Credit Risk in Private Debt Portfolios
Author(s): Mark Carey
Source: *The Journal of Finance*, Vol. 53, No. 4, Papers and Proceedings of the Fifty-Eighth Annual Meeting of the American Finance Association, Chicago, Illinois, January 3-5, 1998 (Aug., 1998), pp. 1363-1387
Published by: Blackwell Publishing for the American Finance Association
Stable URL: http://www.jstor.org/stable/117405
Accessed: 25/11/2009 09:58

THE JOURNAL OF FINANCE • VOL. LIII, NO. 4 • AUGUST 1998

Credit Risk in Private Debt Portfolios

MARK CAREY*

ABSTRACT

Default, loss severity, and average loss rates for a large sample of privately placed bonds are presented and compared with loss experience for publicly issued bonds. The chance of very large portfolio losses is estimated and some determinants of such losses are analyzed. Results show ex ante riskier classes of private debt perform better on average than public debt. Both diversification and the riskiness of individual portfolio assets influence the bad tail of the portfolio loss distribution. Private placements are similar to corporate loans in that both are monitored private debt. The results are thus relevant to management and securitization of private debt portfolios generally.

MORE THAN HALF OF NONFINANCIAL corporate debt finance is issued privately, and decisions of the financial intermediaries that invest in such debt depend importantly on its credit risk. Regulators of such institutions are similarly attentive to portfolio credit risk, and it is an important consideration in the design of securitizations of pools of private debt.

In spite of its importance, little is known empirically about private debt credit risk, especially ex ante risk as a function of portfolio characteristics. Many studies have examined the loan portfolio performance of banks, thrifts, and insurance companies, but the ex ante risk characteristics of the portfolios have been largely unobservable (examples include Jones and King (1995), Berger (1995), and Avery and Berger (1991)). Thus, the results of such studies provide little guidance about the risk posed by an individual portfolio or institution, or even about groups of similar institutions out of sample.

Recently developed formal credit risk management systems have generally finessed this lack of information partly by assuming that private and

*Federal Reserve Board and Society of Actuaries Asset Risk Experience Study Private Placement Subcommittee. This paper is derived from and is an extension of the *1986–1992 Credit Risk Loss Experience Study: Private Placement Bonds*, sponsored by the Society of Actuaries and supervised by the above-mentioned committee. As such, credit for the work accrues to the Society and to others involved in the study as well as myself: Jack Luff, Warren Luckner, Nick Bauer, Mike Zurcher, Kin Tam, Stacy Gill, Serge Goulet, Bill McDonald, Nancy Morse, Mark Poeppelman, Ken Roberts, Brad Scher, George Silos, and Dan Towriss. In addition, thanks to Ned Elton, Darrell Duffie, John Mingo, Ken Singleton, and Arthur Warga for useful conversations and to Margaret Kyle and Mark Hrycay for research assistance. I alone am responsible for any errors, and this paper reflects my views and not necessarily those of the Society of Actuaries, the Board of Governors, or the Federal Reserve System.

public debt perform similarly.[1] In estimating average or expected losses conditional on portfolio composition, private debt instruments are rated on a scale similar to that of the major bond rating agencies, and private debt default risk is assumed to correspond to that of similarly rated, publicly issued bonds. The severity of loss in the event of default is often assumed better for private than for public debt, but published studies of severity are limited. A principal problem with such assumptions is that the credit risk of public debt and monitored private debt might be different.

To limit the chance of unusually large losses (often called "unexpected losses"), portfolio managers traditionally have employed diversification strategies that are in the form of targets or limits. For example, limits are imposed on the share of the portfolio invested in individually risky assets or on exposure to a single borrower. However, little is known empirically about how loss rates in the bad tail of the portfolio loss distribution vary as the limits change. Recently developed risk-management systems take a somewhat different approach, modeling marginal contributions of individual assets to portfolio bad-tail losses. However, such models involve largely untested assumptions about the conditional covariance of losses across assets.

Regardless of the portfolio management method, the equity share of a financial intermediary's capital structure is (in principle) set to provide a buffer sufficient to absorb credit (and other) losses in all but very rare circumstances, thus sharply limiting the risk of institutional failure. Similarly, many collateralized loan obligations (CLOs) involve subordinated tranches large enough to make senior tranches almost free of credit risk.[2]

This paper empirically examines two characteristics of private debt portfolio credit risk loss rate distributions, conditional on individual asset risk and other portfolio characteristics. Estimates of both expected losses and the size of loss rates in the bad tail of the conditional loss distribution are reported. Data are from the combined portfolios of privately placed bonds of about a dozen major insurance companies for the period 1986 to 1992, with default and loss measured at the individual asset level. The data represent about 25 percent of all placements outstanding during the sample period. Ex ante risk of individual assets is measured by both investors' internal risk ratings and regulatory ratings. Results of previously published public-bond default studies are used in comparing expected losses on similarly rated public and private debt.

One important finding is that although average loss rates on private debt are similar to those on public debt for investment-grade assets (those rated BBB or better) and those rated BB, average private portfolio losses are bet-

[1] Examples of such systems include RAROC (risk-adjusted return on capital) systems, CreditMetrics, and CreditRisk+. See James (1996), Gupton, Finger, and Bhatia (1997), and Credit Suisse (1997). Jones and Mingo (1998) offer a generic description of such systems.

[2] CLOs are securitizations of pools of loans. Structures vary widely, but securities with senior and subordinated rights to cash flows from the pool are commonly issued and the subordinated tranches typically absorb credit-related losses up to some specified amount (Carlson and Fabozzi (1992)).

ter for the lower grades, and increasingly so as risk increases. Private default rates are slightly higher than public for the investment grades, but better loss severities have an offsetting effect on average portfolio loss rates. Both private default rates and severities are better for the risky grades, especially B and below. In obtaining these results, expected loss rates are measured as sample average rates for one-year horizons. Default rates are similarly one-year averages and thus may be interpreted as unconditional hazard rates. Loss severities are measured using present values of pre- and post-credit risk event cash flows.

In passing, such results may be interpreted as further evidence that private debt is monitored by lenders (Diamond (1984)). Intuitively, it is appealing that monitoring appears to have a larger effect on loss rates the higher the ex ante risk of the debt.

Loss rates in the bad tail of the portfolio loss distribution, conditional on portfolio characteristics, are estimated by Monte Carlo resampling methods. Simulated portfolios are constructed by drawing assets randomly from the total sample while enforcing diversification targets and limits. This method would provide good nonparametric estimates of loss rate distributions if draws were from the universe of possible assets and outcomes, and if all characteristics used by portfolio managers in selecting assets were modeled. The sample used here is of course not the universe but, as discussed further below, may yield reasonably representative results for the sensitivity of the bad tail to changes in diversification strategy. To my knowledge, this paper represents the first application of this method to estimation of credit risk loss distribution tail mass.

The Monte Carlo estimates imply that both individual asset risk and portfolio diversification strategies influence loss distribution tails. For example, to yield the same probability of credit losses to the senior tranche, a collateralized loan obligation involving a pool of solely below-investment-grade loans would require a subordinated tranche four to five times larger than would a CLO involving only investment-grade loans (when similar diversification rules are imposed). Required equity tranches are three to four times as large for relatively small portfolios as for very large portfolios. Given the limited upside of debt, at least from changes in credit risk, perhaps it is not surprising that it is much harder to diversify a portfolio of individually risky debt than one invested in safer loans.

Such Monte Carlo results are a potentially useful resource in the design of securitizations, in financial intermediary risk management, and in the setting of regulatory capital requirements. However, the bad-tail loss rates reported here are not estimates of the equity finance requirements for all real-world portfolios of private debt. Estimates are most representative for portfolios constructed by a manager who chooses assets randomly from those outstanding while staying within specified diversification limits. In reality, some managers may push the limits, whereas others may stay far away. For example, some may invest in relatively few large assets while still remaining below an exposure-to-one-borrower limit. Such portfolios would probably have fatter loss distribution tails than portfolios with a larger number of smaller assets. Similarly, some portfolio managers may enforce diversification rules not modeled here, and to the extent that such rules are effective their portfolios may pose smaller risk of unusually large losses.

