DRAFT 1.1
Parameter Uncertainty and the Credit Risk of Collateralized Debt Obligations

Erik Heitfield
Division of Research and Statistics
Federal Reserve Board
Washington, DC 20551 USA
202-452-2613
erik.heitfield@frb.gov[1]

May 8, 2008

[1]The views expressed here are those of the author and do not reflect the opinions of the Federal Reserve Board or its staff. This paper is still very much a work in progress. If you would like to cite or distribute it, please contact the author beforehand, as a revised draft will be available shortly.

Abstract

This paper examines the empirical difficulties inherent in assessing the credit quality of collateralized debt obligations (CDOs). Because of the way CDO liabilities are structured, CDO note payouts are sensitive to tail collateral loss events. As a result, in order to assess the likelihood and severity of a CDO note's losses, one needs to know the distribution of losses for each collateral asset, as well as the dependence of losses across collateral assets. In practice, CDO collateral losses are most commonly modeled using normal copulas. I show that for more senior CDO tranches, standard credit risk metrics such as probability of default, expected loss, and conditional expected loss are highly sensitive to model parameters which must be estimated or specified judgmentally by a credit analyst. Given assumptions about the historical data available to a credit analyst, I compute bounds on the accuracy of normal copula parameter estimates and show that in applied settings data constraints are likely to impose severe limitations on an analyst's ability to accurately evaluate CDO tranches. Thus, CDO note credit ratings should generally be viewed as more preliminary and less informative than comparable corporate bond ratings.

1 Introduction

Collateralized debt obligations (CDOs) are structured fixed income securities whose payouts depend on the performance of pools of collateral comprised of corporate bonds or loans, or other structured securities which are themselves backed by underlying collateral pools. The CDO market has grown and evolved considerably over the last decade, but recently CDOs have been implicated in the ongoing financial market turmoil in the US and Europe. As large numbers of highly-rated CDO notes have experienced dramatic credit rating downgrades and/or significant falls in market valuations, investors have become unwilling to hold them, creating liquidity and credit challenges for financial institutions. A number of commentators have pointed out that lack of transparency and complexity of CDO deals pose challenges for market participants seeking to quantify and manage CDO credit risks, and rating agencies that make a business of evaluating the credit quality of debt securities have been criticized for issuing ratings for CDOs that some believe were too optimistic or did not capture the full range of risks associated with these investments.[1]

This paper examines the empirical difficulties inherent in assessing CDO credit performance. I argue that the design of CDOs makes these structures more difficult to evaluate than more traditional types of debt securities with embedded credit risk such as corporate bonds. In assigning credit ratings, rating agencies and others typically focus on a security's probability of default or expected loss. In managing portfolio credit risk, financial institutions and their regulators may make use of more complex risk metrics, such as conditional expected loss or expected shortfall. Because of the way CDO liabilities are structured, CDO note payouts are sensitive to tail collateral loss events. As a result, in order assess the likelihood and severity of a CDO note's losses, one needs to know the distribution of losses for each collateral asset, as well as the dependence of losses across collateral assets. Accurately measuring loss correlations or other higher-order moments of loss distributions for groups of fixed income securities can be exceptionally challenging, particularly when defaults are relatively rare.

The paper is organized as follows. Section 2 briefly surveys prior research on normal copula models commonly used to evaluate the credit risk of CDOs and other structured finance products. Section 3 presents a stripped down version of a normal copula model and shows how it can be used to compute loss distributions for unstructured and structured debt securities. This model is considerably simpler than those typically used to evaluate debt securities. It is intended to capture the salient features of richer copula-based credit models, while allowing one to investigate the role of a small number of key model parameters. Section 4 uses simulations to show how changes in model parameters affect standard metrics of the credit quality of debt securities. These simulations show that risk metrics for more senior CDO tranches are much more sensitive to errors in copula model parameters than risk metrics for more junior CDO tranches or unstructured bonds. Section 5 computes bounds on the precision with which normal copula model parameters can be estimated given assumptions about the historical data available to an analyst. This analysis highlights how both the quantity and character of available data imposes quantifiable limits on an analyst's ability to accurately estimate model parameters. Section 6 uses results from Sections 4 and 5 to

[1] See, for example Mason and Rosner (2007).

simulate the distribution of standard credit risk metrics for different types of structured and unstructured bonds. These results suggest that even when high quality historical data are relatively plentiful, it may be difficult to accurately gauge the credit quality of more senior CDO notes.

2 Literature review

Modeling dependence in realized defaults among groups of credit exposures is critical for portfolio credit risk management. Normal copula models describe default dependence using systems of correlated normal latent credit factors. Copula-based models have become popular over the last decade both because they are computationally tractable, and because they can be derived from the structural corporate debt valuation framework of Merton (1974). Today, normal copula models are used in a broad range of risk management applications. Widely used 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 bank regulatory requirements under the Basel II risk-based regulatory 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[2] and normal copula models are commonly used to price CDO notes (Andersen and Sidenius 2005a).

Despite their popularity, limitations of normal copula models have been well documented. A particular concern is that models capable of fitting observable data under typical credit conditions appear to understate the likelihood of extreme portfolio credit loss events. Numerous authors have proposed extensions or generalizations of the normal copula framework to address this problem. For example, Frey, McNeil and Nyfeler (2001) propose a copula model based on thicker tailed t-distributed latent variables and Andersen and Sidenius (2005b) and Burtschell, Gregory and Laurent (2007) extend the normal copula model to allow for unobserved heterogeneity in latent factor correlations across credit exposures.

Normal copula models and their various extensions depend on vectors of parameters that describe the probability and likely severity of individual credit loss events and the correlation structure of latent credit factors. In applied settings these parameters must be estimated from historical data or specified judgmentally by a credit analyst. Any errors in estimating model parameters will naturally result in miss-measurement of the credit risk associated with individual exposures or portfolios of exposures. Simulation studies by Loffler (2001), Tarashev and Zhu (2007), and Hamerle and Rosch (2005) investigate the sensitivity of portfolio loss measures to errors in estimating copula model parameters.

A small body of research has examined the statistical properties of specific estimators of copula model correlation parameters. Gordy and Heitfield (2002) compare the small sample properties of maximum likelihood and moment-based estimators of credit factor correlation parameters and show how imposing intuitive parameter restrictions can improve estimation

[2] See Jolivet, Lassalvy, Mueller-Heumann and Sieler (2006), Gilkes (2004), and Koo, Cromartie and Vedova (2006) for descriptions of the normal-copula models used by Moody's, S&P, and Fitch respectively.

efficiency. Frey and McNeil (2003) propose maximum likelihood estimators for copula correlation parameters in the presence of non-normal latent credit factors, and Hamerle and Rosch (2005) examine the sensitivity of correlation parameters to miss-specification of the distribution of latent credit factors.

The analysis presented in this paper contributes to the academic literature on copula models in a number of ways. I show how bounds on the accuracy of normal copula model parameters can be determined with minimum assumptions about the actual estimators used. Unlike previous research on the accuracy of copula model parameters, I allow for the possibility that an analyst my have access to data no the credit quality of non-defaulted firms that are useful for estimating model parameters. This extension is important, since in applied settings normal copula correlation parameters are commonly estimated using information on equity returns or imputed asset returns for publicly traded firms or historical ratings transition data for rated bonds.[3] Though others have investigated the effects of copula model specification errors on measures of portfolio credit risk, this paper is the first to rigorously examine how data limitations affect the accuracy of credit risk metrics for structured debt products.

3 The normal copula/beta model for correlated bond credit losses

Throughout this paper, I will use the term "simple" bonds to describe traditional debt claims on corporations or sovereigns. I will use the term "structured" bonds to describe bonds issued by CDOs, whose payouts depend on an underlying collateral pool of simple bonds. This section describe a stripped down version of the normal copula/beta model commonly used to evaluate the credit risk of both simple and structured bonds. The specification used in this analysis allows for cross sectional correlation in realized default rates for simple bonds and stochastic losses for those bonds that default. Correlations in defaults are driven by a single systematic risk factor. For simplicity, I assume that loss rates given default are independent of the systematic factor. It should be noted that both the single systematic risk factor assumption and the assumption that there is no systematic risk in loss given default can be easily relaxed in more applied settings, and indeed, this is commonly done in practice.

3.1 Simple bonds

Under the simplest normal copula framework, bond i defaults during a specified horizon if an unobservable normal latent factor Y_i lies below the default threshold $\Phi^{-1}(\pi)$ where $\Phi^{-1}(\cdot)$ is inverse of the standard normal cumulative density function. The parameter π describes the bond's marginal probability of default. Cross sectional correlation in defaults across pairs of

[3] For example Gupton et al. (1997) describes how equity return correlations are used to infer latent factor correlation parameters used in the CreditMetrics portfolio risk model and Kealhofer (1998) explains how Moodys/KMV uses imputed obligor asset returns to calibrate correlation parameters in its portfolio risk model.

simple bonds i and j arises from correlations in latent credit factors Y_i and Y_j. Let

$$Y_i = \sqrt{\rho}X - \sqrt{1-\rho}E_i$$

where X is a standard normal random factor shared by all bonds, and E_i is a standard normal idiosyncratic factor that is unique for each bond. The parameter ρ lies between zero and one and determines the correlation in credit factors between pairs of bonds. Higher values of ρ imply higher correlation between credit factors, and, by extension, higher correlation in realized defaults across pairs of bonds.

Assume that bond i is a bullet loan that pays $1 + r_i$ at maturity if the obligor does not default, and $(1 + r_i)(1 - \lambda_i)$ if the obligor defaults. λ_i is a random variable describing the realized loss given default of the bond. For corporate bond exposures, this loss rate is often assumed to be drawn from a beta distribution which may or may not depend on the systematic factor(s) that drive asset correlations. As noted above, for simplicity this analysis assume that λ_i is independent of all other random variables. The beta distribution is a two parameter distribution with support on the unit interval that can be fully characterized by a mean parameter μ and a standard deviation parameter σ.[4]

The payout from a one dollar investment in bond i at the terminal date is,

$$V_i = (1 + r_i) - 1\left\{Y_i \leq \Phi^{-1}(\pi)\right\}\lambda_i(1 + r_i)$$

The right-most term is the realized contractual loss per dollar invested. Note that when λ_i is large, this loss rate may exceed 100 percent because of accrued but unpaid interest. Given N homogeneous bonds, the joint distribution of $V_1 \ldots V_N$ is fully described by the parameter vector $\theta = (\pi, \rho, \mu, \sigma)$.

