

A False Sense of Security

To appear in
Recovery Rates and Loss Given Default
Edited by Altman, Edward, Andrea Resti, and Andrea Sironi
Risk Books, 2005

An earlier version of this study appeared in
Risk, August 2003, pp. 63-67

Jon Frye
Senior Economist, Capital Group
Federal Reserve Bank of Chicago
Jon.Frye@chi.frb.org
312-322-5035

The author thanks Doug Stalker for skilled and dedicated research assistance, anonymous referees, and colleagues for providing their time and perspectives. These include Ray Bacon, Nancy Berlad, Adrian D'Silva, Denise Duffy, Matt Foss, David Hamilton, Paul Huck, Dale Klein, Corinne Krincek, Cathy Lemieux, Perry Mehta, Jim Nelson, Tara Rice, and Til Schuermann. If errors remain, they are the author's.

The views expressed are the author's and do not necessarily represent the views of the management of the Federal Reserve Bank of Chicago or the Federal Reserve System.

When forecasting the losses stemming from credit default, most analysts make the same assumption. The assumption is that the degree of loss on an exposure, or loss given default (LGD), varies either independently or not at all. Analysts of individual deals assume independence to develop models to price loans and credit derivatives, and to relate those prices to public ratings. Analysts of portfolio credit risk assume independence when they use credit models such as RiskMetrics CreditManager, KMV Portfolio Manager, and CSFB CreditRisk+.

The assumption of LGD independence has been defended in the past on two grounds. First, independence makes the math more tractable and can sometimes allow quick results in place of lengthy simulations. Second, no convincing evidence could be assembled that systematic LGD risk is real.

Standing opposed to the assumption of LGD independence has been the intuition that LGD is sensitive to an economic downturn. The intuition has two levels. In the first level, recovery on a defaulted obligation stems from the assets of the defaulting party. Those assets are presumed sensitive to systematic risk prior to the default event; logically, the post-default assets are exposed as well. The conclusion is that recovery should be lower, and LGD greater, in a downturn.

The second level of the intuition compares the systematic sensitivity of different kinds of debt instruments. In the abstract, they are “high LGD” and “low LGD” instruments. More concretely, they might be a subordinated bond and a guaranteed senior secured bank loan. Suppose that the average LGDs on these instruments are 50% and 10%, respectively, in a period of low or average default rates. In a downturn, both LGDs should rise, but there is a difference, because LGD can rise no higher than 100%. The LGD of the first kind of debt can rise by 50 points, or by a multiple of two, but the LGD of second kind of debt can rise by 90 points, or by a multiple of ten. Among all LGDs, those that are normally low should be the most sensitive to a downturn.

This study presents evidence that supports both levels of intuition. It contrasts LGD in a period of high default to LGD in a period of low default. For almost all types of debt, LGD appears sensitive to the default rate, and the degree of sensitivity is substantial. Further, debt types that have low LGDs in good years are the most sensitive to the default rate. Not only do low LGD instruments respond more strongly than high LGD instruments, their response outstrips the proportional change in the default rate itself.

Instruments that seem best—senior secured loans, for example—suffer most in a high default period. Therefore, security on a debt instrument does not protect a lender against systematic risk. It provides him a better recovery on average, but in adverse conditions it may expose him to even greater risk; in a phrase, it may give him a false sense of security.

Data and Methods

The data are drawn from Moody's Default Risk Service Database for calendar years 1983-2001. Attention is restricted to issuers domiciled in the U.S. and designated by Moody's as non-financial in both "broad industry" and "specific industry." Default is observed at the issuer level. When a rating is needed for an issuer, it is the issuer-level rating that most recently precedes the calendar year. Default rates are adjusted for one-half the number of issuers that withdraw from public ratings during the year. The series of nineteen annual default rates are shown in Figure 1.