Though tail-loss levels should be interpreted with care, estimates of changes in portfolio tail losses as diversification rules are changed are probably more robust. Here the situation is similar to that of a linear regression in which some independent variables are omitted. If omitted variables are uncorrelated with those included in the regression, parameter estimates for the included variables will not be biased by the omission. By analogy, limits on exposure to firms in a single industry are not modeled in this paper, even though they are often applied in reality, but estimates of the sensitivity of the bad tail to portfolio size are probably not affected too much.

Although private placements differ from bank and nonbank business loans in certain respects, especially in being fixed-rate and somewhat longer in term to maturity, they resemble loans in that they are monitored private debt (see Carey et al. (1993)). Placements are frequently secured and include, for example, restrictive covenants. Thus, this paper's results likely form a more accurate representation of bank loan risk than analyses based mainly on public bond default experience.[3]

Issuance of Rule 144A private placements began to be material only at the end of this paper's sample period and such placements represent an insignificant fraction of the data. In applying the results to 144As, however, it is important to distinguish those 144A placements that are essentially public debt from those that amount to traditional privates. This paper's results are likely to be reasonable estimates of the credit risk of the latter, whereas public bond experience is likely more relevant for the former.

Parts of this study are derived from the Society of Actuaries (1996), which reports results of the first two rounds of a continuing study of private placement credit risk experience. The Society of Actuaries study discusses in detail several matters mentioned only in passing in this paper, such as the relationship between regulatory risk ratings of placements and credit losses. It also presents detailed summary tables of the raw data. To my knowledge, the Society's studies are the first to offer comprehensive loss distribution information for corporate loans. Anecdotally, a few such studies have been done by or for individual financial intermediaries, but the results are proprietary, and the data represent smaller samples of historical experience.

The Society of Actuaries studies make two other contributions to the literature on debt default and loss. As noted, loss severities are based on ratios of net present values of original versus post-default cash flows, and are avail-

[3] On average, the monitoring of private placements appears to be somewhat less intense than that of bank loans, and nonprice terms appear less restrictive in placements, but this may reflect sample selection. Private placement portfolios contain smaller proportions of relatively high-risk assets than do bank loan portfolios. Thus, loss rates and other characteristics may be rather similar when the ex ante risk of loans and placements is similar.

able for every credit risk event. Thus, integrated measures of default and loss are possible. In contrast, public bond loss severities are usually estimated from post-default dealer indicative bid prices, which are not always available. In addition, here it is possible to include restructurings and distress sales in the definition of credit risk events, whereas public bond studies are generally limited to defaults.

For reasons of availability of data, most studies of default and loss on individual debt instruments have involved publicly issued corporate bonds, sovereign debt, or municipals. Examples include Altman (1989), Altman and Saunders (1998), Moody's (1996), Standard and Poor's (1995), Carty (1996), and McDonald and Van de Gucht (1996). Three recent studies of bank loan default rates or loss severities are Altman and Suggitt (1997), Asarnow and Edwards (1995), and Carty and Lieberman (1996). These three studies examine either default or severity, but not both.

In the remainder of the paper, Section I describes the data, Sections II and III describe the methods for measuring default and loss rates and report results, and Sections IV and V describe the Monte Carlo method of estimating loss distribution tails and report results. Section VI summarizes and suggests some avenues for future research.

I. Data

Thirteen major life insurance companies contributed annual portfolio composition and credit risk event (CRE) information for the years 1986 through 1992, although all companies did not contribute data for all years.[4] Contributing companies vary in size, but their aggregate share of all private placements outstanding is approximately 25 percent.[5] Thus, this paper's sample and the statistics drawn from it are probably representative of the universe of placements outstanding during the period. The vast majority of such outstandings were debt obligations of U.S. corporations. Significant volumes of U.S. private issuance by non-U.S. firms are more recent.

Each insurance company contributed information about each individual placement in its portfolio at each year-end during the sample period. Variables for each asset include principal outstanding, asset funding and maturity dates, interest rate coupon and yield, a variety of credit risk ratings, and several other variables. Some variables, such as seniority, secured status, and a direct measure of interest rate spreads, were collected only for

[4] Contributors included Aetna, Great-West Life, John Hancock, Lutheran Brotherhood, Metropolitan, Nationwide, New England Life, Principal Financial, Prudential, SAFECO, Sun Life, TIAA, and Washington Square Capital. Two of the thirteen contributed only for 1990–92, and three only for 1986–89.

[5] Contributing companies' share of total life insurance company private placements ranged from 34 to 44 percent over the years of the sample. In turn, life insurance companies held more than 80 percent of all outstanding private placements (see Carey et al. (1993)). The aggregate sample portfolio amount ranged from about $50 billion in 1986 to about $90 billion in 1992, with a steady uptrend across the years.

1990 through1992. On the whole, the data provide a detailed description of the population of placements owned by each company. Extensive validation and scrubbing of the raw data was performed, especially for placements experiencing CREs.

The ratings measure the ex ante riskiness of each placement and thus are a particularly important element of the data. In principle, current and at-origination insurance company internal and NAIC (regulatory) ratings are available for each asset at each year-end, with the internal risk ratings reported on a scale equivalent to S&P's senior public bond rating scale. Insurance companies regularly risk-rate each bond in their portfolios as part of their monitoring and risk-management procedures. Thus, internal ratings at each year-end should reflect relatively current and informed assessments of credit risk. Unfortunately, many companies' database systems do not maintain a long historical record of internal ratings. Thus, ratings at bond issuance are often not available for placements issued prior to 1985, and such ratings are not examined here.[6]

The 393 CREs in the data include defaults, restructurings designed to avoid or minimize possible credit-related losses, and asset sales with similar motivations. For each placement experiencing a CRE, the data include characteristics (such as the event date and CRE type) and detailed information on the originally contracted schedule of cash flows and the post-CRE schedule. The latter are actual cash flows where known and are insurance company estimates otherwise, with estimates being updated in successive rounds of the study. These data support present-value-based calculations of loss severities, as described below.

II. Methods of Measuring Default and Loss Rates

The loss rate on any pool of debt over any given horizon can be expressed as the product of a value-weighted CRE incidence rate and an average loss severity (loss-given-default rate). In public bond default studies, incidence is almost always defined to include only payment defaults and perhaps distressed exchanges. The latter are relatively infrequent, and other forms of restructuring are even less common because of the structure and diffuse ownership of publicly issued debt. Restructurings of private debt are common, however.

[6] In addition to internal ratings, virtually all private placements in insurance company portfolios are rated by the National Association of Insurance Commissioners' Securities Valuation Office (NAIC SVO). The SVO is essentially a rating agency contained within the U.S. insurance regulatory structure. Its ratings are used in determining various reserving requirements and, in recent years, risk-based capital requirements. Default and loss statistics by NAIC rating appear in the Society of Actuaries (1996). Those statistics are qualitatively similar to the internal-rating statistics presented here, lending credence to the accuracy and consistency of ratings on both scales. The comparability of internal ratings to Standard and Poor's and Moody's ratings is discussed further in Section IV.