3.2 CDOs backed by simple bonds

The simplest types of collateralized debt obligations (CDOs) issue tranches of structured debt securities backed by pools of corporate bonds. The normal-copula/beta model of bond losses can be used to build up a model of CDO tranche credit losses.

Consider a static CDO deal backed by N bonds. Investments are made at the "deal date" and proceeds are distributed to investors at the "terminal data". The value of the collateral pool at the deal date is normalized to 1 and the value of the collateral pool at the terminal date is denoted V_p. c_e of the collateral pool is funded by equity investors at the deal date. The remaining $1 - c_e$ of the pool is funded by a continuum of arbitrarily thin debt tranches. Debt tranches are indexed by $c \in [c_e, 1]$. c is a tranche's attachment point in the CDO capital structure, so higher values of c imply greater seniority. The interest paid to each debt tranche is described by the non-increasing function $r(c)$. At the terminal date, collateral is liquidated and tranche c investors are paid $1+r(c)$ if sufficient funds are available. If V_p is not sufficient to pay all debt investors, tranches are paid according to seniority. If V_p exceeds that needed to pay debt investors, equity investors receive any residual value.

[4] In the statistics literature the beta distribution is most commonly characterized by two shape parameters α and β. I use the less common μ-σ parametrization to make the economic interpretation of model parameters more transparent. It can be shown that $\alpha = (\mu(1-\mu) - \sigma^2)\mu/\sigma^2$ and $\beta = (\mu - 2\mu^2 + \mu^3 - \mu\sigma^2)/\sigma^2$.

Assuming no credit losses, the total value of all debt tranches senior to tranche c is

$$\bar{V}(c) = (1 - c) + R(c)$$

where $R(c) = \int_c^1 r(s)ds$ is the total interest owed to these tranches. The realized value of tranches senior to c is

$$V(c) = \bar{V}(c) - 1\left\{V_p \leq \bar{V}(c)\right\}\left(\bar{V}(c) - V_p\right)$$

The second right-hand term is the value of any realized credit losses for tranches senior to c. Note that the value for a "slice" of the CDO with attachment point c_l and detachment point c_h is $V(c_l) - V(c_h)$. The value of the equity trance is $V_e = V_p - V(c_e)$.

To keep notation simple, this analysis is restricted to CDOs backed by equal-weighted pools of bonds that are homogeneous in the sense that all bonds in the pool share the same parameter vector θ and pay the same interest rate r_p. Let $M = \sum_{n=1}^{N} 1\left\{Y_n \leq \Phi^{-1}(\pi)\right\}$ be a random variable that described the number of bonds in the CDO collateral pool that default by the terminal date, and let $\bar{\lambda}_M = \frac{1}{M}\sum_{j=1}^{M}\lambda_j$ be the average loss given default for those M bonds. The value of the collateral pool at the terminal date is

$$V_p = (1 + r_p) - \frac{M}{N}\bar{\lambda}_M.$$

The random variables M and $\bar{\lambda}_M$ determine V_p. M is a draw from a binomial-normal mixture distribution, and, conditional on M, $\bar{\lambda}_M$ is an average of M independent beta random variables. Neither the marginal distributions of M nor the conditional distribution of $\bar{\lambda}_M$ given M can be expressed in closed form, but both can be computed analytically with high precision. The product of these two distributions is the joint distribution of M and $\bar{\lambda}_M$, which provides all the information necessary to compute the joint distribution of V_p, $V(c)$ (for all c) and V_e.

The distribution of CDO tranche payouts is fully determined by N, r_p, $r(c)$, and the normal copula/beta model parameter vector θ for the collateral pool. Given θ, any number of relevant metrics of the credit risk associated with a CDO note can be computed. The next section examines how three commonly used metrics of credit risk depend on θ.

4 Sensitivity of credit risk metrics to model parameters

This analysis considers three standard metrics of credit quality: probability of default, expected loss, and conditional expected loss. Define the expectation operator $\mathrm{E}[Z]$ as the expected value of the random variable Z whose distribution is determined by θ. Let V be the value of a one dollar investment in a debt security at the terminal date and let r be the contractual interest on that security. The security's probability of default is defined as

$$PD = \mathrm{E}\left[1\left\{V < 1 + r\right\}\right]. \tag{1}$$

PD describes the likelihood of a credit loss, but not the magnitude of the loss. The expected loss

$$EL = (1 + r) - \mathrm{E}[V] \tag{2}$$

summarizes the expected likelihood and the magnitude of a credit loss. Note that EL may exceed 100 percent because both principal and accrued interest may be lost.

PD and EL describe the first moments of a security's loss distribution. In portfolio risk management applications such as economic capital allocation, analysts also require information about a security's marginal contribution to portfolio-wide losses. A number of risk metrics useful for describing the dependence between an individual exposure's credit losses and those of a broader portfolio have been proposed in the risk management literature, and I do not propose to survey them here. This analysis will consider one such measure derived from an asymptotic single risk factor approximation. Gordy (2003) shows that if a portfolio is well diversified and its overall loss rate depends on a single systematic factor $\bar{X}$ then an exposure's marginal contribution to portfolio value-at-risk (VaR) can be determined analytically by calculating its conditional expected loss of the exposure given an adverse draw of the systematic risk factor, The conditional expected loss associated with a qth percentile portfolio VaR measure is

$$EL_q = (1+r) - \mathrm{E}\left[V \mid \bar{X} = \bar{x}_q\right] \tag{3}$$

where $\bar{x}_q$ is the $1-q$th percentile of the stochastic systematic risk factor. Unlike PD and EL, which describe the center of the distribution of V, EL_q describes the tail of this distribution.[5]

Under the normal copula/beta model, EL, PD, and EL_q for both simple and structured bonds are determined by the parameter vector θ. For simple bonds,

$$PD = \pi$$

and

$$EL = (1+r)\pi\mu,$$

If we assume that correlation between the single systematic factor underlying a financial institution's overall asset portfolio $\bar{X}$ and the systematic factor that affects bond default rates X is 50 percent, the conditional expected loss for a simple bond is

$$EL_q = (1+r)\Phi\left(\frac{\Phi^{-1}(\pi) - \sqrt{0.5\rho}\,\bar{x}_q}{\sqrt{1-0.5\rho}}\right)\mu$$

where $\bar{x}_q = \Phi^{-1}(1-q)$.[6]

For simple bonds, PD is determined by the normal copula marginal default probability parameter π, EL depends on both π and the expected loss-given-default μ, and EL_q is a function of π, μ, and the asset value correlation parameter ρ. For structured bonds, simple analytic formulas for PD, EL, and EL_q are not available, but these risk metrics can be computed numerically for any value of θ. In contrast to the case for simple bonds, PD, EL, and EL_q for structured bonds each depend on all four elements of θ.

[5] If a portfolio is not fully diversified or if multiple systematic factors are present (e,g., sector specific factors) then a given exposure's marginal contribution to portfolio VaR depends on all exposures within the portfolio. In these cases, portfolio-specific measures of marginal VaR contributions can be derived. See Pykhtin (2003) for an analytic approach and Heitfield, Burton and Chomsisengphet (2006) for a simulation-based approach to computing marginal VaR contributions in the presence of multiple risk factors.

[6] See Pykhtin and Dev (2002) for detailed analysis of conditional expected loss in the presence of a single portfolio-wide systematic factor and a correlated CDO collateral-pool-specific factor.

Table 1: Credit risk statistics for hypothetical CDO deal backed by 100 mezzanine-rated bonds.

Tranche	Position (%)	Spread (bp)	PD (%)	EL (%)	$EL_{0.95}$ (%)
Junior	3 – 6	224	16.78	11.86	45.88
Jr. Mezz.	6 – 9	84	6.24	4.29	19.64
Sr. Mezz.	9 – 12	34	2.54	1.73	8.04
Senior	12 – 15	14	1.07	0.71	3.07
Sup. Sen.	15 – 100	0	0.43	0.01	0.05
Collateral		56	5.00	3.45	8.14

The true value of the normal copula/beta model parameters, which will be denoted with the subscript "0", cannot be directly observed by a credit analyst. In practice, model parameters are either determined judgmentally or estimated from historical data. Any differences between θ_0 and the value of θ used to compute the expectations in equations (1), (2), and (3) can result in errors in imputed risk metrics.

To illustrate how errors in θ can affect imputed risk metrics, this paper examines two hypothetical CDO deals summarized in Tables 1 and 2. Both CDO deals are backed by homogeneous collateral pools of 100 simple bonds, described in the bottom rows of the tables. In the mezzanine CDO example $\pi_0 = 0.05$, $\rho_0 = 0.20$, $\mu_0 = 0.55$, and $\sigma_0 = 0.35$ and bonds have a maturity of 5 years. In the high-grade CDO example, all normal copula/beta parameters are the same except the default probability parameter, which is set at $\pi_0 = 0.01$. The default probability parameter for the mezzanine and high-grade cases roughly corresponds to the five-year cumulative default rate for corporate bonds rated Baa- and A- respectively by Moody's. The collateral's asset value correlation and LGD parameters are chosen to be broadly consistent with those used in major rating agencies' or regulators' CDO credit risk models.[7] A constant risk-free interest rate of four percent is assumed, and spreads are set so that a risk-neutral pricing model based on θ_0 would value the bonds at par at origination.

As is typical of securitization deals, credit risk metrics for the various tranches of the CDO bear little direct relation to those of the underlying collateral pool but are very sensitive to each tranche's position in the deal capital structure. Tranches lower in the capital structure, which are the first to take losses, have much higher default probabilities, expected losses, and conditional expected losses than more senior tranches.

Figures 1 through 8 show how deviations in each component of θ from θ_0 affects PD, EL, and EL_q. Each line in a panel plots the ratio of a particular risk metric to its true value (i.e., the risk metric computed given $\theta = \theta_0$) as θ changes. For example, the first panel of Figure 1 shows how bonds' implied default probabilities change with π holding all other elements of θ fixed at their true values. All lines cross at the true value of the model parameter in

[7] For example, Basel II uses corporate bond asset correlation parameters ranging from 0.12 to 0.24 depending on firm size, and baseline unsecured loss given default parameter of 0.45 (Basel Committee on Banking Supervision 2004).

Table 2: Credit risk statistics for hypothetical CDO deal backed by 100 high-grade bonds.

