

LOSSCALC V2: DYNAMIC PREDICTION OF LGD

MODELING METHODOLOGY

AUTHORS

Greg M. Gupton

Roger M. Stein

ABSTRACT

LossCalc™ version 2.0 is the Moody's KMV model to predict loss given default (LGD) or ($1 - \text{recovery rate}$). Lenders and investors use LGD to estimate future credit losses. LossCalc is a robust and validated model of LGD for loans, bonds, and preferred stocks for the US, Canada, the UK, Continental Europe, Asia, and Latin America. It projects LGD for defaults occurring immediately and for defaults that may occur in one year.

LossCalc is a statistical model that incorporates information at different levels: collateral, instrument, firm, industry, country, and the macroeconomy to predict LGD. It significantly improves on the use of historical recovery averages to predict LGD, helping institutions to better price and manage credit risk.

LossCalc is built on a global dataset of 3,026 recovery observations for loans, bonds, and preferred stock from 1981-2004. This dataset includes over 1,424 defaults of both public and private firms—both rated and unrated instruments—in all industries.

LossCalc will help institutions better manage their credit risk and can play a critical role in meeting the Basel II requirements on advanced Internal Ratings Based Approach.

This paper describes Moody's KMV LossCalc, its predictive factors, the modeling approach, and its out-of-time and out-of-sample model validation.

© 2005 Moody's KMV Company. All rights reserved. Credit Monitor®, EDFCalc®, Private Firm Model®, KMV®, CreditEdge, Portfolio Manager, Portfolio Preprocessor, GCorr, DealAnalyzer, CreditMark, the KMV logo, Moody's RiskCalc, Moody's Financial Analyst, Moody's Risk Advisor, LossCalc, Expected Default Frequency, and EDF are trademarks of MIS Quality Management Corp.

ACKNOWLEDGEMENTS

The authors would like to thank the numerous individuals who contributed to the ideas, modeling, validation, testing, and ultimately writing of this document. Eduardo Ibarra has provided invaluable research support throughout this project. Nil Demircubuk, Thom King, Jun Zheng, Eugene Naguirner, and Joe McQuown. Particular thanks go to: Jody Rasch, Navneet Arora, and Amnon Levy of Moody's KMV. For his insights into bank loans, we thank Kenneth Emery of Moody's Investors Service. We also received invaluable feedback from Moody's Academic Advisory and Research Committee during a presentation and discussion of preliminary findings, particularly from Darrell Duffie of Stanford University and Alan White of the University of Toronto.

Published by:
Moody's KMV Company

To Learn More:
Please contact your Moody's KMV client representative, visit us online at www.moodyskmv.com, contact Moody's KMV via e-mail at info@mkmv.com, or call us:

FROM NORTH AND SOUTH AMERICA CALL:
1 866 321 MKMV (6568) or 415 874 6000

FROM EUROPE, THE MIDDLE EAST, AND AFRICA CALL:
44 20 7280 8300

FROM ASIA, NEW ZEALAND, AUSTRALIA AND INDIA CALL:
813 3218 1160

TABLE OF CONTENTS

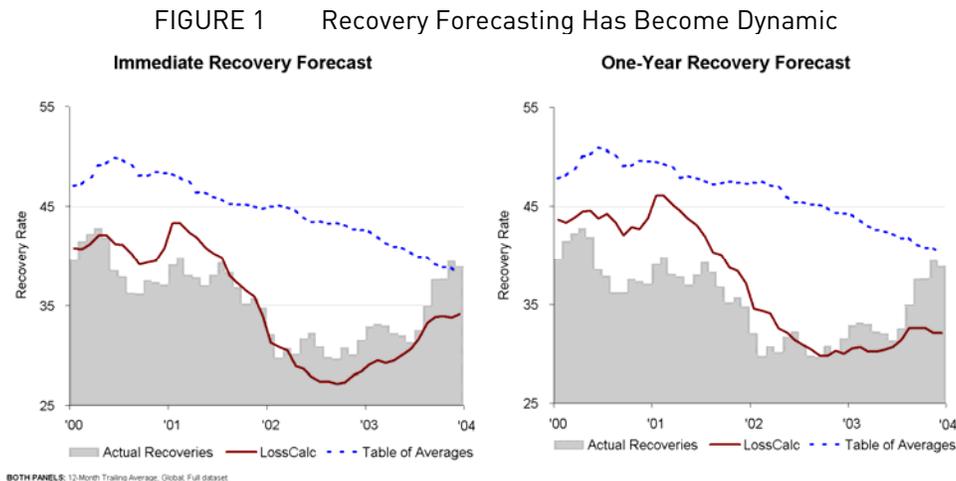
1	OVERVIEW	5
2	THE LGD MODEL: LOSSCALC V2	8
2.1	Overview.....	8
2.2	Time Horizon	8
2.3	Predictive Factors	8
3	FACTORS	9
3.1	Factor Descriptions	9
3.1.1	Collateral and Support	11
3.1.2	Firm-level Information.....	12
3.1.3	Industry.....	13
3.1.4	Macroeconomy	15
3.1.5	Geographic.....	15
4	MODELING FRAMEWORK	16
4.1	Framework	16
4.2	Definition of Loss Given Default.....	17
4.3	Establishing a Loss Given Default Measure	18
4.3.1	Accounting Loss	18
4.3.2	Market Value of LGD	19
4.3.3	Loss Distribution	19
4.4	Mini-modeling	20
4.4.1	Factor Suppression	21
4.5	Modeling and Mapping: Explanation to Prediction.....	21
4.6	Conditional Prediction Interval Estimation	21
4.7	Aligning LossCalc's Output with Banks' Default Definitions	22
5	VALIDATION AND TESTING	22
5.1	Establishing a Benchmark for LossCalc	22
5.1.1	Table of Averages	23
5.1.2	Historical Average	23
5.2	The LossCalc Validation Tests	23
5.2.1	Prediction Error Rates	23
5.2.2	Correlation with Actual Recoveries	24
5.2.3	Prediction of Larger than Expected Losses	25
5.2.4	Relative Performance for Loans and Bonds	26
5.2.5	Reliability and Width of Prediction Intervals	27
5.2.6	Focusing on Loan LGD	28

5.2.7	Other Tests	29
5.2.8	Performance.....	30
6	THE DATASET.....	31
6.1	Historical Time Period Analyzed.....	31
6.2	Scope of Geographic Coverage and Legal Domain	31
6.3	Scope of Firm Types and Instrument Categories.....	32
7	A CASE STUDY.....	32
8	CONCLUSION.....	33
APPENDIX A	BETA TRANSFORMATION TO NORMALIZE LOSS DATA.....	35
APPENDIX B	AN OVERVIEW OF THE VALIDATION APPROACH	36
APPENDIX C	POWER CURVE CONSTRUCTION FOR LGD.....	38
APPENDIX D	RISK NEUTRAL PRICING AND LGD	39
	D.1.Equating Final recovery and Market Price.....	39
	D.2.Uses of Risk Neutral Pricing	39
APPENDIX E	LOSSCALC WEB-SITE OUTPUTS.....	41

1 OVERVIEW

LossCalc Highlights

- LossCalc 2.0 exhibits significant enhancements over LossCalc 1.0 in predicting LGD. LossCalc 2.0 substantially outperforms earlier versions of the model and, dramatically outperforms traditional historical average methods. In particular, LossCalc 2.0:
 - (1) Exhibits lower prediction error and higher correlation with actual losses;
 - (2) Is more powerful at predicting low recoveries—the cases that are most damaging to a portfolio; and
 - (3) Produces narrower prediction bounds.
- LossCalc 2.0 produces estimates of LGD for defaults occurring immediately and those occurring in one year. See Figure 1.
- LossCalc provides estimates of the full distribution of probability weighted predicted values for conducting sensitivity analysis and for use in portfolio applications such as calculating Credit-VaR.
- We have added collateral information to LossCalc’s dataset. This joins with our previous four categories of predictive factors (e.g., debt specific, borrower, industry specific, and macroeconomic characteristics).
- We base Moody’s KMV LossCalc on over 3,000 observations of global recovery values of defaulted loans, bonds, and preferred stock covering the last 23 years. This dataset includes over 1,400 defaulted public and private firms in all industries that are across a wide range of firm and issue sizes.



Using tables of historical recovery averages does not capture the dynamics of actual recovery rates. LossCalc offers both a true forecast plus insight (via its factors) into the drivers of recovery rates.

Why Loss Given Default is Important

Loss Given Default ($1 - \text{recovery rate}$) is an essential input into the process of lending, investing, trading, or pricing of loans, bonds, preferred stock, lines of credit, and letters of credit. Accurate LGD estimates are important for provisioning reserves for credit losses, calculating risk capital, and determining fair pricing for credit risky obligations.

Accurate estimates of LGD are *fundamental* to calculating potential credit losses. They are important because any error in predicting LGD is as damaging as a proportional error in estimating the Expected Default Frequency (EDF), see Gupton and Stein [2001].

$$\text{Expected Credit Loss} = (\text{Expected Default Frequency}) \cdot (\text{Loss Given Default}) \quad (1)$$

Errors in estimating LGD by looking them up in a table of either expert opinions or historical averages can significantly affect estimates of credit losses and the resulting credit risk exposures. Increasing the accuracy of LGD estimates improves the precision of both regulatory and economic capital allocation.

Use of a traditional table look-up approach results in analysis that is backward looking, static, and overlooks changes in LGD behavior. In contrast, LossCalc version 2 addresses the key issues that have prevented institutions from building more accurate LGD models:

- **Lack of Recovery Observations:** There are few financial institutions worldwide with sufficient (both quantity and quality) LGD datasets to fully specify and validate a statistical and predictive LGD model. Institutions are beginning to address this issue in how they collect LGD data, but it still may be several years to a decade before there is sufficient data for many of them to build accurate internal models;¹
- **Complexity of the Recovery Process:** The adjudication of bankruptcy makes it difficult for creditors to predict how their claims will be satisfied. Even where there are legal definitions and guidelines, such as the *Absolute Priority Rule* in North America, the standard is almost never fully applied; and
- **Lack of Empirical Evidence to Determine Predictive Factors:** Even where our clients implement true LGD models (more than mere look-up tables), the predictive factors are typically (1) *static*, (2) *backward looking* (e.g., default rate indices), and (3) *low in power*.

Inaccuracies of Look-up Tables

There is a wide variability in recovery values for instruments even when grouped by debt class and seniority. We commonly see clients with table driven LGD systems that are broken-out by debt type, seniority class, collateral type, loan purpose, business segment, etc. Client LGD tables generally blend their recovery history (which is often brief) with subjective adjustments. These approaches typically lack 1) a time-varying dimension, and 2) a rigorous means of discriminating differences in recovery within any given “cell” of the look-up table (Figure 2).

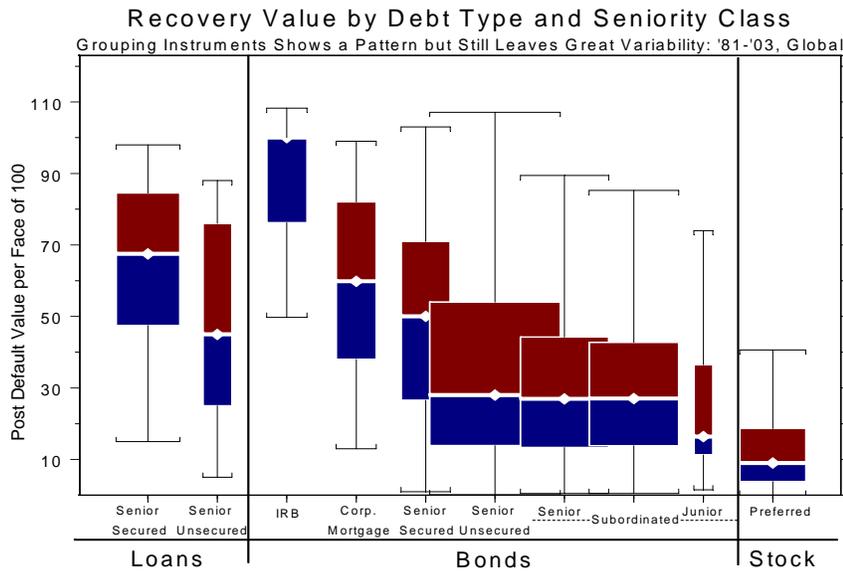
Figure 2 gives the range of recoveries for instruments based on debt-type and seniority. It is similar to annual default studies by Moody’s Investors Service², but for LossCalc we extend it in two ways. First, we separate out Industrial Revenue Bonds from the Sr. Unsecured Bonds group. Second, we separate out Corporate Mortgage Bonds from the Sr. Secured Bonds group. These two groups recover at different levels and our forecast benefits from treating them separately.

The box-and-whisker plot provides a sense of the variability of recoveries in each category. Most of the observations fall within the solid box in each case. For example, even though the median recovery for Senior Unsecured Loans is approximately 45% (horizontal line on the bar), in 50% of the cases, the actual recovery a lender would experience falls in the (wide) range of 25-75% of par. Similar bars for the other seniority classes highlight the wide variability of recoveries even within individual classes.

¹ Moody’s KMV has worked with a number of banks to assist them in specifying data needed for modeling and working to access (commonly) deeply archived paper-based files of historical loan work-out records. Such work is typically highly specialized and requires detailed knowledge of the lending practices of the institution and of the norms of various markets. As a result, for some banks, data is only historically available for a few years.

² See Exhibit #20 in: Hamilton, Gupton and Berhault [2001]

FIGURE 2 Recovery Value by Debt Type and Seniority Class, 1981-2003, Global



The shaded boxes cover the interquartile range with blue extending from the 25th percentile to the median while the red extends from the median to the 75th percentile. White bars mark the medians. Squared brackets mark the substantive data range. The width of each box is proportional to the number of observations within each group. This data comes from Moody's Investors Service, which is the source of LossCalc's recovery observations.

Use of LGD in Basel II

If a bank uses the Advanced IRB approach, then the Basel II accord allows it to use internal models to estimate LGD. While initially a standard LGD allocation may be used (the Foundation Approach), institutions that have adopted the IRB approach for probability of default are being encouraged to use the IRB approach for LGD as well since it gives a more accurate assessment of loss. In many cases, this added precision changes capital requirements.

In order to qualify for the IRB approach:

Basel Committee on Banking Supervision "International Convergence of Capital Measurement and Capital Standards" – June 2004

§ 468

A bank must estimate an LGD for each facility that aims to reflect economic downturn conditions where necessary to capture the relevant risks... cyclical variability in loss severities may be important and banks will need to incorporate it into their LGD estimates. **For this purpose, banks may use ... forecasts based on appropriately conservative assumptions, or other similar methods.**

§ 470

LGD estimates must be grounded in historical recovery rates and, when applicable, must not solely be based on the collateral's estimated market value.

LossCalc can assist in meeting the requirements of LGD specified under Basel II by virtue of both the model's construction and the rigorous testing that Moody's KMV has performed on historical data as far back as 1981. LossCalc includes time-varying factors, which are a Basel requirement and LossCalc uses LGD histories that are longer than the seven year that Basel requires.

2 THE LGD MODEL: LOSSCALC V2

2.1 Overview

LossCalc is a robust and validated model of LGD for loans, bonds, and preferred stock. In developing and testing the model, we used data from Asia (including Australia and New Zealand), Canada, Europe, Latin America, the United States, and the United Kingdom.

LossCalc gives estimates of LGD for defaults occurring immediately and at one year from the time of evaluation.

LossCalc uses predictive information on five different levels to predict debt specific LGD:

- Collateral;
- Instrument;
- Firm;
- Industry; and
- Macroeconomy.

