

**The Past as Prologue:
Lessons Learned from the Crisis in
Mortgage-Backed Structured Securities**

August 2009

Erik Heitfield*
Federal Reserve Board
Erik.Heitfield@frb.gov

* The views expressed here are my own and do not necessarily reflect the opinions of the Federal Reserve Board of Governors or its staff.

1. Introduction

Unexpected deterioration in the credit quality of structured mortgage-backed securities (MBS) and collateralized debt obligations (CDO) backed by mortgage securities played a central role in recent financial market dislocations. As investors became increasingly concerned about the credit risk embedded in MBS and structured finance CDOs, demand for these and other credit products plummeted, creating severe liquidity and credit challenges for many financial institutions. This paper uses a simple normal copula model of portfolio credit risk to argue that credit ratings for structured products did not provide investors with clear signals about the risks embedded in these investments because (i) one-dimensional credit ratings are not typically designed to differentiate among exposures' sensitivity to systematic risk, and (ii) credit ratings do not adequately reflect the uncertainty associated with forecasting tail loss events.

Rating the liabilities of a structured finance deal requires detailed information about the distribution of future losses on the collateral backing the deal. The process of pooling assets and tranching liabilities has the effect of filtering out idiosyncratic risk but concentrating those risks that cannot be diversified away. As a result, holders of structured securities are particularly vulnerable to systematic risk, which leads to correlation in the performance of assets, and model risk, which can cause investors to miss-evaluate the credit performance of many or all assets. Precisely because systematic risk and model risk are not diversifiable in a cross-section of credit exposures, they cannot be measured accurately when available historical data are limited. Credit rating processes designed to rank credit exposures in terms of relative risk at a point in time, or which do not fully account for the effects of systematic and model risk, are not well suited to the purpose of evaluating the collateral backing structured finance products.¹

2. The crisis in mortgage securitization

While the pooling and resale of mortgage assets has been a fixture of U.S. residential housing markets since the 1970s, issuance of “non-agency” residential mortgage backed securities (RMBS) expanded dramatically in the first half of this decade. “Agency” RMBS are pass-through instruments guaranteed by a government sponsored enterprise (typically Freddie Mac or Fannie Mae) and backed by pools of smaller, prime loans that conform to GSE underwriting guidelines. In contrast, non-agency RMBS are structured securities backed by pools of mortgages that are either too large (jumbo loans), or too risky (Alt-A and subprime loans) to receive a GSE guarantee. As shown in Figure 1 annual issuance of non-agency RMBS grew four-fold from just shy of \$300 billion in 2001 to \$1.2 trillion in 2006.

The rapid growth in RMBS issuance, along with the development of standardized documentation for credit default swaps that reference RMBS and other asset backed securities, bolstered the issuance ABS CDOs, which reached a peak of \$308 billion globally in 2006 (see Figure 2). ABS CDOs are structured securities backed by pools of asset-backed securities (ABS) including residential and commercial mortgage-backed

¹ For a more detailed discussion of the issues raised in this paper, see Heitfield (2009).

securities and tranches of other CDOs. A typical ABS CDO may hold cash or synthetic investments in 100 or more asset-backed securities. ABS CDOs are described as “high grade” or “mezzanine” depending on the rating of the collateral held by the CDO. High-grade ABS CDOs generally hold securities rated A- and higher, while mezzanine ABS CDOs are backed by lower rated securities. RMBS collateral held by high-grade CDOs may have higher ratings either because it is backed by higher quality mortgages (e.g., Alt-A rather than subprime), or because it has better credit enhancement (e.g., higher seniority in the RMBS deal structure), or both. As shown in Figure 3, recent vintages of ABS CDOs in general, and mezzanine ABS CDOs in particular, have been heavily invested in RMBS backed by subprime mortgages.²

In 2007, falling house prices and rising default rates led to a dramatic reappraisal of the value of even the most highly rated structured securities backed by residential mortgages. Figure 4 shows the market value (as a percentage of par) for various vintages of the AAA ABX.HE index, which covers the most important AAA-rated RMBS tranches issued over a given half-year interval. By the middle of 2008 AAA RMBS tranches issued in 2007 were trading at half their face value.³

As Figure 5 shows, all classes of mortgage-backed structured securities suffered ratings downgrades. Downgrades of AAA-rated securities are particularly significant because these highly-rated notes comprise the largest share of total RMBS and CDO issuance by value and because investors in these securities expect them to be almost risk free. Standard and Poor’s (2009) has downgraded \$660 billion of the \$2.5 trillion (27 percent) worth of 2005-07 vintage U.S. RMBS tranches it originally rated AAA. Because of their implicit leverage and their concentrated exposure to subprime and other non-agency RMBS, recent vintages of ABS CDOs have experienced even more dramatic rating downgrades. Grouping U.S. RMBS, ABS CDOs and the liabilities of SIV-lites⁴ together, S&P reports that it has downgraded \$1 trillion of \$2.9 trillion (35 percent) in AAA-rated structured securities issued from 2005 to 2007 that were directly or indirectly backed by U.S. mortgages. The other major credit rating agencies have taken similar actions.

It is difficult to overstate the consequences of the collapse of the market for mortgage-backed structured credit products. It forced tiny pension funds and global investment banks alike to take huge losses, precipitated the failure of some of the world’s largest and most venerable financial institutions, and would surely be included on any short list of the most important factors contributing to the global financial crisis. Policymakers and regulators have been highly critical of credit rating agencies. According to the President’s Working Group on Financial Markets (2008), a committee of financial regulators chaired by the Secretary of the U.S. Treasury, “[c]redit rating agencies contributed significantly to the recent market turmoil by underestimating the credit risk of subprime residential-mortgage-backed securities and other structured credit products,

² The data in Figure 3 are reported by Standard and Poor’s (2007a). While these data only reflect deals rated by S&P, similar information published by other rating agencies confirms that recent ABS CDOs are heavily concentrated in subprime RMBS. See also Moody’s (2007) and Fitch (2007).

³ Of course, not all of the change in RMBS market values should be attributed to increases in expected losses. Market-wide liquidity problems and rising risk premiums also contributed adverse valuation changes.

⁴ Similar to CDOs, structured investment vehicles issued structured short term liabilities to fund longer duration investments in fixed income securities including structured securities. While SIVs were managed as ongoing concerns, SIV-lites were set up as one-off, self-liquidating investment vehicles.

notably asset-backed-security CDOs.” Although much attention has focused on the inadequacies of public ratings assigned to the liabilities of structured finance deals, problems were not limited to institutions that relied on these ratings directly. A study commissioned by the Committee on the Global Financial System (2008) reports that, “[w]hile many relatively sophisticated investors...are unlikely to have relied on credit rating agencies to assess the risk of CDO *tranches*, they appear, nonetheless, to have relied on them to assess the risk of the underlying RMBS *collateral*.”