Tranche	Position (%)	Spread (bp)	PD (%)	EL (%)	$EL_{0.95}$ (%)
Junior	1 – 2	183	13.24	9.56	35.17
Jr. Mezz.	2 – 4	57	5.41	2.88	12.84
Sr. Mezz.	4 – 6	15	1.29	0.75	3.35
Senior	6 – 8	5	0.39	0.23	0.90
Sup. Sen.	8 – 100	0	0.12	0.00	0.01
Collateral		11	1.00	0.68	1.92

question.

As these results plainly show, structured bonds are considerably more sensitive to specification errors in each component of θ than simple bonds, and the higher is a bond's position in the CDO capital structure the greater is its sensitivity to parameter errors. Moreover, a risk metric's sensitivity to parameter errors is inversely proportional to its relevance for portfolio risk management; conditional expected loss is more sensitive to errors in θ than expected loss, and expected loss is more sensitive to specification errors than probability of default.

For more senior tranches, risk metrics are remarkably sensitive to model parameters. For example, if a credit risk manager were to use an asset value correlation parameter of 15 percent to evaluate the "Senior" tranche of the high-grade CDO in an environment where the true correlation was 20 percent, she would understate the tranche's default probability and expected loss by over 50 percent, and the tranche's conditional expected loss (equivalent to its marginal capital requirement) by more than 75 percent. Conversely, if she used an asset correlation parameter of 25 percent, she would overstate the tranche's default probability and expected loss by 100 percent and its conditional expected loss by 200 percent.

5 Bounding the accuracy of model parameters

The sensitivity of structured bond risk metrics to θ begs the question, how accurately can model parameters be estimated? The answer, of course, depends on the type and volume of historical data available and the statistical estimator used. Practitioners typically use historical data on the default frequencies of rated bonds to estimate default probability parameters. Factor correlation parameters can also be calibrated from such data but, as Gordy and Heitfield (2002) show, accurately estimating correlation parameters from default rate data requires long data histories and/or strong *ex ante* parameter restrictions. As noted earlier, the latent credit factors of the normal copula model have a structural interpretation as obligor asset returns in a Merton (1974) valuation framework. Leveraging this fact, practitioners often use information on equity returns or imputed asset returns for publicly traded firms as proxies for latent credit factors. Alternatively, some practitioners use information on bond rating transitions (which are far more common than bond defaults) to estimate

Figure 1: Sensitivity of mezzanine CDO risk metrics to errors in the default probability parameter π.

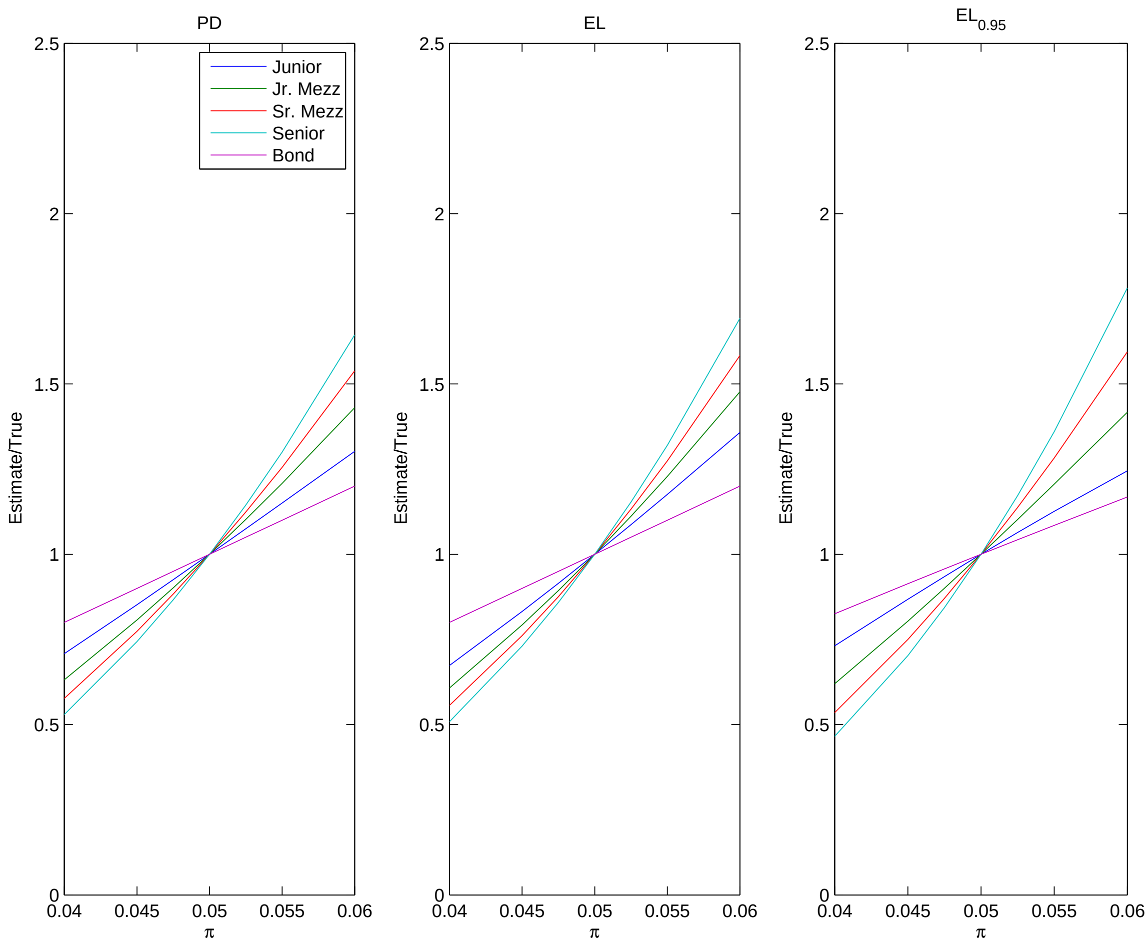

Figure 2: Sensitivity of mezzanine CDO risk metrics to errors in the asset value correlation parameter ρ.

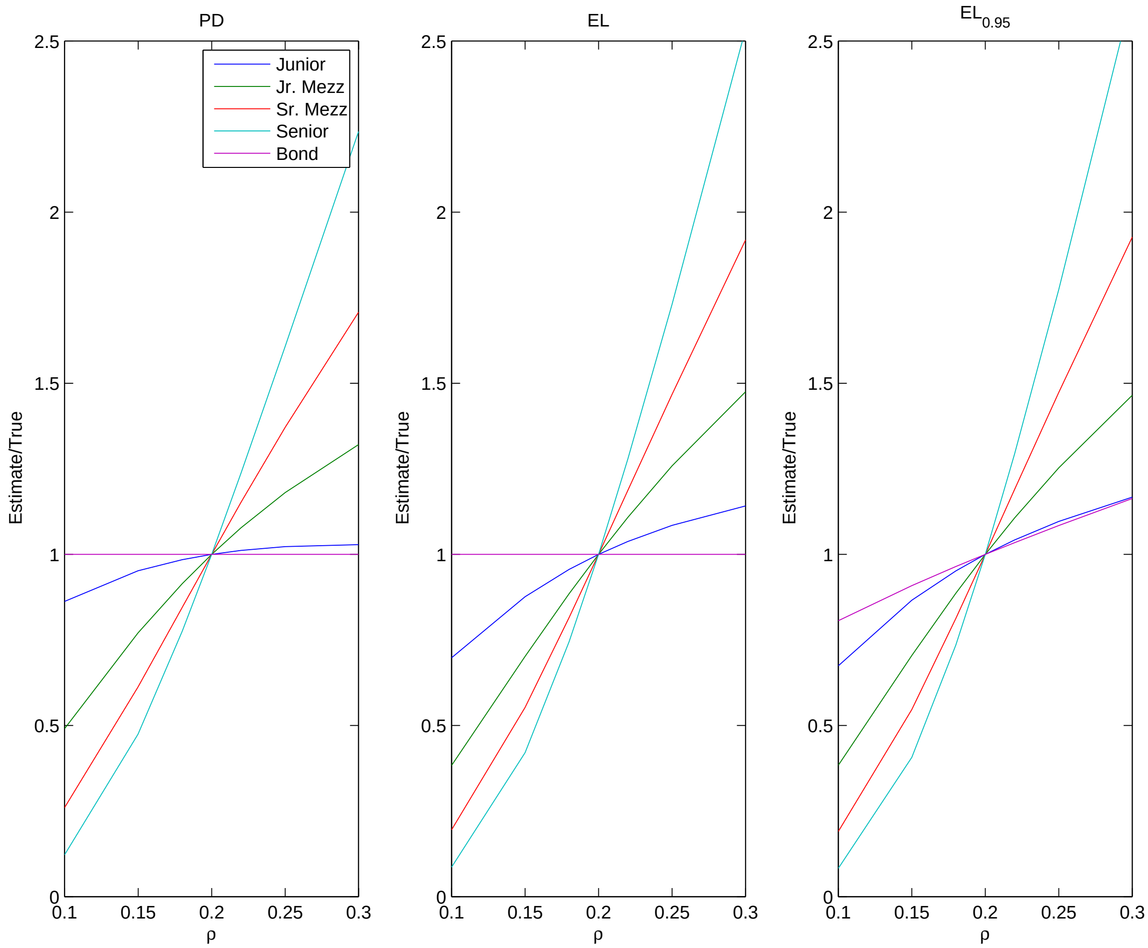

Figure 3: Sensitivity of mezzanine CDO risk metrics to errors in the expected loss given default parameter μ.