2.2 Time Horizon

Because of their dynamic nature, expected recoveries are different depending on the time horizon of the occurrence of a default event. LossCalc captures this by having two risk horizons: “Immediate” and “One-Year”. The Immediate model is fit using data known just before default. We recommend applying it when the risk horizon under analysis is one year or less.³ The One-Year model is fit using data known one year before default. We recommend applying it when the risk horizon is greater than one year. Each time horizon is a separate model. Each model uses the same predictive factors and are modeled and tested in the same way. The only differences are factor lags and weights assigned to each. Lookup-table models of LGD cannot address the different risk horizons due to the inherently *static nature* of a table.

In selecting the time horizon, investors and lenders should match the tenor of the LGD projection to their risk horizon. For example, a Credit-Value-at-Risk calculated over a one-month horizon would seek to account for the credit environment only one month in the future while a Credit-VaR calculated over a two-year horizon would need to project further.

2.3 Predictive Factors

LossCalc v2 uses nine explanatory factors to predict LGD. We organize these nine factors into five broad groups:

- **Collateral & Backing:** In other words, Cash, “All Assets”, Property Plant & Equipment, and support from subsidiaries;
- **Debt-type / Seniority-class:** Loan, bond, and preferred stock/Secured, senior unsecured, subordinate, etc.;
- **Firm Status:** Up to three factors: (1) Leverage adjusted for credit cycles, (2) Relative seniority standing, and (3) Firm specific Moody’s KMV Distance-to-Default (applied when the borrower is a public company);⁴

³ Clients commonly ask why we recommend the Immediate model for a one year risk horizon. Defaults occurring in the coming year will “arrive” uniformly throughout the year. Therefore, at a one-year horizon, there is about a 50/50 split of arriving closer to LossCalc’s Immediate vs. One-Year points of estimation.

⁴ The “Distance-to-Default” is an output of a structural (Merton-type) valuation of credit distress. Simply put, it is the firm’s debt funding measured in units of standard deviations of asset volatility; see Crosbie & Bohn [2003].

- **Industry:** Two factors: (1) Historical average of industry recoveries to set a base level, and (2) Distance-to-Default across many firms aggregated at the industry and regional level to provide a forward-looking indication of the direction of the credit cycle (default rates, etc.). Note that this factor also contributes to forward-looking estimates of the broad economy as detailed in the next point.
- **Macroeconomy / Geography:** Two factors: (1) Regional flags (i.e., Asia, Canada, Europe, Latin America, United Kingdom, and United States), and (2) Distance-to-Default across many firms aggregated at the regional, and industry, level using data local to each region.⁵

These factors exhibit little colinearity (inter-correlation). Each is statistically significant and so they join to produce a more robust prediction of LGD.

The remaining life of an instrument (i.e., tenor the debt would have had if it hadn't defaulted) is *not* predictive of LGD. The relevant timeframe is the risk horizon rather than the instrument's maturity. This finding is consistent with the *Recovery of Par* hypothesis.⁶ LossCalc reports recoveries as a percentage of the par amount of the debt:

$$\text{Recovery Payoff} = (\text{Recovery Rate}) \cdot (\text{Par Amount}) \quad [2]$$

NOTE: Discount and zero coupon instruments take their *accrued* values.

3 FACTORS

LossCalc is a statistical model designed to predict LGD through the inclusion of multiple factors. As outlined above, we designed each factor to capture a specific dimension of LGD behavior.

3.1 Factor Descriptions

Historical averages broken-out by debt type (loan, bond, preferred stock) and seniority class (secured, senior, subordinate, etc.) are important factors in predicting LGD. However, in LossCalc they account for only about 40% of the influence in predicting levels of recoveries. The distinguishing benefit of LossCalc is that it assembles multiple predictive factors at different information levels. This approach is powerful because (1) each factor is predictive, (2) they are uncorrelated with each other (i.e., they speak to very different parts of the puzzle), and (3) they aggregate LGD behavior within a consistent framework.

We group LossCalc's factors into five categories of predictive information as listed in Table 1. These factors exhibit low colinearity (little inter-correlation) and together make a significant and more accurate prediction of LGD. All factors enter both LossCalc forecast horizons (i.e., immediate and one-year) in the same direction.

⁵ A careful reader will notice that we listed a factor with essentially the same name within *both* the "Industry" and "Macroeconomic" categories. In fact, this is just one factor: Distance-to-Defaults across many firms aggregated at the industry-level within each region. We catalog factors this way to underscore the concepts that they address.

⁶ There are various frameworks for expressing the recovery on defaulted debt. Following nomenclature from Schönbucher [2003] they include: Recovery of Treasury (RT), Recovery of Market Value (RMV), and Recovery of Par (RP). Guha [2002] tested these alternatives and found that RP provided the best empirical fit.

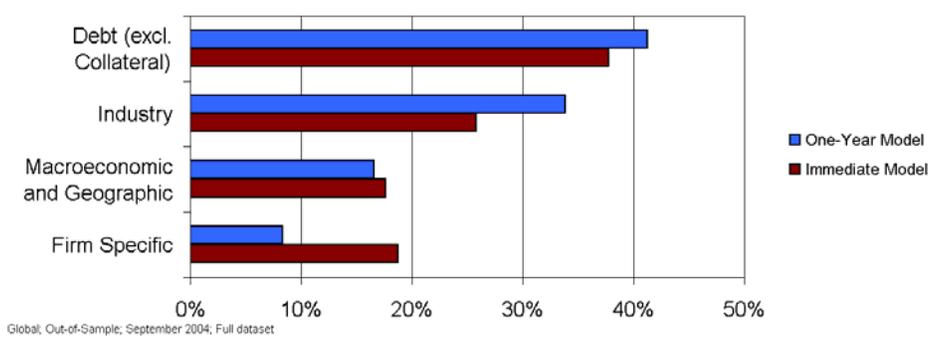
TABLE 1 Explanatory Factors in the LossCalc Models

This is a summary of the factors applied in Moody's LossCalc model to predict LGD. The table highlights the five broad categories of predictive information: collateral, instrument, firm, industry, and broad economic environment. These factors have little inter-correlation and join to make a powerful LGD prediction.

Collateral and Other Support	
The proportion of coverage (of the exposure) by: Cash, "All Assets" or Property, Plant and Equipment. Support from subsidiaries takes a yes/no flag rather than a coverage ratio.	Collateral
Debt Type and Seniority Class	
LGD, controlling for debt-type (loan, bond, and preferred stock) and seniority classes (senior, junior, secured, unsecured, subordinate, etc.).	Historical Averages
Firm-level Information	
Seniority standing of debt within the firm's overall capital structure; this is the <i>relative</i> seniority of a claim. This is different from the <i>absolute</i> seniority stated in Debt Type and Seniority Class above. For example, the most senior obligation of a firm might be a subordinate note if no claim stands above it.	Seniority Standing
Cycle Adjusted Firm Leverage (Gearing): All Corp. Default Rate interacted with the default probabilities directly implied by book leverage	Leverage
The firm's Distance-to-Default (for public firms only)	Firm Distress
Industry	
Historical normalized industry recovery averages after controlling for seniority class.	Industry Experience
The industry's Distance-to-Default (tabulated by country/region)	Industry Distress
Macroeconomic and Geographic	
The country/region's Distance-to-Default (tabulated by industry)	Region Distress
Country/region shifts in mean expectation	Shift

All of the factors in LossCalc are highly statistically significant individually and, in all cases, signs were in the expected direction. Figure 3 below shows the contributions of each broad factor category towards the prediction of the one-year and immediate LGD forecasts. Bars of each color add up to 100%. This shows the relative significance across the predictive factors.

FIGURE 3 Relative Influence of Different Factor Groups in Predicting LGD



This figure shows the normalized marginal effects (relative influence) of each broad grouping of factors. These groupings give intuition as to the source of LossCalc’s power, but there is overlap between groups. For instance, we tabulate industry distance-to-defaults over time and within individual country/regions. Therefore, it influences both the 2nd and 3rd groupings in this graph simultaneously. Note that the influence of collateral is substantive (see Figure 4). Collateral is *not* included here and so it would represent predictive power *in addition to* the factors shown here.

We have calculated the statistics in Figure 3 assuming that collateral was not among the predictive factors. We present the numbers in this way for three reasons:

- Many bankers will use LossCalc to help determine the proper type and amount of collateral to require of the borrower when structuring a loan. Therefore, it is important for the banker to understand his base case.
- The collateral type and amount are commonly changeable when borrower quality deteriorates.
- Portfolio managers may not initially have collateral information electronically available to enter into LossCalc, but will want to understand how LossCalc works.

3.1.1 Collateral and Support

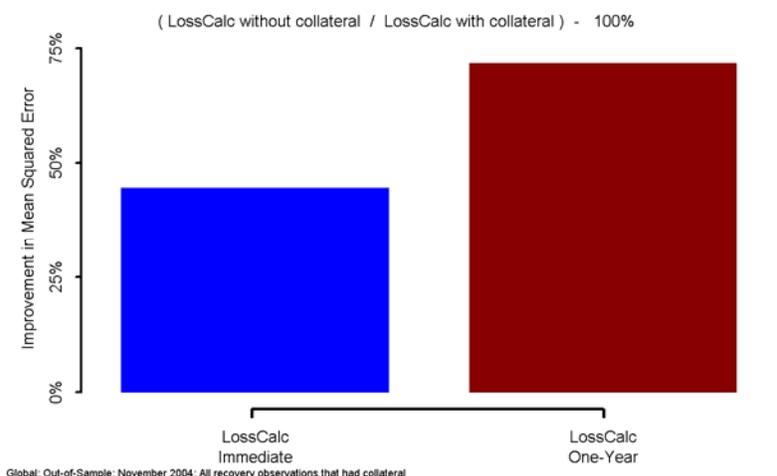
LossCalc allows for six different types of collateral and support for secured instruments. These broad collateral types summarize the more detailed information in the Moody’s Investors Services default database. Of the ten Debt / Seniority types defined in LossCalc, three can be linked with collateral and/or support: (1) Senior Secured Loans, (2) Senior Secured Bonds, and (3) Senior Unsecured Loans. For Senior Unsecured Loans, the only available support is “Subsidiary Support”, which has value for loan recoveries but is technically not a security. The types of collateral and support are:

- **Cash and Marketable Securities Collateral:** This is cash, compensating balances, linked balances, and liquid instruments held on account;
- **Pledge of "All Assets" of the Firm:** All assets in a general designation and may result in enforceability issues. That is taken into account when assessing its impact;
- **Generically Secured, by "Unknown":** This category is used where the specific nature of the collateral is not available. The effect is set at the average across all collateral in LossCalc’s dataset;
- **Property, Plant, and Equipment:** This represents the actual means of production. In the majority of cases it is a physical asset, but in LossCalc it is extended to include other “instruments of production”, for example, airport landing rights;
- **Subsidiary Support:** Refers to any (1) guarantees by subsidiaries, (2) pledged stock of subsidiaries, or (3) pledged and substantively sized key assets of subsidiaries; and
- **Unsecured.**

LossCalc supports multiple selections of collateral backing. This reflects the typical practice of bankers to secure as much collateral as they reasonably can. It also supports Basel II requirements on designating collateral for loans under both the Advanced IRB and Foundation approaches. Mixing different *combinations* of collateral can give a better estimation of recovery. We provide clients with guidance to map 78 types of collateral positions into the LossCalc framework.⁷

When collateral information is available, LossCalc’s forecast accuracy improves. Figure 5 illustrates the accuracy improvement from the collateral information in LossCalc’s dataset. Many clients wish to understand the power of the core LossCalc model. Figures: 6, 9, 11, and 15 do *not* include collateral information and so reflect LossCalc power *before* including collateral.

FIGURE 4 Including Collateral Information Improves LGD Accuracy



This figure shows relative improvement (reduced Mean Squared Error) of LossCalc’s LGD forecasts. We show results for both the Immediate forecast (in blue) and the One-Year forecast (in red). The metric of forecast accuracy is the relative reduction in MSE. For example, when running LossCalc with particular collateral information, its one-year forecast is 72% more accurate compared to LossCalc run selecting merely “Generically Secured by Unknown”. Note that this boost in power from having collateral information is *not* included in Figures: 6, 9, 11, or 15.

3.1.2 Firm-level Information

LossCalc considers three types of information for the firm: (1) credit distress of the firm, (2) standing in the capital structure, and (3) cycle-adjusted leverage ratio.

- **Firm-specific Distance-to-Default** measures credit distress for public firms. A firm that suffers more severe distress is more likely to have a higher LGD. The measure uses the firm’s capital structure as well as information from the equity markets to produce a timely and efficient signal of market participants’ collective belief of the firm’s future prospects and the market value of the firm’s assets. The Distance-to-Default as used in the MKMV public-firm model calculates the market value of the firms’ assets and compares that to the book value of its liabilities, thus addressing the “coverage” aspect of the leverage ratio and doing a more accurate job.

This measure is only available for publicly traded firms.

- **Standing within the firm’s capital structure** is a debt’s *relative* seniority (i.e., are there claimants who stand more senior at the time of default). For example, preferred stock is the lowest seniority class short of common stock, but it might hold the *highest* seniority rank within a particular firm that has no funding from loans or bonds. In addition, some more senior class of debt may mature thus revealing another (lower) debt class to be “most senior”.

⁷ These 78 collateral descriptions reduce to 15 unique combinations within LossCalc. This is because we judge different types of collateral to have the same benefit for LGD.

The ordering that LossCalc applies for this tabulation is the same as shown in Figure 2. Exceptions are “Industrial Revenue Bonds” and “Corporate Mortgage Bonds”. These two debt classes are outside of LossCalc’s relative seniority scheme. They are not influenced by (and do not influence) any other seniority class that may be in firm’s capital structure at the time of default.

Alternative measures not selected were “the *amount* of debt that stands more senior” or “the *proportion* of total liabilities that is more senior”. The two main reasons are:

Resolution Procedure. In bankruptcy proceedings, a junior claimant’s ability to extract concessions from more senior claimants is not proportional to its claim size. Junior claimants can force the full due process of a court hearing and so have a practical veto power on the *speediness* of an agreed settlement.⁸

Availability of Data. Claim amounts at the time of default are not the same as original borrowing/issuance amounts. In many cases, borrowers partially pay down their debt before maturity through amortization schedules (for loans) and sinking funds (for bonds). In addition, data on the exposure at default may be unavailable for firms that pursue multiple funding channels. Requiring such an extensive detailing of claims before being able to make *any* LGD forecast would be onerous.

- **Firm leverage** (or gearing) is how much asset value is available to cover the liabilities of the firm. Since asset values may rapidly decline as a firm approaches bankruptcy, a Heckman adjustment⁹ is used. This adjustment takes into account that the leverage ratio is itself predictive of the event of default (as companies’ leverage increases, their default probability increases). We do not apply the leverage ratio in the case of secured debt or financial industries.¹⁰
Firm leverage has more of an impact on LGD during periods of economy-wide credit distress. When relatively many firms in the economy are defaulted, then a firm’s leverage tends to count against it more than during less distressed times. We address this relationship by interacting the leverage ratio with the *Global All Corporates Default Rate* as published by Moody’s Investors Service.

3.1.3 Industry

Industry is an important indicator of recovery rates. Most institutions use recovery averages broken out by industry to refine historical estimates of LGD.¹¹ This approach implies that industry characteristics do not change over time.

This approach does not capture the industry-level variability in recovery rates *across time* and across industry. For example, some sectors, such as the Telephone industry, enjoy periods of prolonged superior recoveries, but fall well below average recoveries at other times. A constant industry bump up or notch-down cannot capture such behavior.

The two upper panels of Figure 5 show the mean recoveries for different industries and the distribution around those recoveries compared to the entire population. They show a low recovery of Business Services relative to high recoveries in Gas Utilities. In addition, industry recovery distributions change over time. The Telephone industry is a recent example. The bottom panels of Figure 5 show the shift in Telephone recoveries.

For many years, the phone industry was mature in its technology and somewhat like a utility with redeployable assets and above average recoveries. However, as wireless technology emerged, the asset base in the industry faced a quicker

⁸ We tested this on a sub-population selected to have fully populated claim amount records. The best predictor of recoveries, both univariately and in combination with a core set of LossCalc factors was a simple flag of Who-Has-Highest-Standing. We tested many alternatives, such as amount or proportion of “cushion” and amount or proportion of “overhead” as well as certain transformations such as logarithms.

