 1 on page 113. If the straight line were wholly above the curved line, its expected loss would be too high. Therefore, the straight line and the curved line cross. If parameter v takes a positive value, as is likely, the lines cross twice. Therefore, a calibration of (6.1) would be likely to produce a straight line that is shallower than (2.3) at the left and steeper than (2.3) at the right. If this calibration were statistically significant, the verdict would be that (2.3) is too steep in some places and too flat in others.

Such an answer is not without interest, but we address a simpler question. If the LGD function of (2.3) does not adequately represent the data, we want to know whether a better function is steeper or flatter. Therefore our alternatives have a special feature: when the parameter takes a particular value, the alternative becomes identical to (2.3). When the parameter takes a different value, the alternative becomes steeper or flatter than (2.3). For all values of the parameter, the mathematical expectation of loss is equal to the value of parameter EL. When we test against such an alternative, we are testing for a difference in slope alone. Although the slope of the LGD–default relationship is not the only aspect of systematic LGD risk that is important, it has first-order importance.

Summarizing, we create alternatives that

- contain one more parameter than (2.3),
- collapse to (2.3) when the parameter takes a specified value,
- are steeper or flatter than (2.3) otherwise,
- produce the same value of EL irrespective of the parameter value.

Alternative A takes the following form, using (for convenience) the substitution $\mathrm{EL} = \mathrm{PD}\ \mathrm{ELGD}$:

$$\mathrm{LGDA} = \mathrm{ELGD}^{a}\, \Phi\left[\Phi^{-1}[\mathrm{DR}] - \frac{\Phi^{-1}[\mathrm{PD}] - \Phi^{-1}[\mathrm{EL}/\mathrm{ELGD}^{a}]}{\sqrt{1-\rho}}\right] \Big/ \mathrm{DR} \tag{6.2}$$

The additional parameter in alternative A is symbolized by a. If a takes the value zero, alternative A becomes identical to (2.3). If a takes the value 1.0, the function collapses to ELGD; in other words, when $a = 1$, alternative A becomes a model in which LGD is a constant in the asymptotic portfolio.

Figure 3 on the next page illustrates alternative A for five values of a. If parameter a takes a negative value, alternative A is steeper than (2.3). If parameter a takes a value greater than 1.0, alternative A is negatively sloped. Thus, alternative A can represent an entire spectrum of slopes of the LGD–default relationship: equal to, steeper than or flatter than the null hypothesis.

Irrespective of the value of ρ or the decomposition of EL into PD and ELGD, the expectation of loss equals the value of the parameter EL:

$$\begin{aligned} E[\mathrm{DR}\ \mathrm{LGDA}] &= E\left[\mathrm{ELGD}^{a}\, \Phi\left[\Phi^{-1}[\mathrm{DR}] - \frac{\Phi^{-1}[\mathrm{PD}] - \Phi^{-1}[\mathrm{EL}/\mathrm{ELGD}^{a}]}{\sqrt{1-\rho}}\right]\right] \\ &= \mathrm{ELGD}^{a}\, \frac{\mathrm{EL}}{\mathrm{ELGD}^{a}} \\ &= \mathrm{EL} \end{aligned} \tag{6.3}$$

Thus, the value of a affects the relationship between LGD and default but has no effect on EL.

We use alternative A to challenge the null hypothesis, but it is also an approximation of other LGD models that might be used instead. Appendix B compares alternative A with the LGD model of Pykhtin (2003) and finds that the approximation is quite good

FIGURE 3 Alternative A for five values of a.

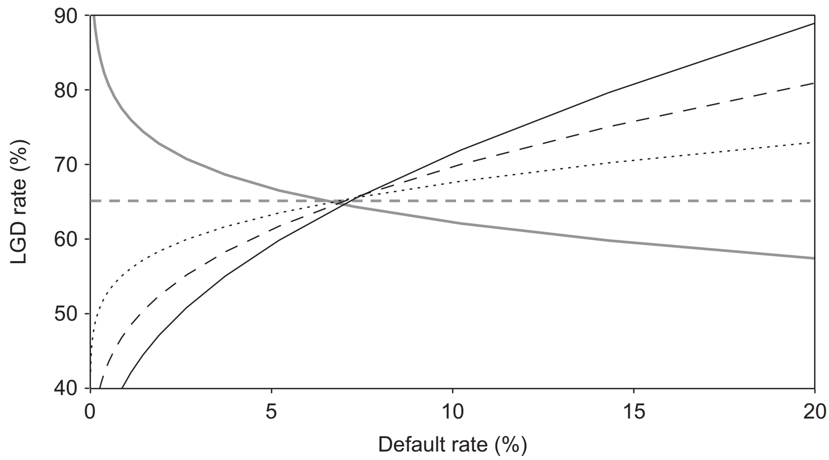

Solid black line: $a = -2$. Dashed black line: $a = -1$. Dotted black line: $a = 0$ (null hypothesis (2.3)). Dashed gray line: $a = 1$. Solid gray line: $a = 2$.

when the value of a is zero. This is exactly the case that survives statistical testing. Therefore, although we do not test explicitly against Pykhtin's model, we believe that we test against a model that is quite similar.

Alternatives B and C have forms similar to alternative A. In fact, the three alternatives can be identical to each other when applied to a single homogeneous cell. However, the assumption that parameter b (or c) is uniform across several cells is different from the assumption that parameter a is uniform across the cells. For this reason, alternatives B and C are defined here and applied in Section 8:

$$\mathrm{LGDB} = \mathrm{PD}^b\, \Phi\left[\Phi^{-1}[\mathrm{DR}] - \frac{\Phi^{-1}[\mathrm{PD}] - \Phi^{-1}[\mathrm{EL}/\mathrm{PD}^b]}{\sqrt{1-\rho}}\right] \Big/ \mathrm{DR} \tag{6.4}$$

$$\mathrm{LGDC} = \mathrm{EL}^c\, \Phi\left[\Phi^{-1}[\mathrm{DR}] - \frac{\Phi^{-1}[\mathrm{PD}] - \Phi^{-1}[\mathrm{EL}/\mathrm{EL}^c]}{\sqrt{1-\rho}}\right] \Big/ \mathrm{DR} \tag{6.5}$$

When parameter b (or c) takes the value 0, alternative B (or C) becomes identical to (2.3). When the parameters take other values, the associated LGD–default relationships become steeper or flatter. In any case, the mathematical expectation of loss equals the parameter EL.