To compare the behavior of LGD in two sub-periods, this study defines the "high default years" somewhat arbitrarily as years where the default rate exceeds 4%. These years are 1990, 1991, 2000, and 2001. "Low default years" are then 1983-1989 and 1992-1999. (The sample period does not include 2002 and the costly telecom defaults by Worldcom, NTL, Intermedia, Nextel, and others.) For convenience the sub-periods are often referred to simply as "bad years" and "good years."

LGD is measured as 100 minus the debt instrument price observed by Moody's. Prices are observed as the bid-side average two to eight weeks after the default event, and may be viewed as contemporaneous, market-based estimates of the loss to lenders. The size of the data sample is shown in Table 1.

An alternative to the market price measure of LGD would be the discounted value of all the cash flows that occur until the obligation is finally settled. This approach is preferred by many analysts as a more reliable measure of LGD, for four principal reasons. First, the cash flows that ultimately result are known with certainty, while the market price of a newly defaulted instrument is an uncertain forecast. Second, the market for the defaulted asset may be illiquid or difficult to observe, such as for certain loans. Third, the market price of newly defaulted assets may be temporarily depressed. Forth, the holder of a defaulted loan may be a bank or other institution that does not account on a market value basis, whether the assets are defaulted or non-defaulted.

Other analysts prefer market price measures of LGD on conceptual and practical grounds, for reasons of their own. First, the cash flows that ultimately result from a defaulted asset are only half the picture. Those cash flows must be compared across time using the appropriate discount rate. It may take years to receive the cash flows stemming from a defaulted instrument. After the fact, the appropriate discount rate is not known, but must be chosen to reflect the risk that was born. The risk in a defaulted asset exceeds risk in a non-defaulted asset, other things equal. But risk differs substantially among defaults, and only an intimate knowledge of the default can guide the choice of the appropriate rate. Worse, as a default is being resolved, risk changes as time passes and expectations become more definite. Therefore, the "ultimate cash flow" method depends critically on an unknown and variable discount rate that is difficult to estimate for a particular situation. In this regard, an advantage of market price LGD is that knowledgeable participants bundle their discount rate estimates into the market price of the defaulted instrument along with their cash flow estimates.

Second, the market price is the same for all holders of the defaulted asset. By contrast, the cash flow approach depends to a degree on the choices made by the individual debt holder. A holder may choose to simply exit the credit (and to accept the market price for the defaulted instrument). If the holder continues its investment in the defaulted asset, it may pursue its interests by any of several strategies and intensities of pursuit. Every choice affects the net cash flows ultimately received.

Third, ultimate cash flows are collected only by institutions that survive to collect them. If there are periods during which the market underprices defaulted assets (and this assertion is far from proved), institutions must survive those periods. Specifically, if a market value assessment were to show an institution to be insolvent, it is likely that investors would refuse to fund it or that supervisors would refuse to let it continue to operate. This would be the case whether or not the institution chooses to perform the market valuation itself.

Over and above these considerations, market price indicators of LGD are consistently observed and continuously available throughout the sample period. Two caveats are in order regarding rated loans. First, a smaller proportion of loans are rated, as compared to bonds. Second, defaulted loans are less likely to have an observed price than are defaulted bonds. Either could introduce a bias, and in neither case is the direction or magnitude known with certainty.

An important choice is to study LGD at a granular level, by noting the “debt type” designated by Moody’s. The debt type designation generally repeats language from the bond or loan documents. Since that language is rather specific, there are many debt types, and the issues sharing the same debt type designation are rather similar to each other. Table 2 provides detail on the distribution of LGDs among debt types. If an issuer defaults on more than one issue having the same debt type, the market prices are averaged at the outset to avoid double counting. This reduces the 1184 defaulted instrument prices appearing in Table 1 to the 960 LGDs appearing in Table 2.

Empirical analysis

Having in place the data, this section compares for each debt type the average LGD in high default years to average LGD in low default years.¹ Forty-nine of the 121 debt types experience default in both sub-periods. These debt types contain 859 of the 960 LGDs, and they are the subjects of the empirical analysis.