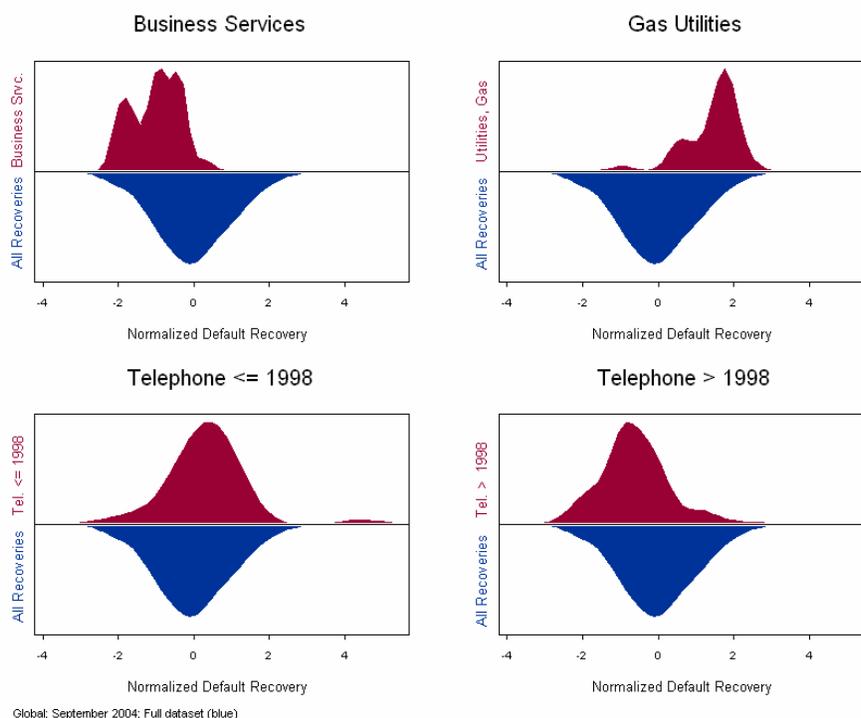
⁹ Heckman [1979] describes the potential bias caused by applying a factor that is itself predictive of being in the conditioned state of the world. In this case, since leverage is a strong predictor of being in default, then the leverage ratio needs to have a bias adjustment to properly predict loss *given* default.

¹⁰ Specifically, this applies to “Senior Secured Loans”, “Senior Secured Bonds” and “Corporate Mortgage Bonds” as well as Lessors, Real Estate, Banks and S&Ls, and Finance not otherwise classified. Secured claims look principally to the security for satisfaction and the leverage of financial firms is far more problematic to judge than for corporate and industrials.

¹¹ See Altman & Kishmore [1996] and Ivorski [1997] for broad recovery findings by industry and Borenstein & Rose [1995] for a single industry (airlines) case study.

obsolescence cycle, and a significant proportion of once hard assets shifted to *assets* like “air rights”. Recoveries fell and the prior industry-level distribution did not prove to be predictive.¹²

FIGURE 5 Industries Recover Differently and Are Not Static



In all panels, we contrast the variability of individual industry recoveries (in red) against the entire LossCalc dataset (in blue). All recoveries are after LossCalc’s normalization (see § 4.3.3), which is visible in the “bell” shape (Normal distribution shape) of the blue areas. The *y*-axes scales show relative frequency. The *x*-axes are the number of standard deviations away from average a recovery was after considering debt type. The top two panels contrast Business Services vs. Gas Utilities industries. The bottom two panels illustrate the Telephone industry change.

To address industry differences across time, we first organize the dataset into the sixty-two specific industries defined by Moody’s KMV.¹³ We then produce two different measures of industry behavior:

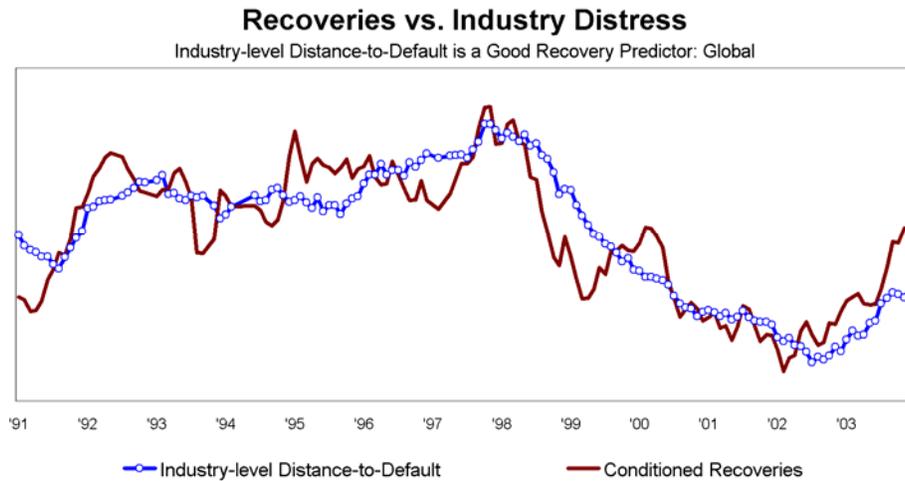
- the industry’s historical recovery experience, and
- the aggregated Distance-to-Default across all firms in that industry (and region).

We recompute both of these measures on a monthly basis and observe significant time variation over the span of our dataset. The Distance-to-Default is the more dynamic of these two factors. As shown in Figure 6, these measures generally exhibit stable predictive behavior across the industries and the country/regions we have examined.

¹² The difference in the two Telephone mean recoveries shown in the bottom panels have a *z*-statistic of 8.65 indicating near zero likelihood of occurring by chance. Statistical significance is also made possible by our large (356 observations) sample size in this industry.

¹³ A mapping of SIC codes to these industry groups is available upon client request.

FIGURE 6 Recoveries are Lower in Distressed Industries



In LossCalc, we aggregate Distance-to-Default measures separately for each industry-region, but for illustration here, we aggregate into *one* index only (blue line) the Distance-to-Default measures that coincide with an LGD observation. We contrast this factor against averaged recovery rates (red line). These are LossCalc’s normalized recoveries controlled for seniority class.

3.1.4 Macroeconomy

Distance to Default: Because recoveries have positive and significant inter-correlation within bands of time, we use our Distance-to-Default measure to bring macroeconomic information into LossCalc. The positive correlation between EDF and LGD, in an economic downturn, would lengthen the tail of a portfolio loss distribution thus raising economic capital assessments.

The Distance-to-Default replaces a number of traditional macro indicators that analysts use as proxies for the credit cycle, including the Moody’s Bankruptcy Bond Index (MBBI)¹⁴ and changes in Leading Economic Indicators.

Trailing 12-month All Corporate Default Rate (global): We use this LossCalc factor to adjust each firm’s book leverage for the credit cycle. It replaces the *U.S.* Speculative-grade Trailing 12-month Default Rate used in version 1. Moody’s Investors Service publishes both indices monthly. This substitution is more appropriate for a global model.

We had considered an alternative factor, which was a set of sector-level default rated indices. We rejected this because these indices turn out to be far too noisy to be useful in prediction. Since defaults are rare and can occur in clumps, default rate averages are volatile indicators when aggregated across small populations such as an industry sector. In contrast, our Distance-to-Default statistic is a continuous measure available for 25,000 firms globally. Distance-to-Default also has the advantage of being *forward* looking.

3.1.5 Geographic

LossCalc is an international model, as the LGD observations come from a global data set. We also source LossCalc’s predictive factors locally by country / region, so the predictive factors used in the model vary depending on the location of the borrower. Factors applied in Europe, for example, are not weighted, mixed, or supplemented in any way with the factors applied in any other country / region – including the United States. Because we have detailed credit information on 25,000 firms worldwide, we have been able to construct granular and powerfully predictive indices broken out by (1) country/region, (2) industry, and (3) across time.

¹⁴ Hamilton & Berthault [2000]

Legal Differences

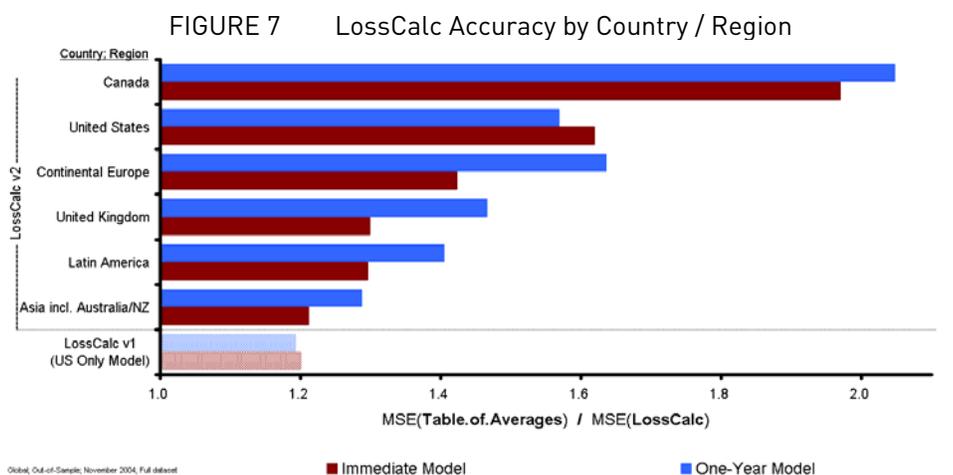
While legal differences in bankruptcy play a part in recovery, fundamental economics drives firm values and recoveries in a predictable way regardless of the particular legal jurisdiction. A simplified view of the bankruptcy process is universal for all countries and has two distinct steps.

- First, there is a determination of the aggregate firm-level value of the defaulted entity. This will be the total worth available to satisfy all the claimants of the firm.
- Second, there is a determination as to how this firm value will be divided-up among the claimants.

Country-by-country differences in bankruptcy law, though important, only affect this second step.

We address the legal issues among countries by having certain fields active only for certain countries. For example, a guiding principle in the United States is the *Absolute Priority Rule*. The Canadian experience closely emulates the US even though strict written law differs. Outside North America, however, various considerations beyond an absolute rule of priority come into play. To address this, we only apply the “Most Senior Debt” factor in the United States and Canada.

By properly addressing both “steps” in the LGD process, LossCalc captures a significant portion of the variability across countries and regions.



Shown here is the improvement in the Mean Squared Error performance measure of LossCalc relative to a Table of Averages. We show two risk horizons for each of six country/regions that LossCalc covers. A horizontal bar with no length (i.e., simply marking the 1.0 value) would say that LossCalc did as well as a Table of Averages. Thus, LossCalc outperformed in all cases. LGD predictions were most accurate in Canada. All countries / regions outperformed the US-only version 1 of LossCalc (see ghosted bars at bottom).

4 MODELING FRAMEWORK

4.1 Framework

LossCalc is a data intensive, empirically driven, statistical model that adheres to economic principles. The broad steps in this framework are transformation, modeling, and mapping.

- **Transformation:** Rather than taking the simple levels of predictive factors, we transform the raw data or “mini-model” in it to create more powerful univariate factors. For example, we find that leverage is more damaging to recoveries during downturns in the credit cycle. We thus interact leverage by the “Global All Corporate Default Rate”.

- **Modeling:** Once we have transformed individual factors and converted them into mini-models, we aggregate these using familiar multivariate regression techniques.
- **Mapping:** We statistically map the model output to historical LGD.

4.2 Definition of Loss Given Default

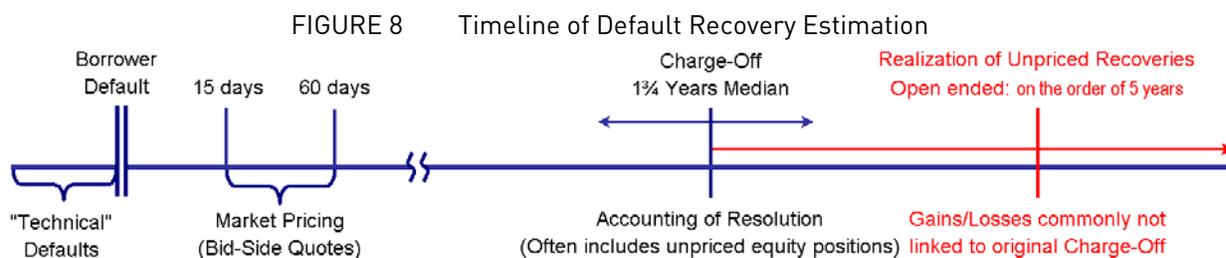
LossCalc defines recovery ($1-LGD$) on a defaulted instrument as its market value approximately one-month after default.¹⁵ The model uses security-specific bid-side market quotes¹⁶ rather than broker-created “matrix” prices¹⁷.

Moody’s KMV chose to use price observations approximately one month after default for three reasons:

- The time period gives the market sufficient time to assimilate new post-default corporate information;
- The time period is not so long after default that market quotes become too thin for reliance; and
- This period best aligns with the goal of many investors to trade out of newly defaulted debt.

This definition of recovery value avoids the practical difficulties associated with determining the post-default cash flows of a defaulted debt or the identification and value of instruments provided in replacement of the defaulted debt. The very long resolution times in a typical bankruptcy proceeding (commonly 1.25 to 5 years) compounds these problems. Broker quotes on defaulted debt provide a more timely recovery valuation than requiring the completion of court ordered resolution payments, potentially several years after the default event. Market quotes are commonly available in the period 15 to 60 days after default. We exclude securities with no price or an unreliable price from the dataset.

Figure 8 shows the timing of price observation of recovery estimates and the ultimate recovery of the claims.



This diagram illustrates the timing of the possible observation of recovery estimates. Recovery from a default is an extended process rather than a culminating event. After default, the market prices the expectation of anticipated recoveries. These are most liquid 15 to 60 days post-default. Some 1 3/4 years later, half of defaults have been charged-off the accounting books. Recoveries at that point include cash, new debt extensions, equity in the (emerging) borrower, etc. Equities may not trade and may not have a market price. Eventually, all receipts in satisfaction of the default realize a value, but this is typically not traceable back to the original borrower.

The Relationship of Market Pricing and Ultimate Realization: There have been several studies of the market’s ability to price defaulted debt efficiently.¹⁸ These studies do not always show statistically significant results (typically due to small sample sizes), but they consistently support the market’s efficient pricing of ultimate realization. At different times,

¹⁵ The date of default is not always well defined. As an example, bankers commonly write loan covenants with terms that are more sensitive to credit distress than those of bond debentures. Thus, different debt obligations of a single defaulted firm may technically default on different dates. The vast majority of securities in our dataset have quotes within the range of 15 to 60 days after the date assigned to initial default of the firm’s public debt. Importantly, our study found no distinction in the quality or explicability of default prices across this 45-day range.

¹⁶ Contributed by Goldman Sachs, Citibank, BDS Securities, Loan Pricing Corporation, Merrill Lynch, Lehman Brothers, and (newly added LoanX and LPC).