The fourth alternative is designated as alternative E. It assumes that both loss and default have Vasicek distributions, but that the correlation relevant to loss, e, differs from the correlation relevant to default, ρ. Substituting the Vasicek CDF and inverse

FIGURE 4 Alternative E for three values of e.

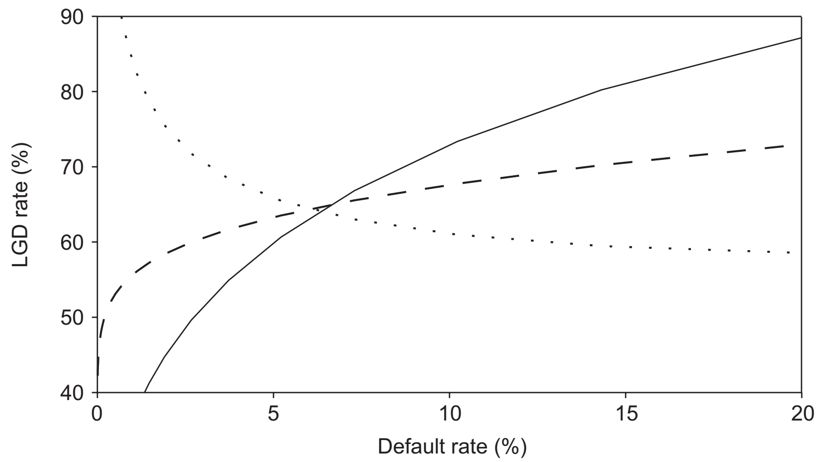

Solid line: $e = 19\%$. Dashed line: $e = 14.51\%$. Dotted line: $e = 10\%$.

CDF into (2.2):

$$\text{LGDE} = \left[\frac{\sqrt{\rho}\Phi^{-1}[\text{EL}] - \sqrt{e}(\Phi^{-1}[\text{PD}] - \sqrt{1-\rho}\Phi^{-1}[\text{DR}])}{\sqrt{\rho}\sqrt{1-e}} \right] \Big/ \text{DR} \qquad (6.6)$$

When $e = \rho$, (6.6) becomes identical to (2.3). Figure 4 illustrates this when e takes the value 14.51%. Relative to that, greater values of e make the function steeper and lesser values of e make the function flatter or negatively sloped. For any value of e, the mathematical expectation of loss equals the mean of the Vasicek loss distribution, which is parameter EL.

This section introduces four alternative LGD models. Each alternative contains an extra parameter that can allow LGD to be more (or less) sensitive to DR than the null hypothesis. The additional parameter has no effect on the expectation of loss. Therefore, the alternatives focus solely on the question of the slope of the LGD–default relationship. Later sections use the alternatives in statistical challenges to (2.3).

7 TESTING CELLS SEPARATELY

This section performs tests on the twenty-five cells, one cell at a time. Each cell isolates a particular Moody's rating and a particular seniority. Each test calibrates (4.5) twice: once using (2.3) and once using an alternative LGD function. The likelihood ratio statistic determines whether the alternative produces a significant improvement.

Judged as a whole, the results to be presented are consistent with the idea that (2.3) does not misstate the relationship between LGD and default.

As with most studies that use the likelihood ratio, it is compared with a distribution that assumes an essentially infinite, “asymptotic” data set. The statistic itself, however, is computed from a sample of only twenty-seven years of data. This gives the statistic a degree of sampling variation that it would not have in the theoretical asymptotic data. As a consequence, tail observations tend to be encountered more often than they should be. This creates a bias toward finding statistical significance. This bias strengthens a finding of no significance, such as that produced here.

Most risk managers are currently unable to calibrate all the parameters of a loss model using maximum likelihood estimation (MLE). A scientific finding that is valid only when MLE is employed would be useless to them. Instead, we calibrate mean parameters along the lines followed by practitioners. Our estimator for PD in a cell is the commonly used average annual default rate. Our estimator for EL is the average annual loss rate. ELGD is the ratio of EL to PD.

In the case of ρ, we find MLE to be more convenient than other estimators. (The next section checks the sensitivity of test results to the estimate of ρ.) We begin with the MLE found by maximizing the following expression of ρ within each cell:

$$\mathrm{LnL}_{\rho}[\rho] = \sum_{t=1983}^{2009} \log\left[\int_0^1 f_{\mathrm{DR}}[\mathrm{DR}_t]\,\mathrm{DR}^{d_t}(1-\mathrm{DR})^{N_t-d_t}\binom{N_t}{d_t}\,\mathrm{d}\,\mathrm{DR}\right] \tag{7.1}$$

where $f_{\mathrm{DR}}[\cdot]$ is the PDF of the Vasicek density with parameters $\widehat{\mathrm{PD}}$ and ρ. Consistent with the assumptions made in developing (2.3), this value of ρ is assumed to be valid for the loss distribution as well, except in the case of alternative E.

The parameter σ measures the random dispersion of an individual LGD around its conditionally expected value. This is needed to calibrate the distribution of loss for the finite portfolio, but σ has no role in the asymptotic LGD function of (2.3). From this perspective, σ is a “nuisance” parameter. To estimate it, we consider every cell-year in which there are two or more LGDs. In each such cell-year we calculate the unbiased estimate of the standard deviation. The dispersion measured around the data mean is less than the dispersion around any other number, including the conditional expectation. Therefore, the average standard deviation, 20.30%, should represent an underestimate of σ and should understate the contribution of purely random causes.

These parameters (PD, ELGD, ρ and σ) are the only ones required under the null hypothesis. The alternative has the extra parameter that controls the slope of the LGD–default relationship, and that parameter is estimated by MLE. Thus, the only parameter informed by the loss data in the context of the loss model is the additional parameter of the alternative hypothesis. This is believed to bias the test toward finding a statistically significant result, and this bias strengthens the findings of no statistical significance.

Table 2 on the next page shows summary statistics, parameter estimates and test statistics for each cell. The test statistics are stated as the difference between the maximum log likelihood using the alternative and the log likelihood using the null. Twice this difference would have the χ^2 distribution with one degree of freedom in the asymptotic portfolio. Differences greater than the 5% critical value of 1.92 are noted in bold. The test statistics for alternatives B and C would be identical to those presented for alternative A.

Along the bottom and on the right of Table 2 on the next page are averages. The overall averages at the bottom right corner contain the most important fact about credit data: it is low in quantity. The average cell has only thirty-one defaults, which is about one per year. Since defaults cluster in time, the average cell has defaults in only nine years, and only these years can shed light on the connection between LGD and default.