For each of the forty-nine debt types, Figure 2 shows the average LGD in good years and the average LGD in bad years. The area of each bubble is proportional to the total number of LGDs in the debt type. Thus, the smallest bubbles represent debt types having two LGDs (one in the good years and one in the bad years), and the largest bubble represents a debt type having 111 LGDs.

¹ All averages are simple, or “default weighted”, averages, rather than dollar weighted.

Bubbles representing debt types in unsecured or subordinated seniorities appear mostly in the top right part of the diagram, indicating that they tend to have the greatest LGDs in both the good years and the bad years. Senior secured loans tend to appear in the bottom left. The figure contains three auxiliary lines. One is a constant 30% of par in bad years. The second is a 45-degree line that represents equality between bad year LGD and good year LGD. The third is a regression line to be described.

The diagram has two striking features. First, only two debt types have bad year average LGDs below 30%, and those two debt types contain very few defaults. Managers of portfolios of any of the debt types, including managers of the senior secured loans, would have been well served to prepare for LGDs of at least 30%, irrespective of performance in the good years.

A second feature is that most of the bubbles appear above the 45° line. This means that for most debt types, LGD is greater during the high default years than during the low default years. This feature contradicts what would be expected, if there were no systematic effect. Under the null hypothesis of no systematic variation of LGD, a bubble would appear above the 45° line with probability at most 0.5. Of the forty-nine bubbles, assuming independence, thirty-nine or more would appear above the line with probability at most 0.00002. In fact, thirty-nine bubbles do appear above the line. This decisively rejects the null hypothesis that LGD is independent of high default years.

This result supports the first level of the intuition. It is consistent with the expanding literature regarding systematic LGD variation surveyed by Allen and Saunders 2002, which includes Frye 2000b, Altman, Brady, Resti and Sironi 2001, and Hu and Perraudin 2002.

Acharya, Bharath, and Srinivasan 2004 find that if an industry is in a distressed state, defaulting firms within the industry provide less than normal recovery. This introduces the potential that the previous test result could have come about if a small number of industries contribute a large number of LGDs that are highly sensitive to the default rate.

A series of non-parametric tests guards against this possibility. The results are shown in Table 3. The first row shows the test described above: the 859 LGDs are members of 49 debt types, of which 39 have greater average LGD in bad years. The second row eliminates the industry that contributes the greatest number of LGDs, which is the oil industry. When oil industry defaults are removed, 816 LGDs remain. These are members of only 48 debt types, of which 38 appear above the 45° line. The associated test significance equals 0.00003. Each subsequent row of Table 3 removes one more industry, ordered by the number of LGDs contributed. The process continues until nineteen of the original 105 industries are removed. At that point, less than half the original LGDs remain. In none of the twenty tests is significance greater than one basis point. Thus, it appears that LGD sensitivity is pervasive among industries as well as among debt types.

Figure 2 also supports the second level of intuition. The overall impression is that across the board, bad year LGD exceeds good year LGD by approximately the same amount, 15% or so. Expressed as a proportion, of course, low LGDs tend to be more severely affected.

This overall impression can be checked by a simple model in which the 859 LGDs are regressed on fifty dummy variables, one for each debt type and one to indicate the bad years.

$$LGD_{ij} = aBAD + \sum_{j=1}^{49} b_j DT_j (1 + cBAD) + e_{ij}$$

In this regression, the least squares estimate of a equals 0.171 with an associated t statistic of 10.7. This is strong evidence that LGD rises in high default years, and it agrees with the non-parametric test results. But the estimate of c (-0.014) lacks statistical significance. Thus, the regression line in Figure 2 has a slope of very slightly less than 1.0 and represents a shift not significantly different from a uniform shift affecting all debt types equally. This test does not challenge the overall impression that the LGDs of different debt types have similar "add-on" responses to the bad years.

Practical Significance

One way to assess the practical significance of LGD variation is to compare LGD variation to the variation of the default rate. Figure 3 shows grade-specific default rates in good years and bad years.² The logarithmic axes emphasize percentage differences. The solid line represents a default rate equal to twice the good year rate.