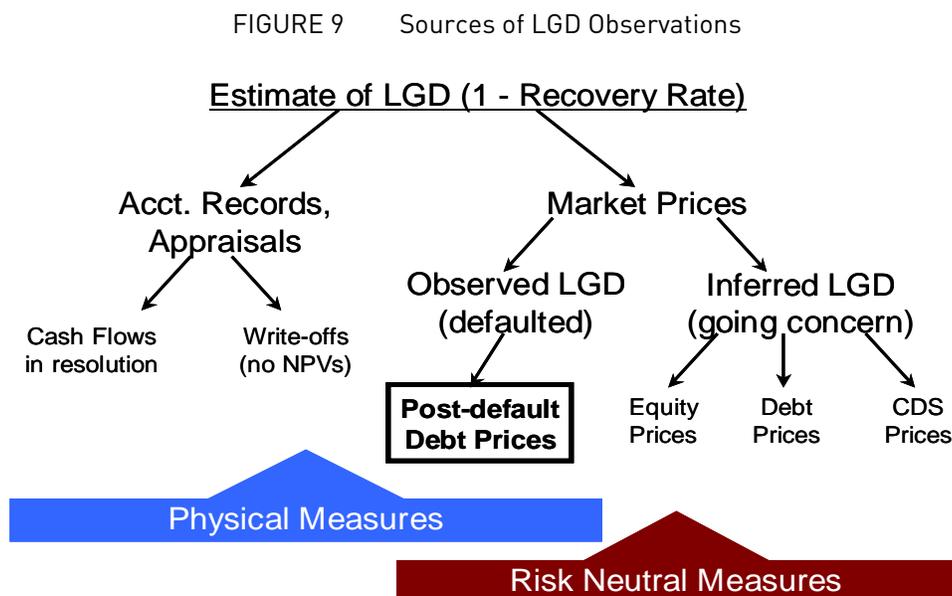
¹⁷ Matrix prices are tables specified across maturity, credit grade, and instrument type, without consideration of the specific issuer. As such, they are inferior.

¹⁸ See Eberhart & Sweeney [1992], Wagner [1996], Ward & Griepentrog [1993], Altman, Gande & Saunders [2004].

Moody's has studied recovery estimates derived from both bid-side market quotes and discounted estimates of resolution value. We find, consistent with outside academic research that these two valuation approaches reconcile with an appropriate risk adjusted valuation. See Appendix D for more information.

4.3 Establishing a Loss Given Default Measure

There are several ways to measure LGD. The two key measures are Economic (Risk Neutral) and Accounting (Physical Measure). Figure 9 compares these two measures.



There are several ways of observing a LGD (or recovery rate). They each have pros and cons. We can directly observe the accounting records so they are “Physical Measures”. Market valuations represent fair market value so they are “Risk Neutral Measures”.¹⁹ The dependent variable for LossCalc falls in the middle since it is both “physical” (a defaulted asset can be liquidated for cash when the price is observed) and “risk neutral” (the market price is a risk neutral estimate of the present value of all cash flows during resolution). LossCalc uses bid-side quotes about one month post default as the recovery rate.

4.3.1 Accounting Loss

Accounting measures of LGD are typically available only for loans since banks source these from their accounting records. This information is also scarce because typical bank systems have not captured recovery cash flows nor have they tracked the collateral realizations electronically; collecting this information requires a manual extraction from paper files. Accounting loss is usually calculated using one of two techniques:

- **Cash flows from a default resolution:** Using cash flows requires the full profile of when the creditor, typically a lending bank, receives economic value. In order to calculate the Credit-VaR, the institution could calculate a net present value of these recovery cash flows.
- **Charge-off amounts:** An alternative estimate of accounting-LGD is the write-off value when recovery efforts are exhausted or when they choose to abandon a debt.

¹⁹ In mathematical finance, a *risk-neutral measure* is today's fair (i.e. arbitrage-free) price of a security, which is equal to the discounted expected value of its future payoffs. The measure is so called because, under the measure, all financial assets in the economy have the same expected rate of return, regardless of the asset's “riskiness”. This is in contrast to the *physical measure* - i.e. the actual prices where (typically) more risky assets (those assets with a higher price uncertainty) have a greater expected rate of return than less risky assets.

Cash Flows from Default Resolution

The value at resolution is often subjective since equity, rights, and warrants received as payment in resolution commonly have no market price. One study, Hamilton [1999], found that 15% of the value of the recoveries for Senior Secured Loans came in the form of equity of the defaulted firm (see Figure 8). Since these payments with equity interests (e.g., common stock, preferred, and warrants) commonly do not trade, their value will be unrealized and unclear for years. When these equity values are eventually realized/known (often well past the write-off date), it would be *atypical* for a bank's accounting system to track flows back to the original charge-off.

When we assist clients in databasing their own institution's LGD histories, we have always found it necessary to examine archived paper records. The full tracking of default resolution realized values (cash flows) has been far more informative than sourcing simply the accounting write-offs.

Charge-Off Amounts

Charge-off amounts represent the amount that an institution writes off its books when the loan becomes non-performing. The problem with using charge-offs is that they can occur substantially before the final resolution. Also, bank accounting records do not link charged-off accounts to what a bank eventually received in the true ultimate recovery resolution. Charge-offs are essentially the financial institution's best guess of the ultimate recovery, without having market forces to help determine its correctness.

4.3.2 Market Value of LGD

The LGD of an instrument can be inferred from market prices (right-hand branch of Figure 9). The market-based values represent fair pricing and are "Risk-Neutral Measures." There are two alternative approaches:

- **Implied Loss Given Default:** This estimation typically requires the initial step of estimating the default likelihood of the firm (given either a "structural" or "reduced form" model), and then picking the LGD to best reconcile the market price of the debt or CDS with its model valuation. This approach is most suited for securities valuation.
- **Direct Observations of Market Value of LGD:** For investors in liquid credits, the market price directly represents their realizable recovery. Market liquidity is good at about one month post default as many investors are disallowed from holding defaulted assets and so trade their positions to a group that specializes in defaulted debt.

Market prices of defaulted debt are valuable and are not merely reflections of supply and demand in the distressed debt market. By its construction, the Moody's Bankrupt Bond Index would capture price variability caused by supply/demand. The fact that an MBBi predictive factor dropped-out as insignificant when Industry Distance-to-Default entered our model is strong evidence that presumed imbalances in supply and demand do not drive market prices of defaulted debt. Thus, market prices are efficiently stated and predictive of future recoveries. Fundamental credit information drives defaulted debt prices rather than the ebbs and flows of supply and demand.

The loan market prices are also robust, based on empirical evidence found by Altman, Gande & Saunders [2004]. These authors find that, "...the loan market is informationally more efficient than the bond market around loan default dates and bond default dates." LossCalc's estimates benefit from efficient prices.

4.3.3 Loss Distribution

LossCalc uses the observed post-default market valuation for the recovery ratio. These recoveries, however, are not normally distributed. An alternative that better approximates the prices in our dataset is the Beta distribution.

Using a Beta transformation, we convert the market prices to be normally distributed, see blue areas in Figure 5. A Beta distribution is suited to LGD analysis because it ranges from 0 to 1, corresponding to 100% loss or zero loss, but is not restricted to being symmetrical. Application of a Beta distribution is robust across many LGD models and datasets.²⁰

4.4 Mini-modeling

Before the final modeling, we assessed the individual variables on a stand-alone (univariate) basis. We do this to both boost the univariate power of factors and to deal with interactions between factors. Examples include:

- The relationship of leverage and corporate default rates;
- Historical LGD and industry level; and
- Industry-level distance-to-default.

Leverage and corporate default rates: Higher leverage typically leads to worse recoveries. This effect is more pronounced during times when many firms in the economy are defaulting.²¹ In order to capture this impact, we interact leverage with corporate default rates rather than using it directly. The result is that in times of stress, higher leveraged companies will be affected more than lower leveraged firms will. However, the influence of the default rate on leverage plateaus beyond a certain limit. In order to reflect this, we combine leverage and the default rate so that the factor has a reduced marginal effect above a certain bound.

Historical LGD and industry: Across time, there are differences in industry recoveries. Some industries, such as service firms, may have “softer” assets compared to others such as a Gas Utility (top panels of Figure 5); however, these effects are not constant over time. Thus, recoveries in one industry may be higher than those in another during one phase of the economy, but lower than the other industry in a different economic environment. To measure this, we constructed industry indices of LGD. To fully utilize all of our LGD observations, we had to state recoveries from any given seniority class on a common basis. We achieved this by normalizing recovery observations—each within its own debt type—and standardizing these to have a mean of zero and a standard deviation of one. These numbers all have a consistent interpretation.

Industry-level Distance-to-Default (DD): By its construction, the industry Distance-to-Default is a potent measure of credit distress in an industry. It is forward looking, powerful, and consistent across industries, countries, and time. We construct our industry-DD factor by tabulating the X^{th} percentile of the firm DDs in each industry. Our factor is linear in its response except that we dampen the response near its limits. Since the primary buyer of defaulted assets is another firm in the defaulter’s industry, they will be better buyers if they themselves are not credit distressed. In credit distressed industries:

- Firms are more likely to be retrenching to stave-off a default of their own.
- There are more likely multiple defaulted firms at the same time thus flooding the market with defaulted assets.
- Firms have fewer resources with which to buy defaulted assets.
- Industry growth is likely slow limiting new project opportunities that would benefit from asset purchases.

By building mini-models, the factors are better stand-alone predictors of LGD. Both mini-models and stand-alone factors are univariate measures. We do this level of modeling before assembling an “overall” multivariate model.

²⁰ See recent research: Gordy & Jones [2002], Ivanova [2004], Onorota & Altman [2003], Pesaran, et.al. [2003], Singh [2003], and Tasche [2004].

²¹ See Gupton & Stein [2002].

4.4.1 Factor Suppression

The model drops certain factors in certain cases when it would not make economic sense to apply them. For example, although *leverage* is one of the nine predictive factors in the LossCalc model, it is not included for financial institutions.²² These are typically highly leveraged with lending and investment portfolios having very different implications than an industrial firm's plant and equipment.

Similarly, we do not consider leverage when assessing secured debt.²³ The recovery value of a secured obligation depends primarily on the value of its collateral rather than broad recourse to general corporate assets.

4.5 Modeling and Mapping: Explanation to Prediction

The modeling phase of the LossCalc methodology involves statistically determining the appropriate weights to use to combine the transformed variables and mini-models. The combination of the predictive factors is a linear weighted sum, derived using regression techniques without an intercept term. The model takes the additive form:

$$\hat{r} = \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + \dots + \beta_k x_k \quad (3)$$

where the x_i are either the transformed values or mini-models, the β_i , are the weights and \hat{r} is the normalized recovery prediction. \hat{r} is in "normalized space" and it is not yet in "recovery rate space." Therefore, the final step is to apply inverse Beta-distribution transformations. See Appendix A for more details.

4.6 Conditional Prediction Interval Estimation

Prediction Intervals (PI) provide a range around the prediction within which the actual value should fall a specified percentage of the time. There is typically a high variability around the estimates of recovery rates produced by tables. See Figure 2.

LossCalc provides a conditional estimate of the prediction interval (i.e., 5% and 95% bounds) on the recovery prediction that is much tighter than that of tables. Risk managers can use the PI's within a portfolio Credit-VaR model. The *width* of this prediction interval provides information about both (1) the ability to realize the mean expectation of the prediction, and (2) the inherent uncertainty of the recovery rate process. It does *not* describe the precision of mean recovery forecast.

A 90% prediction interval around the predicted mean value is the range (bounded by an upper bound and lower bound) in which the realized value will fall 90% of the time. Therefore, we only expect the realized value to be below the lower bound or above the upper bound, 10% of the time.

Although standard parametric tools can produce an (in-sample) estimate of the prediction intervals, these estimates are relatively wide. The LossCalc PI prediction approach produces narrower ranges of prediction. To achieve this heightened precision, we use a combination of quintile regression²⁴ and analytic results. Quantile regression provides a convenient means for estimating the quantiles of the recovery distributions conditional on LossCalc predictive factors. As a second step, we used quasi-analytic approximations to convert these back to Beta distributions. This prediction approach takes into account debt type, seniority class, and other LossCalc factors.

The resulting prediction interval, where prediction is of the post-default prices, has implications for the assessment of risk-neutral quantities. We discuss this briefly in Appendix D.

²² Industries that do not participate in the firm leverage factor include: Banks and S&Ls, Finance N.E.C., Lessors, and Real Estate.

²³ Seniority classes that do not participate in the firm leverage factor include: Senior Secured Loans, Senior Secured Bonds, and Corporate Mortgage Bonds.

²⁴ See, Fitzenberger, et.al. [2002].

4.7 Aligning LossCalc's Output with Banks' Default Definitions

In a certain percent of cases, banks experience no loss when a firm defaults on a loan. For example, a loan may be 90 days past due but are restructured (cured). This means that it never goes through “workout” and (depending upon the bank’s systems) it may or may not be recorded as a default with (likely) a 100% recovery. We find that it is more typical for banks to *not* capture these in their records of historical LGD. There is the potential of a resulting mismatch in the effective default definition used to calculate default probabilities versus the definition used to calculate LGD. We refer to this difference as the “cure” rate.²⁵ In order to bring the calculation of default into line with that used in LGD, banks can apply a cure rate to LossCalc’s loan LGDs. This adjustment aligns the recovery rates to reflect the percentage of defaulted loans that have no losses. This adjustment applies only to loans since bonds are publicly reported and do not have this same sort of “cure” process.

To make this adjustment for a specific institution, we first need to determine what a bank’s own *cure rate* is. Banks’ cure rates differ and Moody’s KMV cannot predict ahead of time what this might be for any given institution, although our experience with clients and bank data suggests that cure rates are commonly between 20% and 50%. As a result, LossCalc assumes a 100% recovery on the (bank specified) cured portion and then applies its statistical forecast of LGD to the remaining ($1 - \text{cure rate}$) portion. This is consistent with LossCalc’s dataset, which represents “loss given a loss.” Mathematically,

$$\text{Aligned Recovery} = (\text{Cure Rate} \cdot 100\%) + (1 - \text{Cure Rate}) \cdot \text{LossCalc Recovery Forecast} \quad (4)$$

5 VALIDATION AND TESTING

The primary goals of validation and testing are to:

- Determine how well a model performs;
- Ensure that a model has not been overfit and that its performance is reliable and well understood; and
- Confirm that the modeling approach, not just an individual model, is robust through time and credit cycles.

To validate the performance of LossCalc, we have used the approach adopted and refined by Moody’s KMV termed *walk forward* validation. It involves fitting a model on one set of data from one time period and testing it on a subsequent period. We then repeat this process, moving through time until we have tested the model on all periods up to the present. Thus, we never use data to test the model that we used to fit its parameters and so we achieve true out-of-time and out-of-sample testing with an efficient use of the data. We can also assess the behavior of the modeling approach over various economic cycles. Walk forward testing is a robust methodology that accomplishes the three goals set out above. See Figure 19 for an illustration of the process.

Model validation is an essential step to credit model development. We must perform rigorous and robust testing to guard against unintended errors. For example, the same model may get different performance results on different datasets, even when there is no specific selection bias in choosing the data. To facilitate comparison, and avoid misleading results, we use the same dataset to evaluate LossCalc and competing models.

Sobehart, Keenan, and Stein [2000] describe the walk-forward methodology more fully. Appendix B gives a brief overview of the approach.

5.1 Establishing a Benchmark for LossCalc

The standard practice in the market is to estimate LGD by some historical average. There are many variations in the details of how these averages are constructed: long-term versus moving window, by seniority class versus overall, dollar weighted versus simple (event) weighted. We chose two of these methodologies as being both representative and broadly

²⁵ The lowest cure rate we have seen is 20%.

applied in practice. We then use these traditional approaches as benchmarks against which to measure the performance of the LossCalc models.

5.1.1 Table of Averages

The majority of financial institutions use a table look-up to estimate LGD. This often reflects expert opinion as to what LGD *ought* to be, but more commonly LGD look-up tables list historical average LGDs either from the institution's own experience (usually based on write-offs rather than economic loss) or (not uncommonly) taken from rating agency recovery studies. A leading source of this type of agency recovery table is in Moody's Annual Default Studies. With Moody's sizable dataset, it represents a high quality implementation of this classic approach.

In order to compare LossCalc performance against the most widespread model (look-up tables) we used Moody's Investors Services tables since they reflect economic loss. Annually, Moody's publishes LGD averages that show recoveries by debt type and seniority class.

5.1.2 Historical Average

Some institutions use a simple historical average recovery rate as their recovery estimate. Therefore, as a second hurdle (and we believe this represents a naive LGD model), we also tabulated the overall recovery rate across all instruments (the "Historical Average").