Not only is the data sparse, it has a low signal-to-noise ratio: the random variation of LGD, measured by $\sigma = 20.30\%$, is material compared with the magnitude of the systematic effect and the number of LGDs that are observed. A data set such as that used here, spanning many years with many LGDs, provides the best opportunity to see through the randomness and to characterize the degree of systematic LGD risk.

In Table 2 on the next page, there are two cells with log likelihood pickups greater than 1.92: loans to B2-rated firms and senior subordinated bonds issued by C-rated firms. This does not signal statistical significance because many tests are being performed. If twenty-five independent tests are conducted, and if each has a size of 5%, then two or more nominally significant results would occur with probability 36%. Of the two nominally significant results, one cell is estimated steeper than (2.3) and one cell is estimated flatter than (2.3). Nothing about this pattern suggests that the LGD function of (2.3) is either too steep or too flat.

Considering all twenty-five cells including the twenty-three cells that lack nominal significance, there is about a 50:50 split. About half the cells have an estimated LGD function that is steeper than (2.3) and about half have an estimated LGD function that is flatter than (2.3). A pattern like this would be expected if the null hypothesis were correct.

Summarizing, this section performs statistical tests of the null hypothesis one cell at a time. Two cells produce nominal significance, which is an expected result if the null hypothesis were correct. Of the two cells, one cell has an estimated LGD function that is steeper than the null hypothesis and the other cell has an estimated LGD function that is flatter than the null hypothesis. Of the statistically insignificant results, about half the cells have an estimated LGD function that is steeper than the null hypothesis and the other half have an estimated LGD function that is flatter than

TABLE 2 Basic statistics, parameter estimates and test statistics by cell. [Table continues on next two pages.]

		Senior secured loans		Senior secured bonds		Senior unsecured bonds		Senior subordinated bonds		Subordinated bonds		Averages	
Ba3 rating													
EL (%)	D	0.2	4	0.7	3	0.4	6	0.8	9	0.9	26	0.6	10
PD (%)	N	0.6	616	2.1	179	0.8	703	1.2	525	1.5	874	1.1	579
ELGD (%)	D years	42	3	33	3	49	4	63	6	64	8	55	5
ρ (%)	N years	7.6	14	1.0	26	27.5	27	5.6	26	7.9	21	11.8	23
Firm PD (%)	Firm D	0.7	5	2.1	3	1.2	9	1.2	9	1.7	31	1.3	11
a	Δ LnL	−9.00	0.37	1.45	0.01	2.07	0.17	4.75	0.47	1.16	0.05	0.09	0.21
e (%)	Δ LnL	20.4	0.49	1.0	0.00	11.9	0.23	2.8	0.43	7.1	0.04	8.6	0.24
B1 rating													
EL (%)	D	0.2	9	0.2	2	1.0	13	1.4	22	1.3	38	0.8	17
PD (%)	N	0.8	1332	0.6	205	1.8	757	1.9	909	2.5	756	1.5	792
ELGD (%)	D years	28	5	29	2	53	10	74	10	51	10	54	7
ρ (%)	N years	14.4	14	1.0	27	1.0	27	5.2	25	8.8	25	7.9	24
Firm PD (%)	Firm D	1.8	25	0.8	3	2.3	17	2.0	24	3.0	45	2.1	23
a	Δ LnL	0.82	0.04	−5.46	0.00	−14.28	0.29	1.89	0.07	−0.36	0.01	−3.48	0.08
e (%)	Δ LnL	12.3	0.02	2.3	0.00	3.4	0.29	4.4	0.07	9.8	0.02	6.5	0.08

TABLE 2 Continued.

		Senior secured loans		Senior secured bonds		Senior unsecured bonds		Senior subordinated bonds		Subordinated bonds		Averages	
B2 rating													
EL (%)	D	0.4	24	2.5	4	2.5	45	2.1	36	3.2	35	1.5	29
PD (%)	N	1.2	1842	5.8	168	4.1	826	3.0	740	5.7	325	2.8	780
ELGD (%)	D years	36	10	43	4	60	14	69	8	57	11	55	9
ρ (%)	N years	5.0	14	56.9	26	11.6	27	8.3	21	12.1	26	9.9	23
Firm PD (%)	Firm D	3.0	61	6.1	5	4.9	51	3.2	40	6.1	39	3.8	39
a	Δ LnL	−1.80	0.11	1.76	1.17	−2.35	0.50	1.24	0.05	−0.68	0.04	−0.37	0.38
e (%)	Δ LnL	6.9	0.11	29.6	1.07	16.5	0.68	7.4	0.03	13.8	0.04	14.9	0.39
B3 rating													
EL (%)	D	0.4	19	2.9	11	3.5	45	4.2	33	6.5	58	2.6	33
PD (%)	N	1.4	1374	7.3	218	6.7	813	10.5	382	10.1	371	5.3	632
ELGD (%)	D years	29	6	40	8	52	17	40	14	64	13	48	12
ρ (%)	N years	21.3	14	38.3	26	13.9	27	9.4	22	11.7	21	18.0	22
Firm PD (%)	Firm D	4.1	47	8.1	13	8.5	58	10.9	36	12.0	73	7.3	45
a	Δ LnL	**3.56**	**2.54**	2.55	1.40	−0.58	0.07	−2.72	1.08	0.81	0.15	0.73	1.05
e (%)	Δ LnL	**3.1**	**2.40**	9.7	1.23	15.7	0.08	18.3	1.18	9.9	0.14	11.4	1.01

TABLE 2 Continued.