Ten Moody's rating grades have defaults in both good years and bad years. (Rating grades Caa and B3 are nearly indistinguishable in Figure 4.) Nine of the ten grades appear above the 45-degree line, indicating greater default rates in bad years than in good years. The overall pattern is that the data clusters around the line that reflects a doubling of the default rates. As a first approximation, the default rate in bad years is about twice the rate in good years. The effect is no more powerful on low default grades than on other grades.

On the LGD side, Figure 4 shows the LGD data from Figure 2 on logarithmic axes. A solid line again represents twice the good year rate. Instead of clustering along that line, there is a difference from left to right. At the left, all debt types appear above the "2x" line, and at the right they appear below it. Thus, "approximate doubling" does not describe the effect of bad years on LGD rates. Low LGD rates increase by more than doubling, and high LGD rates tend to increase, but by less than doubling. Another difference between the figures is that Figure 4 appears to have greater sampling noise

² The default rate for a rating equals the number of defaults in the sub-period divided by the number of firm-rating-years.

than Figure 3. That is principally because there are more debt types than rating grades, and fewer defaults than issuers.

A surprising fact emerges: bad years have a stronger effect on low LGDs than on the default rate itself. This is especially surprising because “bad years” are *defined* as high default years. Default rates respond to the bad years by definition. Nonetheless, the low LGD rates respond by a greater proportional amount.

When the LGD data is aggregated to less granular level, it is possible to average away most of the effect. Considering the eleven senior secured loan debt types shown in Figure 2, five of them are well above the 45° line, two of them are well below the line, and four are near the 45° line. Table 4 shows the result when this diversity is aggregated. Then, LGD rises only from 32 to 35 as a response to bad years. At this level of aggregation, the overall impression would be that senior secured loans are largely immune to systematic risk. This impression would have been a poor guide to managing risk in the bad years, as we have seen.

A granular analysis reveals the strength and pervasiveness of systematic LGD risk (the first level of intuition). Only a granular analysis can distinguish the effects of bad years on different types of debt (the second level of intuition). Less granular approaches ignore differences between debt types and the potential for defaults to shift between them. Mindful of this potential, risk managers generally prefer data at the most granular level. These granular distinctions appear to be especially important in the analysis and forecasting of LGD.

Implications for risk management

The strong and pervasive rise of LGD in bad years worsens the loss in portfolios of credit-risky assets. LGD sensitivity therefore should affect the pricing of credit risky assets, including loans, bonds, and credit derivatives. These implications are beyond the scope of this study. Instead, we look at the tools risk managers use to discern risk: stress testing, credit risk modeling, and economic capital functions.

Managers stress test to envision performance under adverse circumstances. All too often, the circumstances include only adverse default rates and not simultaneously adverse LGD rates. Stress testing with adverse LGD is probably most important for banks, which normally expect LGDs in the lower range. Among the various types of bank loans, stress testing with adverse LGD is apt to be especially important for senior secured or asset-based lending.

A credit risk model can be viewed as a collection of credit scenarios united by a probability distribution. The data presented here suggest that when the economy enters an unfavorable state, LGD rates tend to rise. As a first approximation, the rise appears to be uniform from debt type to debt type. A successful credit risk model is apt to permit such a pattern to occur.

An economic capital allocation function acts as a summary of a credit risk model. The function values might appear in the “grid” used by some institutions to assign capital to transactions. Grid dimensions usually include the borrower’s probability of default (PD), the lending facility’s expected LGD, and possibly other variables such as maturity, domicile, and product type.

A capital function associated to a first generation capital model uses average historical or expected LGD as a multiplier. In such a function, debt types with low historical LGDs enjoy proportionally low levels of capital. The models assume LGD independence, but as shown, LGD is not independent of the default rate. As a consequence, many capital functions fail to account for systematic LGD risk, they assign relatively too little capital to low LGD deals, and they provide too much incentive for risk takers to find ways to reduce LGD.