5.2 The LossCalc Validation Tests

Since LossCalc produces an estimate of an *amount* of recoveries, it seeks to fit a continuous variable. Thus, the diagnostics we use to evaluate its performance reflect this.

Since 1992 is the end of the first half of our dataset, we fit LossCalc to data from 1981-1992 and then forecast 1993 recoveries. We saved these out-of-time and out-of-sample predictions in a result set. We then fit LossCalc to 1981-1993 data, forecast 1994, and saved those to the result set. Following this walk-forward procedure, we constructed a validation *result set* containing over 1,851 observations, representing 915 default events (some firms defaulted more than once and were counted each time) from Moody's extensive database from 1993 to 2004. This result dataset was over 60% of the total observations in the full dataset, and is a representative sampling of rated and unrated public and private firms, in all industries, and country/regions. See Appendix B for more details.

5.2.1 Prediction Error Rates

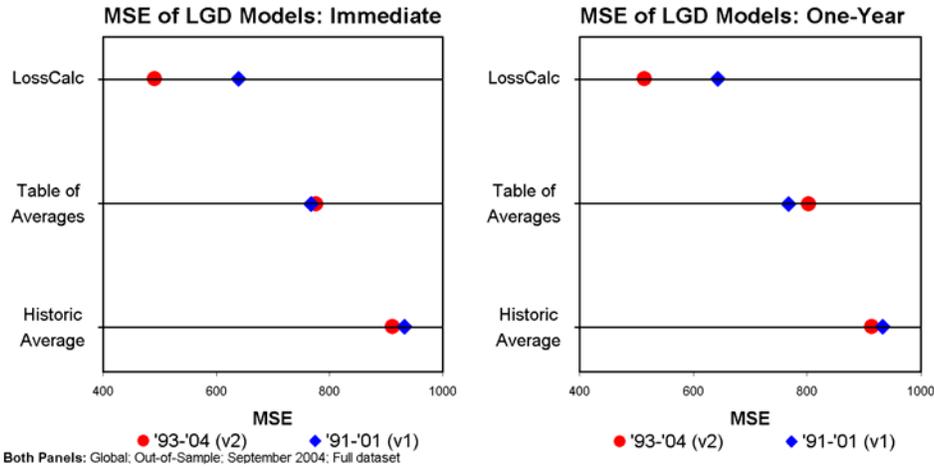
As a first measure of performance, we examined the error rate of the models. This is measured with an estimate of the mean squared error (MSE) of each model. The MSE is calculated as:

$$MSE = \frac{\sum (r_i - \hat{r}_i)^2}{n-1}, \tag{5}$$

where r_i and \hat{r}_i are the actual and estimated recoveries, respectively, on security, i . The variable, n , is the number of securities in the sample.

Models with lower MSE have smaller differences between the actual and predicted values and thus they predict actual recoveries more closely. Thus, better performing models in Figure 10 below will have their symbols further to the left. LossCalc outperforms a Table of Averages by a large margin. For comparison, we have included the comparable accuracy measures from version 1 of LossCalc. Version 1 statistics were tabulated on a purely US dataset and so are not directly comparable to version 2. Indeed this may understate the performance differential since the version 2 dataset is more heterogeneous: including both US and non-US firms.

FIGURE 10 Mean Squared Error (MSE) of LossCalc Models and Other Alternative Models



This figure shows the out-of-time and out-of-sample MSE for LossCalc, the Table of Averages, and the Historical Average. Note that better performance is towards the left hand side in each panel, which is the opposite of Figure 11. It is clear that in both the immediate and one-year prediction, LossCalc has smaller error in comparison with the two alternative models.

TABLE 2 Mean Squared Error (MSE) of LGD Prediction Accuracy across Models and Horizons
Here we list the specific out-of-sample MSE values illustrated in Figure 10.

Out-of-Sample	Immediate MSE		One-Year MSE	
	1991-2001	1993-2004	1991-2001	1993-2004
Historical Average	933.2	910.8	933.3	913.4
Table of Averages	767.3	775.8	767.3	802.4
LossCalc (v1 & v2)	639.1 (v1)	490.6 (v2)	643.0 (v1)	514.0 (v2)

It is useful to think in terms of the *difference in error rates* between two models as a type of “savings.” For example, over 10% of the time, the “savings” in absolute²⁶ error rate is greater than 28% of par value.

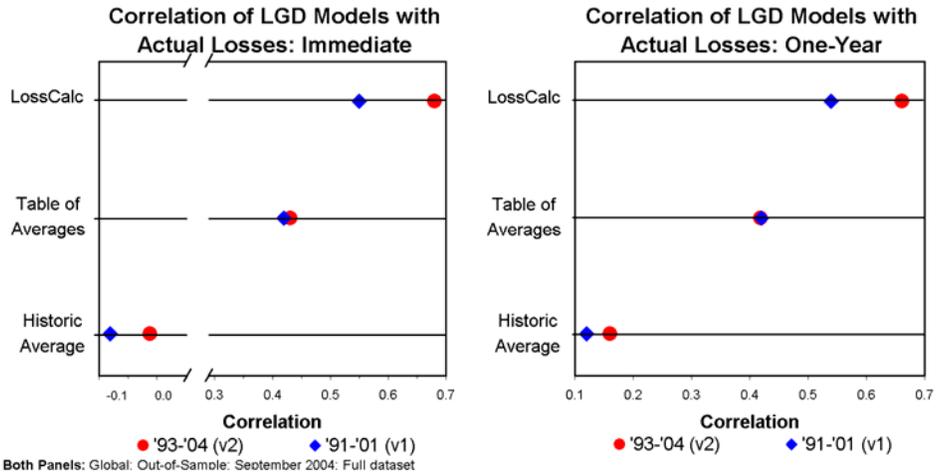
5.2.2 Correlation with Actual Recoveries

The correlation of the various models’ predictions with the actual loss experience is an important measure of a model’s usefulness. Correlation is measured by plotting the actual recovery with the predicted recovery to see how well the two estimates track one another. In a model with high correlation, the model will give high estimated recoveries when actual recoveries are high and low estimated recoveries when actual recoveries are low. The less deviation there is from this relationship, the higher the model’s correlation

Figure 11 below, shows the correlation of predicted versus actual for the three models. We also report this in tabular form in Table 3. The Historical Average actually exhibits a *negative* correlation with actual recovery experience, meaning it tended to show high recoveries when actual recoveries are low and vice versa. This is because it is a moving average and tends to reflect what has happened previously, rather than forecasting future movements.

²⁶ Measured as $(|Table\ Error| - |LossCalc\ Error|)$

FIGURE 11 Correlation of LossCalc Models and Alternatives with Actual Recoveries



This figure shows the out-of-time and out-of-sample correlation for LossCalc, the Table of Averages, and the Historical Average. Note that better performance is towards the right hand side of this graph, which is the opposite of Figure 10. It is clear that over both the immediate and one year horizons, LossCalc has better correlation in comparison with the two alternative models.

TABLE 3 Correlation of LGD Prediction Accuracy across Models and Horizon
Listed here are the specific out-of-sample correlation values that we show in Figure 11.

Out-of-Sample	Immediate Correlation		One-Year Correlation	
	1991-2001	1993-2004	1991-2001	1993-2004
Historical Average	-0.13	-0.06	-0.13	-0.09
Table of Averages	0.42	0.43	0.42	0.42
LossCalc (v1 & v2)	0.55 (v1)	0.68 (v2)	0.54	0.66 (v2)

Models with higher correlation have smaller differences between the actual and predicted values and thus they predict actual recoveries more closely. Better performing models in Figure 11 above will have their symbols further to the right. LossCalc outperforms a Table of Averages by a large margin.

5.2.3 Prediction of Larger than Expected Losses

MSE and correlation are parametric measures of model performance that are sensitive to model *calibration*, which refers to the overall level of model outputs. An alternative and *nonparametric* measure of model performance focuses on a model’s ability to discriminate “good” outcomes from “bad” outcomes; also referred to as a model’s *power*. Therefore, when assessing the accuracy of an LGD model, how well it ranks orders higher-than-average losses versus lower-than-average losses is important. A financial institution’s ability to correctly predict when losses are going to be higher than the average is key. To measure the power of an LGD model, we designed the following test:

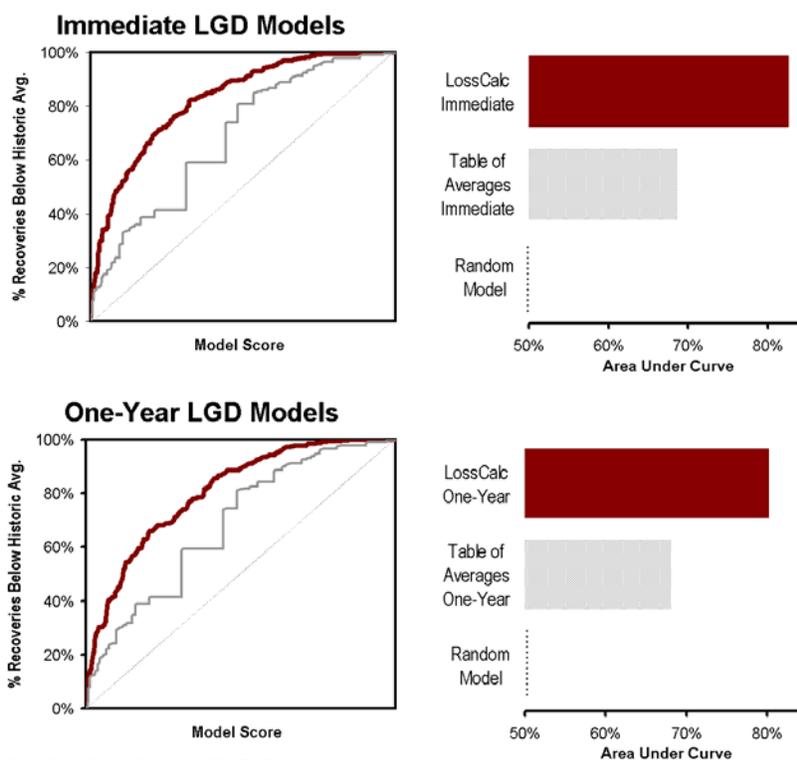
- Each record is labeled to determine if the actual loss experienced was greater or less than the historical mean loss for all instruments to date (i.e., the Historical Average was first referenced in §5.1.2).
- The out-of-sample predictions for each model were ordered based on the largest predicted loss to smallest predicted loss.
- We calculated how well the model ranked these larger than average losses with the lower than average losses.
- The ranking ability of each model was captured using standard power tests.

For details on this test design, see Appendix C.

If a model was powerful at ranking losses, the above average loss predictions of the model would correspond to the above average actual results. The way the power curve was constructed, this would result in the curve for a good model being bowed out towards the Northwestern corner of the chart. The power curve of a perfect model would rise straight up the y-axis extend straight across the top of the chart yielding 1.0 area under the curve. The random model would be a 45° line.

Figure 12 shows both the Power Curves at left and the area under the curves at right—the larger the area under the curves, the more accurate the model. Both the Table of Averages and LossCalc models perform better than random at differentiating high- and low-loss events, but the LossCalc models outperform the Table of Averages by a considerable margin.²⁷ This relationship persists over both the immediate and one-year horizons. The comparison of areas under the curves confirms this observation.

FIGURE 12 Power in Predicting Higher than Expected LGD



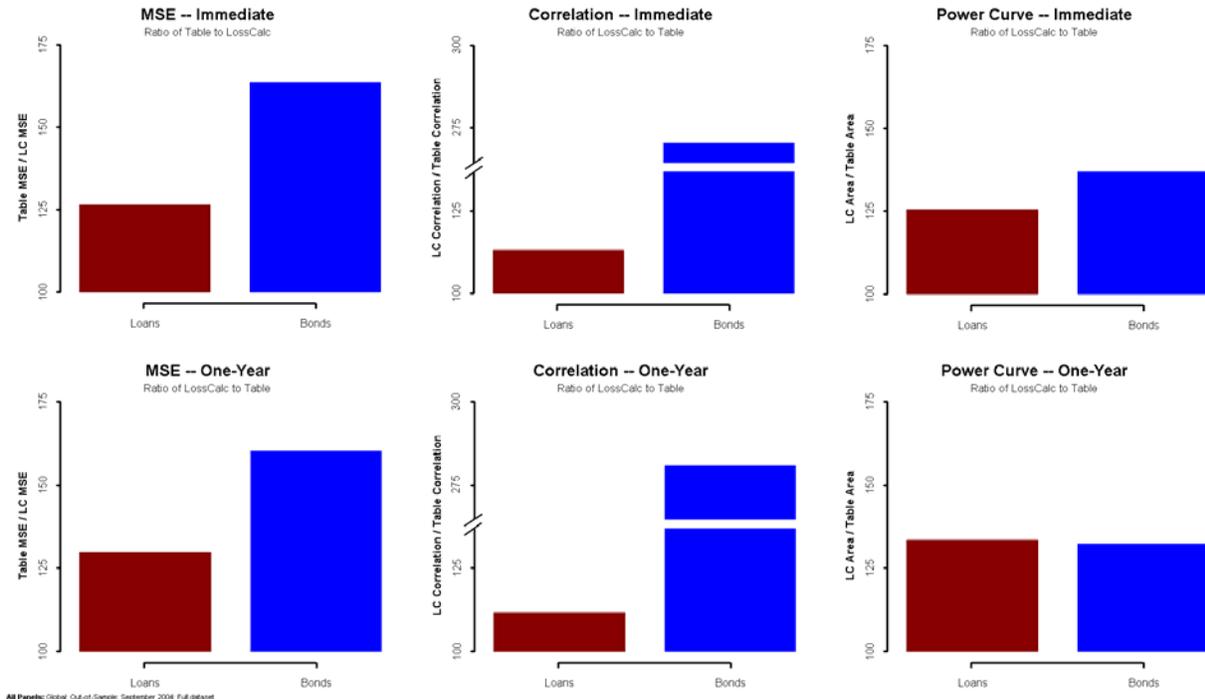
This figure shows the superiority in predictive power of LossCalc relative to the Table of Averages. It is clear that LossCalc has significantly greater power at both the immediate and one-year horizons, (i.e., its power curve is more towards the upper left corner and the area under the curve is greater).

5.2.4 Relative Performance for Loans and Bonds

In Figure 13, we show the relative performance of LossCalc versus the common practice of a Table of Averages with respect to three measures of relative performance: (1) Mean Squared Error (MSE), (2) Correlation, and (3) Area Under a Power Curve. We show results for both loans and bonds. By all three measures of performance, LossCalc increases predictive performance significantly. However, we find that it is important to examine multiple dimensions of performance to gain the greatest understanding of model performance.

²⁷ Note that “plateaus” can occur in the Power Curves indicating that multiple instruments received the same LGD prediction. This mostly affects the Table of Averages model where all instruments of a particular seniority class (for a particular year) receive the identical LGD prediction. In principle, if one was a “bad” and the other was not, the ordering of these becomes important; although testing indicates that, the differences would not change the overall conclusions.

FIGURE 13 LossCalc Performance Increase Over a Table of Averages



This figure shows the relative increase in performance for both loans and bonds. We show three metrics of performance: (1) Mean Squared Error (MSE), (2) Correlation, and (3) Area Under a Power Curve. In all panels, a bar of zero height would indicate that LossCalc did only as well as a Table of Averages. Therefore, as shown, in every scenario LossCalc outperforms a Table of Averages. As a specific example, we show the improvement in “Area Under a Power Curve” with the realized LGD between a table and LossCalc (at the one-year horizon) in the bottom right chart. LossCalc is at “133” so it is 33% better than a Table of Averages for bonds. See Appendix C for an explanation of a Power Curve.