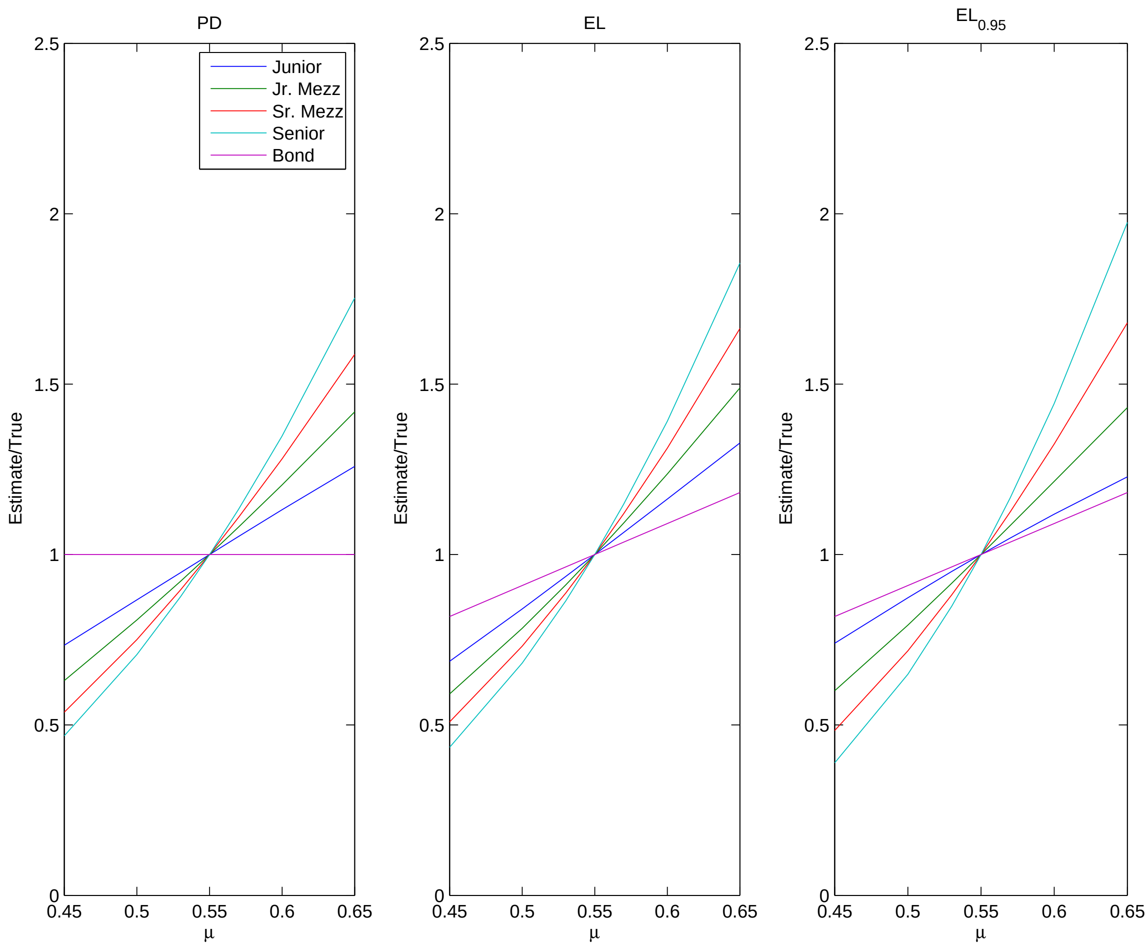

Figure 4: Sensitivity of mezzanine CDO risk metrics to errors in the loss-given-default volatility parameter σ.

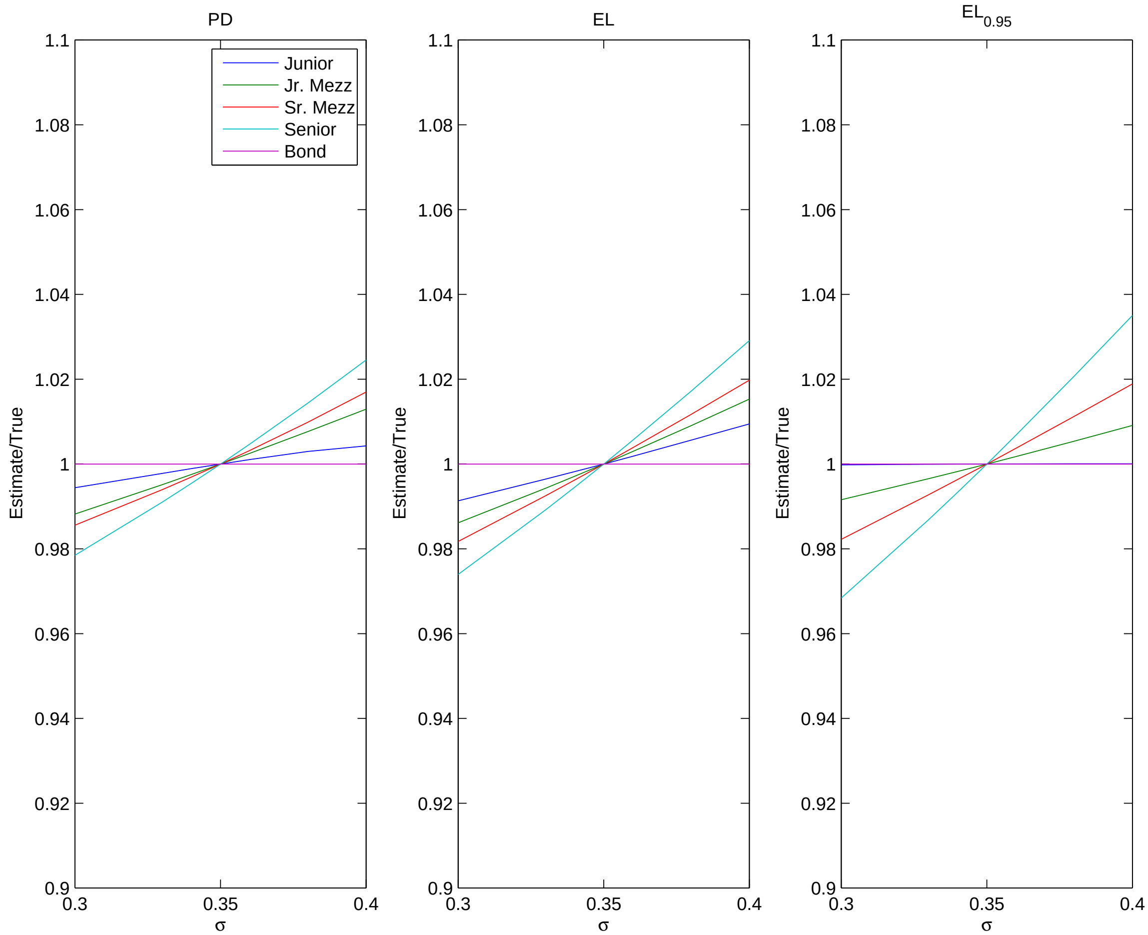

Figure 5: Sensitivity of high-grade CDO risk metrics to errors in the default probability parameter π.

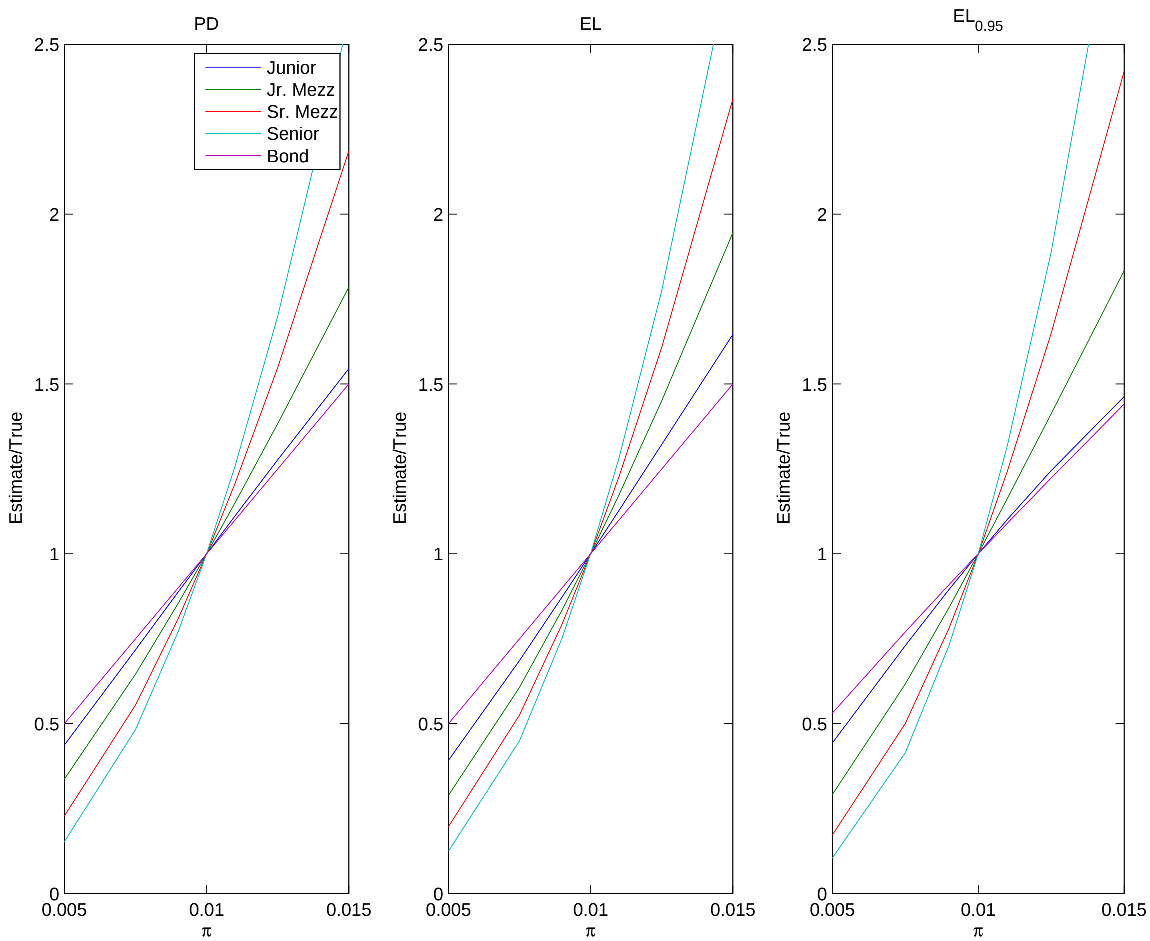

Figure 6: Sensitivity of high-grade CDO risk metrics to errors in the asset value correlation parameter ρ.