To date, models are not readily available to estimate the stress LGD appropriate for capital. As an interim work-around, rather than multiplying by historically observed or expected LGD, practitioners could multiply by a function of average historical or expected LGD. This study suggests a simple add-on approach might be effective. Another possibility would be to multiply by a fractional power of expected LGD, rather than by LGD itself. The amount of the add-on or the value of the fractional exponent would be chosen to produce a level of LGD appropriate for an economic downturn; the downturn would probably be more severe than the ones contributing data to this study.

Conclusion

Not only do loss given default rates rise at the same time as the default rate, the risk is especially great for low LGD debt types. The “at-risk” group contains the forms of debt that are usually the most highly regarded, including senior secured bank loans.

The sensitivity of LGD works against managers of risky assets, because it increases loss in high default periods. This increase in loss, especially for low LGD debt types, has implications for stress testing, for capital models, for capital functions, and for the pricing of credit risky assets.

To date, most of the work on credit risk has focused solely on the systematic variation of the default rate. The simultaneous variation of LGD has been largely ignored. Much work remains to bring this blind spot into focus.

Jon Frye is Senior Economist in the Capital Group at the Federal Reserve Bank of Chicago. The author would like to thank Doug Stalker for skilled and dedicated research assistance, and to thank referees and colleagues for providing their time and perspectives, including Ray Bacon, Nancy Berlad, Adrian D'Silva, Denise Duffy, Matt Foss, David

Hamilton, Paul Huck, Dale Klein, Corinne Krincek, Cathy Lemieux, Perry Mehta, Jim Nelson, Tara Rice, and Til Schuermann. If errors remain, they are the author's.

The views expressed are the author's and do not necessarily represent the views of the management of the Federal Reserve Bank of Chicago or the views of the Federal Reserve System.

References

Acharya, Viral, Sreedhar Bharath, and Anand Srinivasan, 2004, Understanding the Recovery Rates on Defaulted Securities, Working paper,

<http://libweb.london.edu/facresnew/docs/recoveryrates.pdf>

Allen, Linda and Anthony Saunders, 2002, "A Survey of Cyclical Effects in Credit Risk Measurement Models", NYU Stern School of Business, Finance Working Paper No. FIN-02-018,

<http://www.stern.nyu.edu/fin/workpapers/papers2002/pdf/wpa02018.pdf>

Altman, Edward I., Brooks Brady, Andrea Resti and Andrea Sironi, 2003, "The Link between Default and Recovery Rates: Theory, Empirical Evidence and Implications", NYU Salomon Center Working Paper Series # S-03-4, forthcoming, Journal of Business, Winter 2005,

http://pages.stern.nyu.edu/~ealtman/Link_between_Default_and_Recovery_Rates.pdf

Frye, Jon, 2000a, "Collateral Damage", *Risk*, April, pages 91-94,

http://www.chicagofed.org/banking_information/credit_portfolio_research.cfm

Frye, Jon, 2000b, "Depressing Recoveries", *Risk*, November, pages 108-111,

http://www.chicagofed.org/banking_information/credit_portfolio_research.cfm

Gupton, Greg and Roger Stein, 2002, "LossCalcTM: Moody's Model for Predicting Loss Given Default (LGD)", Moody's Investors Service,

http://www.moodyskmv.com/research/whitepaper/losscalc_methodology.pdf

Hu, Yen-Ting and William Perraudin, 2002, "The Dependence of Recovery Rates and Defaults", Working Paper,

<http://wp.econ.bbk.ac.uk/respap/documents/research/recov15.pdf>

Moody's Investors Service, 2002, Default Risk Service Database

Table 1. Counts of ratings, defaults, and LGDs		
	Low default years	High default years
Issuer-rating years	22,129	7,366
Issuer defaults	381	369
Bonds with LGD	535	544
Loans with LGD	32	73