		Senior secured loans		Senior secured bonds		Senior unsecured bonds		Senior subordinated bonds		Subordinated bonds		Averages	
C rating													
EL (%)	D	2.1	88	5.1	49	6.9	125	19.0	58	4.2	12	6.0	66
PD (%)	0 N	5.6	956	9.8	449	11.5	914	25.1	288	5.8	183	10.2	558
ELGD (%)	D years	38	10	52	17	60	21	76	14	73	6	59	14
ρ (%)	N years	16.9	14	12.3	27	8.2	27	16.4	16	11.2	19	12.9	21
Firm PD (%)	Firm D	23.2	178	13.2	62	14.3	149	27.7	68	10.0	18	18.3	95
a	Δ LnL	−1.26	0.35	−1.49	0.38	0.01	0.00	**−6.58**	**2.78**	0.99	0.03	−1.66	0.71
e (%)	Δ LnL	23.3	0.47	17.5	0.51	8.1	0.00	**31.1**	**3.20**	10.1	0.02	18.0	0.84
Averages													
EL (%)	D	0.6	29	2.9	14	3.0	47	3.6	32	2.4	34	2.1	31
PD (%)	N	1.8	1224	6.1	244	5.3	803	5.6	569	3.9	502	3.9	668
ELGD (%)	D years	35	7	47	7	57	13	65	10	61	10	55	9
ρ (%)	N years	12.8	14	19.5	26	12.1	27	7.8	22	9.5	22	11.8	22
Firm PD (%)	Firm D	5.9	63.2	7.6	17.2	6.6	56.8	6.0	35.4	4.8	41.2	6.1	43
a	Δ LnL	−1.53	0.68	−0.24	0.59	−3.03	0.21	−0.28	0.89	0.39	0.06	−0.94	0.49
e (%)	Δ LnL	13.2	0.70	12.0	0.56	11.1	0.26	12.8	0.98	10.1	0.05	11.9	0.51

EL, PD and ρ are estimates as discussed in the text. $\mathrm{ELGD} = \mathrm{EL}/\mathrm{PD}$. "$D$" denotes the number of defaults in the cell, counting within all twenty-seven years. "N" denotes the number of firm-years of exposure in the cell, counting within all twenty-seven years. "D years" denotes the number of years that have at least one default. "N years" denotes the number of years that have at least one firm exposed. "NomD" denotes the number of nominal defaults (including where the resulting loss is unknown). "NomPD" denotes the average of annual nominal default rates. a and e are the MLEs of the parameters in alternatives A and E. "Δ LnL" denotes the pick-ups in $\mathrm{LnL}_{\mathrm{loss}}$ provided by alternative A or E relative to the null hypothesis. Statistical significance at the 5% level is indicated in bold. Along the right and bottom margins, averages EL, PD and ρ are weighted by N; other averages are unweighted.

the null hypothesis. The pattern of results is of the type to be expected when the null hypothesis is correct. This section provides no good evidence that (2.3) either overstates or understates LGD risk.

8 TESTING CELLS IN PARALLEL

This section tests using several cells at once. To coordinate the credit cycle across cells, we assume that the conditional rates are connected by a comonotonic copula. Operationally, the conditional rate in every cell depends on a single risk factor. All cells therefore provide information about the state of this factor.

We begin by analyzing the five cells of loans taken together. There are 6120 firm-years of exposure in all. The cell-specific estimates of EL and PD are equal to those appearing in Table 2 on page 126. The average of the standard deviation of loan LGD provides the estimate $\sigma = 23.3\%$. We estimate $\rho = 18.5\%$ by maximizing the following likelihood in ρ:

$$\mathrm{LnL}_{\rho} = \sum_{t=1996}^{2009} \log\Bigg[\int_0^1 \prod_{i=1}^{5} \left(\frac{\Phi^{-1}[\widehat{\mathrm{PD}}_i] + \sqrt{\rho}\,\Phi^{-1}[q]}{\sqrt{1-\rho}}\right)^{d_{t,i}} \times \left(1 - \frac{\Phi^{-1}[\widehat{\mathrm{PD}}_i] + \sqrt{\rho}\,\Phi^{-1}[q]}{\sqrt{1-\rho}}\right)^{N_{t,i}-d_{t,i}} \binom{N_{t,i}}{d_{t,i}}\,\mathrm{d}q\Bigg] \tag{8.1}$$

Part (a) of Table 3 on the next page shows the estimates of the parameters and pickups of LnL that result. The estimates of a, b, c and e suggest a steepness that is slightly less than the null hypothesis, but none of the alternatives comes close to the statistically significant pickup of $\Delta\,\mathrm{LnL} > 1.92$. For the five cells of loans taken together, the null hypothesis survives testing by each of the four alternatives.

Turning to the twenty cells of bonds, some firms have bonds outstanding in different seniority classes in the same year. Of the total of 10 585 firm-years of bond exposure, 9.0% have exposure in two seniority classes, 0.4% have exposure in three classes and 0.1% have exposure in all four classes. This creates an intricate dependence between cells rather than independence. Assuming that this degree of dependence does not invalidate the main result, part (b) of Table 3 on the next page shows parameter values suggesting steepness slightly greater than the null hypothesis. Again, none of the alternative models come close to statistical significance and the null hypothesis survives testing.

When all loans and bonds are considered together, 16.0% of firm-years have exposure to two or more classes. Analyzing these simultaneously produces the parameter estimates in part (c) of Table 3 on the next page. Once again, the alternative models remain far from statistically significant and the null hypothesis survives testing.

TABLE 3 Testing cells in parallel.

(a) Loans only
$\sigma = 23.3\%$, $\rho = 18.5\%$

Parameter	Estimate	Δ LnL
a	0.01	0.00
b	0.19	0.31
c	0.11	0.19
e	0.158	0.28

(b) Bonds only
$\sigma = 19.7\%$, $\rho = 8.05\%$

Parameter	Estimate	Δ LnL
a	−0.43	0.28
b	−0.03	0.03
c	−0.03	0.06
e	0.085	0.10

(c) Loans and bonds
$\sigma = 20.3\%$, $\rho = 9.01\%$

Parameter	Estimate	Δ LnL
a	−0.81	1.28
b	−0.10	0.41
c	−0.09	0.55
e	0.102	0.76

The foregoing tests use maximum likelihood estimates of ρ. Risk managers take estimates of ρ from various sources. These include vended models, asset or equity return correlations, credit default swaps, regulatory authorities and inferences from academic studies. All of these sources are presumably intended to produce an estimate of the statistical parameter that appears in a Vasicek distribution relevant for an asymptotic portfolio.[2] Still, it is natural to ask whether a different value of ρ would lead to a different conclusion about the statistical significance of the alternative hypotheses.

To investigate this, we repeat the analysis of Table 3 for the collection of loan cells. In each repetition, we assume a value of ρ. Based on that, we calculate the

[2] Frye (2008) discusses the difference between the correlation in a statistical distribution and the correlation between asset returns that is often used as an estimator.

FIGURE 5 Log likelihood, all loans, as a function of ρ.