3. Securitization concentrates exposure to systematic risk

Why did so many structured credit products perform so poorly relative to expectations? To a great extent, the potential for significant, contemporaneous drops in the market value and credit quality of large numbers of highly rated structured securities is a foreseeable consequence of the pooling and tranching of credit risk inherent in securitization. The securitization process tends to filter out idiosyncratic credit risk while concentrating systematic risk. Highly rated structured securities are designed to perform well under a wide range of economic circumstances; however, unlike more traditional fixed income securities such as corporate bonds, these securities are also very sensitive to systematic risk. In the parlance of finance, structured credit products are “high beta” instruments.

Credit ratings are typically designed to reflect a credit obligation’s probability of default or expected loss, not its exposure to systematic risk. Thus, a bond issued by a strong company in a not particularly cyclical industry might receive the same AAA credit rating as the senior tranche of an RMBS deal backed by subprime mortgages. Both ratings may be “accurate” in the sense that both securities might share the same very low *unconditional* default probability. Nonetheless, under conditions of severe systematic stress, the structured credit product can be expected to perform much more poorly than the corporate debt obligation.⁵

3.1 A simple normal copula model

To illustrate how securitization leverages exposure to systematic risk, I will use a very simple version of a normal copula model. Normal copula models are ubiquitous in risk management applications. Industry standard portfolio evaluation tools developed by The RiskMetrics Group (Gupton, Finger, and Bhatia 1997) and Moody’s/KMV (Kealhofer 1998) can be interpreted as variants of normal copula models (Li 2000). Normal copula models are used to compute risk-based regulatory capital requirements under the Basel II bank capital accord (Basel Committee on Banking Supervision 2004). Moody’s, Standard and Poor’s, and Fitch, the three largest bond rating agencies, rely on normal copula models to develop credit ratings for CDOs and normal copula models are commonly used to price CDO notes (Andersen and Sidenius 2005). Because the model presented here is intended for pedagogical purposes, I will make simplifying assumptions wherever helpful, though it should be noted that the models used in more applied settings are considerably richer than the one presented in this paper.

⁵ The converse is also true. Under benign macroeconomic conditions the corporate debt obligation would be more likely to default than the structured security.

Under the simplest normal copula model of loss dependence, the default status of credit obligation i at a one-year-ahead assessment horizon depends on the realization of a continuous index of the obligor's credit quality denoted Y_i . For corporate loan obligations this variable is commonly interpreted as a measure of the firm's return on assets over the assessment horizon. If the realized value of Y_i lies below a critical default threshold then the obligation defaults. For simplicity I assume that the loss-given-default per dollar exposure (LGD_i) is nonrandom. The unconditional probability of default of obligation i (PD_i) is simply the probability that the realized value of Y_i lies below the default threshold. So if Y_i is a standard normal random variable, then the loss per dollar exposure associated with the obligation is

$$L_i = 1\{Y_i \leq \Phi^{-1}(PD_i)\} \cdot LGD_i$$

where $1[A]$ is an indicator function that is equal to one if A is true and zero if A is false and $\Phi^{-1}(p)$ is the inverse of the standard normal cumulative distribution function.

In normal copula models of this sort, dependence in losses across obligors is modeled by assuming that Y_i depends on a number of underlying risk factors, some of which are shared in common by groups of obligors and some of which are unique to individual obligors. The presence of common factors that influence groups of obligors leads to correlations in losses across portfolio exposures. Under the simplest possible specification, Y_i is the weighted sum of two independent standard normal factors: a systematic risk factor X that affects all obligations and an idiosyncratic factor U_i that is unique to obligation i :

$$Y_i = X\sqrt{\rho} + U_i\sqrt{1-\rho}$$

The factor loading, ρ , is a parameter that lies between zero and one and describes the sensitivity of Y_i to the systematic factor X .

By construction, the unconditional default probability for obligation i is PD_i . One way of thinking about this unconditional default probability is that it is our assessment of the likelihood that obligation i will default, assuming we know nothing about the realized value of the systematic factor X . If X were observed the credit assessment could be refined. The conditional default probability of obligation i given X is

$$CPD_i(X) = \Phi\left(\frac{\Phi^{-1}(PD_i) - X\sqrt{\rho}}{\sqrt{1-\rho}}\right).$$

If we think about a large number of obligations that share the same PD and factor loading, $CPD(X)$ should closely approximate the realized default rate over a one-year horizon. This default rate will change from year to year. While the long-run average annualized default rate for a pool of loans should be equal to their common PD, the realized default rate in any given year could be substantially higher or lower depending on the realized value of X . The higher is ρ , the greater is the volatility of annual default rates relative to the long-run average.

Given this simple framework, it is straightforward to compute the probability that a credit obligation will default under adverse systematic stress conditions. Since a lower value of X implies a higher conditional default probability, an intuitive adverse systematic stress

scenario is one in which the systematic risk factor takes on a value less than some low percentile $\Phi^{-1}(1-q)$. We call this the q -th percentile or worse stress scenario. It is easy to show that under this stress scenario a credit exposure with an unconditional default probability of PD will default with probability

$$PD^q = \frac{\Phi_2(\Phi^{-1}(PD), \Phi^{-1}(1-q); \sqrt{\rho})}{1-q}$$

where $\Phi_2(v_1, v_2; r)$ is bivariate CDF for a standard normal random double with correlation parameter r .

3.2 *Pooling assets does not mitigate systematic risk*

Structured credit products depend on the performance of large pools of credit obligations. Commercial mortgage backed securities may reference several dozen loans while traditional CDOs may reference hundreds of bonds. Securitizations of retail loans such as residential mortgages, credit cards, or auto loans commonly reference several thousand individual credit obligations. Resecuritizations, such as ABS CDOs, might directly reference a hundred structured securities, but since each of these securities is itself backed by hundreds or thousands of loans, resecuritizations can indirectly reference many thousand credit obligations.

Basic probability theory predicts that a large pool of loans will perform differently under stress conditions from an individual loan, even if all loans in the pool have similar unconditional default probabilities and expected losses. Pooling reduces uncertainty about future losses, but pools of loans are actually more sensitive to systematic risk. To see this, consider a pool of N loans that are homogeneous in the sense that each has the same unconditional probability of default, loss given default and factor loading. Let $\bar{L} = \frac{1}{N} \sum_i L_i$ denote the loss rate for the pool of loans. The expected loss for one loan is

$$E[L_i] = PD \cdot LGD$$

and, likewise, the expected loss rate for the loan pool is

$$E[\bar{L}] = \frac{1}{N} \sum_i E[L_i] = PD \cdot LGD.$$

Thus, if we were rating investments to an expected loss standard, we would assign a single loan the same credit rating as a pool of N such loans even though the distribution of losses for the two investments are obviously quite different.