Any Table of Average’s predicts recoveries by segregation (bucketing) debts and assigning each segment a static recovery value. Across time, the only variability in the Table of Averages comes from the changes in re-tabulated averages from one year to the next. Thus, the Table of Averages focuses primarily on the *between-group* (debt type and seniority class) variability rather than the *within group* (firm, industry, and macroeconomic) variability in recoveries. This is a fundamental weakness of *any* table look-up approach. In contrast, LossCalc uses additional information, which allows it to incorporate both within- and between-group variability more completely.

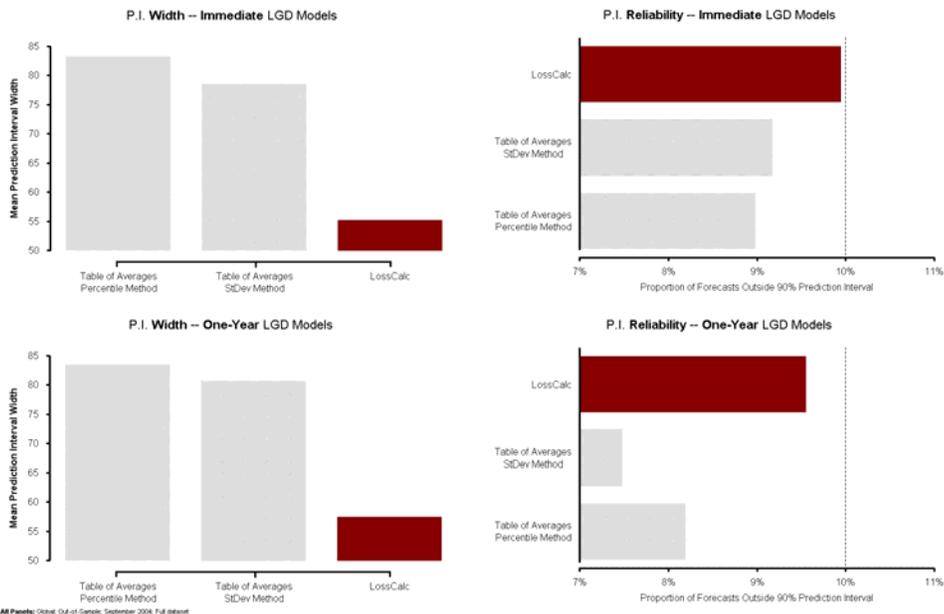
5.2.5 Reliability and Width of Prediction Intervals

To test the soundness of the prediction intervals produced by the models, we examined two dimensions: width and reliability. The average *width* of a prediction interval (PI) provides information about the precision and efficiency of the estimate. In general, a narrow PI is good because it gives less uncertainty to potential LGD realizations and allows an institution to assess losses or allocate capital more efficiently. However, too narrow a PI may not *reliably* reflect the actual observed PI in the future.

To examine these two issues, we generated PIs for each model, calibrated them in sample, and tested them out-of-sample and out-of-time. We show the average widths of these PIs in the left hand panels of Figure 14. Then we tested out-of-sample and out-of-time the number of cases in which the actual observed losses exceeded the predicted interval.

We examined several methods for PI prediction. Some of these required the use of actual prediction errors from previous periods. These tests were on about 500 observations for the one-year horizon and close to 600 for the immediate tests.

FIGURE 14 The Width and Reliability of Model Prediction Intervals



These panels show the Prediction Interval Widths (left two panels) and Reliabilities (right two panels) for both immediate predictions (top two panels) and for defaults that may occur one-year from now (bottom two panels). In all cases, we compare the LossCalc performance relative an alternative Table of Averages model. We find that LossCalc’s Prediction Intervals are both narrower and more reliable than other alternatives.

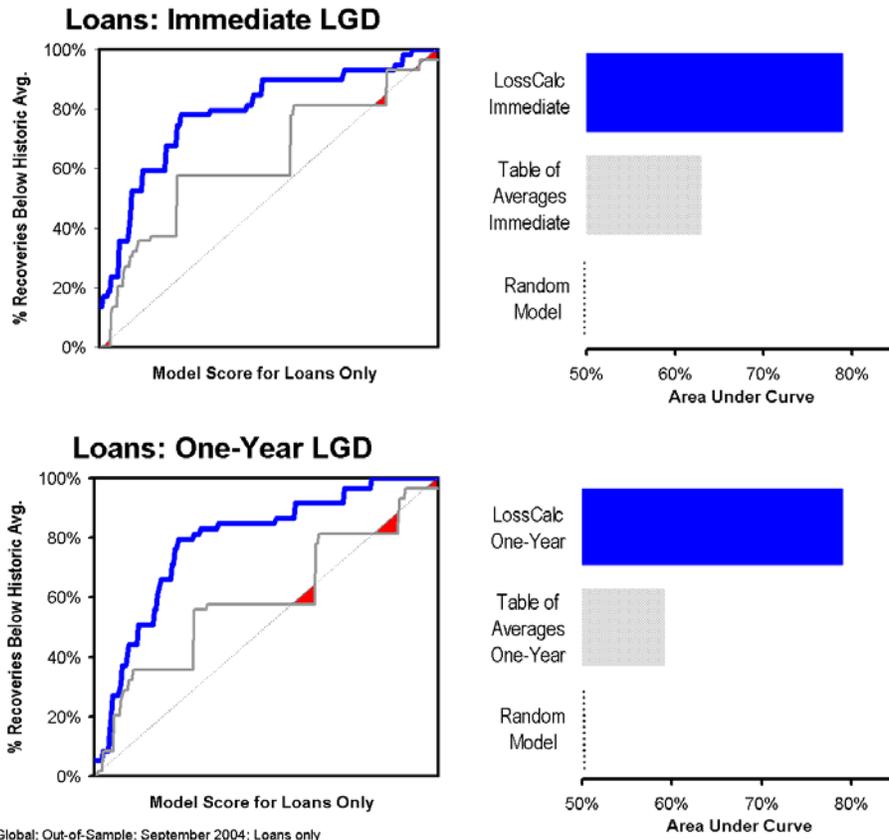
Figure 14 above, shows that LossCalc’s PIs are both more precise (narrower) and more reliable in out-of-sample and out-of-time testing when compared to two possible approaches that we applied to the Table of Averages.

5.2.6 Focusing on Loan LGD

We find that the LossCalc predictive factors work powerfully and consistently in all subpopulations of our dataset including loans. Naturally, loans are special because they are far more commonly collateralized relative to bonds. We know that collateral information is important and improves LossCalc’s accuracy significantly. See Figure 4.

Figure 15 shows predictive power for exclusively loan LGDs within LossCalc’s dataset. LossCalc exhibits significantly more power to rank order loan LGDs when compared to a Table of Averages model, which is the most common internal bank model that we see our clients using. This test is significant because it runs LossCalc without the benefit of collateral information. LossCalc achieves the power shown in Figure 15 from our *other* predictive factors, which are at the firm, industry, macroeconomic, and geographic levels.

FIGURE 15 Loans: Power in Predicting Higher than Expected Loan LGD



Global: Out-of-Sample: September 2004; Loans only

For just the loans in our dataset, this figure shows the predictive power of LossCalc relative to the Table of Averages. LossCalc has significantly greater power at both the immediate and one-year horizons. Its power curves (blue lines in left panels) are more towards the Northwest corner thus capturing more area under the curve (right panels). What was surprising to us was that there were some strata at which a Table of Averages technique actually produced LGD predictions that ranked loans no better than randomly, i.e., where the gray line (in left panels) falls below the diagonal dotted line (areas in red).

A surprising finding, illustrated in Figure 15, was that there were some strata at which a Table Of Averages technique would actually produce LGD predictions that ranked loans no better than randomly, i.e., where the gray line falls below the diagonal dotted line of the random model. On investigation, we find that this unfortunate outcome is possible for any model that assigns an identical model score to a large group of instruments. Within that grouping (i.e., within an individual “cell” of the Table Of Averages), there exists no power to discriminate one LGD from another. Therefore, within that cell it is a random model.

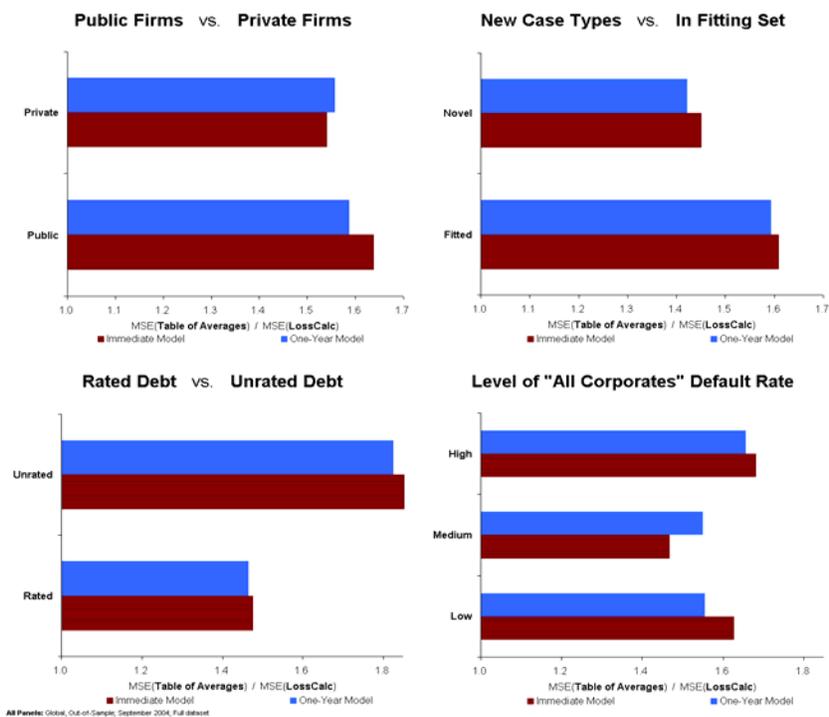
5.2.7 Other Tests

We have tested the LossCalc model extensively. We show a sampling of these tests in Figure 16 below. We have performed all of these tests on our full global dataset, out-of-time and out-of-sample, as of September 2004. The top left panel shows that LossCalc is accurate for both private and public borrowers even though one of the model inputs is our public-firm EDF technology (i.e., Distance-to-Defaults).

The top right panel shows that LossCalc can predict LGD for new situations not included in its fitting dataset. We did this by testing the model on industry and seniority class that were not present within the “fitting set” of data.

The bottom two panels show that for unrated instruments (on the left) and for periods of above/below average default rate environments, LossCalc performs better than tables.

FIGURE 16 LossCalc has Undergone Extensive Testing



These figures examine LossCalc’s predictive accuracy for four different types of divisions across our dataset. By construction, if these bars have any length to the right of the “1.0” point, then LossCalc outperforms a Table of Averages. LossCalc performs well for both Public and Private firms with only slightly diminished accuracy for Private firms (top left panel). LossCalc’s performance was only somewhat reduced for cases (Industry/Seniority combinations) that were unrepresented in its fitting set (top right panel). LossCalc performed relatively better during *unusual* times or *odd* cases. Shown here are two examples of this: (1) the case of unrated debt (bottom left panel), and (2) the case of either above or below average default rates (bottom right panel).

5.2.8 Performance

LossCalc is a better predictor of LGD than the traditional methodologies of historical averages segmented by debt type and seniority. By “better,” we mean that LossCalc:

- Estimates have significantly *lower error* as might be stated by mean squared error or standard deviation. See Figure 10.
- Estimates have significantly *more correlation* with actual outcomes. This means they have better tracking of both high and low recoveries. See Figure 11.
- Makes *far fewer large errors*. See Figure 12.
- Provides *better discrimination between instruments of the same type*. For example, the model provides a much better ordering (best to worst recoveries) of bank loans than historical averages. See Figure 13.
- Over 10% of the time, shows *improvement in predictive accuracy (reduction in error) greater than 28% of the actual true LGD*.
- Version 2.0 has *tighter prediction intervals* than other approaches. Therefore, a portfolio manager has more certainty of LGD forecasts.

6 THE DATASET

We developed LossCalc v2 on 3,026 *global* observations of LGD for defaulted loans, bonds, and preferred stock occurring between Jan-1981 and Dec-2003. The dataset includes 1,424 defaulted public and private firms in all industries. In US\$ equivalents, the issue sizes range from \$370 thousand to \$4.6 billion, with a median size of about \$125 million. The median firm size (sales at annual report before default) was \$660 million, but ranged from zero up to \$48 billion. (Since neither the size of the debt nor firm size is predictive of recovery, these size statistics reflect only the nature of the dataset.)

The dataset used to develop and test LossCalc is Moody's Investors Service proprietary default and recovery database (used in Moody's annual default study) plus other types of data such as financial statements, credit indices and firm, industry, and credit market information made available with the 2002 merger with KMV.

The default data is secondary market pricing of defaulted debt as quoted one month after the date of default. We use debt-issue specific market quotes that are *not* "matrix" prices, which are broad broker-created tables keyed off maturity, credit grade, and instrument type with no particular consideration of the specific issuer.

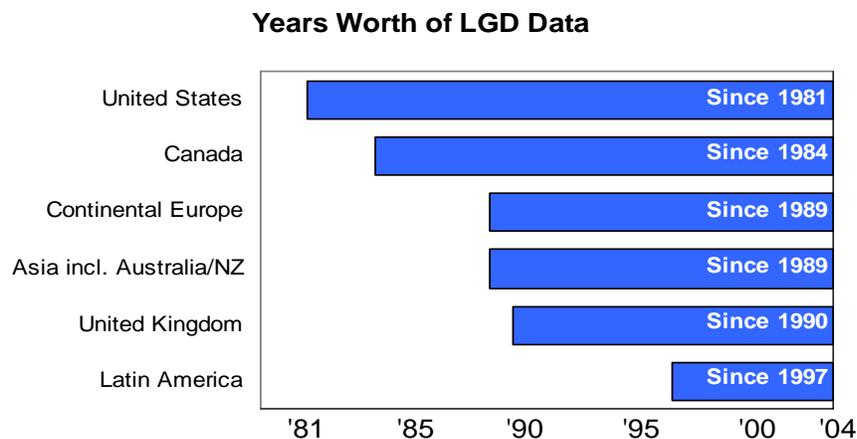
6.1 Historical Time Period Analyzed

LossCalc uses recovery observations since January 1981 so that it covers two full economic cycles. We did this because our research found that the credit cycle was a strong determinant of recoveries. We also require financial statement data, which are only reliably available since 1981. For both of these reasons, the period we used was Jan-1981 to the present.

6.2 Scope of Geographic Coverage and Legal Domain

Bankruptcy laws vary across legal domains, for example; UK law strongly seeks to protect creditors. French law does not recognize the priority of claims of specifically identified security. Some domains allow *creditors* to file a petition for insolvency.²⁸ LossCalc maintains predictive power across countries and regions by specifying a consistent framework that uses locally sourced predictive factors. We construct these factors on our dataset of 25,000 firms worldwide.

FIGURE 17 Depth of LGD Data by Country / Region



The LGD dataset for LossCalc extends across countries and regions. Although data collection has been in place the longest in North America, all regions have at least the 7-year minimum dataset as prescribed by the Advanced IRB requirements.

²⁸ See West, & de Bodard [2000a,b,c] and Bartlett [1999].

Bankruptcies in all countries are a two-step process of (1) determining the fundamental economic value of the firm that is available to satisfy creditors, and (2) determining the allocation of those monies to creditors. We find that the first step drives a significant proportion of LGD variability. This is why we find that our predictive factors perform so well worldwide. See Figure 7.

6.3 Scope of Firm Types and Instrument Categories

Our dataset includes three broad debt instrument types: (1) bank loans, (2) public bonds, and (3) preferred stock. We have organized loans broadly into two seniority classes: “senior secured”, which are the more numerous, and “senior unsecured.” Public bonds subdivide into seven seniority classes:

- Industrial revenue bonds;
- Corporate mortgage bond;
- Senior secured;
- Senior unsecured;
- Senior subordinated;
- Subordinated; and
- Junior subordinated.²⁹

In addition, since Medium Term Note programs (MTNs) are typically a large number of small issues, we consolidate the many individual issues that an obligor may have into a single recovery observation per default event. Otherwise, LossCalc would over-weight MTNs. The recovery rate realized for this proxy observation is the simple average of all the individual issuances. It happens that there is never large variability in the recovery rates across issues within a single MTN program.