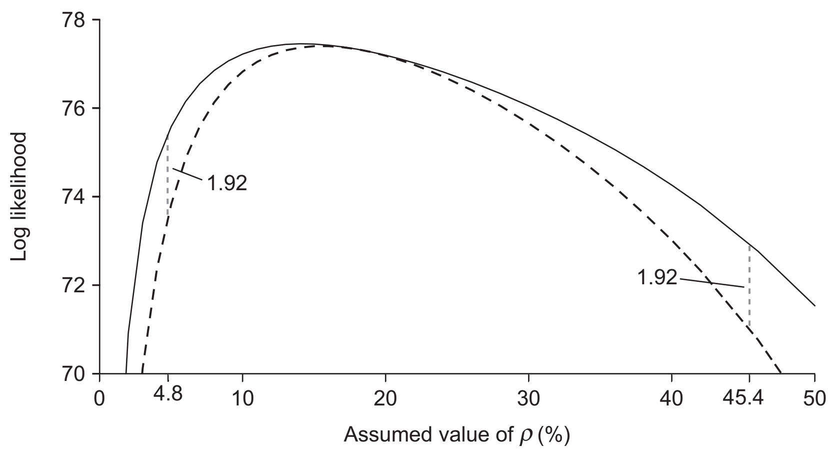

Solid line: maximum log likelihood under alternative A. Dashed line: log likelihood under null hypothesis.

log likelihood under the null hypothesis and under alternative A. A significant result would be indicated by a difference in log likelihoods greater than 1.92.

Figure 5 displays the results. The dashed line shows the log likelihood of loss under the null hypothesis, while the solid line depicts the maximum log likelihood of loss under alternative A. When ρ equals 18.5%, the two are nearly equal, as already shown in Table 3 on the facing page. When ρ takes a different value, the two log likelihoods tend to differ. However, the difference between them never exceeds 1.92 for ρ in the range 4.8% to 45.4%. It is likely that any estimate of correlation for Moody's-rated loans would be in this range. Therefore, the null hypothesis appears robust with respect to the uncertainty in the estimate of correlation.

The results of a statistical test depend on every technique used. For example, our estimator of PD is the average annual default rate. Estimation of PD using maximum likelihood, by contrast, would take into account both the default rates and the numbers of exposures. One must always be aware that other techniques, as well as other alternative hypotheses or other data sets, could lead to other conclusions. An exhaustive check is impossible. It seems more important that existing credit loss models are tested for statistical significance rather than attempting to anticipate every possibility.

The tests of this section employ three sets of cells: all loans, all bonds or all instruments including both loans and bonds. This allows many or all cells to contribute information about the unobserved systematic risk factor. None of the tests produce evidence to suggest that the null hypothesis seriously misstates LGD risk. With respect to the collection of loans, the conclusion is shown to be robust over a range of possible values of correlation.

9 APPLICATIONS AND INCENTIVES

This section discusses the practical implementation and the practical effects of the LGD model. The model can be included within a standard credit loss model without much difficulty. Outside the model, it can be used to obtain scenario-specific LGDs. If the LGD model were used to establish the capital charges associated to new credit exposures, new incentives would result.

A standard credit loss model could use (2.3) to determine conditionally expected LGD. Estimates of parameters PD and EL (or ELGD) are already part of the credit model. The value of ρ has little impact on the LGD–default relationship. A practical estimator of ρ might be a weighted average of an exposure's correlations with other exposures.

Some credit loss models work directly with unobserved factors that establish the conditional expectations, and these models would have the conditionally expected default rate, DR, readily available. Other credit models only have the simulated default rate available. On each simulation run, these models could place the portfolio default rate within a percentile of its distribution, and use that percentile to estimate the DR of each defaulted exposure in the simulation run. An LGD would be drawn from a distribution centered at the conditionally expected rate. This approximation is expected to produce reasonable results for the simulated distribution of loss. Every exposure would have LGD risk, and portfolio LGD would be relatively high in simulation runs where the default rate is relatively high.

Outside a credit loss model, risk managers might want to have an estimate of expected LGD under particular scenarios. An important scenario is one in which DR has a tail realization. In a tail event, there would be many defaults, and individual LGDs should average out quite close to the conditionally expected LGD rate.

In the bad tail, conditionally expected LGD is greater than ELGD. Figure 6 on the facing page shows the difference at the 99.9th percentile. Functions for six different exposures are illustrated. Based on its PD, each exposure has ρ taken from the Basel II formula.

An exposure with PD $= 10\%$ is illustrated by the solid black line in Figure 6 on the facing page. If ELGD were equal to 10%, LGD in the 99.9th percentile would equal $(10\% + 12\%) = 22\%$, which is more than twice the value of ELGD. If ELGD were equal to 20%, LGD in the 99.9th percentile would equal $(20\% + 16\%) = 36\%$. The figure shows that the LGD function extracts a premium from exposures having the low-ELGD, high-PD combination. Relative to systems that ignore LGD risk,

FIGURE 6 99.9th percentile LGD minus ELGD.

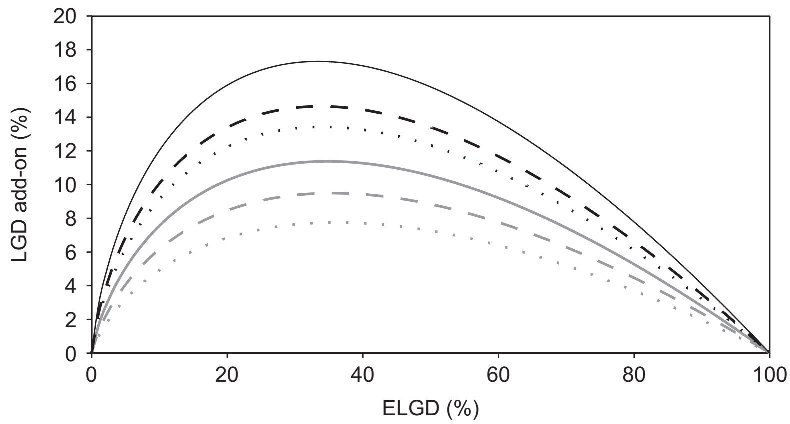

Solid black line: PD = 10%, ρ = 12.1%. Dashed black line: PD = 3%, ρ = 14.7%. Dotted black line: PD = 1%, ρ = 19.3%. Solid gray line: PD = 0.3%, ρ = 22.3%. Dashed gray line: PD = 0.1%, ρ = 23.4%. Dotted gray line: PD = 0.03%, ρ = 23.8%.

this discourages exposures that have exhibited low historical LGD rates and favors exposures that have low PD rates.