Suppose we ask, what is the likelihood that losses on an investment will exceed some critical threshold? For a single loan, this probability is simply its PD. For a pool of loans, realized losses are drawn from a normal-binomial mixture distribution and loss exceedance probabilities must be computed numerically. The top panel of Table 1 reports loss exceedance probabilities for hypothetical loan pools of different sizes assuming each loan has a PD of ten percent, an LGD of 45 percent, and a factor loading of 0.2. Observe that the larger the pool, the lower are its tail loss exceedance probabilities. In this sense, diversification reduces uncertainty.

But now suppose we ask, what is the likelihood that losses on an investment will exceed some critical threshold *given severe systematic stress conditions*? The lower panel of Table I documents loss exceedance probabilities for pools of loans, conditional on a 98-th percentile or worse adverse draw of the systematic risk factor X . Obviously, all loss exceedance probabilities increase under stress conditions, but what is most striking is that exceedance probabilities for diversified investments are not particularly smaller, and are often much larger, than those for undiversified investments. Pooling loans can reduce the overall likelihood of tail loss events, but diversification offers little protection against tail loss events that are driven by adverse *systematic* shocks.

3.3 *Tranching liabilities leverages exposure to systematic risk*

Senior tranches of structured finance deals only default when the losses on deal collateral exceed the value of all subordinated tranches.⁶ As a result, senior tranches are highly sensitive to the tail of the distribution of collateral losses. Consider a hypothetical structured finance deal in which N homogeneous loans are funded by two debt tranches: a junior tranche which funds proportion C of deal assets and a senior tranche that funds the remaining $1 - C$ of assets. The probability that the senior tranche will incur a credit loss is simply the unconditional probability that the pool loss rate \bar{L} exceeds the attachment point C . As we have just shown, diversification can lower this unconditional exceedance probability without substantially reducing the likelihood that the senior tranche will default under systematic stress conditions.

Table 2 compares conditional default probabilities and expected losses under stress conditions for four hypothetical senior structured finance securities and a whole loan. Collateral loans backing the structured securities are assumed to have the same PDs, LGDs, and factor loading as those described in Table 1. Structured securities that reference better diversified collateral pools are given less credit enhancement so that all four securities have an unconditional default probability of about 90 basis points.⁷ The whole loan is assumed to have a 90bp default probability and a factor loading of 0.2. All the credit exposures in Table 2 would receive the same PD-based credit ratings despite their very different performance under systematic stress conditions.

To be sure, greater credit enhancement provides added protection for structured finance investors both on average and in times of severe systematic stress, but the benefits of greater seniority are attenuated under stress conditions. If the collateral backing a structured finance deal consists of a very large number of homogenous credit exposures⁸ the realized pool loss rate is well approximated by $\bar{L} = \text{CPD}(X) \cdot \text{LGD}$. In this case, the unconditional default probability for a senior tranche with attachment point C is

⁶ This is a bit of an over-simplification, since, depending on the details of the a deal's cash flow waterfall, some interest and principal payments may be directed to junior tranches before senior tranches are fully paid off. Typically, this is only permitted when realized credit losses on collateral assets fall below pre-specified benchmarks so that senior tranches are protected.

⁷ Some minor variation in the unconditional default probability arises from the fact that for finite N and nonrandom LGD, \bar{L} is a discrete random variable.

⁸ For more formal treatments of asymptotically fine grained portfolios in a single systematic factor model, see Gordy (2003) and Vasicek (1998).

$$PD_C = \Pr[\bar{L} > C] = \Phi\left(\Phi^{-1}(PD)\frac{1}{\sqrt{\rho}} - \Phi^{-1}\left(\frac{C}{LGD}\right)\sqrt{\frac{1-\rho}{\rho}}\right).$$

Under a q-th percentile or worse stress scenario, the tranche default probability is

$$PD_C^q = \begin{cases} \frac{1}{1-q} \Phi\left(\Phi^{-1}(PD)\frac{1}{\sqrt{\rho}} - \Phi^{-1}\left(\frac{C}{LGD}\right)\sqrt{\frac{1-\rho}{\rho}}\right) & \text{if } \Phi^{-1}(PD)\frac{1}{\sqrt{\rho}} - \Phi^{-1}\left(\frac{C}{LGD}\right)\sqrt{\frac{1-\rho}{\rho}} < \Phi^{-1}(1-q) \\ 1 & \text{if } \Phi^{-1}(PD)\frac{1}{\sqrt{\rho}} - \Phi^{-1}\left(\frac{C}{LGD}\right)\sqrt{\frac{1-\rho}{\rho}} \geq \Phi^{-1}(1-q) \end{cases}$$

The three panels in Figure 6 show how tranche seniority affects unconditional and stress scenario tranche default probabilities given low-, medium-, and high-quality collateral. The green and red solid lines show that increasing the tranche attachment point lowers both the unconditional tranche default probability and the stress scenario tranche default probability. The dashed line in the figure shows the stress condition default probability for a whole loan whose unconditional default probability is set to match the unconditional tranche default probability. Comparing the solid and dashed red lines reveals that senior structured securities have much higher default probabilities under severe stress conditions than whole loans with the same unconditional default probabilities.

Ratings designed to reflect unconditional default probabilities or expected losses can fail to predict credit performance under severe stress conditions. Furthermore the liabilities of a structured finance deal are usually finely tuned to achieve the best possible distribution of credit ratings given the collateral backing them. If a deal is backed by higher quality or better diversified collateral, the deal's senior tranche can garner a high credit rating with less credit enhancement (i.e., a lower attachment point) than otherwise. As a result, higher-quality or better diversified collateral does not necessarily imply less risk for investors in AAA-rated structured securities and may actually expose these investors to higher losses under stress conditions. In light of this analysis, it is perhaps not surprising that so many highly rated structured securities performed so poorly when subjected to a nationwide fall in house prices and escalating mortgage default rates.

4. Modest errors in rating collateral imply large errors in rating structured securities

Because credit losses on structured securities can be highly sensitive to systematic shocks, a great deal of historical data on the performance of collateral is needed to accurately rate them. Moreover, it is not sufficient that such data cover a large number of individual collateral exposures; it must also cover a broad range of macroeconomic conditions. This section uses a simple Monte Carlo simulation and the preceding copula model to illustrate how modest errors in measuring the default rates of collateral assets can translate into large errors in the measured default rate for a senior structured security.