A. Incidence Rates

Four basic methods of measuring CRE incidence rates have appeared in the literature: average one-year default rates (Hickman (1958)), average multi-year cohort rates (Moody's (1996), Standard and Poor's (1995)), mortality rates (Altman (1989)), and probabilities estimated using a hazard model (Carty (1996), McDonald and Van de Gucht (1996)).[7] Apart from the inclusion of restructurings and distress sales, this study's CRE incidence rates are most similar to average one-year default rates: A pool of assets exposed to loss (with an associated total exposed amount of principal) is formed for each year, with the incidence rate by number (amount) being the ratio of the number (principal amount) of assets experiencing CREs to the number (amount) of assets in the pool.[8] One-year rates may be weight-averaged across years of the study or examined individually for time patterns. Pools may include all placements or may differ by asset characteristics, for example by rating, supporting analysis of the relationship between bond characteristics and loss rates.

A principal goal of this study is comparison of incidence and loss rates on private placements to those on publicly issued bonds. Public bond statistics are drawn from Moody's (1996) and Standard and Poor's (1995) studies, but their methods differ in certain details from the methods used here. Most importantly, public bond incidence rates are for defaults only, and public incidence rates are computed on an issuer rather than an asset basis. For example, a firm with twenty bonds outstanding would be counted only once in the Moody's and Standard and Poor's studies. As described in more detail below, for comparability, results are presented both for all CREs and for defaults only, and on both an asset and an issuer basis.

B. Loss Severity and Economic Loss Rate

The loss severity (loss-given-default rate) for a given pool of assets is measured as the sum of dollar losses on pool assets experiencing CREs (or defaults where the analysis is limited to defaults) divided by the sum of principal outstanding on those assets at the time of the CREs. Dollar losses are the difference in net present values of original and revised contract cash flows, multiplied by the ratio of principal outstanding to the present value of original cash flows. The latter ratio is applied in order to place dollar losses on a book value basis for comparability with the denominator of the severity

[7] Average cohort rates are conceptually similar to average one-year rates, but the horizon over which default may occur may be greater than one year. Mortality rates differ in that cohorts are formed based on year of issuance, and marginal mortality rates are computed differently. Hazard models permit simultaneous evaluation of the effects of many characteristics on the likelihood of default.

[8] Where an asset was acquired or matured during a year, it was counted as half a unit of exposure and the principal amount exposed was set to half the amount outstanding at the end or beginning of the year, respectively.

measure.[9] Discount rates applied to each cash flow in making present value calculations are from the Treasury spot yield curve existing on the CRE date plus interest rate spreads appropriate to the risk rating of the assets.[10]

The overall economic loss rate for a given pool is the sum of dollar economic losses on pool assets experiencing CREs divided by total pool principal exposed to loss. This is identical to the product of the pool loss severity and the pool CRE incidence rate by amount.

As noted, a virtue of this study's data and methods is the existence of loss severity estimates for each asset experiencing a CRE, permitting a partitioning of all of CRE incidence, loss severity, and overall loss rates by asset characteristics. For publicly issued bonds, loss severities are typically estimated using dealer indicative bid (or ask) prices soon after default (and perhaps factoring in one lost coupon payment). Even indicative prices are not available for all defaulted bonds, however, so public bond loss rates must be estimated by multiplying default rates by an average loss severity estimate, typically around 60 percent.

Each method has strengths and weaknesses. This study's method is dependent on the quality of contributing companies' estimates of post-CRE cash flows, and also does not take into account all administrative costs of workouts.[11] Post-default public bond indicative bid prices should take into consideration at least some administrative costs, but sample selection biases may arise from the use of severities for only those bonds for which indicative bids were available.

III. Results for Default, Loss Severity, and Portfolio Loss Rates

A. Aggregate Private Placement Experience Over Time

Average annual percentage loss rates appear in the first column of Table I. In 1986, for example, the aggregate of participating insurance company portfolios lost approximately 32 cents per $100 invested. The simple average of these annual rates is 0.24 percent for the 1986 to 1989 period and 0.49 percent for 1990 to 1992; over all sample years, the weighted-average loss rate was 0.37 percent. The jump in loss rates in the early 1990s is sensible given the recession that occurred and the large volume of defaults in corpo-

[9] Loss severities for individual CREs are measured as the difference of net present values divided by the present value of original contract cash flows.

[10] Present values of cash flows omit the value of embedded options. For publicly issued bonds, the most important are usually call options held by the borrower. Some of the bonds in this paper's sample were callable, but beginning in the mid-1980s, newly issued private placements generally included prepayment penalties that compensated the lender not only for any change in the general level of interest rates since issuance but also as a practical matter for changes in the borrower's market interest rate spread. Thus, the value of unmodeled embedded options was probably small for most CREs.

[11] The cash flow estimates appear pretty good, however—companies contributing to the 1990–92 round of the Society of Actuaries study provided revisions to some cash flow estimates for 1986–89 CREs, and in general the revisions were not large.

Table I

Average Incidence, Severity and Loss Rates Over Time

The average portfolio loss rate is the percentage loss per dollar exposed for the total sample during the given year. Credit risk event (CRE) incidence rates are like default rates but include restructurings and distress sales in the definition of CRE. Incidence rate by number is the percentage of exposed assets experiencing CREs, whereas incidence by dollar amount is the percentage of dollar exposure experiencing CREs. Loss severity is the percentage of exposed dollars lost on assets experiencing CREs. One minus loss severity would be an average recovery rate.

Year	Average Portfolio Loss Rate (%)	CRE Incidence Rate (%)		Average Loss Severity (%)
		By Number	By Dollar Amount	
86	0.32	0.68	0.79	41
87	0.21	0.77	0.87	24
88	0.17	0.42	0.43	40
89	0.26	0.48	0.65	40
90	0.35	0.60	0.75	47
91	0.66	1.30	1.54	43
92	0.47	1.08	1.77	26
All	0.37	0.74	1.04	36

rate debt markets in general. Standard and Poor's (1995) reports an approximate doubling of public bond default rates across the same periods.

Economic loss rates rose because incidence rates rose, not because loss severities were worse during 1990 to 1992. The middle columns of Table I (CRE Incidence Rate) display incidence rates computed both as the number of assets experiencing CREs relative to the total number and as the dollar volume experiencing CREs relative to the total amount exposed. Incidence by dollar amount is higher than that by number in each year (though often not by much), indicating that assets experiencing CREs had larger than average dollar amounts outstanding. Incidence rates approximately doubled during 1990 to 1992 by both measures. In contrast, average loss severities, shown in the last column, were close to 40 percent in most years and averaged 36 percent. They show no time trend.

The apparent low variance of average loss severities from year to year does not imply that severities for individual bonds are highly predictable. In fact, they are widely distributed and thus rather unpredictable, as shown in Figure 1, which displays the distribution for all CREs. The distribution is bimodal, with one mode occurring at severities between zero and 10 percent and another at losses between 90 and 100 percent. That is, the most likely individual outcomes of a CRE involve losing nothing or very little, or else everything or almost everything.

Moody's (1996) finds that public bond default loss severities are also widely distributed, but their distribution is unimodal. The two publicly available studies of bank loan loss severities offer conflicting evidence. Carty and Lieberman (1996) find a unimodal distribution with more than half of loan loss

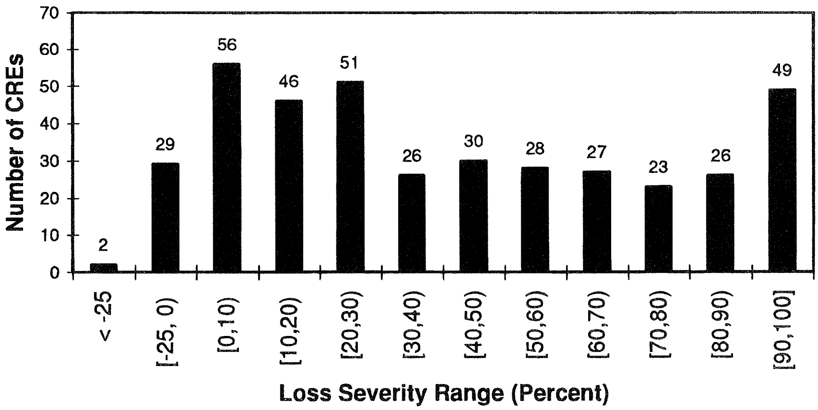

Figure 1. Distribution of Loss Severities, All CREs. This figure displays the number of credit risk events (CREs) falling in each range of loss severity. Individual CRE loss severity is measured as the difference between the present values of original and revised contract cash flows divided by the present value of original contract cash flows.

severities at 30 percent or less, whereas Asarnow and Edwards (1995) find a bimodal distribution more similar to that in Figure 1. Methods and data differ across the studies and some samples are small, so it is difficult to draw firm conclusions, but on the whole it appears that bank loan loss severities are a little smaller on average than those for private placements, which are in turn much smaller than public bond severities (which average 60 percent).