In Figure 6, the conditional LGD rate depends on both parameters: PD and ELGD. That traces back to the derivation of the LGD function. If the LGD function had no sensitivity to PD, the credit loss distribution would have three parameters rather than two. Thus, the idea that the distribution of credit loss can be seen only two parameters deep with existing data has a very practical risk management implication.

If this approach to LGD were used to set capital charges for extensions of credit, new incentives would result. The asymptotic loss distribution has two parameters, EL and ρ. Assuming that the parameter ρ is uniform across a set of exposures, two credit exposures having the same EL would have the same credit risk. The capital attributed to any exposure would primarily be a function of its EL. The EL, rather than the breakdown of EL into PD and ELGD, would become the primary focus of risk managers. This would produce an operational efficiency and also serve the more general goals of credit risk management.

10 CONCLUSION

If credit loss researchers had thousands of years' worth of data, they might possess a detailed understanding of the relationship between the LGD rate and the default rate. However, only a few dozen years' worth of data exist. Logically, it is possible that this data is too sparse to allow careful researchers to distinguish between theories. This possibility motivates this paper.

This study begins with simple statistical models of credit loss and default, and infers LGD as a function of the default rate. Using a long and carefully observed data set, this function is tested but it is not found to be too steep or too shallow. It produces greater LGD rates with greater default rates. It uses only parameters that are already part of credit loss models; therefore, the LGD function can be implemented as it is. It can also be subject to further testing. By far the most important tests would be against the portfolio credit loss models now running at financial institutions. If those models do not have statistical significance against (2.3), they should be modified to improve their handling of systematic LGD risk.

APPENDIX A: ANALYSIS OF THE LOSS GIVEN DEFAULT FUNCTION

This appendix analyzes (2.3). It can be restated using the substitution $\text{EL} = \text{PD ELGD}$:

$$\text{LGD} = \Phi\left[\Phi^{-1}[\text{DR}] - \frac{\Phi^{-1}[\text{PD}] - \Phi^{-1}[\text{PD ELGD}]}{\sqrt{1-\rho}}\right] \Big/ \text{DR} \tag{A.1}$$

The parameters combine to form a single value that we symbolize by k:

$$k = \frac{\Phi^{-1}[\text{PD}] - \Phi^{-1}[\text{PD ELGD}]}{\sqrt{1-\rho}}, \qquad \text{LGD} = \frac{\Phi[\Phi^{-1}[\text{DR}] - k]}{\text{DR}} \tag{A.2}$$

LGD functions differ from each other only because their parameter values produce different values of k. We refer to k as the LGD risk index.

Figure A.1 on the facing page illustrates the LGD function with a base case (E_0) and with three contrasting cases. The base case has the parameter values $\text{ELGD} = 32.6\%$, $\text{PD} = 4.59\%$ and $\rho = 14.51\%$. These produce the value $k = 0.53$. Each contrasting case doubles one of the three parameters: ELGD, PD or ρ. Case E_1 is the same LGD function as the Vasicek case illustrated in Figure 1 on page 113.

Along each LGD function, LGD rises moderately with default. In their nearly linear regions from 5% to 15%, LGD rises by slightly less than 10% for each of the illustrated exposures.

Loss given default lines cannot cross, because the LGD risk index k acts in a similar way to a shift factor. Comparing the three contrasting cases, E_1 is the most distant from E_0. That is because the unconditional expectation, ELGD, has the most effect on k; not surprisingly, ELGD is the most important variable affecting the conditional expectation, LGD. Next in terms of importance is PD, which has partially offsetting

FIGURE A.1 LGD functions for four exposures.

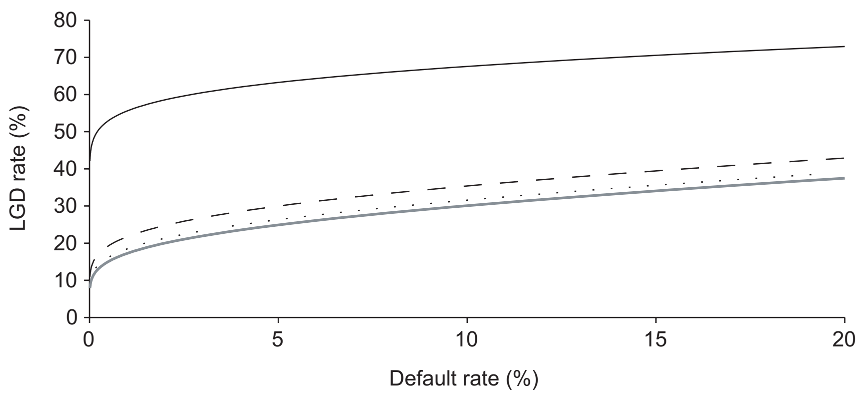

Solid black line (E_1): ELGD = 65.1%, PD = 4.59%, ρ = 14.5%, EL = 3%, k = 0.21. Dashed black line (E_0): ELGD = 32.6%, PD = 4.59%, ρ = 14.5%, EL = 1.5%, k = 0.53. Dotted black line (E_3): ELGD = 32.6%, PD = 4.59%, ρ = 29%, EL = 1.5%, k = 0.58. Solid gray line (E_2): ELGD = 32.6%, PD = 9.18%, ρ = 14.5%, EL = 3%, k = 0.60.

influences on the numerator of k. Least important is the value of ρ. This is useful to know because the value of ρ might be known within limits that are tight enough, say, 5% to 25% for corporate credit exposures, to put tight bounds on the influence of ρ.

In general, an estimate of PD tends to be close to the average annual default rate. (Our estimator of PD is in fact exactly equal to the average annual default rate.) An estimate of ELGD, however, tends to be greater than the average annual LGD rate. The latter is sometimes referred to as "time-weighted" LGD, since it weights the portfolio average LGDs that are produced at different times equally. By contrast, an estimate of ELGD is "default-rate-weighted". This tends to be greater than the time-weighted average, because it places greater weight on the times when the default rate is elevated, and these tend to be times when the LGD rate is elevated. As a consequence, the point (PD, ELGD) tends to appear above the middle of a data swarm.