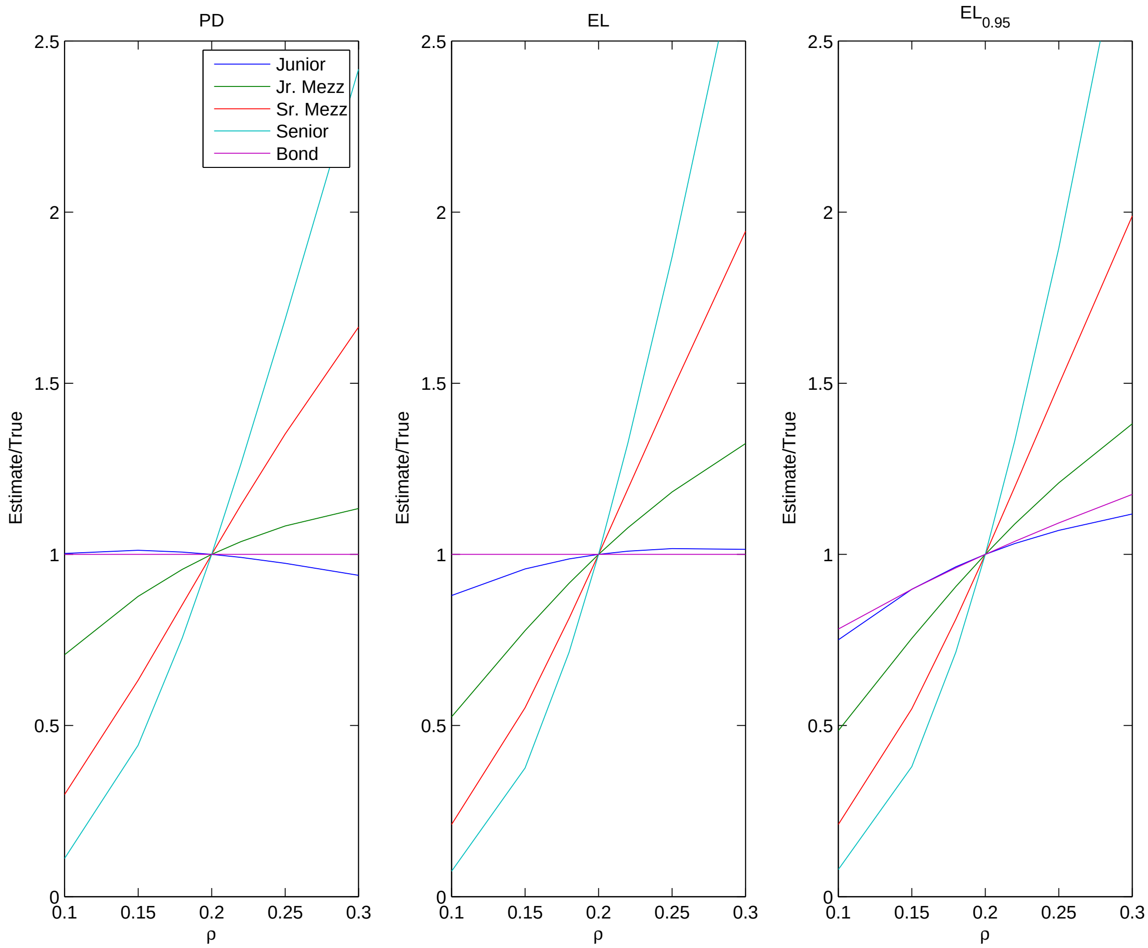

Figure 7: Sensitivity of high-grade CDO risk metrics to errors in the expected loss given default parameter μ.

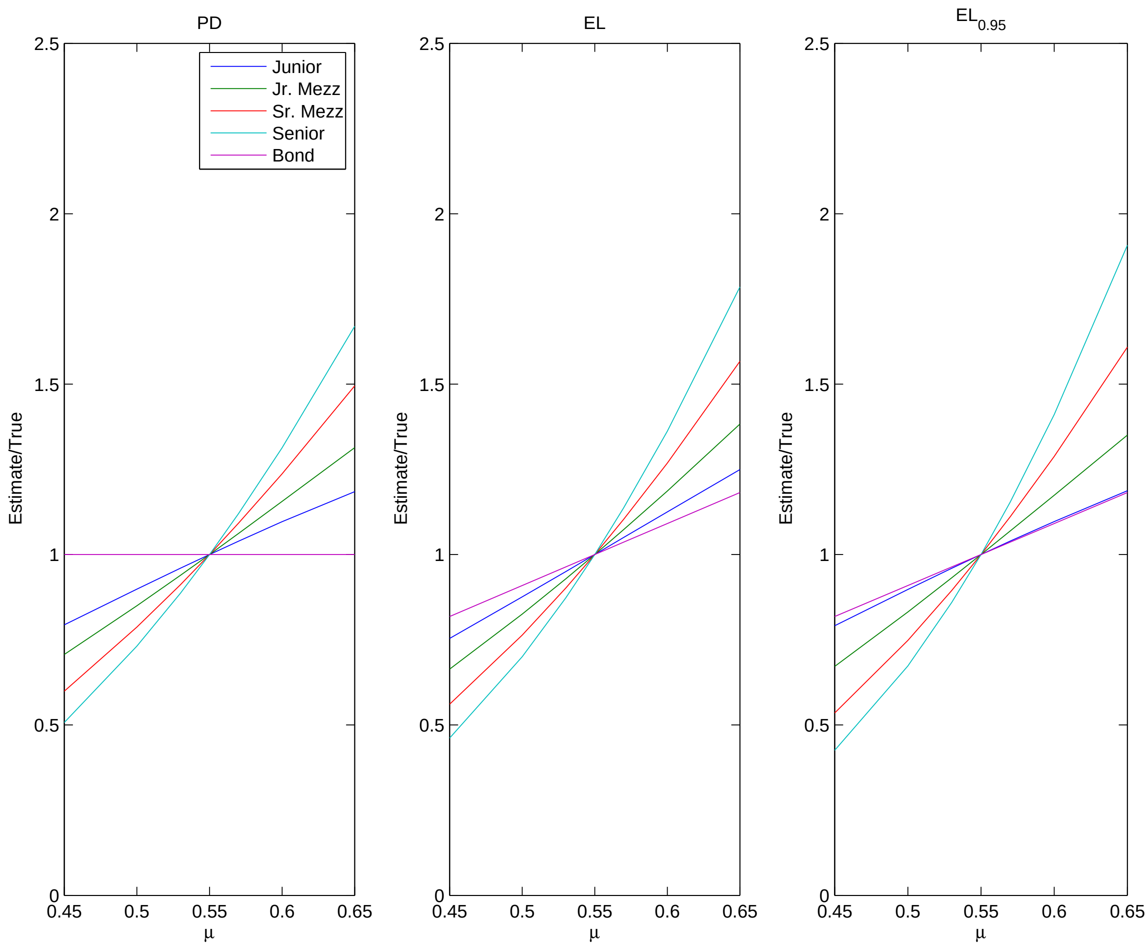

Figure 8: Sensitivity of high-grade CDO risk metrics to errors in the loss-given-default volatility parameter σ.

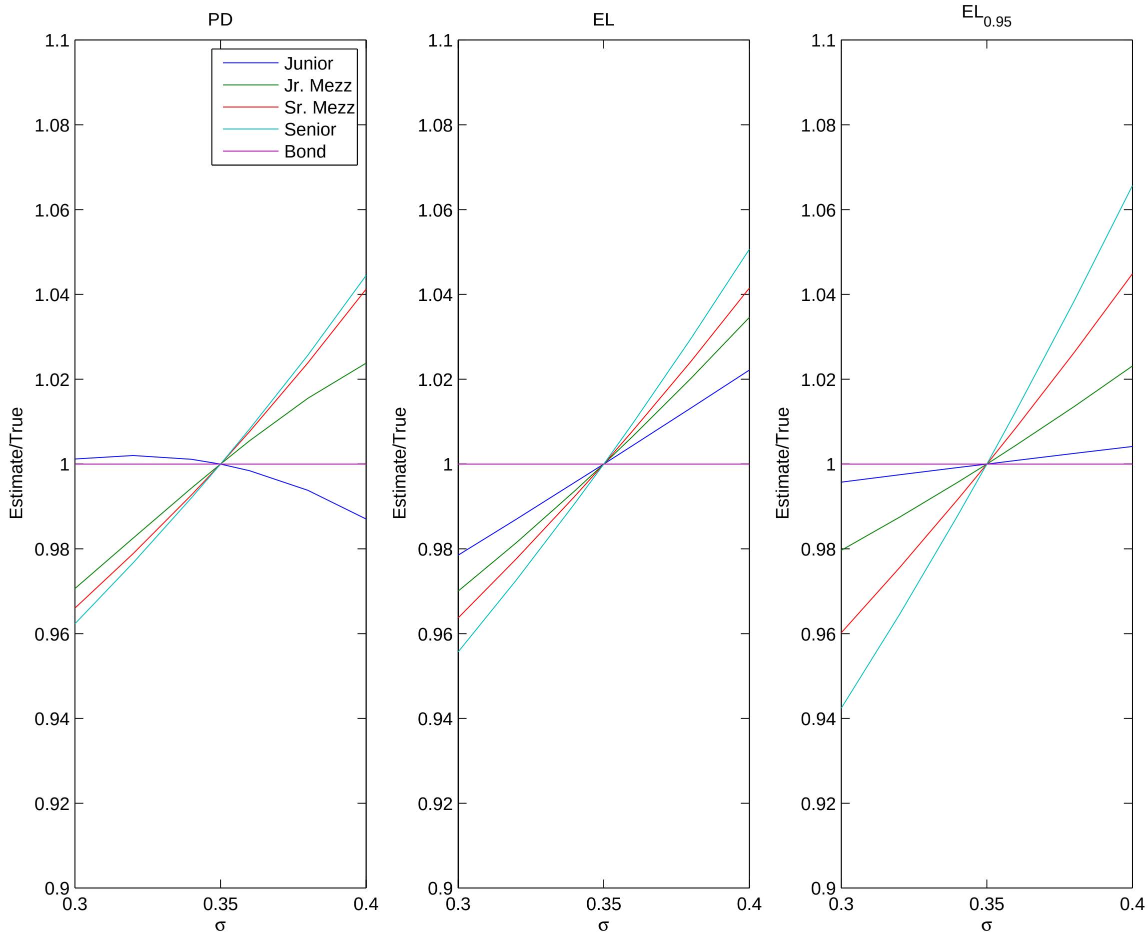

factor correlations. Loss given default parameters are typically estimated from data on the market price of traded debt shortly after a default event or ultimate recoveries on defaulted bonds or loans.

In this section, I present a stylized data generating process (DGP) capable of describing the range of data that might reasonably be available to a credit analyst. I assume that the normal copula/beta model is correctly specified in the sense that the DGP is consistent with this model given a true parameter vector θ_0. Using the DGP, I derive lower bounds on the sampling variance of any unbiased estimator of θ_0. This allows one to investigate how the quality and character of available historical data affects the accuracy of normal copula parameter estimates.