In the previous section we showed that if the senior tranche of a structured finance deal is backed by a well-diversified pool of homogeneous credit exposures the tranche default probability can be computed directly when the PD, LGD and ρ of the deal collateral are known. But what if some of these parameters must be estimated? To keep things simple, let us make the heroic assumption that both LGD and ρ are known and that only PD must be estimated from historical data. Assume we observe T annual cohorts of data on the credit performance of an effectively infinite number of homogeneous loans. In the Appendix, I show that under these assumptions it is possible to construct a unique

minimum variance unbiased estimator (MVUE) of the loans' unconditional default probability (PD) and any tranche default probability (PD_C).

Table 3 and Figure 7 describe the distribution of MVUE PD estimates given 5-, 10-, and 30-year simulated data histories, assuming that the true collateral default probability is ten percent.⁹ Notice that there is significant dispersion among these estimated default probabilities even though we have assumed that the number of exposures in each cohort is so large that the observed default rate for a cohort is arbitrarily close to its conditional expectation given X . Because systematic risk leads to correlation in defaults across obligors, a long time-series of data – not a large cross-section – is needed to accurately estimate the unconditional default probability.

Table 4 and Figure 8 shows the distribution of MVUE PD_C estimates for a senior tranche with a true default probability of 91bp (this is the same tranche described in the last row of Table 2). Even if 30 years of data on a very large number of credit exposures are available for estimating the collateral default probability, the dispersion in estimated default probabilities for senior structured securities is exceptionally large. A 90 percent confidence interval would cover a range from about one-third to double the true default probability. When fewer years of data are available the sampling error is even larger.

To put these results in perspective, Table 5 and Figure 9 describe the distribution of PD estimates for a whole loan whose true default probability matches that of the senior structured security. Notice that the third line in Table 4 is similar to the first line in Table 5. In other words, about six times as much historical data are required to rate the structured security with the same precision as a whole loan with the same true default probability.

The results illustrated in this simple Monte Carlo exercise are quite general. The Appendix provides analytic expressions for the sampling variances of best unbiased PD and PD_C estimators given any configuration of model parameters. Figure 10 compares the standard deviation of the MVUE of a tranche default probability with that of the MVUE of a whole loan default probability with the same actual PD as the tranche.¹⁰ In all cases, the standard deviation of the tranche default probability estimator is more than double that of the whole loan default probability estimator. The difference is greatest for tranches with more credit enhancement (higher attachment points) or backed by higher-quality (lower-PD) collateral. Evidently, the safest structured securities are the hardest to rate relative to whole loans of comparable quality.

This analysis illustrates how relatively modest uncertainty about the credit quality of collateral assets can translate into very large uncertainty about the credit quality of structured liabilities backed by those assets. Note that, if anything, the simplifying assumptions used in this analysis tend to minimize the uncertainty associated with evaluating structured credit products. Factor loadings and LGDs are notoriously difficult

⁹ In each simulation, T pseudorandom realizations of X were drawn and used to compute T values of $CPD(X)$. 100,000 Monte Carlo iterations were run. To my knowledge, the MVUE PD estimators used in these simulations have not been applied in practice. In assessing the credit quality of whole loans, practitioners commonly use long-run averages of rating-cohort default rates. These estimators are unbiased PD estimators under standard assumptions, but when systematic risk is present, they are not efficient.

¹⁰ This analysis assumes a factor loading of 0.2, an LGD of 45 percent, and 10 years of historical data. Similar results obtain under alternative parameter configurations.

to pin down and ample empirical research finds that much more flexible copula models are needed to capture observable features of real-world data.¹¹

5. Looking forward

Moody's, Standard and Poor's, and Fitch use the same ordinal rating scales to rank traditional corporate debt securities and structured securities. Even as the crisis in mortgage-backed structured securities was unfolding, the agencies argued that their ratings should be treated as comparable across debt classes. For example, S&P (2007b) asserted:

Our ratings represent a uniform measure of credit quality globally and across all types of debt instruments. In other words, an 'AAA' rated corporate bond should exhibit the same degree of credit quality as an 'AAA' rated securitized debt issue.

The analysis presented here suggests that the uncertainty associated with ratings for high quality structured products is likely to be much greater than that associated with high-quality corporate debt securities. Ratings designed to capture a credit exposure's unconditional default probability or expected loss provide little insight into how a structured security might perform under systematic stress conditions. Furthermore estimated default probabilities and expected losses for senior structured securities are likely to be highly sensitive to modest and foreseeable errors in assessing the credit quality of underlying collateral.

Financial regulators (President's Working Group, 2008) have urged rating agencies to "make changes to the credit rating process that could clearly differentiate ratings for structured products from ratings for corporate and municipal securities" and "work with investors to provide the information investors need to make informed decisions about risk, including measures of the uncertainty associated with ratings and of potential ratings volatility." In response, rating agencies have taken steps to better convey the uncertainty associated with structured finance ratings. Moody's now provides sensitivity scores alongside its structured finance ratings that are designed to provide users with better information about how ratings are affected by model assumptions and the credit quality of underlying collateral. S&P publishes regular forecasts of the performance of structured credit products under a range of systematic stress scenarios. Disclosures such as these may help to increase market confidence in the agencies' structured finance rating processes.

However, ultimate responsibility for risk management rests not with rating agencies, but with those institutions that elect to invest in structured securities. Even in the absence of model riskb similarly rated credit exposures may perform very differently under systematic stress conditions, so one-dimensional credit ratings cannot be treated as sufficient summaries for credit quality. To the extent that securitization is used to fund retail and corporate debt going forward, risk-management systems designed primarily with loan-level underwriting in mind need to be adapted to better reflect undiversifiable risks.

¹¹ Heitfield (2009) examines the effects of errors in estimating factor loadings and loss-given default on CDO risk metrics.

Appendix: Minimum Variance Unbiased Estimators for Default Probabilities when the Number of Bonds in a Cohort is Very Large but the Number of Cohorts is Fixed

The PD estimators used in this paper are derived under the following assumptions.

- Individual whole loan defaults are described by a one-factor normal copula model. All whole loans are homogeneous in the sense that they share the same unconditional default probability (PD) and factor loading (ρ). The loss rate per default (LGD) is nonrandom and constant across loans.
- Realized default rates ($F_1 \dots F_T$) are observed for T cohorts of whole loans. Each cohort is arbitrarily large, so that

$$F_t = \Phi\left(\frac{\Phi^{-1}(\text{PD}) - X_t \sqrt{\rho}}{\sqrt{1-\rho}}\right)$$

where ($X_1 \dots X_T$) are IID standard normal random variables.

- PD is not directly observable, but ρ and LGD are known.

A.1 Estimating the default probability for a whole loan

I will first construct the unique minimum variance unbiased estimator for PD. Observe that $E[F_1] = \text{PD}$ so a single observed default rate is an unbiased (though very inefficient) estimator of PD. The Lehmann-Scheffe Theorem (see Bickel and Doksum, 1977, pg. 122) provides a means of constructing a unique minimum variance unbiased estimator (MVUE) for a function of a model parameters. To apply this result, we need to find a complete sufficient statistic for the unknown parameter of the model.