The LGD function passes close to the point (PD, ELGD). This can be seen by inspection of (A.1). In the unrealistic but mathematically permitted case that $\rho = 0$, if DR = PD, then LGD = ELGD. In other words, if $\rho = 0$, the LGD function passes exactly through the point (PD, ELGD). In the realistic case that $\rho > 0$, the LGD function passes lower than this. In Figure A.1, function E_1 passes through (4.59%, 62.9%), which is 2.3% lower than (PD = 4.59%, ELGD = 65.1%). Function E_2 passes through (9.18%, 29.4%), which is 3.2% lower than (PD = 9.18%, ELGD = 32.6%).

For a given combination of PD and ELGD, the “drop” (the vertical difference between the point (PD, ELGD) and the function value) depends on ρ; a greater ρ produces a greater drop. (On the other hand, greater ρ allows the data to disperse further along the LGD function. This is the mechanism that keeps EL invariant when ρ becomes greater.) The amount of the drop can be placed within limits that are easy to calculate. If ρ takes the value of 25%, the drop is between 4% and 7% for all PD less than 50% and all ELGD between 10% and 70%. If ρ takes the value of 4%, the drop is less than 1% for all PD and all ELGD. For all values of parameters that are likely to be encountered, the LGD function tends to pass slightly lower than the point (PD, ELGD).

The LGD function of (2.3) is strictly monotonic. Figure A.2 on the facing page illustrates this for seven exposures that share a common value of PD (5%) and a common value of ρ (15%), but differ widely in ELGD.

Because both the axes of Figure A.2 on the facing page are on a logarithmic scale, the slopes of the lines in the figure can be interpreted as elasticities, which measure responsiveness in percentage terms. The elasticity of LGD with respect to DR is defined as:

$$\eta_{\mathrm{DR}}\,\mathrm{LGD} = \frac{\partial\,\mathrm{LGD}}{\partial\,\mathrm{DR}}\frac{\mathrm{DR}}{\mathrm{LGD}} \tag{A.3}$$

Looking at the range $1\% < \mathrm{DR} < 10\%$, the slope is greater for lines that are lower; that is, the elasticity of LGD with respect to DR is high when ELGD is low. Thus, when default rates rise, the biggest percentage changes in LGD are likely to be seen in low-ELGD exposures.

By extension, Figure A.2 on the facing page represents the entire range of LGD functions that can arise. Each of the LGD functions illustrated could apply to infinitely many other exposures that have parameters implying the same value of k.

APPENDIX B: ALTERNATIVE A AND PYKHTIN’S LOSS GIVEN DEFAULT MODEL

A solid theoretical model of LGD is provided by Pykhtin (2003). This appendix discusses Pykhtin’s model and illustrates that alternative A is similar to it. In fact, alternative A can be thought of as an approximation to Pykhtin’s model, if the slopes are low or moderate. Therefore, although we do not test directly against Pykhtin’s model, this suggests that we test against a very similar alternative.

Pykhtin’s LGD model depends on a single factor that can be the same one that gives rise to variation of the default rate. Adapting Pykhtin’s original notation and reversing the dependence on Z, there are three parameters that control the relationship between

FIGURE A.2 LGD functions: PD = 5%, ρ = 15% and seven values of ELGD.

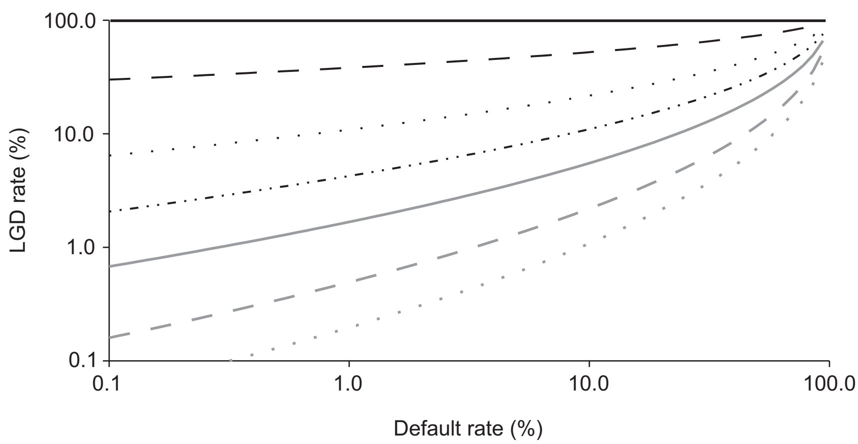

Solid black line: ELGD = 100%, k = 0.0. Dashed black line: ELGD = 50%, k = 0.34. Dotted black line: ELGD = 20%, k = 0.74. Dot-dashed black line: ELGD = 10%, k = 1.01. Solid gray line: ELGD = 5%, k = 1.26. Dashed gray line: ELGD = 2%, k = 1.57. Dotted gray line: ELGD = 1%, k = 1.78.

$\mathrm{LGD_{Pyk}}$ and the standard normal factor Z:

$$\mathrm{LGD_{Pyk}} = \Phi\left[\frac{-\mu/\sigma + \rho_{\mathrm{LGD}} Z}{\sqrt{1-\rho_{\mathrm{LGD}}^2}}\right] - \exp[(\mu + \sigma^2(1-\rho_{\mathrm{LGD^2}})/2) - \sigma\rho_{\mathrm{LGD}} Z] \times \Phi\left[\frac{-\mu/\sigma + \rho_{\mathrm{LGD}} Z - \sigma(1-\rho_{\mathrm{LGD}}^2)}{\sqrt{1-\rho_{\mathrm{LGD}}^2}}\right] \tag{B.1}$$

Pykhtin's three parameters are μ, σ and ρ_{LGD}. Roughly stated, these 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, respectively. Obviously, this is a model of the dynamics of collateral; LGD is determined as the outcome of those dynamics. If there is very little collateral, LGD takes a high value and there is very little for the model to do. Thus, the contribution of the model is most apparent when ELGD is low.

Pykhtin's LGD model can be combined with Vasicek's (2002) default model, which relates the rate of default to the unobserved risk factor Z:

$$\mathrm{DR} = \Phi\left[\frac{\Phi^{-1}[\mathrm{PD}] + \sqrt{\rho} Z}{\sqrt{1-\rho}}\right], \qquad Z = \frac{\sqrt{1-\rho}\,\Phi^{-1}[\mathrm{DR}] - \Phi^{-1}[\mathrm{PD}]}{\sqrt{\rho}} \tag{B.2}$$

FIGURE B.1 Alternative A and variants of the Pykhtin model.