Assume that an analyst observes T cohorts which each contain N simple bonds. All bonds in all cohorts share the same underlying model parameter vector θ_0. Within a cohort, all bonds are sensitive to the same systematic risk factor X, but systematic risk factors are assumed to be independent across cohorts. Such data could arise, for example if one observed information on T non-overlapping cohorts of bonds over time.

For each bond i in cohort t, the analyst observes the following information *ex post*. First the analyst observes an indicator variable D_{it} which is equal to one if the bond defaults (i.e, if $Y_{it} \leq \Phi^{-1}(\pi_0)$) and zero otherwise. Second, if bond it defaults, the analyst observes the realized loss rate given default λ_{it}. Finally, if the bond does not default, the analyst observes a noisy signal of the bond's realized latent credit factor Y_{it}, denoted Y^*_{it}. Y^*_{it} is a weighted sum of Y_{it} and a standard normal error term U_{it} which is assumed to be independent of all other variables in the model:

$$Y^*_{it} = \sqrt{\omega}\left(\sqrt{1-\psi}Y_{it} + \sqrt{\psi}U_{it}\right).$$

The parameter ψ lies on the unit interval and captures the relative amount of noise in the signal Y^*_{it}. In the limiting case where $\psi = 0$, Y^*_{it} is perfectly correlated with the realized credit factor Y_{it}. At the other extreme, where $\psi = 1$, Y^*_{it} provides no information about Y_{it}. In this case Y^*_{it} can be safely ignored by the analyst, and inference about θ must be based solely on information about whether or not bonds have defaulted. ω controls the scale of Y^*_{it}. The nuisance parameters ω and ψ are not of direct interest to the analyst, but they must be estimated in order to use information from Y^*_{it} to estimate π and ρ.[8] For the remainder of this paper we will redefine θ to include these nuisance parameters.

In applied settings, our stylized assumption that historical data can be grouped into independent, homogeneous cohorts could be easily generalized in useful ways. For example, a more realistic DGP might allow for correlations across cohort-specific systematic factors since in practice cohorts may well be overlapping in time. Within cohort heterogeneity, particularly with respect to π_o, could also be accommodated as in Gordy and Heitfield (2002). This would allow one to investigate the costs and benefits of pooling historical data across different types of credit exposures. However, for our current purposes, relatively stylized DGP assumptions are useful, because they allow us to investigate how broad features of available data affect parameter estimates.

[8] An exception occurs in the limiting cases where $\psi_0 = 0$ or $\psi_0 = 1$. When $\psi_0 = 0$ there is no uncertainty about the information content of Y^*_{it} and hence no need to estimate ψ_0. In this case it is still necessary to estimate the scale parameter ω_0 When $\psi_0 = 1$, Y^*_{it} is irrelevant to inference of θ, so there is no need to estimate either ψ or ω.

Conditional on X, the likelihood of observing the vector $(D_{it}, Y_{it}^*, \lambda_{it})$ can be expressed as the product of a marginal likelihood for (D_{it}, Y_{it}^*) given X which depends on π and ρ and a conditional likelihood for λ_{it} given D_{it} that depends on μ and σ:

$$f\left(D_{it}, Y_{it}^*, \lambda_{it}|X;\theta\right) = f\left(D_{it}, Y_{it}^*|X;\pi,\rho,\psi,\omega\right) f\left(\lambda_{it}|D_{it};\mu;\sigma\right) \tag{4}$$

The first right-hand term is

$$\begin{aligned} f\left(D_{it}, Y_{it}^*|X;\pi,\rho,\psi,\omega\right) = & \left[\Phi\left(\frac{\Phi^{-1}(\pi)-\sqrt{\rho}X}{\sqrt{1-\rho}}\right)\right]^{D_{it}} \\ & \times\left[\phi\left(\frac{(Y_{it}^*/\sqrt{\omega})-\sqrt{1-\psi}\sqrt{\rho}X}{\sqrt{S}}\right)\right]^{1-D_{it}} \\ & \times\left[\Phi\left(-\frac{S\Phi^{-1}(\pi)-\psi\sqrt{\rho}X-(1-\rho)\sqrt{1-\psi}\left(Y_{it}^*/\sqrt{\omega}\right)}{\sqrt{S}\sqrt{\psi(1-\rho)}}\right)\right]^{1-D_{it}} \end{aligned} \tag{5}$$

where $S = (1-\psi)(1-\rho)+\psi$. See the appendix for a derivation of (5). The second term in (4) is the likelihood of the loss rate λ_{it}, which is only observable if $D_{it} = 1$. It can be written

$$f\left(\lambda_{it}|D_{it};\mu,\sigma\right) = \left[\beta\left(\lambda_{it}\right)\right]^{D_{it}}$$

where $\beta(z)$ is the PDF for a beta-distributed random variable with mean and standard deviation parameter μ and σ respectively. The joint likelihood for cohort t is

$$L_t(\theta) = \int\prod_{i=1}^{N} f\left(D_{it}, Y_{it}^*|x;\pi,\rho\right)\phi(x)dx\prod_{i=1}^{N} f\left(\lambda_{it}|D_{it};\mu,\sigma\right)$$

Given the likelihood of the DGP, the well known Cramer-Rao lower bound defines the minimum covariance matrix of any asymptotically unbiased estimator for θ_0. The Fischer information matrix for our DGP is

$$I\left(\theta_0\right) = \mathrm{E}\left[\frac{\partial^2 \ln L_t(\theta_0)}{\left(\partial\theta\right)^2}\right].$$

If an estimator $\hat{\theta}$ is asymptotically unbiased, then the difference between its covariance matrix and $\frac{1}{T}\left(I\left(\theta_0\right)\right)^{-1}$ is a positive semi-definite matrix. Among other things, this implies that the variance of each element of $\hat{\theta}$ is at least as large as the corresponding diagonal element of the average of the inverse information matrix. If we posit a value of θ_0, the Cramer-Rao bound can be computed directly for any combinations of T, N, and ψ_0.

Because the marginal distribution of D_{it} does not depend on μ and σ and the conditional distribution of λ_{it} does not depend on π, ρ, ψ, and ω, the information matrix for θ is block diagonal with respect to (π,ρ,ψ,ω) and (μ,σ).[9] As a result, we can investigate the accuracy of these two subvectors separately. For the remainder of this paper we will focus on the first subvector, though the methods described here are equally applicable to analysis of μ and σ.

Figures 9 and 10 show how the minimum standard deviation of π and ρ vary as N, T, and ψ change. Several useful results can be gleaned from this analysis.

[9]Note that this result does not hold for the more general case where loss given default is not independent of the systematic factor X.

- As one might expect sampling errors for π and ρ tend to be higher relative to the true parameter values for populations with low default rates. It is more difficult to estimate the frequency and volatility of low probability events with precision.

- The figures dramatically illustrate the value of observing high quality historical data on the credit performance of those bonds that do not default. Comparing results for $\psi_0 = 1.0$ and $\psi_0 = 0.0$ we see that observing credit factors for non-defaulted firms reduces standard errors by about two-thirds for the mezzanine bond population and by about three-quarters for the high-grade bond population. Indeed, for high-grade bond populations, estimating ρ with reasonable precision virtually requires some type of credit factor data. Interestingly, these results also show that the accuracy of default probability parameters can be greatly improved by incorporating information on latent credit factors. Given that default probability parameters are commonly calibrated using only long-run default frequencies, this fact does not appear to be widely appreciated.

- Cohort size is more important when defaults are rare (Figure 10) than when they are relatively common (Figure 9). For example, increasing N from 50 to 200 reduces the standard deviation of ρ by about one-third for the mezzanine bond population, but by over half for the high-grade bond population. For the high-grade population increasing N and T have about the same effect on sampling errors. To see this, compare the points for $N = 50$ and $T = 40$ with those for $N = 200$ and $T = 10$).

6 Distribution of credit risk metrics

Section 4 shows how standard measures of credit risk depend on estimated normal copula model parameters, and Section 5 shows how the accuracy of parameter estimates depends on the data available to an analyst. This section combines these results to examine how the distribution of PD, EL and EL_q are affected by the characteristics of the data generating process.

Given the sampling distribution for an estimator $\hat{\theta}$ of θ_0, the distribution of PD, EL, and EL_q for simple and structured bonds can be estimated using Monte Carlo simulation. Unfortunately, while the Fischer information inequality allows one to bound the sampling variance of unbiased estimators, it provides no additional insights into the small sample properties of such estimators. Hence, in order to simulate the distribution of PD, EL, and EL_q, we need to make some assumptions about the sampling distribution of $\hat{\theta}$.

π and ρ, are bounded between zero and one, so the natural assumption that $\hat{\theta}$ is drawn from a multivariate normal distribution is inappropriate. One can circumvent this problem by reparameterizing the normal copula model in such a way that model parameters are defined over the entire real line. Let $\tilde{\pi} = \Phi^{-1}(\pi)$ and $\tilde{\rho} = \Phi^{-1}(\rho)$. $(\tilde{\pi}, \tilde{\rho})$ has support $\Re^2$, and the sampling variance for a minimum variance unbiased estimator of this vector can be derived by inverting the Fischer information matrix for the reparameterized likelihood function.

Figure 9: Minimum standard deviation of unbiased normal copula parameter estimators for mezzanine bonds under various data generating processes ($\pi_0 = 0.05$, $\rho_0 = 0.20$).