The significant fraction of negative severities displayed in Figure 1 is atypical of public bonds (31 of 393 assets experiencing CREs had recoveries greater than 100 percent). The present value of post-CRE private placement cash flows can exceed the pre-CRE present value mainly because the post-CRE coupon rate, the amount to be repaid, or the amortization schedule differ from pre-CRE values. For example, after a workout or restructuring many of the revised cash flows for an asset might occur earlier than the originally scheduled cash flows and, after discounting, the revised cash flows might therefore have a larger net present value than the original cash flows. Negative-CRE cash flow data were audited especially closely, and thus in general such CREs likely reflect genuine economic gains to the investor rather than data errors.

B. Experience by Credit Risk Rating

Experience by most recent internal rating (that is, the contributing insurance company's rating of each asset as of the start of each portfolio year) is summarized in Table II. Average incidence and portfolio loss rates were low for assets with the equivalent of investment-grade ratings (AAA through BBB) during the period 1986 to 1992 but rose steeply in the speculative grades. No clear relationship between loss severities and ratings is evident, however. Severities are between 24 and 39 percent except for the AA grade

Table II

Loss Experience by Credit Rating

The average portfolio loss rate is the percentage loss per dollar exposed for all placements with the given rating at the beginning of the experience year, averaged across all years 1986–92. Credit risk event (CRE) incidence rates are like default rates but include restructurings and distress sales in the definition of CRE. Incidence rate by number is the percentage of exposed assets experiencing CREs, whereas incidence by dollar amount is the percentage of dollar exposure experiencing CREs. Loss severity is the percentage of exposed dollars lost on assets experiencing CREs. One minus loss severity would be an average recovery rate. n.c. means no CREs occurred for assets rated AAA.

Most Recent Rating	Average Portfolio Loss Rate (%)	CRE Incidence Rate (%)		Average Loss Severity (%)
		By Number	By Dollar Amount	
AAA	n.c.	n.c.	n.c.	n.c.
AA	0.03	0.06	0.04	76
A	0.02	0.10	0.09	24
BBB	0.24	0.48	0.71	33
BB	1.50	2.76	3.89	39
B	2.16	3.74	5.66	38
<B	4.36	4.38	7.91	55
Unknown	0.42	0.83	1.29	32
All	0.37	0.74	1.04	36

(where the number of CREs is small and the average likely noisy) and for the riskiest rating (<B). The number of CREs in the <B category is large enough (39) to make the average severity estimate (55 percent) credible.

The start-of-year rating is missing for about 29 percent of the placement-experience-years, but the missing values do not appear to impart a substantial selection bias. Focusing on the last two rows of Table II, average incidence and portfolio loss rates for missing-rating observations are a little higher than the overall average, and average severities are a bit lower. Thus, it appears that the average ex ante risk of observations with missing ratings is a bit higher than that of the rest of the sample, but it is very unlikely that such observations are predominantly high-risk assets.[12]

C. Comparison with Public Bond Experience

Comparison of private placement loss experience with results in public bond default studies is complicated by differences in methods. To promote comparability of results, private placement incidence rates are reestimated

[12] Two patterns of missing ratings appear in the data for individual insurance companies: similar small fractions of missing observations in each year, perhaps indicative of assets not yet rated or of minor database retrieval problems in constructing a data submission; and large fractions missing in particular years, almost certainly symptomatic of database retrieval problems. There are only two contributing insurance companies for which ratings are always missing. In the aggregate, placements with missing ratings are distributed fairly evenly across years.

in a manner as similar as possible to the methods of Moody's (1996) and Standard and Poor's (1995) studies. Most importantly, the volume of CREs is computed for defaults only, and computations are on an issuer rather than an asset basis. For example, CREs on multiple placements issued by the same firm are counted as a single event by this alternative method, and multiple placements outstanding for the same issuer at the beginning of a year are counted as a single unit of exposure. Such adjustments affect only CRE incidence rates by number, not overall private placement loss rates.[13] The public bond default rates are weighted-average one-year rates for 1986 to 1992, and are comparable to the adjusted private placement incidence rates by number of assets. Public bond portfolio loss rates are estimated as the product of default rates from Standard and Poor's (1995) study and a constant loss severity of 60 percent, which is very close to the overall average severity reported by Moody's (1996).[14]

Incidence and loss rates by most recent internal ratings for private and agency ratings for publicly issued bonds appear in Table III. Panel A shows private placement statistics when all CREs are included, Panel B shows statistics when only defaults are included, and Panel C shows public bond default and loss rates. Private incidence rates are of course smaller when only defaults are included, and proportionately rather substantially so (by about 15 to 30 percent in general). Loss rates are only somewhat smaller, however, because the restructurings that are omitted from Panel B have lower average severities than private defaults (24 percent versus 45 percent; see Society of Actuaries (1996) for detailed statistics by CRE type).

Private placement incidence rates are higher than public bond default rates for all but the B and <B grades, and perhaps the BB grade. Comparison of Panels B and C shows that for the investment grades the incidence rate differences are in the range 0.09 to 0.29 percentage points, which is absolutely rather small but proportionately substantial. For the BB category, the

[13] In addition, a few asset-years of data for certain contributing companies were omitted from the analysis because placement issuers could not be identified. Placements were counted as exposed to loss only if they were outstanding at the beginning of a calendar year, and three CREs were dropped because they occurred during the calendar year of funding and thus would not have been counted as exposed in Standard and Poor's (1995) and Moody's (1996) studies. Where the same placement was in the portfolios of multiple contributing insurance companies and was rated differently, the less risky rating was used if the difference was a single grade, but the rating was set to "unknown" for larger differences (few such alterations were required). The effect of these adjustments on loss rates may be seen by comparing loss rate column values in Table II and Panel A of Table III. Where the effect is nonnegligible, private loss rates are increased, thus biasing against a finding that private debt performs better than publicly issued bonds.

[14] The estimated economic losses are based on Standard and Poor's rather than Moody's ratings for three reasons. First, Standard and Poor's (1995) study includes experience for bonds rated CCC whereas Moody's (1996) includes only those rated B or better. CCC experience appears to represent a significant portion of overall public bond losses. Second, Moody's sample includes a number of sovereign and structured finance securities and thus is less comparable to the private placement sample. Finally, Standard and Poor's ratings are nominally pure default ratings and thus are perhaps more suitable for combination with loss severities.

Table III

Public versus Private Loss Experience by Rating

All incidence, default, and loss rates are in percent. This table compares incidence and loss rates by rating for all credit risk events (CREs) and for those involving only defaults with default and loss rates on publicly issued bonds over the same time period, as reported by Standard and Poor's (1995) and Moody's (1996). Exposure, CREs, and incidence by number are counted on an issuer basis for comparability with public bond statistics (earlier results are on an individual asset basis). Incidence rate by number is the percentage of exposed assets experiencing CREs. The average portfolio loss rate is the percentage loss per dollar exposed for all placements with the given rating at the beginning of the experience year, weight-averaged across years. The loss rate for public bonds is estimated by multiplying the S&P default rate by a loss severity rate of 60 percent. n.c. means no CREs for AAA assets. n.r. means not reported.