Let

$$Z_t = \Phi^{-1}(F_t) \sqrt{\frac{1-\rho}{\rho}}.$$

Note that since we have assumed that ρ is known, Z_t can be computed directly from the observed default rate for cohort t. By construction, each Z_t is a standard normal random variable with mean

$$\mu = \frac{\Phi^{-1}(\text{PD})}{\sqrt{\rho}}$$

and unit variance, and all Z_t are independent of one another. It is well known that the sample average of a vector of IID normal random variables is a complete sufficient statistic for the population mean when the variance is known, so $\bar{Z} = \frac{1}{T} \sum_t Z_t$ is a complete sufficient statistics for μ . Since $\text{PD} = \Phi(\mu \sqrt{\rho})$, and F_1 is an unbiased estimator of PD, the Lehmann-Scheffe Theorem implies that the minimum variance unbiased estimator for PD is $E[F_1 | \bar{Z}]$. We can derive an explicit expression for this conditional expectation. Let W be a standard normal random variable that is independent of all Z_t , then

$$\begin{aligned}
E[F_1 | \bar{Z}] &= E\left[\Phi\left(Z_1 \sqrt{\frac{\rho}{1-\rho}}\right) \mid \bar{Z}\right] \\
&= E\left[\mathbb{1}\left\{W < Z_1 \sqrt{\frac{\rho}{1-\rho}}\right\} \mid \bar{Z}\right] \\
&= \Pr\left[W - Z_1 \sqrt{\frac{\rho}{1-\rho}} \leq 0 \mid \bar{Z}\right]
\end{aligned}$$

Using the fact that $Z_1 | \bar{Z} \sim N(\bar{Z}, \frac{T-1}{T})$, it follows immediately that the unique MVUE for PD is

$$\hat{P}(\bar{Z}) = \Phi\left(\bar{Z} \sqrt{\frac{\rho}{1-\rho/T}}\right).$$

The variance of this estimator has a closed form expression. Observe that $\bar{Z} \sim N(\mu, \frac{1}{T})$ and let W_1 and W_2 be two independent standard normal random variables.

$$\begin{aligned}
\text{Var}[\hat{P}(\bar{Z})] &= E\left[\Phi\left(\bar{Z} \sqrt{\frac{\rho}{1-\rho/T}}\right)^2\right] - \text{PD}^2 \\
&= E\left[\mathbb{1}\left\{W_1 \leq \bar{Z} \sqrt{\frac{\rho}{1-\rho/T}}\right\} \cdot \mathbb{1}\left\{W_2 \leq \bar{Z} \sqrt{\frac{\rho}{1-\rho/T}}\right\}\right] - \text{PD}^2 \\
&= \Pr\left[\left(W_1 - \bar{Z} \sqrt{\frac{\rho}{1-\rho/T}}\right) \sqrt{\frac{T-\rho}{T}} \leq 0, \left(W_2 - \bar{Z} \sqrt{\frac{\rho}{1-\rho/T}}\right) \sqrt{\frac{T-\rho}{T}} \leq 0\right] - \text{PD}^2 \\
&= \Phi_2\left(\Phi^{-1}(\text{PD}), \Phi^{-1}(\text{PD}); \frac{\rho}{T}\right) - \text{PD}^2.
\end{aligned}$$

A consistent estimator for the sampling variance of $\hat{P}(\bar{Z})$ can be computed by plugging the estimated PD into the right-hand expression.

A.2 Estimating the default probability for a tranche backed by a large number of homogeneous bonds

A corollary to the Lehmann-Scheffe Theorem (Patil and Wani, 1964) provides a method of constructing the MVUE of a distribution function when complete sufficient statistics for the parameters of that distribution are available. Adapting this result to the current application provides the unique MVUE of PD_C . The loss rate for an arbitrarily large pool of homogeneous bonds is

$$F_t \cdot \text{LGD} = \Phi\left(Z_t \sqrt{\frac{\rho}{1-\rho}}\right) \cdot \text{LGD}.$$

Substituting this expression for the collateral loss rate into the definition of PD_C (page 7) and rearranging terms yields

$$\text{PD}_C = \Pr\left[Z_t > \Phi^{-1}\left(\frac{C}{\text{LGD}}\right) \sqrt{\frac{1-\rho}{\rho}}\right].$$

We have already shown that \bar{Z} is a complete sufficient statistic for the distribution of Z_t and that $Z_t | \bar{Z} \sim N(\bar{Z}, \frac{T-1}{T})$. By Patil and Wani's corollary, the MVUE of PD_C is

$$\begin{aligned}
\hat{P}_C(\bar{Z}) &= \Pr\left[Z_t > \Phi^{-1}\left(\frac{C}{\text{LGD}}\right) \sqrt{\frac{1-\rho}{\rho}} \mid \bar{Z}\right] \\
&= \Phi\left(\left(\bar{Z} - \Phi^{-1}\left(\frac{C}{\text{LGD}}\right) \sqrt{\frac{1-\rho}{\rho}}\right) \sqrt{\frac{T}{T-1}}\right).
\end{aligned}$$

As with whole loan, we can derive an analytic expression for sampling variance of the MVUE of the tranche default probability:

$$\begin{aligned}
\text{Var}[\hat{P}_C(\bar{Z})] &= E\left[\Phi\left(\left(\bar{Z} - \Phi^{-1}\left(\frac{C}{LGD}\right)\sqrt{\frac{1-p}{\rho}}\right)\sqrt{\frac{T}{T-1}}\right)^2\right] - (PD_C)^2 \\
&= E\left[1\left\{W_1 \leq \left(\bar{Z} - \Phi^{-1}\left(\frac{C}{LGD}\right)\sqrt{\frac{1-p}{\rho}}\right)\sqrt{\frac{T}{T-1}}\right\} \cdot 1\left\{W_2 \leq \left(\bar{Z} - \Phi^{-1}\left(\frac{C}{LGD}\right)\sqrt{\frac{1-p}{\rho}}\right)\sqrt{\frac{T}{T-1}}\right\}\right] - (PD_C)^2 \\
&= \Pr\left[W_1\sqrt{\frac{T-1}{T}} - \bar{Z} \leq -\Phi^{-1}\left(\frac{C}{LGD}\right)\sqrt{\frac{1-p}{\rho}}, W_2\sqrt{\frac{T-1}{T}} - \bar{Z} \leq -\Phi^{-1}\left(\frac{C}{LGD}\right)\sqrt{\frac{1-p}{\rho}}\right] - (PD_C)^2 \\
&= \Phi_2\left(\Phi^{-1}(PD)\sqrt{\frac{1}{\rho}} - \Phi^{-1}\left(\frac{C}{LGD}\right)\sqrt{\frac{1-p}{\rho}}, \Phi^{-1}(PD)\sqrt{\frac{1}{\rho}} - \Phi^{-1}\left(\frac{C}{LGD}\right)\sqrt{\frac{1-p}{\rho}}; \frac{1}{T}\right) - (PD_C)^2
\end{aligned}$$

Table 1: This table shows loss rate exceedance probabilities for less diversified and more diversified loan pools. Each pool contains N loans, each of which has a default probability of 10 percent, a loss given default of 45 percent, and a factor loading of 0.20. Diversification reduces unconditional tail exceedance probabilities. Under systematic stress conditions, all exceedance probabilities increase and diversification does not reduce the likelihood of tail losses.