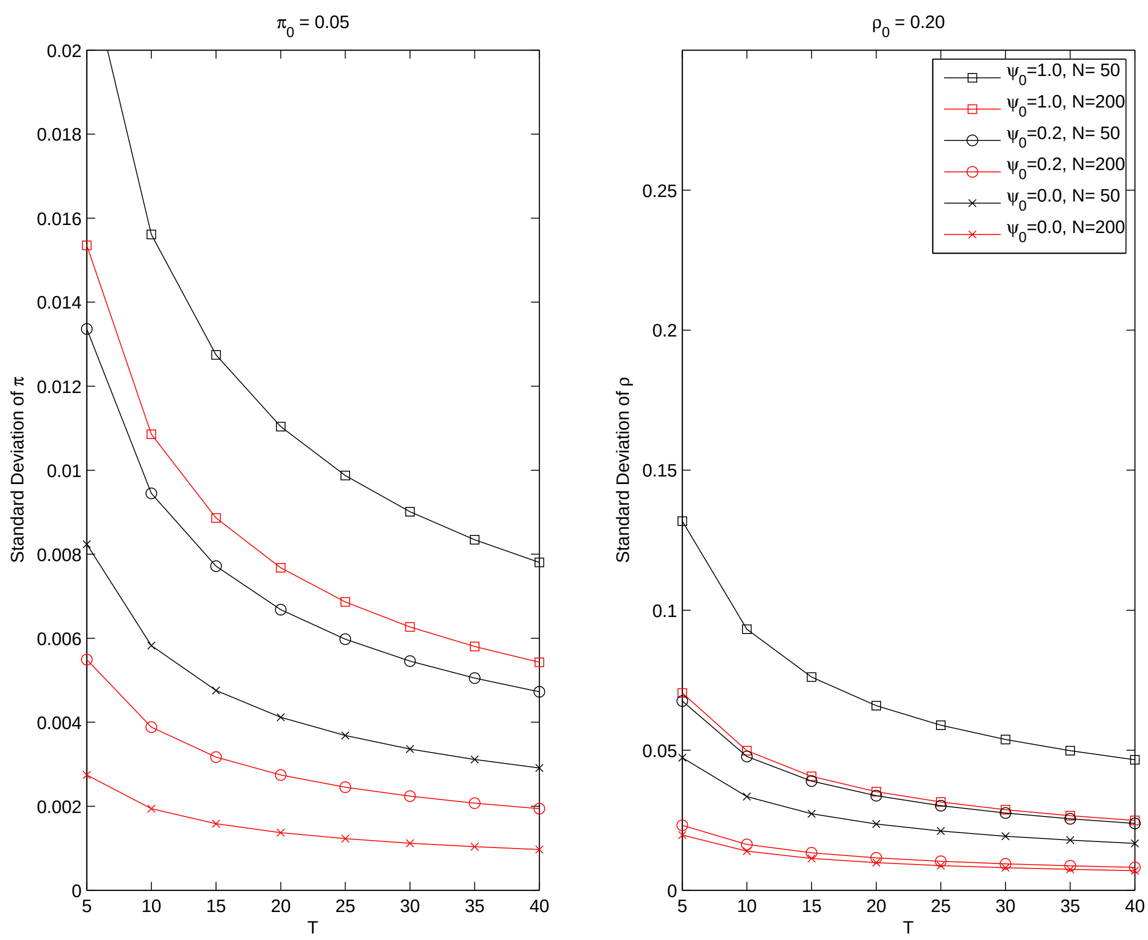

Figure 10: Minimum standard deviation of unbiased normal copula parameter estimators for high-grade bonds under various data generating processes ($\pi_0 = 0.01$, $\rho_0 = 0.20$).

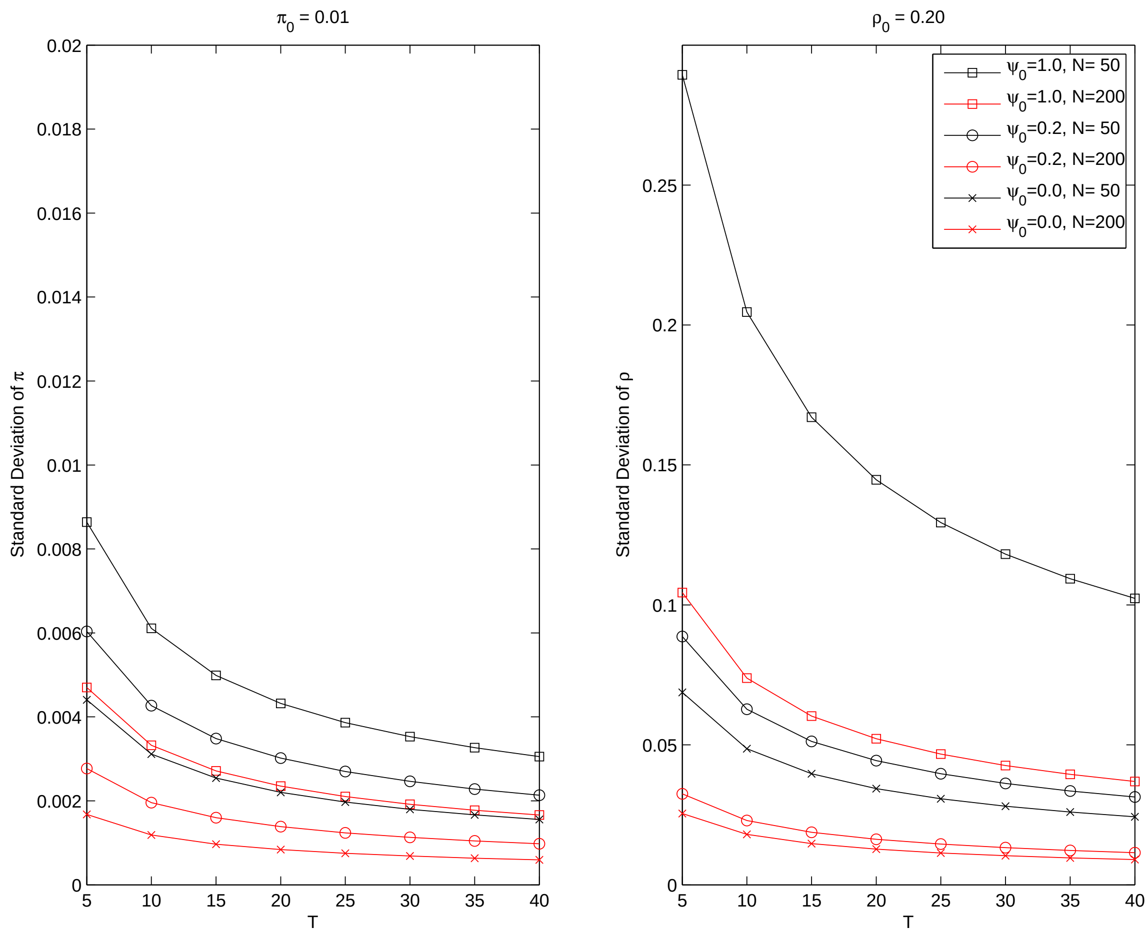

The sampling distributions of PD, EL, and EL_q are simulated as follows. For each Monte Carlo iteration I draw a value of $(\tilde{\pi}, \tilde{\rho})$ from a bivariate normal distribution with mean $(\tilde{\pi}_0, \tilde{\rho}_0)$ and covariance matrix equal to the Cramer-Rao bound. Using this parameter value, I compute implied PD, EL, and EL_q for simple bonds and CDO tranches. Two-thousand Monte Carlo iterations are run for each of four hypothetical data generating processes:

- a "wide" panel with partially observable credit factors ($N = 200$, $T = 10$, $\psi = 0.2$);
- a "long" panel with partially observable credit factors ($N = 50$, $T = 40$, $\psi = 0.2$);
- a "wide" panel with unobservable credit factors ($N = 200$, $T = 10$, $\psi = 1.0$); and
- a "long" panel with unobservable credit factors: ($N = 50$, $T = 40$, $\psi = 1.0$).

In each case, simulations are run for both the mezzanine and the high-grade CDO deals described in Section 4.

Simulation results are reported in Tables 3 and 4 and Figures 11 through 18. Several conclusions can be drown from these simulations.

- Standard measures of credit risk such as probability of default, expected loss, and conditional expected loss for more senior CDO tranches have significantly larger sampling errors than those for more junior tranches or unstructured bonds.
- Standard measures of credit risk for CDOs backed by high-grade securities are likely to have larger sampling errors than those backed by mezzanine-grade securities.
- Long, narrow data panels generally produce more accurate risk metrics than short, wide panels with comparable numbers of observations, particularly when no data on latent credit factors are available.
- Proxy data for latent credit factors can significantly improve the accuracy of risk metrics for both unstructured and structured securities. Notably, such data significantly improves the accuracy of PD ane EL estimates for unstructured securities, even though these credit risk metrics do not depend on asset correlation parameters.

Overall, the results presented here suggest that statements about the credit quality of senior structured securities should be viewed with considerable skepticism, particularly when those securities are backed by high-grade collateral. Normal copula models provide a valuable tool for ranking the relative risks of similar classes of structured securities, but they may be less useful for making comparisons across different types of credit products.

Table 3: 90% confidence intervals for risk metrics of a mezzanine CDO.

	PD (%)			EL (%)			$EL_{0.95}$ (%)		
Tranche	True	5th	95th	True	5th	95th	True	5th	95th
$N = 50$, $T = 40$, $\psi_0 = 0.2$									
Junior	16.78	12.59	20.91	11.86	8.15	15.70	45.88	33.53	57.24
Jr. Mezz.	6.24	3.80	9.02	4.29	2.40	6.68	19.64	11.17	29.61
Sr. Mezz.	2.54	1.24	4.38	1.73	0.76	3.22	8.04	3.45	14.99
Senior	1.07	0.41	2.18	0.71	0.24	1.56	3.07	0.97	7.14
Collateral	5.00	4.27	5.81	3.45	2.95	4.01	8.14	6.73	9.72
$N = 200$, $T = 10$, $\psi_0 = 0.2$									
Junior	16.78	12.19	21.51	11.86	7.81	16.28	45.88	32.21	58.61
Jr. Mezz.	6.24	3.54	9.42	4.29	2.17	6.97	19.64	10.14	30.81
Sr. Mezz.	2.54	1.11	4.57	1.73	0.68	3.35	8.04	3.03	15.63
Senior	1.07	0.36	2.31	0.71	0.21	1.69	3.07	0.82	7.69
Collateral	5.00	4.18	5.93	3.45	2.89	4.10	8.14	6.59	9.91
$N = 50$, $T = 40$, $\psi_0 = 1.0$									
Junior	16.78	10.47	24.11	11.86	6.70	18.22	45.88	28.14	63.40
Jr. Mezz.	6.24	2.95	10.75	4.29	1.72	8.00	19.64	8.02	34.86
Sr. Mezz.	2.54	0.81	5.38	1.73	0.47	4.12	8.04	2.06	19.13
Senior	1.07	0.22	2.94	0.71	0.12	2.18	3.07	0.47	9.93
Collateral	5.00	3.79	6.48	3.45	2.61	4.48	8.14	6.15	10.60
$N = 200$, $T = 10$, $\psi_0 = 1.0$									
Junior	16.78	10.38	28.56	11.86	7.10	20.59	45.88	30.97	64.42
Jr. Mezz.	6.24	3.54	10.98	4.29	2.33	7.63	19.64	10.86	32.14
Sr. Mezz.	2.54	1.24	4.69	1.73	0.77	3.21	8.04	3.44	14.89
Senior	1.07	0.42	2.12	0.71	0.25	1.48	3.07	0.99	6.71
Collateral	5.00	3.50	7.22	3.45	2.42	4.99	8.14	6.44	10.25

Table 4: 90% confidence intervals for risk metrics of a high-grade CDO.