	Panel A: Private Placement Statistics by Issuer, All CRE Types			Panel B: Private Placement Statistics by Issuer, Defaults Only			Panel C: Public Bond Statistics		
							Default Rate (%)		Average Loss Rate (%)
Most Recent Rating	Number of CREs	Incidence by Number (%)	Average Portfolio Loss Rate (%)	Number of CREs	Incidence by Number (%)	Average Portfolio Loss Rate (%)	S&P	Moody's	S&P basis
AAA	0	n.c.	n.c.	0	n.c.	n.c.	0.00	0.00	0.00
AA	3	0.13	0.03	2	0.09	0.02	0.00	0.04	0.00
A	10	0.21	0.01	7	0.15	0.01	0.05	0.00	0.03
BBB	47	0.70	0.26	37	0.55	0.22	0.32	0.26	0.19
BB	53	3.55	1.53	36	2.41	1.14	1.39	2.42	0.83
B	44	5.54	2.66	36	4.53	2.27	6.62	9.34	3.97
<B	23	5.45	5.53	17	4.03	3.77	23.71	n.r.	14.23
Unknown	104	1.41	0.53	81	1.10	0.44	n.r.	n.r.	n.r.
All	284	1.13	0.39	216	0.86	0.31	1.94	1.49	1.16

private placement default rate is about 1 percentage point higher than that computed from Standard and Poor's (1995) study but about the same as that from Moody's (1996) study. The private default rate is 2 to 5 percentage points lower for B, and nearly 20 percentage points less for bonds rated less than B.[15]

Even though default rates are higher on privates than publics in the investment grades, loss rates are strikingly similar, with better loss severities offsetting the effect of higher private default rates.[16] For the speculative grades, both default and loss rates are better for privates (with the caveats for the BB category already noted), implying substantially better overall performance of private debt.

The better overall performance of private placements might be ascribed to differences between companies' internal rating schemes and the major rating agencies' scales—perhaps contributing insurance companies are more conservative than the agencies, for example placing assets in the B category that the agencies would rate as less risky. The pattern of results does not support such a hypothesis, however. If contributors were systematically more likely than the agencies to place less risky assets in the riskier grades, private default and loss rates in the less risky grades should be substantially better than public bond rates, whereas in fact they are a little worse or about the same.

It is more likely that monitoring by private lenders is responsible for the better performance of private debt in the riskier grades. The literature on private debt argues that lender monitoring and the presence of associated nonprice terms such as covenants and collateral are the main distinguishing characteristics of private debt (Diamond (1984), Berlin and Mester (1992)). Such monitoring is likely to be more important and more intensive for riskier debt, and thus to add more value in terms of improved loss rates.

The results seem to imply that monitoring has little effect for low-risk, investment-grade debt. However, the incremental value of monitoring is probably understated in this paper's results because statistics are average one-year rates. Many originally investment-grade bonds that end up in default

[15] Assessment of the statistical significance of the differences is difficult. If default is viewed as a binary random variable that is distributed identically within each rating class, all differences are statistically significant in that they exceed two standard deviations (except default-only private incidence and Moody's public default rate for the BB rating). Moody's and Standard and Poor's results also differ significantly by this criterion. However, the identical-distribution assumption is unrealistic, especially for the lower ratings (a BB− likely differs significantly in default probability from a BB+).

[16] As noted previously, the rating is missing for about 29 percent of private exposure units. Such missing values have different implications for public–private comparisons than the omission of unrated and especially very low-rated bonds (the CC and C grades) from the public-bond statistics. The unknown-rating privates are included in overall loss experience estimates, so such values are representative of all private experience to the extent the portfolios of the companies contributing to this study form a representative sample. However, the omission of some low-rated bonds from public bond statistics biases overall public loss rates downward. Thus, relative private experience is likely even better than shown.

first migrate to the junk grades, and loss experience is better there, so an analysis involving horizons longer than one year would likely show private loss rates to be better in the investment grades as well as the lower grades.

IV. Diversification and the Risk of Unusually Large Portfolio Losses

Events in the bad tail of the portfolio loss distribution can cause insolvency of financial intermediaries or, similarly, losses on senior tranches of collateralized loan obligations. In principle, sufficient equity is placed in the capital structures of such entities to absorb tail-event losses up to some probability cutoff, and senior liabilities of the intermediary or CLO are then priced to reflect the remaining far-out-in-the-tail credit risk.[17] Because equity finance is more expensive than debt, portfolio managers have incentives to diversify in order to limit the equity required. This section and the next analyze the sensitivity of the bad tail to various portfolio characteristics. In the empirical work, the size of losses at high percentiles of the distribution are compared as ex ante portfolio characteristics are varied.

Standard factor-model logic provides a helpful framework for thinking about the empirical strategy. Suppose that default probability and loss severity for any given asset at a given time are functions of asset characteristics, systematic factors, and idiosyncratic factors. Diversification involves limiting the loss exposure of a portfolio to individual systematic factors, thus limiting the likelihood that a bad factor realization will be associated with very large portfolio losses. Such diversification might be achieved by investing only in loans with characteristics that sharply limit loss severities (cash collateral, for example), but historically the most common methods have involved variations in portfolio size and limits on exposure to a single borrower (that borrower's condition would be a systematic factor), or to a single industry, country, etc. Recently developed formal credit risk management systems often estimate cross-asset loss correlations in measuring diversification, using loadings on common factors to model such correlations. However, data upon which such estimates can be based are usually lacking, and thus untested parametric assumptions about loss correlations are embedded in the structure of the models.

This paper avoids assumptions about parametric forms, instead reporting nonparametric estimates of the size of losses in the bad tail, conditional on portfolio composition. For each set of portfolio parameters, 50,000 portfolios are drawn randomly from the database without replacement until the specified portfolio size is reached. Drawn assets are rejected for inclusion if they fail to satisfy the parameters for the given exercise. For example, most exercises specify a target percentage of the portfolio to fall in each risk-rating

[17] In practice, implicit or explicit government or third-party guarantees and regulatory capital standards can complicate real-world capital structure and pricing outcomes, but such considerations are ignored in this paper. Also ignored are the information asymmetries that prevent financial intermediaries and CLOs from being financed entirely with senior debt.

category. A BBB-rated asset would be rejected for inclusion if sufficient assets with that rating had already been drawn, even if the total simulated portfolio was not yet filled. Portfolio parameters have three components: a portfolio size, a set of fractions of the portfolio to be drawn from assets with each internal risk rating, and a limit on the percentage of the portfolio to be lent to a single borrower.

Empirical losses are computed for each drawn portfolio, and the frequency distribution of such losses forms an estimate of the relevant loss distribution. For example, after rank-ordering the portfolios by loss rate, the loss rate for the 49,950th portfolio would represent the estimated equity share of portfolio finance sufficient to absorb all losses with probability 0.999. Although cross-asset loss correlations are not estimated explicitly, such correlations are implicit in the experience database and thus the effects of realistic correlations are reflected in the estimates.

For any given simulated portfolio, the draw is in two stages: one of the seven experience years in the database is drawn, and individual assets with experience in that year are then drawn to fill the portfolio. Using experience from multiple years for a given simulated portfolio would tend to understate tail losses because the results of different realizations of systematic factors would be unrealistically combined.[18]

The Monte Carlo method's strength is its lack of assumptions about functional forms of distributions, about covariances among asset returns, and about the relationship between default and loss severity. Its two principal weaknesses flow from the (unrealistic) implicit assumption that the Society of Actuaries database represents the universe of possible loss experience. One such weakness is that a form of survivorship bias may be present in the results, arising because data were contributed only by insurance companies that survived the sample period. I believe such bias to be small, for two reasons: first, almost all of the large life insurance companies that failed during the past ten years did so because of bad commercial real estate investments. To my knowledge, none failed due to losses on private placements. Second, as shown below, the average loss rate on a portfolio with characteristics similar to those of large commercial banks is about 1 percent, which is similar to the long-term average loss rate on bank loans. Other results shown below imply a relationship between average and far-tail loss rates, so as long as the database is representative of average losses, no first-order survivorship bias seems likely.