	Number of Loans	Loss Exceedance Probability		
		$\bar{L} > 10\%$	$\bar{L} > 15\%$	$\bar{L} > 20\%$
Unconditional	1	10.0	10.0	10.0
	25	12.1	3.4	0.8
	50	9.7	3.1	0.7
	100	9.5	2.5	0.6
	∞	9.1	2.3	0.5
Conditional on 98 th Percentile or Worse Systematic Shock	1	41.2	41.2	41.2
	25	96.2	72.2	32.3
	50	98.8	81.1	31.2
	100	99.9	84.6	29.4
	∞	100.0	100.0	24.3

Table 2: This table compares default probabilities for a whole loan and four structure credit products. Each credit exposure is constructed to have roughly the same 90 basis point unconditional default probability. The structured credit products are likely to perform much worse than the whole loan under systematic stress conditions and senior tranches backed by better diversified collateral pools are likely to perform worse than those backed by more concentrated pools.

Number of Loans in Collateral Pool	Senior Tranche Attachment Point	Unconditional Senior Tranche Default Probability	Stress Condition Senior Tranche Default Probability*
Whole Loan	n.a.	0.90	7.86
25	20.0	0.85	32.25
50	19.0	0.89	38.88
100	18.5	0.90	42.95
∞	18.0	0.92	45.82

* Stress condition is defined as a 98th percentile or worse systematic shock.

Table 3: This table summarizes the distribution of simulated T-year average default rates for arbitrarily large cohorts of credit obligations. Each credit exposure has a default probability of 10 percent and a factor loading of 0.20. Even though the cohorts are arbitrarily large, there is considerable dispersion in estimates of unconditional default rates, particularly when T is small.

T	Mean	Standard Deviation	5 th Percentile	95 th Percentile
5	10.01	3.56	5.02	16.57
10	10.01	2.52	6.30	14.51
30	10.00	1.44	7.78	12.49

Table 4: This table summarizes the distribution of simulated default probabilities for a senior structured tranche backed by a well-diversified collateral pool assuming that a T-year average default rate is used as a proxy for the pool’s true expected default rate. Loans in the collateral pool have a 10 percent default probability, a 45 percent loss-given-default and a factor loading of 0.20. The senior tranche has an attachment point of 20 percent and a true default probability of 91 basis points. The dispersion in estimated tranche default probabilities is extremely large compared to the true parameter, even when the collateral default rate is estimated over a 30-year time span.

T	Mean	Standard Deviation	5 th Percentile	95 th Percentile
5	0.92	1.43	0.03	3.50
10	0.92	0.91	0.12	2.66
30	0.92	0.47	0.34	1.81

Table 5: This table summarizes the distribution of simulated T-year average default rates for arbitrarily large cohorts of credit obligations. Each credit exposure has a default probability of 92 basis points and a factor loading of 0.20. Comparing the results of this table with those of Table 4, we see that the dispersion in the estimated default probabilities for a whole loan is substantially lower than that for a senior structured security with the same true default probability.

T	Mean	Standard Deviation	5 th Percentile	95 th Percentile
5	0.92	0.52	0.31	1.92
10	0.92	0.36	0.44	1.59
30	0.92	0.20	0.62	1.28

Figure 1

Annual Issuance of Non-agency Residential Mortgage Backed Securities





Figure 2

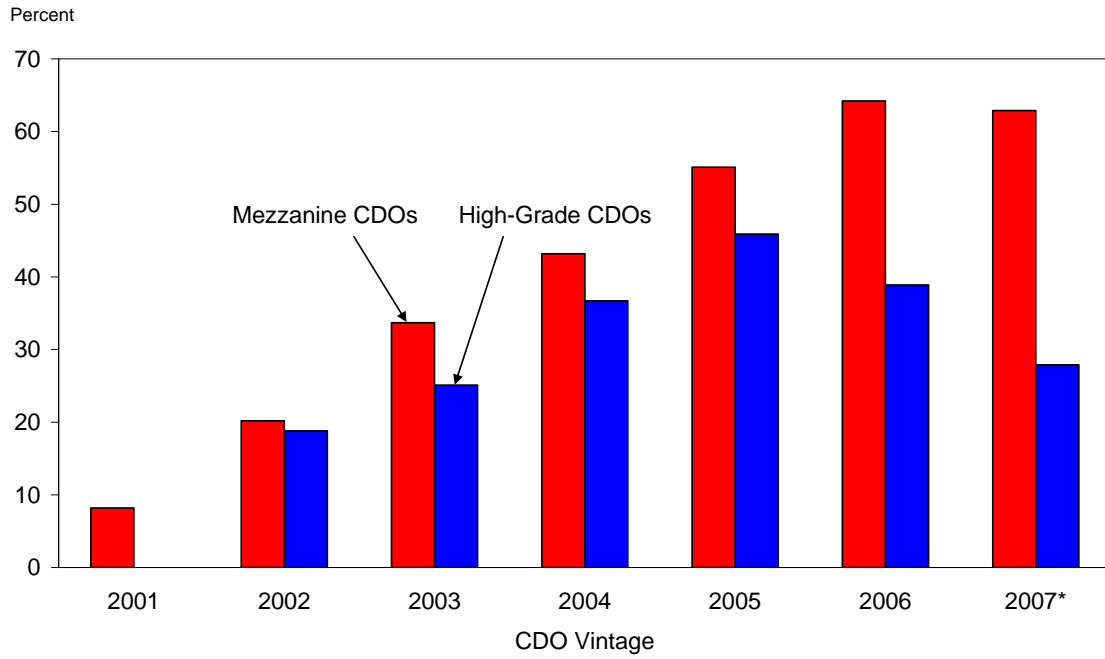
Quarterly Issuance of CDOs Backed by Structured Finance Collateral



Source: Securities Industry and Financial Markets Association

Figure 3


Subprime RMBS Share of ABS CDO Collateral, by CDO Vintage



* 2007 vintage includes deals completed through September
Source: Standard and Poor's