Our dataset includes wide range of both public and private firms. The median firm size (sales at annual report before default) was \$660 million, but ranged from zero up to \$48 billion.

7 A CASE STUDY

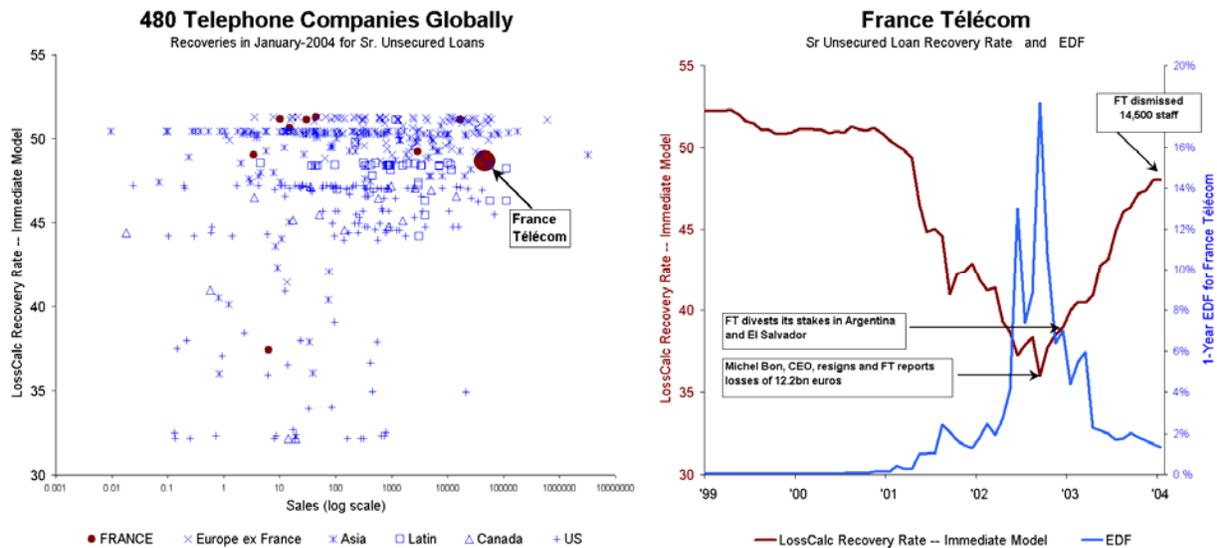
The implications of a fully specified dynamic LGD like LossCalc can be dramatic for analysts who are more accustomed to Table of Averages type LGD models. It is worth giving a brief example to highlights LossCalc’s behavior.

France Télécom is a stable, utility-like, firm that expanded aggressively into wireless communication through a series of acquisitions. These acquisitions included: British NTL (£8B; Jul’99), CTE Salvador (US\$275M; Sep’99), Jordan Telecom (US\$500M; Jan’00), Germany’s MobilCom (3.7B; Mar’00), British Orange (£26.5B; May’00), as well as acquiring Telecom Argentina through a consortium. By May 2002, NTL had filed for bankruptcy, while that summer MobilCom was near bankruptcy. By the end of 2002, France Télécom’s 70B of total debt was three times the value of the company.

Although France Télécom was half state owned, the French government did not move to rescue MobilCom. Europe had been seeking to privatize state run firms. There were also the two recent American Telecom examples of WorldCom’s bankruptcy filing in January 2002 and Global Crossing filing in July 2002.

²⁹ See Mann [1997] for a discussion of the common use and legal treatment of differing seniority classes and Stumpp, et.al. [1977] for a detailing of bank loan structure and collateral.

FIGURE 18 Telecom Recoveries and Specifically France Télécom



Shown here is the Telecom industry both at a point in time, left panel, and one of its members across time, right panel. The left panel shows recoveries (for Senior Unsecured Loans) during January 2004, for 480 firms in all country / regions. Note that even this homogeneous group can have a wide range of possible recovery expectations with the top recoveries nearly double the bottom recoveries. The right panel shows one firm, France Télécom, over a five-year period from Jan'99 through Jan'04. Its recovery dramatically varies during this period. Extreme swings in credit distress, as evidenced by the Moody's KMV Expected Default Frequency (EDF), drove this result.

Figure 18 illustrates a dramatic example of LGD zigzagging across time. More commonly, shifts at the industry/region level are what cause LossCalc changes in LGD over time. What is most clear is that there is no economic reason to believe that a security's LGD should be static across time. Given our knowledge of what drives differences in LGD, it becomes only a question of how best to utilize that knowledge to understand, moderate, and manage a portfolio's risk profile.

8 CONCLUSION

The issue of prediction horizon for LGD is one that has gotten little attention due to the largely static nature of the dominant historical average approach. This implicitly ignores the effects of the credit cycle and other time-varying environmental factors. LossCalc, by its dynamic nature, allows for a much more exact specification of LGD horizon and produces estimates on an immediate and one-year horizon.

LossCalc performs better than common alternatives such as an overall historical average of LGD or a table of average LGDs in categories. LossCalc's performance is superior in both out-of-sample and out-of-time prediction error and correlation of the predictions with actual recovery experience. The model is also better at identifying recoveries that are lower than historical average methods and has fewer large errors.

Combining the results of Moody's KMV LossCalc for recovery rates with EDF values allows a risk manager to estimate the expected credit losses of an instrument, given knowledge of the exposure amount. This is a much more complete portrayal of credit risk than probability of default alone.

Estimation of the average (mean expected) losses due to credit is commonly used to:

- Calculate provisions to a reserve for credit losses;
- Set collateral requirements that are industry sensitive and dynamic with the economy;
- Establish minimum pricing levels at which new credit exposures to an obligor may be undertaken;

- Support Basel compliance towards the Advanced IRB approach;
- Judge relative pricing of credit risky instruments such as corporate bonds or credit default swaps; and
- Calculate portfolio Credit-VaR such as in PortfolioManager™.

LossCalc's estimate of the downside prediction interval of LGD (a 90% two-tailed prediction interval) gives greater meaning, consistency, and guidance to stressing recovery rate estimates and for provisioning loss reserves.

LossCalc represents a robust and validated global model of LGD for the debt types it covers. We believe that this is a productive step forward in answering the call for rigor that the BIS has put forth in the Basel II Accord.

APPENDIX A BETA TRANSFORMATION TO NORMALIZE LOSS DATA

To create an approximately normally distributed dependent variable from the raw observations of recovery, we first confirmed that defaulted debt valuations were approximately Beta distributed. There is no theoretical reason that this is the “correct” shape of the defaulted debt prices, but previous studies have concluded that its characteristics make the Beta a reasonable description of the empirical shape.

Beta distributions are described in this case by an upper and lower bound and by two shape parameters, α and β . They are usually bounded between zero and one, where the mean can be any value strictly within this range. For LossCalc we generalize this distribution to accommodate the rare, but non-trivial cases where recoveries can range somewhat above 1.0. The conversion of the Beta distributed recovery values to a more normally distribution dependent variable is explicitly defined as follows:

$$\text{Dependent Variable} = Y_i = N^{-1} \left[\text{Betadist} (\text{RecovRt}_i, \alpha_d, \beta_d, \text{Min}, \text{Max}_d) \right]$$

where

$N^{-1} \equiv$ The Inverse of the Normal Cumulative Distribution

$\text{RecovRt} = \min (\text{Max} - \varepsilon, \text{observed recovery rate}) \quad \varepsilon = \text{some small value}$

$\alpha_d =$ The Beta Distribution's *center* parameter (6)

$\beta_d =$ The Beta Distribution's *shape* parameter

$\text{Min} =$ Set to zero in all cases

$\text{Max}_d =$ Set to 1.1 for bonds, but otherwise is 1.0

$d = \{ \text{loans, bonds, preferred stock} \}$

We use the sub-notation “ d ” to emphasize that LossCalc fits each debt type to its own distribution.

Thus, much of the distributions of our three separate asset classes can be captured by specifying only two shape parameter values: the α and the β of each Beta-distribution. There are various ways of fitting the distribution parameters. It is also possible, through algebraic manipulation, to specify the Beta distribution that simply matches the mean and standard deviation, which are functions of the shape and boundary parameters.

Mathematically, a Beta distribution is a function of Gamma distributions. With the lower bound, Min , fixed at zero, the distribution is as follows:

$$\text{Beta} (x, \alpha, \beta, \text{Min} = 0, \text{Max}) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \left(\frac{x}{\text{Max}} \right)^{\alpha-1} \left(1 - \frac{x}{\text{Max}} \right)^{\beta-1} \left(\frac{1}{\text{Max}} \right) \quad (7)$$

The shape parameters can be derived in a variety of ways. For example, the following give them in terms of population mean and standard deviation.

$$\alpha = \frac{\mu}{\text{Max}} \left[\frac{\mu \cdot (\text{Max} - \mu)}{\text{Max} \cdot \sigma^2} - 1 \right] \quad \text{and} \quad \beta = \alpha \left[\frac{\text{Max}}{\mu} - 1 \right] \quad (8)$$

Conversely, given Beta distribution parameters, it is straightforward to calculate the mean and standard deviation.

$$\mu = \text{Max} \cdot \left(\frac{\alpha}{\alpha + \beta} \right) \quad \text{and} \quad \sigma = \text{Max} \cdot \sqrt{\frac{\alpha \cdot \beta}{(\alpha + \beta)^2 \cdot (1 + \alpha + \beta)}} \quad (9)$$

APPENDIX B AN OVERVIEW OF THE VALIDATION APPROACH

To validate the performance of LossCalc, we have used the approach adopted and refined by Moody's KMV which is also used to validate the RiskCalc models of default prediction. The approach, termed *walk forward* validation, is a robust means for ensuring that:

- Models have not been “over-fit” to the data;
- Future performance can be well understood; and
- The modeling approach, as well as any individual model produced by it, is robust through time and credit cycles.

We give only a brief overview of the methodology here; a fuller description is detailed in Stein [2002].³⁰

Controlling for “Over Fitting”: Walk-Forward Testing: In order to avoid embedding unwanted sample dependency, we have found it useful to develop and validate models using some type of out-of-sample, out-of-universe, and out-of-time testing approach on panel or cross-sectional datasets.³¹ However, such an approach can generate false impressions about reliability of a model if done incorrectly. “Hold out” testing can sometimes miss important model problems, particularly when processes vary over time, as credit risk does.³² Dhar & Stein [1997] suggest a framework for framing these issues as well as providing a more detailed discussion and some examples from finance.

We designed our testing approach to test models in a realistic setting that emulates closely the practical use of these models. The trading model literature often refers to this procedure as “walk-forward” testing.

The walk-forward procedure works as follows:

1. Select a year, for example, 1992.
2. Fit the model using all the data available on or before the selected year.
3. Once the model's form and parameters are established for the selected period, generate the model outputs for all of the instruments available during the following year (in this example 1993). Note that these are out-of-time and generally out-of-sample (though there are rare cases of firms defaulting more than once).
4. Save the prediction as part of a *result set*.
5. Now move the window up one year (e.g.: to 1993) so that all the data through that year can be used for fitting and the data for the following year can be used for testing.
6. Repeat steps (2) to (5) adding the new predictions to the result set.

Collecting all the out-of-sample and out-of-time model predictions produces a set of model performances. We use this *result set* to rigorously validate the performance of the model in more detail.

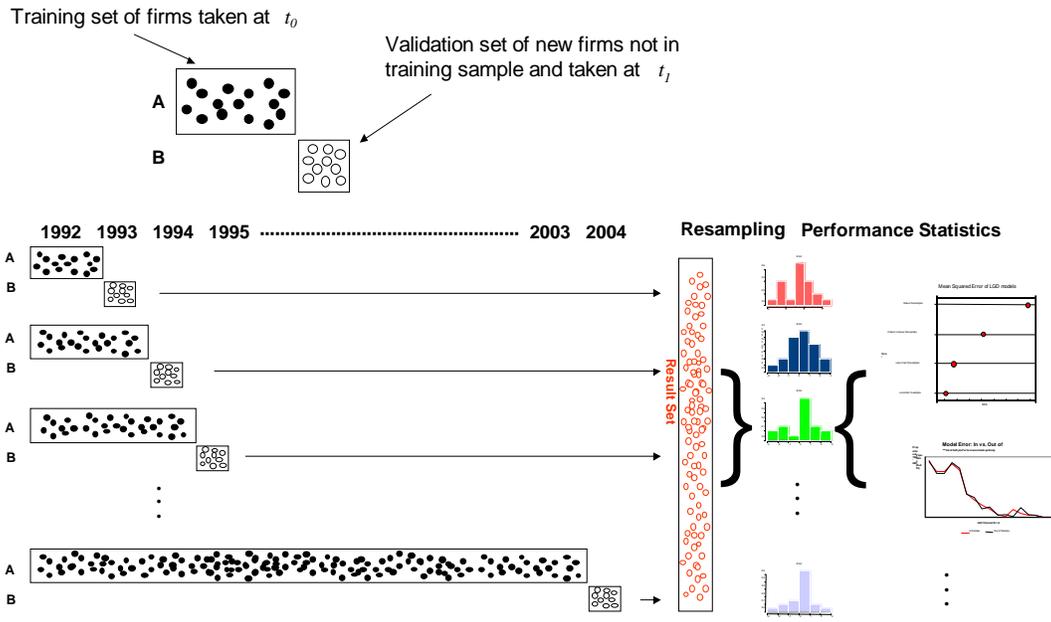
Note that this approach closely simulates how Moody's KMV and our clients actually use LossCalc in practice. Each year, the model is refit and used to predict recoveries one year hence. We outline the walk-forward validation process in Figure 19.

³⁰ Much of what follows was adapted from Sobehart, Keenan & Stein [2000] and Stein [2002].

³¹ A panel dataset contains observations over time on many individuals. A cross sectional dataset contains one observation on many individuals.

³² See, for example, Mensah [1984].

FIGURE 19 Validation Methodology: End-to-End



We fit a model using a sample of historical recovery data and test the model using data on new recoveries one-year later (upper portion of exhibit). Dark circles represent data for fitting the model and white circles represent validation data. We perform “walk-forward testing” (bottom left) by fitting the parameters of a model using data through a particular year, and testing on data from the following year, and then stepping the whole process forward one year.

Note that this approach has two significant benefits. First, it allows us to get a realistic view of how a particular model would perform over time. Second, it allows us to leverage to a higher degree the availability of data for validating models. Unless otherwise noted, all results presented in this study are from this type of out-of-sample and out-of-time walk-forward testing.

APPENDIX C POWER CURVE CONSTRUCTION FOR LGD

In order to determine how well a model works, the power curve looks to determine how well the model predicts high losses and orders the LGD outcomes from high losses to low losses. If an LGD is greater than the historical LGD, this is a “bad” loss (similar to a company defaulting in a default model). If the actual default is less than or equal to the historic average, then it is considered a “good” outcome (similar to a non-default in a default model).

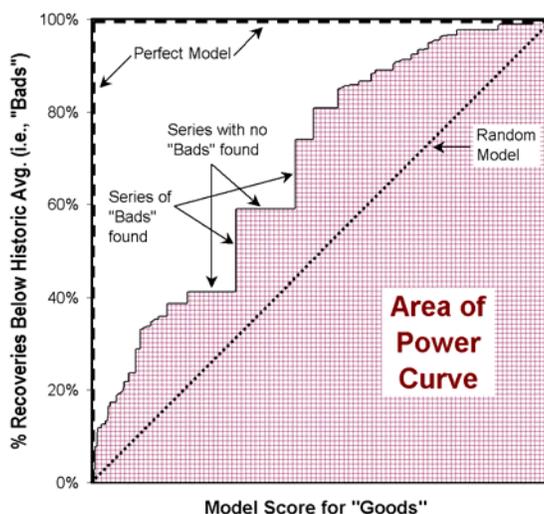
- *good*: Actual LGD \leq Historical Average
- *bad*: Actual LGD $>$ Historical Average

This technique will label half the population as “goods” and half as “bads.” Since this labeling is a function only of the actual and the average, not the specific models under evaluation, the same set of labels will apply to any alternative models we picture in the graph.