Table 2. Distribution of LGDs by seniority and debt type	
Grouping by seniority or debt type	Count
Senior secured seniority	
Guaranteed senior secured revolving credit facility	32
Senior secured notes	12
Guaranteed senior secured term loan B	12
44 additional debt types	117
Total senior secured LGDs	173
Senior unsecured seniority	
Senior notes	87
Guaranteed senior notes	59
Notes	21
34 additional debt types	102
Total senior unsecured LGDs	269
Senior subordinated seniority	
Senior subordinated notes	66
Guaranteed senior subordinated notes	66
Senior subordinated debentures	10
13 additional debt types	19
Total senior subordinated LGDs	161
Subordinated and junior subordinated seniorities	
Convertible subordinated debentures	82
Subordinated debentures	79
Senior subordinated debentures	46
32 additional debt types	150
Total Subordinated and junior subordinated LGDs	357
All seniorities	
Total number of LGDs	960

Table 3. Non-parametric test significance, excluding successive major industries					
Industry Excluded	Counts				Test
	LGDs	Bubbles	Above 45°	Below 45°	Significance
None	859	49	39	10	0.00002
Oil	816	48	38	10	0.00003
Telecommunications	783	47	38	9	0.00001
Automotive Parts	755	46	37	9	0.00002
Airlines	727	44	36	8	0.00001
Healthcare Services/Equipment	700	41	36	5	0.00000
Motion Pictures	673	41	36	5	0.00000
Textiles	649	41	36	5	0.00000
Steel	627	41	36	5	0.00000
Computers/Peripherals	605	41	36	5	0.00000
Apparel	584	41	36	5	0.00000
Retail-Grocery Chain	563	40	36	4	0.00000
Machinery	542	39	34	5	0.00000
Restaurants/Fast Food	522	39	33	6	0.00001
Food/Soft Drinks	504	37	32	5	0.00000
Home Building	487	37	32	5	0.00000
Casinos	470	34	29	5	0.00002
Broadcasting	454	32	28	4	0.00001
Retail-Discount/Variety	438	32	27	5	0.00006
Entertainment	422	31	26	5	0.00010

Table 4. Average LGDs at non-granular levels		
	Average LGD	
	Good Years	Bad Years
All 859 LGDs in Figure 2	54.5%	67.8%
88 Senior secured loan LGDs	31.8%	34.8%