A second, more important problem is that the period 1986 to 1992 naturally does not capture the full range of possible combinations of systematic loss factor realizations. For example, the United States did not experience a depression, and many specific industries did not experience severe distress. Thus, the sensitivity of loss distribution tails to the portfolio fraction loaned

[18] The same asset is prevented from appearing twice in a drawn portfolio (it may appear in the data more than once for a given year by virtue of being on the balance sheets of multiple insurance companies).

to firms in *specific* industries cannot realistically be estimated, and all tail mass estimates are understated in that the effects of very bad systematic events are not modeled.

To some extent, this problem is insurmountable by any method. Data covering a very long period would mitigate but not eliminate it. Parametric modeling approaches are similarly problematic because historical data must be used to estimate model parameters. The approach taken here is to emphasize the need for care in interpretation of results, and also to conduct some sensitivity analysis. As a practical matter, recessions would appear to be the most important systematic factor, and the Society of Actuaries' database does include a recession in which business loan losses as measured by public bond defaults were historically large. Reported below is a comparison of estimates of tails for portfolios drawn only from nonrecession years and tails based on recession years' experience.

As noted previously, a third weakness of this paper's method is that the *levels* of estimated loss rates at each loss distribution percentile may not be realistic because all of the diversification strategies employed by real-world portfolio managers may not be modeled. The levels of tail loss rates are most representative of portfolios constructed by managers that choose assets randomly, while imposing limits as modeled.[19]

V. Empirical Sensitivity of the Loss Distribution to Portfolio Characteristics

A. Main Results

Table IV reports Monte Carlo results for two base sets of portfolio parameters. Both involve a portfolio size of $1 billion and a limit on loans to a single obligor of 3 percent of the drawn portfolio, or $30 million. For reference, looking across portfolio years for individual insurance companies contributing data, the maximum observed exposure to a single borrower for single insurance companies' portfolio-years has a median of 2.1 percent, a mean of 3 percent, and a maximum of 12.2 percent (the next-highest instance is about 10 percent). Thus, a one-borrower limit of 3 percent represents a reasonable contraint from the standpoint of actual practice. This limit is imposed throughout the analysis unless otherwise noted. Observations in the database with nonmissing risk ratings total $310 billion of total exposure, so any drawn $1 billion portfolio represents a small fraction of the eligible observations.[20]

[19] Some risk management systems use covariances among debt issuer stock returns as a proxy for individual debt instrument loss covariances. However, because a substantial fraction of issuers of private debt has no publicly traded equity outstanding, as a practical matter such methods are infeasible for most private debt portfolios.

[20] About 4 percent of database observations cannot enter a drawn portfolio in this particular exercise because they represent more than $30 million of exposure and thus by themselves violate the one-borrower limit. Such observations represent about 30 percent of total dollar exposure.

Table IV

Loss Rate Distribution for Base Portfolio Characteristics

This table reports Monte Carlo estimates of portfolio loss rates at the mean and various percentiles of the credit loss rate distribution, conditional on simulated portfolio size and the proportions of assets rated below investment grade (<BBB) and investment grade. An exposure-to-one-borrower limit of 3 percent of the portfolio size was enforced in building simulated portfolios. $1b indicates portfolios were $1 billion in size. Results in each row are based on 50,000 simulated portfolios. The typical private placement portfolio as measured by this paper's data has 13 percent of assets rated <BBB. The typical large U.S. commercial bank business loan portfolio has 50 percent of assets rated <BBB.

Portfolio Characteristics		Simulated Loss Rates (%)						
			At Loss Distribution Percentiles:					
Percentage Rated < BBB	Size ($billions)	Mean	95	97.5	99	99.5	99.9	99.95
13	$1b	0.33	1.42	1.78	2.29	2.53	3.09	3.41
50	$1b	1.05	2.94	3.42	4.03	4.45	5.40	5.89

The first row of Table IV reports loss rates at high percentiles when the drawn portfolio has fractions in each rating category that are representative of the aggregate private placement portfolio: 1 percent rated less than B, 4 percent B, 8 percent BB, 35 percent BBB, and 52 percent in the A or better grades, which were lumped together in the Monte Carlo exercises. The estimated mean of the loss rate distribution is 0.33 percent, appropriately close to the overall value reported in Table I. The loss at the 99th percentile is 2.29 percent, and at the 99.9th percentile is 3.09 percent.

The second row of Table IV reports percentiles for risk-rating proportions that I believe to be typical of the commercial loan portfolios of large U.S. commercial banks: 3 percent rated less than B, 12 percent B, 35 percent BB, 30 percent BBB, and 20 percent rated A or better. The mean loss for such portfolios is about three times more than for the typical private placement portfolio, whereas bad tail losses at the 99th and 99.9th percentiles are somewhat less than twice as large. Intuitively, although loss covariances among individual assets are important in determining tail mass, the mix of individual asset risk levels is also important because the effect of covariances is magnified for riskier assets. As an extreme example, the loss correlation among assets in a portfolio of U.S. Treasury debt is near one, but the very low individual default risk of such assets makes the tail of a Treasury portfolio loss distribution very thin in spite of the high correlation.[21]

[21] To place results in context, suppose a designer of a CLO involving a pool of private placements, constructed to conform to the characteristics of the exercise reported in the first row of Table IV, wished to achieve a BBB/Baa senior public debt rating on the senior tranche. Results in Moody's (1996) and Standard and Poor's (1995) imply that to support such a rating, the junior or equity tranche(s) must be sufficient to absorb all losses for about 99.75 percent of

Table V

Loss Rate Distribution as the Percentage of Portfolio Assets Rated below Investment Grade Is Varied

This table reports Monte Carlo estimates of portfolio loss rates at the mean and various percentiles of the credit loss rate distribution, conditional on a simulated portfolio size of $1 billion and the given percentage of portfolio assets rated below investment grade (<BBB). An exposure-to-one-borrower limit of 3 percent of the portfolio size was enforced in building simulated portfolios. Results in each row are based on 50,000 simulated portfolios. Portfolio fractions in individual rating grades are proportional to the fractions in the base private placement portfolio.

Portfolio Characteristics		Simulated Loss Rates (%)						
			At Loss Distribution Percentiles:					
Percentage Rated < BBB	Size ($billions)	Mean	95	97.5	99	99.5	99.9	99.95
0	$1b	0.11	0.74	0.93	1.40	1.53	2.03	2.20
10	$1b	0.27	1.28	1.63	2.04	2.46	2.98	3.23
20	$1b	0.47	1.74	2.20	2.61	2.91	3.68	3.88
30	$1b	0.66	2.20	2.59	3.02	3.38	4.10	4.42
40	$1b	0.87	2.60	3.02	3.60	4.02	4.97	5.21
50	$1b	1.06	2.94	3.44	4.07	4.52	5.56	6.00
60	$1b	1.26	3.34	3.90	4.54	5.03	6.20	6.61
70	$1b	1.47	3.70	4.27	5.02	5.59	6.68	7.11
80	$1b	1.67	4.04	4.62	5.35	5.92	7.14	7.71
90	$1b	1.88	4.48	5.10	5.97	6.57	7.72	8.40
100	$1b	2.08	4.84	5.50	6.35	7.00	8.45	8.77

Tables V and VI offer different glimpses of the influence on loss distribution tails of a portfolio's mix of individual asset risk ratings. In Table V, the fraction of the portfolio rated below investment grade (BB or riskier) is varied between 0 and 100 percent.[22] The losses at each bad tail percentile increase monotonically with the proportion below investment grade, with an all-junk portfolio featuring tail losses between four and six times larger than a no-junk portfolio. Table VI reports results for smaller portfolios ($500 million) composed entirely of assets with a particular risk rating. Again, an increase in expected losses (a worsening of rating) is associated with an increase in the size of losses at each bad tail percentile.