The Power Curve estimates the accuracy of the model in ranking the recoveries from bad to good. To do this, the model scores are ranked from highest LGD to lowest. This ranks the model’s observations from what the model rated as bad to good. Each model observation is compared in order of model score (LGD) against the actual bad or good result. If a result was “bad” and the model predicted this, then the curve moves up one point on the y -axis. If the model predicted a result was bad and it was good, this does not increase the model power so it moves one point along the x -axis. If a result is good and the model identifies it as good, the curve moves again across the x -axis. The result of this process is that if a model correctly identifies all the “bads” it will move vertically up the y -axis, and the remainder would be goods that would move it horizontally across the x -axis. The result would be a “perfect model” shown in Figure 20.

If a model is random, the goods and bads will be evenly mixed throughout the sample so the line would step-up at a 45-degree angle. For an individual model, its power is given by how well it ranked bad and good results. The sooner it correctly identifies bads, the sooner and higher it would go on the y -axis, and the more area the curve would cover.

FIGURE 20 Interpretation of an LGD Power Curve



A Power Curve for LGD graphically portrays the ability of an LGD model to discriminate between “good” LGDs (those below the Historical Average) versus “bad” LGDs (those above the historical average). The more accurately an LGD model assigns its lowest (worst) scores to cases that truly are poor, the more rapidly the power curve will ascend. A perfect model will cover the entire area. A random model will have no discriminatory power and will be a diagonal line. The area under the power curve (shaded) is a nonparametric measure of LGD model power.

APPENDIX D RISK NEUTRAL PRICING AND LGD

D.1. Equating Final recovery and Market Price

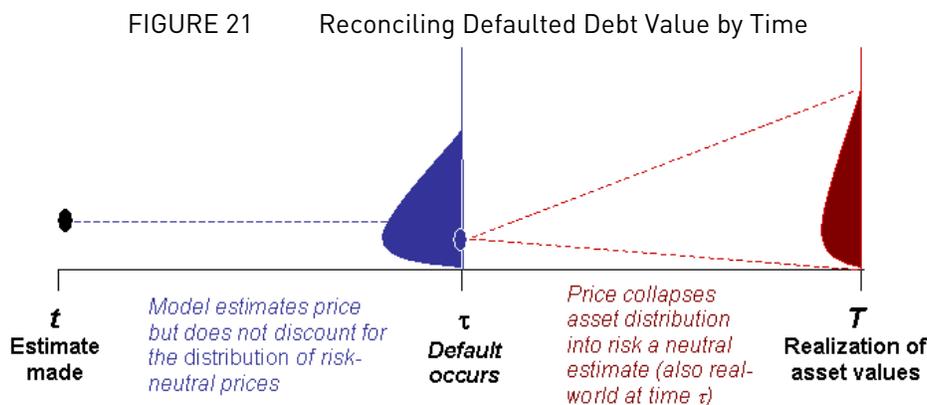
The ultimate recovery of a defaulted instrument can be related back to an observable market price by using risk neutral pricing analysis.

Default recoveries are uncertain and their ultimate realized value will be across a range of possible outcomes. If these possible outcomes are discounted back to the expected default date, that will give a range of possible default prices. The market estimates the most likely value of these possible outcomes, discounting for the uncertainty. This is the price at default.

D.2. Uses of Risk Neutral Pricing

Risk neutral LGD pricing has two uses. First, if the default price is known, then the ultimate realized value is implied.

Second, for market valuation of assets, the default price prediction can be discounted back to give the risk neutral recovery rate. This, combined with the risk neutral probability of default, can be used to value an instrument.³³



Starting at the far right, the value of a defaulted debt instrument is known when the defaulted borrower emerges from bankruptcy as a restructured entity, liquidates, merges, etc. This may take many different values (the red distribution). At the time of default, τ , the market seeks a risk neutral value (price) by discounting the expected distribution of these possible realizations. Using LossCalc, we can place ourselves earlier in time, at t . There a distribution of possible prices that may result from a default, time τ , as represented by the LossCalc prediction interval (the blue distribution). LossCalc does not apply a discount at time t for this uncertainty. Not pictured here is the accounting charge-off event, which may occur before time T , see Figure 8.

Market price p_τ is an ex-ante, risk neutral expectation, $r_{\tau|T}^*$ of ultimate realized value, at time τ , the time of default:

$$p_\tau = r_{\tau|T}^* \tag{10}$$

- At time t , $t < \tau$, the model estimate of market price at some future time is an estimate of what the risk neutral expectation will be at time τ :

³³ Portions of this section are based on Stein [2004].

$$\hat{y}_{t,t<\tau} = E_t[p_\tau] = E_t[r_{\tau|T}^*] \quad (11)$$

- However, for valuation, we should also discount for the uncertainty about the future value of $r_{\tau|T}^*$:

$$\hat{y}_{t,t<\tau}^* = E_t^*[r_\tau^*] \quad (12)$$

APPENDIX E LOSSCALC WEB-SITE OUTPUTS

Output Parameter	Definition
LossCalc Recovery Rate:	This is the mean recovery rate expectation. This and all other model outputs are shown for both an “Immediate” horizon and a “1-Year” horizon. The Immediate figures reflect the expectation stemming from a default that occurs “tomorrow.” The 1-Year figures reflect the expectation stemming from a default that occurs one year from now. Note that LossCalc generates “physical” as opposed to “risk-neutral” measures. At the Immediate horizon, these two are the same.
Lower, Upper 50%-Prediction Interval:	This range is expected to cover half of future LGD realizations. So, 50% of future recovery rate realizations can be expected to be outside this interval. Note, due to the bounded range of possible recoveries, that this interval is not necessarily centered on the mean recovery rate expectation. Note also that this does not reflect the more familiar, <i>standard error of the prediction estimate</i> , $\hat{\sigma}_e$, which is much smaller by its definition.
Lower, Upper 90%-Prediction Interval:	This range is expected to cover 90% of future LGD realizations. So, 10% of future recovery rate realizations can be expected to be outside this interval. Note, due to the bounded range of possible recoveries, that this interval is not necessarily centered on the mean recovery rate expectation. Note also that this does not reflect the more familiar, <i>standard error of the prediction estimate</i> , $\hat{\sigma}_e$, which is much smaller by its definition.
Median:	This is the median recovery rate forecast rather than the mean expectation (see above). It is modeled in its own right rather than being derived from, say, the Beta distribution described below.
Historical Average Recovery Rate:	This is a simple average of all recovery rates (instrument weighted) that have occurred for a particular debt-type and seniority class. This is meant to represent the publicly available information that would be available to benchmark recoveries within a debt-type and seniority class. Except for differences in compilation period, these figures should resemble numbers from the historical research of Moody’s Investors Service.
LGD (1 - LossCalc Recovery Rate):	This is simply one (1) minus the Recovery Rate generated by LossCalc.
Portfolio Manager™ LGD Parameter (k):	This parameter is used within the Portfolio Manager™ methodology to describe the shape of the recovery distribution: $k = \alpha + \beta + 1$
Beta Distribution Parameter Alpha (α):	One of two parameters that define a Beta distribution; α is sometimes referred to as the “location” parameter.
Beta Distribution Parameter Beta (β):	One of two parameters that define a Beta distribution; β is sometimes referred to as the “scale” parameter. (Alternative texts refer to $[\beta - \alpha]$ as being the “scale.”)

Output Parameter	Definition
Beta Distribution Variance (σ^2):	<p>This is a measure of the width or dispersion of the Beta distribution. Over a "unit interval", (i.e., recovers from 0% to 100%), the maximum range of the variance can extend from 0 to 25%. It can be calculated from the α and β parameters directly as:</p> $\frac{\alpha\beta}{(\alpha + \beta)^2 (\alpha + \beta + 1)}$

REFERENCES

1. Altman, E. I., A. Gande and A. Saunders (2004). "Informational Efficiency of Loans Versus Bonds: Evidence from Secondary Market Prices." New York University and Vanderbilt University.
2. Altman, E. I. and V. M. Kishmore (1996). "Almost Everything you Wanted to Know about Recoveries on Defaulted Bonds." *Financial Analysts Journal*. November/December: 56-62.
3. Bartlett, F. (1999). "Regimes, Recoveries and Loan Ratings: the importance of insolvency legislation." Fitch/IBCA - Loan Products.
4. Basel Committee on Banking Supervision (2001). "The New Basel Capital Accord." Bank for International Settlements.
5. Borenstein, S. and N. L. Rose (1995). "Do Airlines in Chapter 11 Harm Their Rivals?: Bankruptcy and Pricing Behavior in U.S. Airline Markets." National Bureau of Economic Research.
6. Crosbie, P. J. and J. R. Bohn (2003). "Modeling Default Risk." Moody's KMV.
7. Davydenko, S. A. and J. R. Franks (2004). "Do bankruptcy codes matter? A study of defaults in France, Germany, and the UK." London Business School.
8. Dhar, V. and R. Stein (1998). "Finding Robust and Usable Models with Data Mining: Examples from Finance." *PCAI*.
9. Eberhart, A. C. and R. J. Sweeney (1992). "Does the Bond Market Predict Bankruptcy Settlements?" *Journal of Finance*. Vol. XLVII. 3: 943-980.
10. Fischer, T.C.G. and J. Martel (2003). "The Firm's Reorganization Decision: Empirical Evidence from Canada." Wilfrid Laurier University, and Université de Cergy-Pontoise.
11. Fitzenberger, B. Koenker R. and Machado, J.A.F. (2002), *Economic Applications of Quantile Regression*, Physica-Verlag Heidelberg, New York.
12. Gordy, M. and D. Jones (2002). "Capital Allocation for Securitizations with Uncertainty in Loss Prioritization." Federal Reserve Board.
13. Guha, R. (2002). "Recovery of Face Value at Default: Theory and Empirical Evidence." London Business School.
14. Gupton, G. M., C. C. Finger and M. Bhatia (1997). "CreditMetrics™ - Technical Document." J.P. Morgan.
15. Gupton, G. M., D. Gates and L. V. Carty (2000). "Bank Loan Loss Given Default." Moody's Risk Management.
16. Gupton, G. M. and R. M. Stein (2001). "A Matter of Perspective." *Credit*. Vol. 2.9: 22-23.
17. Gupton, G. M. and R. M. Stein (2002). "LossCalc™: Model for Predicting Loss Given Default (LGD)." Moody's Investors Service.
18. Hamilton, D. T., et.al. (2004). "Default & Recovery Rates of Corporate Bond Issuers," Moody's Investors Service.
19. Hamilton, D. T. (1999). "Debt Recoveries for Corporate Bankruptcies," Moody's Investors Service.
20. Hamilton, D. T. and A. Berthault (2000). "The Investment Performance of Bankrupt Corporate Debt Obligations: Moody's Bankrupt Bond Index 2000." Moody's Investors Service.
21. Hamilton, D. T., G. M. Gupton and A. Berthault (2001). "Default and Recovery Rates of Corporate Bond Issuers: 2000." Moody's Risk Management.
22. Heckman, J. J. (1979). "Sample Bias as a Specification Error." *Econometrica*. Vol. 47. 1: 153-161.
23. Ivanova, V. (2004). "LGD-Rating for a Portfolio of Retail Loans." University of Oxford.
24. Jarrow, R. (2001). "Default Parameter Estimation Using Market Prices." *Financial Analysts Journal*. Vol. 57. 5: 75-92.
25. Longhofer, S. D. (1997). "Absolute Priority Rule Violations, Credit Rationing, and Efficiency." *Journal of Financial Intermediation*. Vol. 6. 3: 249-267.
26. Longhofer, S. D. and C. T. Carlstrom (1995). "Absolute Priority Rule Violations in Bankruptcy." *Economic Review*: 21-30.
27. Mann, R. J. (1997). "Explaining the Patterns of Secured Credit." *Harvard Law Review*. Vol. 110. 3: 625-683.

28. Mella-Barral, P. (1999). "The Dynamics of Default and Debt Reorganization." *The Review of Financial Studies*. Vol. 12. 3: 535-578.
29. Mensah, Y. M. (1984). "An Examination of the Stationarity of Multivariate Bankruptcy Prediction Models: A Methodological Study." *Journal of Accounting Research*. Vol. 22. 1: 380-395.
30. Onorota, M. and E.I. Altman (2003). "An Integrated Pricing Model for Defaultable Loans and Bonds." City University London and New York University.
31. Pesaran, M.H., T. Schuermann, B.J. Treutler and Weiner (2004). "Macroeconomic Dynamics and Credit Risk: A Global Perspective." University of Cambridge.
32. Schmit, M. and J. Stuyck (2002). "Recovery Rates in the Leasing Industry." Leaseurope.
33. Schönbucher, P. J. (2003). *Credit Derivatives Pricing Models*. John Wiley & Sons.
34. Singh, M. (2003), "Recovery Rates from Distressed Debt - Empirical Evidence from Chapter 11 Filings, International Litigation, and Recent Sovereign Debt Restructurings", IMF.
35. Sobehart, J., S. Keenan and R. Stein (2000). "Validation methodologies for default risk models." *Credit*: 51-56.
36. Sobehart, J. R., S. C. Keenan and R. M. Stein (2000). "Benchmarking Quantitative Default Risk Models: A Validation Methodology." Moody's Risk Management.
37. Sobehart, J. R. and R. M. Stein (2000). "Moody's Public Firm Risk Model: A Hybrid Approach To Modeling Short Term Default Risk." Moody's Risk Management.
38. Sprent, P. (1998). *Data Driven Statistical Methods*. Chapman-Hall/CRC Press.
39. Stein, R. M. (2002). "Benchmarking Default Prediction Models: Pitfalls and Remedies in Model Validation." Moody's KMV.
40. Stein, R. M. (2004) "Discussion Acharya *et al*", The Moody's NYU Innagural Credit Risk Clonference, *Recent Advances in Credit Risk Research*, New York, May 19.
41. Stumpp, P., T. Marshela, M. Mulvaney, et.al. (1997). "A Sense of Security: Moody's Approach to Evaluating Bank Loan Structure and Collateral." Moody's Investors Service.
42. Swank, T. A. and T. H. Root (1995). "Bonds in Default: Is Patience a Virtue?" *Journal of Fixed Income*. June: 26-31.
43. Tasche, D. (2004). "The Single Risk Factor Approach to Capital Charges in Case of Correlated Loss Given Default Rates." Bundesbank.
44. Unal, H., D. Madan and L. Guntay (2001). "A Simple Approach to Estimate Recovery Rates with APR Violation from Debt Spreads." University of Maryland.
45. Wagner III, H. S. (1996). "The Pricing of Bonds in Bankruptcy and Financial Restructuring." *The Journal of Fixed Income*. June: 40-47.
46. Ward, D. J. and G. L. Griepentrog (1993). "Risk and Return in Defaulted Bonds." *Financial Analysts Journal*. May/June: 61-65.
47. West, M. and E. de Bodard (2000). "Bankruptcy and Ratings: A Leveraged Finance Approach for Europe." Moody's Investors Service.
48. West, M. and E. de Bodard (2000). "Bankruptcy and Ratings: A Leveraged Finance Approach for Europe - Part II: France." Moody's Investors Service.
49. West, M. and E. de Bodard (2000). "Bankruptcy and Ratings: A Leveraged Finance Approach for Europe - Part III: Germany." Moody's Investors Service.
50. Xu, P. (2004). "Bankruptcy Resolution in Japan: Corporate Reorganization vs. Civil Rehabilitation." University and RIETI.