Figure 1. Default rate, Moody's U.S. non-financial issuers




Figure 2. LGD rates in two sub-periods, by debt type

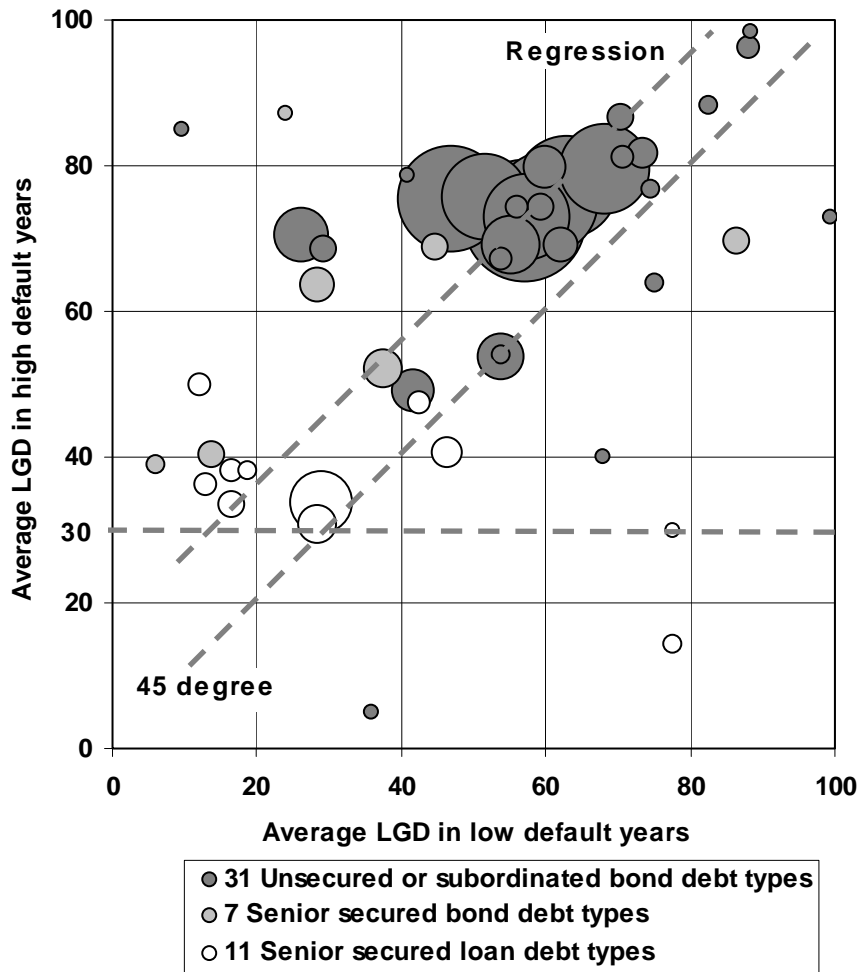


Figure 3. Default rates in two sub-periods, by rating

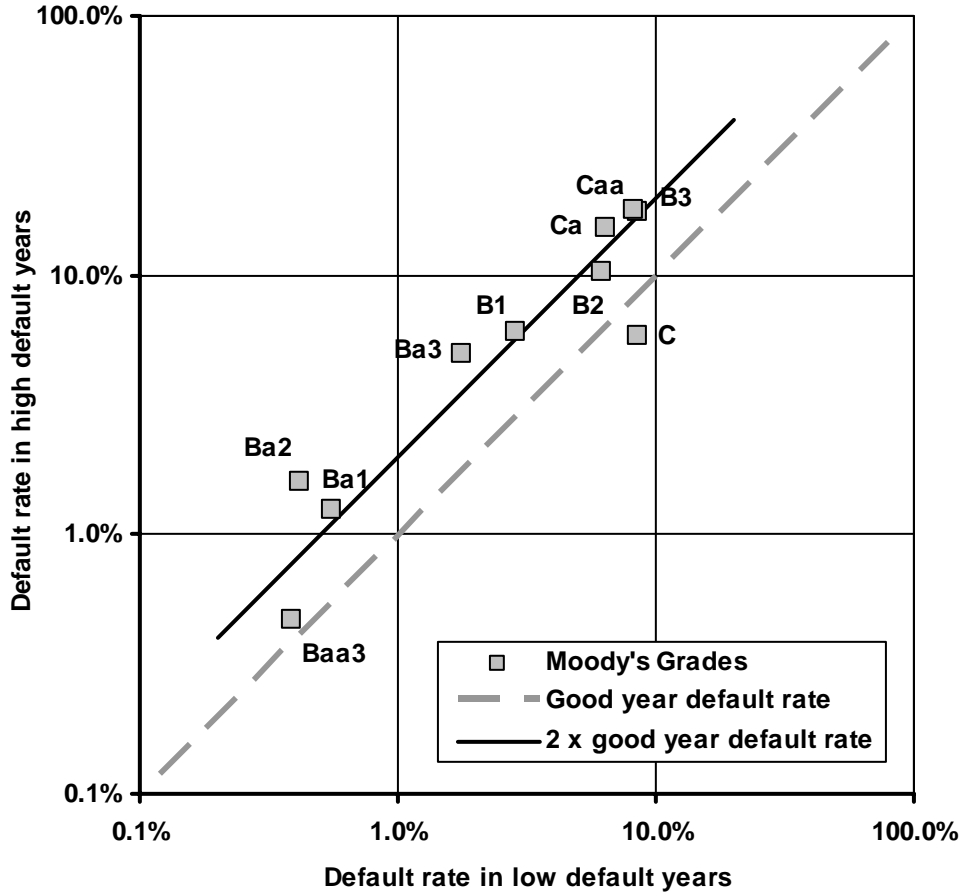


Figure 4. LGD rates in two sub-periods, by debt type