Other means of diversification can substitute to at least some extent for less individually risky assets. Larger portfolios are likely to be better diversified than smaller portfolios, even where the loan-to-one-borrower limit is

outcomes. Table IV implies that such a junior tranche must, at a minimum, represent between 2.5 and 3.1 percent of the financing of the CLO (which are the loss rates at the 99.5th and 99.9th percentiles). A typical pool of bank loans likely would require a larger junior tranche, at least 4.4 to 5.4 percent of CLO financing. However, as noted, levels of tail losses should be interpreted with caution in applications involving real-world portfolios.

[22] The relative proportions of below-investment-grade assets rated BB, B, or <B are kept the same as the proportions in the reference private placement portfolio in Table IV (and similarly for investment grade proportions).

Table VI

Loss Rate Distribution When All Portfolio Assets Have Given Rating

This table reports Monte Carlo estimates of portfolio loss rates at the mean and various percentiles of the credit loss rate distribution, conditional on a simulated portfolio size of $500 million and all portfolio assets falling in the specified rating grade. An exposure-to-one-borrower limit of 3 percent of the portfolio size was enforced in building simulated portfolios. Results in each row are based on 50,000 simulated portfolios.

Portfolio Characteristics		Simulated Loss Rates (%)						
			At Loss Distribution Percentiles:					
All Assets Rated	Size ($millions)	Mean	95	97.5	99	99.5	99.9	99.95
AAA,AA,A	$500m	0.05	0.37	0.82	1.50	1.64	1.85	2.15
BBB	$500m	0.19	1.17	1.48	1.87	2.09	2.57	2.79
BB	$500m	1.68	5.78	6.72	7.70	8.10	8.72	8.83
B	$500m	2.21	6.71	7.76	8.90	9.13	9.90	9.96

the same, because larger portfolios likely contain more assets. Table VII reports results for portfolios constructed with the standard private placement proportions in each grade, but of increasing size. A large portfolio of $15 billion features losses at each percentile one-third to one-quarter as large as a small portfolio of $500 million. The same qualitative result obtains when proportions in each grade are the same as the "bank" portfolio discussed previously (not shown).

B. Bad Years versus Good Years

As noted previously, a weakness of the Monte Carlo method is that a full range of possible systematic factor realizations is not represented in the data. However, both intuition and the results in Tables I and V suggest that the state of the economy is the single most important systematic factor influencing diversified debt portfolio loss rates. Table I shows that average or expected loss rates were much higher during the recession of 1990 to 1992, and highest during 1991. Table V implies that the chance of very large portfolio losses increases with mean or expected losses, and thus the risk of very large losses is probably higher during a recession than at other times. Although the data encompass only one recession, examination of the sensitivity of tail loss amounts to its occurrence provides a sense of the likely impact of recessions that are more or less severe.

Table VIII reports results of Monte Carlo exercises in which draws are from subsets of sample years. Apart from the years involved in draws, the exercises are similar to those reported in the first and last rows of Table V—the top panel of Table VIII is for simulated portfolios that include only investment-grade assets, the bottom panel features all below-investment-grade assets. The first row of each panel reports estimates of tail loss rates

Table VII

Loss Rate Distribution as Portfolio Size Varies

This table reports Monte Carlo estimates of portfolio loss rates at the mean and various percentiles of the credit loss rate distribution as portfolio size varies. Portfolio sizes are in $billions except that $500m indicates $500 million. The fraction of each simulated portfolio rated below investment grade (<BBB) is 13 percent, the same as the base private placement portfolio. An exposure-to-one-borrower limit of 3 percent of the portfolio size was enforced in building simulated portfolios. Results in each row are based on 50,000 simulated portfolios.

Portfolio Characteristics		Simulated Portfolio Loss Rates (%)						
			At Loss Distribution Percentiles:					
Percentage Rated < BBB	Size	Mean	95	97.5	99	99.5	99.9	99.95
13	$500m	0.31	1.55	2.00	2.61	2.91	3.82	4.21
13	$1b	0.33	1.42	1.78	2.29	2.53	3.09	3.41
13	$2b	0.32	1.10	1.31	1.57	1.71	2.13	2.30
13	$3b	0.33	0.97	1.14	1.42	1.60	1.99	2.15
13	$4b	0.35	0.94	1.18	1.54	1.66	2.02	2.11
13	$5b	0.35	0.90	1.14	1.34	1.49	1.69	1.76
13	$10b	0.35	0.76	0.85	0.94	1.02	1.16	1.24
13	$15b	0.35	0.67	0.73	0.82	0.87	0.96	0.99

in good years, defined as 1986 to 1989, the second row has estimates for the bad years 1990 to 1992, and the third row is based only on the year with the worst average loss rate, 1991.

Focusing first on below-investment-grade portfolios (the bottom panel), the estimates confirm that the state of the economy is an important determinant of bad-tail losses. Loss rates at the high percentiles are a bit more than 50 percent larger in the case of the worst observed recession factor realization (1991) than during the good years. This increase is smaller than that for average losses, which more than double. Looking across the columns, in the good years about a 6 percent equity share of portfolio finance is sufficient to make the portfolio insolvency probability very small. But in the worst year, a 6 percent equity share is consistent with an insolvency probability above 0.05, which is large. In the 1991-only results, a 10 percent equity share is needed to make the insolvency probability very small.

Tail-loss risk for investment-grade portfolios behaves remarkably differently, as shown in the top panel of Table VIII. Although the loss at the 95th and 97.5th percentiles is about twice as high in the worst year as in the good years, the difference at the higher percentiles is quite modest. The implication is that losses on investment-grade assets (and portfolios) are associated more with idiosyncratic factors than with systematic factors.[23] Thus, it appears the relatively short seven-year coverage of the database probably im-

[23] However, reductions in the market values of private debt that are associated with rating downgrades, which are ignored in this paper, might well be correlated with the state of the economy.

Table VIII

Loss Rate Distribution When Monte Carlo Draws Are from Good versus Bad Years

This table compares Monte Carlo estimates of portfolio loss rates at the mean and various percentiles of the credit loss rate distribution when Monte Carlo draws are limited to the "good" years 1986–89, the "bad" years 1990–92, and the "worst" year 1991. All drawn portfolios are $1 billion in size. The top and bottom panels, each with three rows, report results when all simulated portfolio assets are investment grade and below investment grade (rated <BBB), respectively. An exposure-to-one-borrower limit of 3 percent of the portfolio size was enforced in building simulated portfolios. Results in each row are based on 50,000 simulated portfolios.

Portfolio Characteristics		Simulated Portfolio Loss Rates (%)						
			At Loss Distribution Percentiles:					
Percentage Rated < BBB	Years Used in Monte Carlo	Mean	95	97.5	99	99.5	99.9	99.95
0	Good: 86–89	0.09	0.53	0.74	1.40	1.46	1.98	2.14
0	Bad: 90–92	0.15	0.87	1.26	1.45	1.59	2.22	2.28
0	Very Bad: 91	0.16	0.91	1.40	1.54	1.67	2.28	2.36
100	Good: 86–89	1.73	4.18	4.63	5.11	5.43	5.91	6.05
100	Bad: 90–92	2.53	5.59	6.31	7.19	7.82	8.95	9.33
100	Very Bad: 91	3.76	6.68	7.30	8.04	8.55	9.72	10.19

parts little bias to Monte Carlo estimates for investment-grade portfolios. For riskier portfolios, the extent to which tail loss amounts are misstated depends on one's prior about the severity and frequency of periods of credit losses worse than the 1990 to 1992 period.