Figure 4

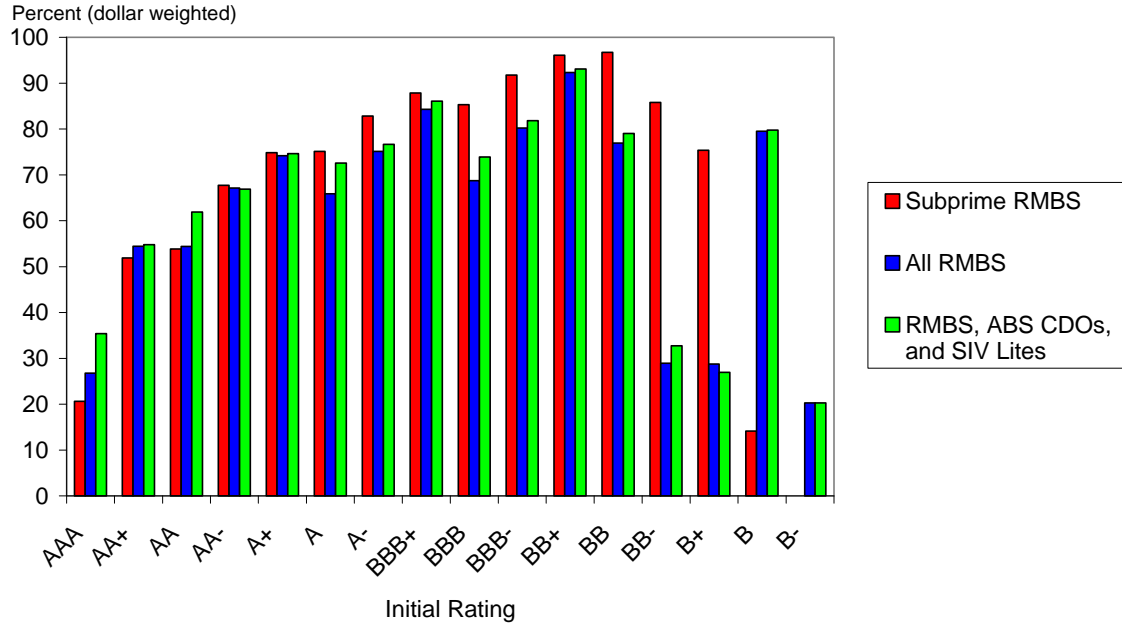
AAA ABX.HE Indexes



Source: MarKit

Figure 5

Downgrade Rates for 2005-07 Vintage
U.S. Mortgage-Backed Structured Securities



Source: Standard and Poor's

Figure 6

Unconditional and Stress Scenario Tranche Default Probabilities Compared with the Stress Condition Default Probability for a Comparably-Rated Whole Loan.


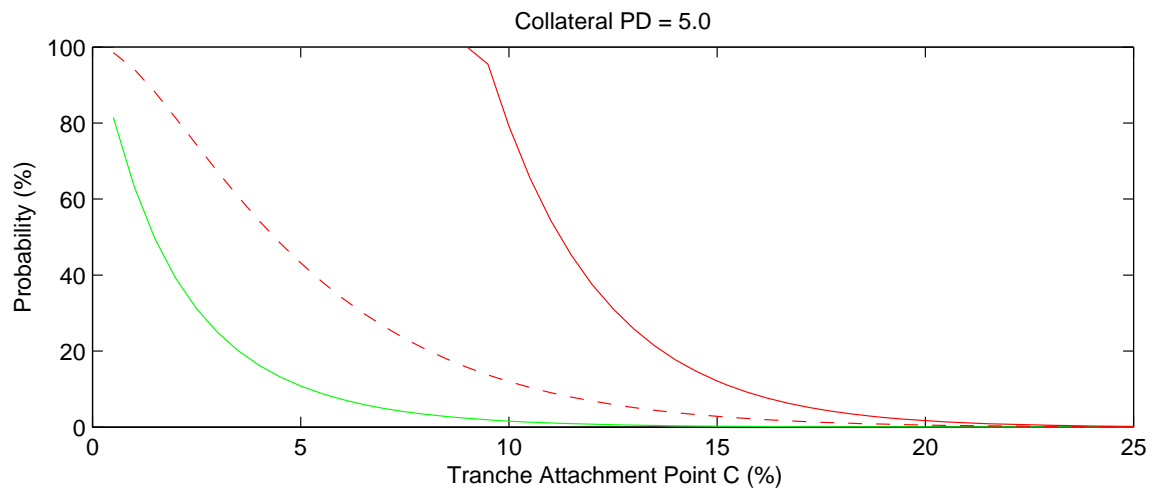
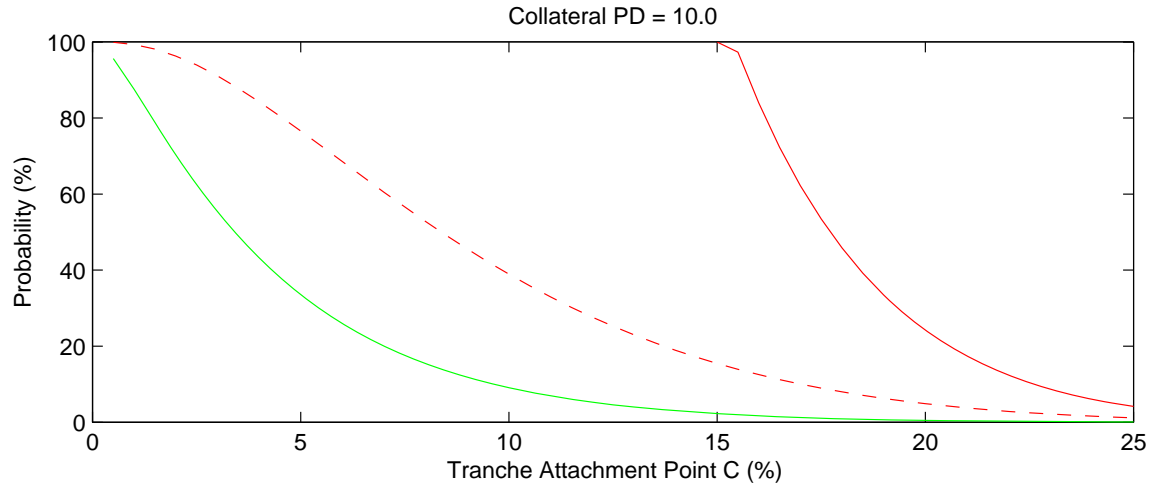


Figure 7
Distribution of Estimated Default Probabilities for a
Whole Loan with a True Default Probability of Ten Percent




Figure 8
Distribution of Estimated Default Probabilities for a Senior
Structured Security with a True Default Probability of 90 Basis Points