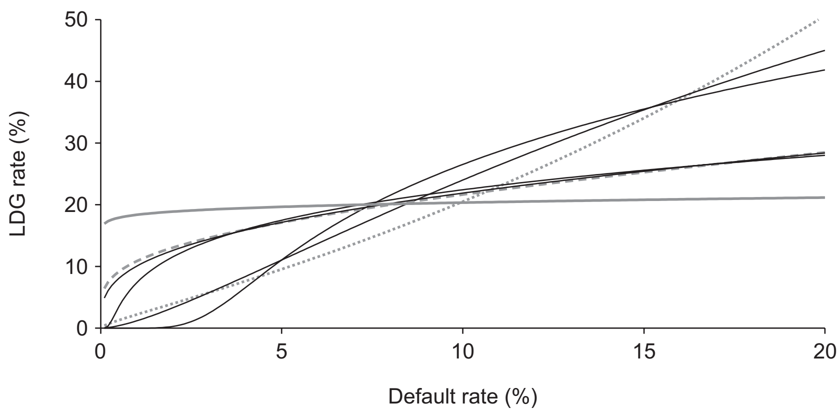

Solid gray line: alternative A with $a = 0$. Dashed gray line: alternative A with $a = 0.87$. Dotted gray line: alternative A with $a = -3$. Near these lines are the variants of the Pykhtin model detailed in Table B.1 on the facing page.

The expression for Z can be substituted into (B.1) to produce a relationship between LGD and default. In this relationship, LGD is a monotonic increasing function of DR that approaches the limits of zero and one as DR approaches the same limits.

Pykhtin's LGD model could be used to test the null hypothesis of (2.3). To produce the correct value of EL, the parameter values must obey the following restriction:

$$\mathrm{EL} = \int_{-\infty}^{\infty} \mathrm{LGD}_{\mathrm{Pyk}}[z]\, \mathrm{DR}[z] \phi[z]\, \mathrm{d}z \tag{B.3}$$

Maximizing (4.6) using the LGD function of (B.1) and subject to the constraint expressed by (B.3) is believed to require a substantial commitment to numerical optimization of what is apt to be a weakly identified model. In the much simpler distribution of loss for the asymptotic portfolio, Frye (2010) finds that the Pykhtin parameters interact strongly with each other and produce an optimum with limiting behavior; that is, to produce the maximum likelihood one of the parameters must be allowed to tend toward negative infinity.

Rather than test directly against Pykhtin's model, we test against alternative A and other alternatives. We compare the two LGD models for a low-ELGD credit exposure: $\mathrm{PD} = 5\%$, $\mathrm{ELGD} = 20\%$, $\mathrm{EL} = 1\%$ and $\rho = 15\%$. In alternative A, this specification leaves only the value of parameter a undetermined. In the Pykhtin model, one of the three LGD parameters is implied by (B.3), leaving two undetermined.

Figure B.1 illustrates the comparison at three distinct levels of LGD risk: low, medium and high. The low level of LGD risk produces an almost-constant LGD

TABLE B.1 LGD functions in Figure B.1 on the facing page.

(a) Low LGD risk

Alternative A a	Pykhtin model μ	σ	ρ_{LGD}
0.867	–0.220	0.100	0.100

(b) Medium LGD risk

Alternative A a	Pykhtin model μ	σ	ρ_{LGD}
0.000 (null hypothesis)	0.294	0.950	0.230
	–0.169	0.075	0.950

(c) High LGD risk

Alternative A a	Pykhtin model μ	σ	ρ_{LGD}
–3.000	0.550	0.950	0.590
	–0.044	0.235	0.950

function. The medium level is consistent with (2.3). In the high level, the LGD functions are steep and varied. At each level of LGD risk, the gray lines in Figure B.1 on the facing page represent alternative A. Each line in Figure B.1 on the facing page produces expected loss equal to 1%. The parameter values of the LGD functions are shown in Table B.1.

When LGD risk is low, the LGD–default relationship is nearly flat at 20%. This is true of both Pykhtin's model ($\sigma = \rho_{\mathrm{LGD}} = 10\%$) and of alternative A ($a = 0.867$). The two lines appear as nearly constant functions almost indistinguishably in the rendering of Figure B.1 on the facing page.

Two variants of Pykhtin's model are compared with the LGD model of (2.3), which is alternative A with $a = 0$. The extra parameters of Pykhtin's model introduce only some nuance into the shape of the relationship, even though the parameter values are large (either σ or ρ_{LGD} equals 95%).

Comparing with the high-risk case when parameter a equals -3, the nuance of Pykhtin's model is clear. Economically, the borrower posts considerable collateral (μ is relatively high), but the collateral is subject to both great systematic risk and to great idiosyncratic risk. The shapes produced by the Pykhtin model are different from the shape of alternative A and somewhat different from each other. Therefore, if the slope of the LGD function were found to be this steep, the nuance provided by the Pykhtin model might make a significant contribution relative to alternative A.

To summarize this illustration, alternative A is a close approximation of the Pykhtin model when LGD risk is low or moderate, but the two models differ when LGD risk is high. Since the level of LGD appearing in the Moody's data appears to be moderate, the null hypothesis, alternative A with $a = 0$, is not rejected by the tests. We believe that we have tested the LGD function against an alternative that is very much like Pykhtin's model. We do not claim that we have shown that Pykhtin's model would demonstrate a lack of statistical significance if the LGD function were tested against it. That test is left for future research.

REFERENCES

Altman, E. I., and Karlin, B. (2010). Special report on defaults and returns in the high-yield bond and distressed debt market: the year 2009 in review and outlook. Report, New York University Salomon Center, Leonard N. Stern School of Business.

Frye, J. (2000). Depressing recoveries. *Risk* **13**(11), 108–111.

Frye, J. (2008). Correlation and asset correlation in the structural portfolio model. *The Journal of Credit Risk* **4**(2), 75–96.

Frye, J. (2010). Modest means. *Risk* **23**(1), 94–98.

Gordy, M. (2003). A risk-factor model foundation for ratings-based bank capital rules. *Journal of Financial Intermediation* **12**(3), 199–232.

Gupton, G., Finger, C., and Bhatia, M. (1997). CreditMetrics. Technical Document, JP Morgan, New York.

Pykhtin, M. (2003). Unexpected recovery risk. *Risk* **16**, 74–78.

Pykhtin, M., and Dev, A. (2002). Analytical approach to credit risk modeling. *Risk* **15**(3), 26–32.

Vasicek, O. (2002). Loan portfolio value. *Risk* **15**(12), 160–162.