C. Other Robustness Checks

Changes in portfolio size have the same qualitative effect as reported in Table VII when Monte Carlo draws are from 1991 only, and when portfolios include only below-investment-grade assets. For example, comparing $1 and $2 billion portfolios when draws are from 1991 alone and 13 percent of assets are in the junk grades, losses at the 99.9th percentile are 3.62 and 2.45 percent, respectively (not shown in the tables). The difference represents a 32 percent reduction in tail loss rate. When draws are from all years and 100 percent of assets are in the junk grades, 99.9th percentile loss rates are 8.45 and 6.03 percent, respectively, a 29 percent reduction. In Table VII, where portfolios have 13 percent of assets below investment grade, the $1 and $2 billion 99.9th percentile loss rates are 3.09 and 2.13 percent, a 31 percent reduction. Thus, although the levels of tail loss rates differ in each case, the proportional effect of changes in portfolio size appears relatively insensitive to individual asset risk.

In passing, one implication is that, other things equal, large financial intermediaries would appear to have a comparative advantage relative to small intermediaries with respect to the (costly) equity share of riskier private debt portfolio finance. This is because the absolute magnitude of the reduction in the equity share when portfolio size increases is larger for riskier portfolios. However, other factors are involved as well in risk posture decisions, and in this paper's data there is no clear tendency of larger insurance companies to invest larger fractions of their private debt portfolios in below-investment-grade placements.

No sample selection bias appears to flow from the exclusion from the Monte Carlos of assets with a missing credit rating. As noted previously, the pattern of missing ratings across contributing insurance companies and time implies they are associated mainly with random insurance company database retrieval problems, not with any systematic tendency by the insurance companies to avoid rating either safe or risky assets. As shown in Table II, average loss rates are similar for rated and unrated assets. Comparable Monte Carlo exercises for rated and unrated assets also yield similar estimates of loss rates at the high percentiles (not shown in the tables).

Limits on exposure to a single borrower appear to be an important element of diversification, but an evaluation of their impact on bad-tail loss rates indicates this topic should be left mostly to future research. Duplicating the exercise reported in the first row of Table IV ($1 billion portfolio, 13 percent junk) but with no exposure limit increases the 99th percentile loss only slightly, from 2.29 percent to 2.42 percent, whereas at the 99.95th percentile the loss rate jumps much more, from 4.29 percent to 6.24 percent. A form of survivorship bias causes this granularity: Because all contributing insurance companies presumably impose exposure limits, there are few credit risk events in the data involving very large exposures. Large exposures might be simulated by multiplying all magnitudes in the database by some factor greater than one, but the data suggest there may be a relationship between asset size and loss rate, and thus such a simple strategy might yield misleading results. The apparent size-loss relationship may proxy for relationships between loss rates and amortization or time since asset issuance. Thus, appropriate grossing-up of exposure to evaluate one-borrower-limits awaits development of a hazard model involving such factors as well as asset size.

VI. Concluding Remarks

This paper and related studies present the first publicly available integrated empirical evidence on the default and loss performance of private corporate debt where ex ante credit risk is observable. In the investment grades, private debt performs as well or better than public debt. Private debt performs better in the riskier grades, consistent with monitoring having substantial value.

Results obtained using a new method of estimating loss rates in the bad tail of the conditional portfolio credit loss distribution imply that both the riskiness of individual portfolio assets and factors related to diversification have important implications for portfolio management, for the capital structure of financial intermediaries, and for security design. Although such a result is intuitive, methods for measuring the impact on credit losses of different portfolio management strategies have been lacking.

Many related subjects appear worthy of future research. Perhaps most importantly, a study of conditional loss distribution properties at horizons longer than one year would be useful. Given appropriate data, this paper's Monte Carlo method might be applied to publicly issued bond portfolios. Also useful would be evaluations of the quality of the various parametric and semiparametric approximations of credit loss distributions that are embedded in the formal risk management models currently used by many financial institutions.

REFERENCES

Altman, Edward I., 1989, Measuring corporate bond mortality and performance, *Journal of Finance* 44, 909–922.

Altman, Edward I., and Anthony Saunders, 1998, Credit risk measurement: Developments over the last 20 years, *Journal of Banking and Finance* 22, forthcoming.

Altman, Edward I., and Heather J. Suggitt, 1997, Default rates in the syndicated bank loan market: A mortality analysis, Working paper, New York University Salomon Center S-97-23.

Asarnow, Elliot, and David Edwards, 1995, Measuring loss on defaulted bank loans: A 24-year study, *The Journal of Commercial Lending* 77, 11–23.

Avery, Robert B., and Allen N. Berger, 1991, Risk-based capital and deposit insurance reform, *Journal of Banking and Finance* 15, 847–874.

Berger, Allen N., 1995, The relationship between capital and earnings in banking, *Journal of Money, Credit, and Banking* 27, 432–456.

Berlin, Mitchell, and Loretta J. Mester, 1992, Debt covenants and renegotiation, *Journal of Financial Intermediation* 2, 95–133.

Carey, Mark, Stephen Prowse, John Rea, and Gregory Udell, 1993, The economics of private placements: A new look, *Financial Markets, Institutions and Instruments* 2(3), 1–66.

Carlson, John H., and Frank J. Fabozzi, 1992, *The Trading and Securitization of Senior Bank Loans* (Probus Publishing: Chicago).

Carty, Lea V., 1996, An empirical investigation of default risk dynamics, Ph.D dissertation, Columbia University.

Carty, Lea V., and Dana Lieberman, 1996, *Special Report: Defaulted Bank Loan Recoveries* (Moody's Investors Service, New York), November.

Credit Suisse, 1997, *CreditRisk+: A Credit Risk Management Framework* (Credit Suisse Financial Products: London).

Diamond, Douglas W., 1984, Financial intermediation and delegated monitoring, *Review of Economic Studies* 51, 393–414.

Flannery, Mark, 1994, Debt maturity and the deadweight costs of leverage: Optimally financing banking firms, *American Economic Review* 84, 320-31.

Gupton, Greg M., Christopher C. Finger, and Mickey Bhatia, 1997, *CreditMetrics—Technical Document* (J.P. Morgan, New York).

Hickman, W. Braddock, 1958, *Corporate Bond Quality and Investor Experience* (Princeton University Press, Princeton, NJ).

James, Christopher, 1996, RAROC-based capital budgeting and performance evaluation: A case study of bank capital allocation, Working paper, University of Florida.

Jones, David S., and Kathy King, 1995, The implementation of prompt corrective action: An assessment, *Journal of Banking and Finance* 19, 491–510.

Jones, David, and John Mingo, 1998, Industry practices in credit risk modeling and internal capital allocations: Implications for a models-based regulatory capital standard, Working paper, Federal Reserve Board.

Kwan, Simon, and Willard T. Carleton, 1995, The role of private placement debt issuers in corporate finance, Working paper, Federal Reserve Bank of San Francisco.

McAllister, Patrick H., and John J. Mingo, 1996, Bank capital requirements for securitized loan pools, *Journal of Banking and Finance* 20, 1381–1405.

McDonald, Cynthia G., and Linda M. Van de Gucht, 1996, The default risk of high-yield bonds, Working paper, Louisiana State University.

Moody's Investors Service, 1996, *Special Report: Corporate Bond Defaults and Default Rates 1938–95* (New York), January.

Society of Actuaries, 1996, *1986–92 Credit Risk Loss Experience Study: Private Placement Bonds* (Schaumberg, IL).

Standard and Poor's, 1995, Corporate defaults level off in 1994, in *Creditweek* (New York), May 1.