Figure 9
Distribution of Estimated Default Probabilities for a
Whole Loan with a True Default Probability of 90 Basis Points

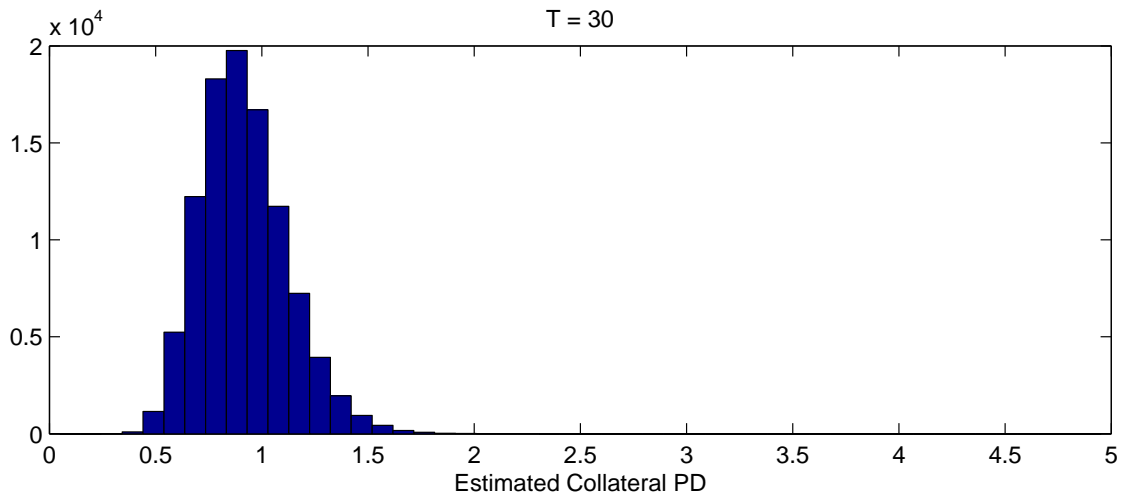

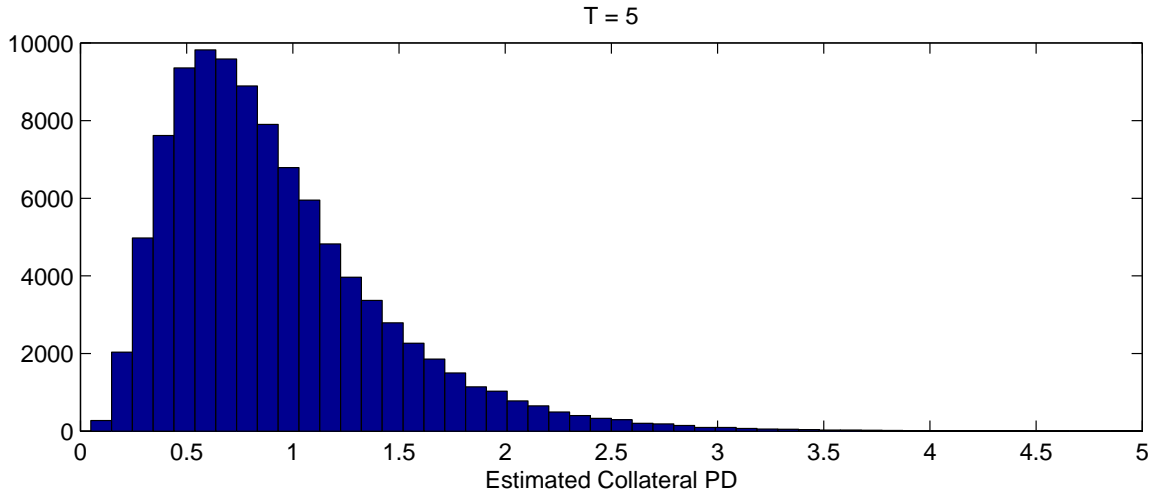



Figure 10

Ratio of the Standard Deviation of the Best Unbiased Estimator of the Default Probability for a Structured Bond and a Whole Loan with the Same True PD



Works Cited

- Andersen, L., and J. Sidenius (2005): "CDO Pricing with Factor Models: Survey and Comments," *Journal of Credit Risk*, 1(3), 71–88.
- Basel Committee on Banking Supervision (2004): *International Convergence of Capital Measurement and Capital Standards: A Revised Framework*, Bank for International Settlements, Basel.
- Bickel, P. and K. Doksum (1977): *Mathematical Statistics: Basic Ideas and Selected Topics*, Holden-Day, Oakland.
- Committee on the Global Financial System (2008): "Ratings in structure finance: what went wrong and what can be done to address shortcomings?" *CGFS Papers*, 32.
- Fitch (2007): "Rating Stability of Fitch-Rated Global Cash Mezzanine Structured Finance CDOs with Exposure to U.S. Subprime RMBS," *DerivativeFitch Structured Credit Special Report*.
- Gordy, M. (2003): "A risk factor model foundation for ratings-based bank capital rules." *Journal of Financial Intermediation*, 12, 199-232.
- Gupton, Greg, Christopher C. Finger, and Mickey Bhatia. 1997. *CreditMetrics – Technical Document*, J. P. Morgan and Co., New York.
- Heitfield, E. (2009): "Parameter Uncertainty and the Credit Risk of Collateralized Debt Obligations". Working paper (Available at SSRN: <http://ssrn.com/abstract=1190362>).
- Kealhofer, S. (1998): "Portfolio Management of Default Risk," *KMV Corporation Discussion Paper*.
- Li, D. X. (2000): "On Default Correlations: A Copula Function Approach," *Journal of Fixed Income*, 9, 43–54.
- Moody's (2007): "The Impact of Subprime Residential Mortgage-Backed Securities on Moody's-Rated Structured Finance CDOs: A preliminary Review," *Moody's Structured Finance Special Comment*.
- Patil, G. and Wani, J. (1964): "Minimum variance unbiased estimation of the distribution function admitting a sufficient statistic," *Annals of the Institute of Statistical Mathematics*, 18(1), 39-47.
- President's Working Group on Financial Markets (2008): "Policy Statement on Financial Market Developments," Policy statement.
- Standard and Poor's (2007a). "The Effect of Mortgage Market Stress on U.S. CDO Ratings in Third-Quarter 2007," *Standard and Poor's CDO Spotlight*.
- Standard and Poor's (2007b): "Principles-Based Rating Methodology for Global Structured Finance Securities," *Standard and Poor's Discussion paper*.
- Standard and Poor's (2009): "Transition Study: Structured Finance Rating Transition and Default Update as of Feb. 27, 2009," *Standard and Poor's RatingsDirect*.
- Vasicek, O. A. 1991. "Limiting loan loss probability distributions." *KMV Corporation Discussion Paper*.