Tranche	PD (%)			EL (%)			$EL_{0.95}$ (%)		
	True	5th	95th	True	5th	95th	True	5th	95th
$N = 50$, $T = 40$, $\psi_0 = 0.2$									
Junior	13.24	8.84	18.59	9.56	5.93	14.32	35.17	23.20	48.79
Jr. Mezz.	5.41	3.05	8.78	2.88	1.47	5.17	12.84	6.55	22.26
Sr. Mezz.	1.29	0.54	2.69	0.75	0.28	1.70	3.35	1.17	7.90
Senior	0.39	0.12	1.02	0.23	0.06	0.67	0.90	0.22	2.88
Collateral	1.00	0.71	1.39	0.68	0.48	0.94	1.92	1.37	2.66
$N = 200$, $T = 10$, $\psi_0 = 0.2$									
Junior	13.24	8.43	18.31	9.56	5.55	14.10	35.17	21.30	49.18
Jr. Mezz.	5.41	2.72	8.80	2.88	1.21	5.26	12.84	5.43	23.26
Sr. Mezz.	1.29	0.41	2.91	0.75	0.20	1.91	3.35	0.84	8.92
Senior	0.39	0.07	1.16	0.23	0.03	0.79	0.90	0.12	3.43
Collateral	1.00	0.69	1.38	0.68	0.46	0.93	1.92	1.29	2.74
$N = 50$, $T = 40$, $\psi_0 = 1.0$									
Junior	13.24	6.40	24.24	9.56	4.45	18.16	35.17	18.35	55.23
Jr. Mezz.	5.41	2.25	10.93	2.88	1.00	6.24	12.84	4.36	26.39
Sr. Mezz.	1.29	0.30	3.32	0.75	0.13	2.08	3.35	0.56	9.61
Senior	0.39	0.04	1.25	0.23	0.02	0.89	0.90	0.06	3.75
Collateral	1.00	0.54	1.70	0.68	0.37	1.15	1.92	1.16	2.99
$N = 200$, $T = 10$, $\psi_0 = 1.0$									
Junior	13.24	7.52	20.74	9.56	5.14	15.61	35.17	19.64	51.39
Jr. Mezz.	5.41	2.51	9.52	2.88	1.03	5.74	12.84	4.55	24.96
Sr. Mezz.	1.29	0.28	3.19	0.75	0.12	2.22	3.35	0.51	10.47
Senior	0.39	0.03	1.49	0.23	0.01	1.06	0.90	0.05	4.59
Collateral	1.00	0.63	1.51	0.68	0.42	1.02	1.92	1.23	2.89

Figure 11: Distribution of estimated mezzanine CDO risk metrics given a long data panel with well observed credit factors ($N = 50$, $T = 40$, $\psi = 0.2$).

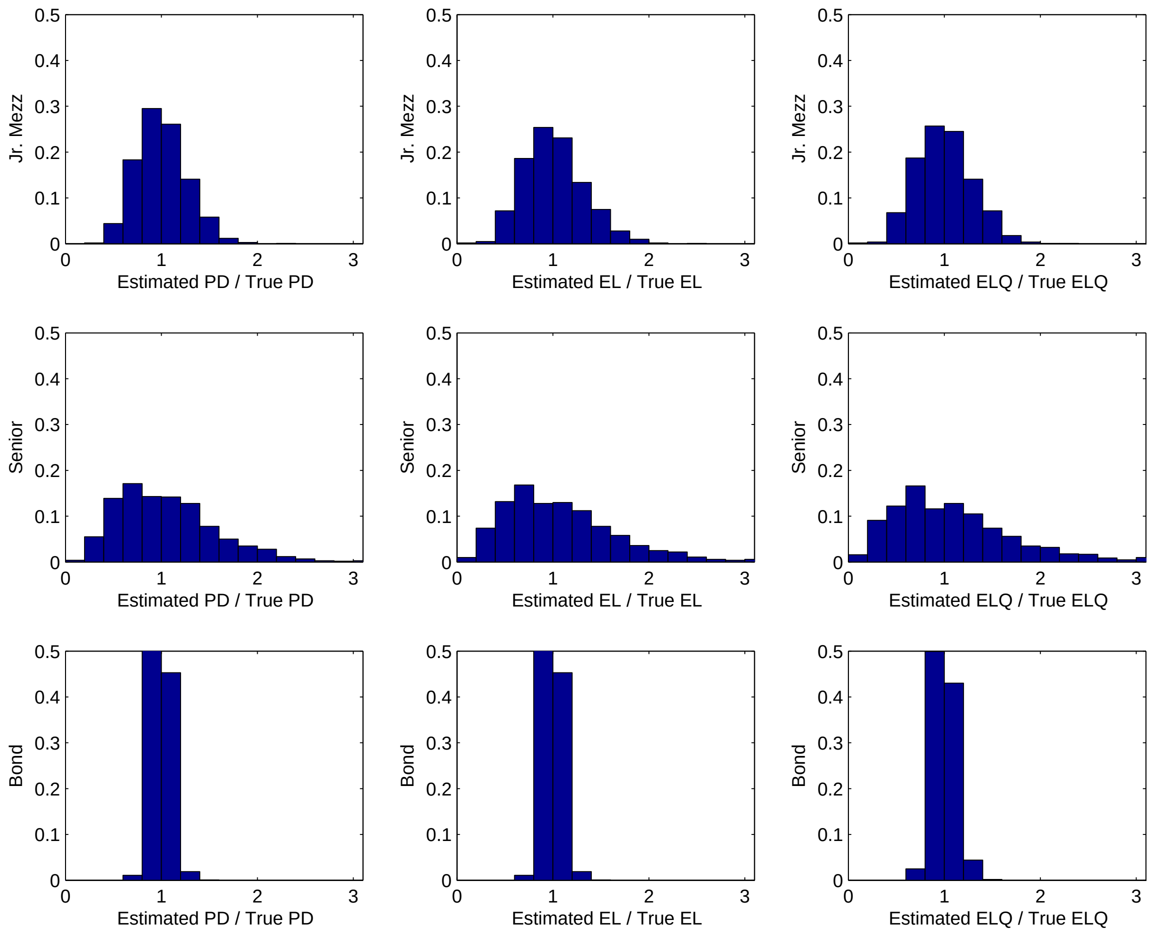

Figure 12: Distribution of estimated mezzanine CDO risk metrics given a wide data panel with well observed credit factors ($N = 200$, $T = 10$, $\psi = 0.2$).

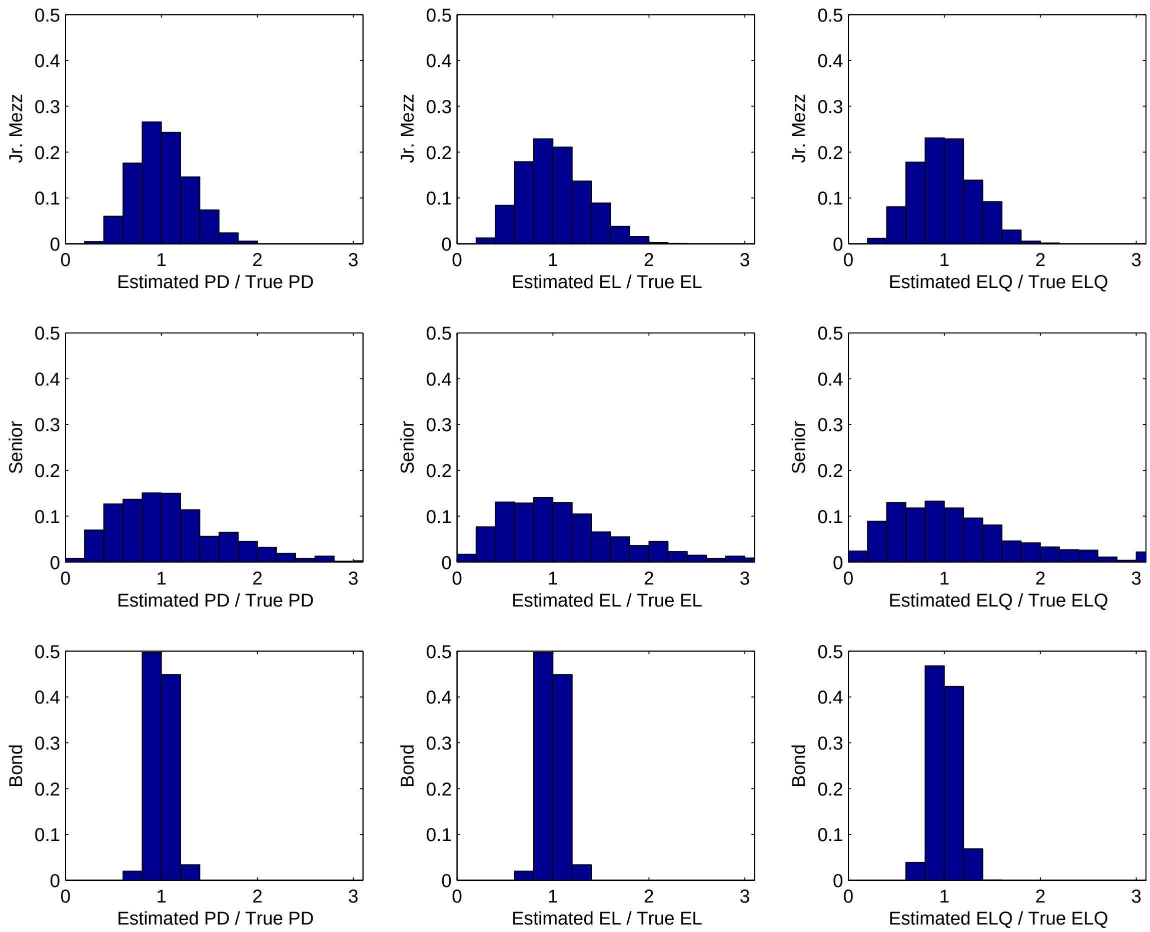

Figure 13: Distribution of estimated mezzanine CDO risk metrics given a long data panel with unobserved credit factors ($N = 50$, $T = 40$, $\psi = 1.0$).

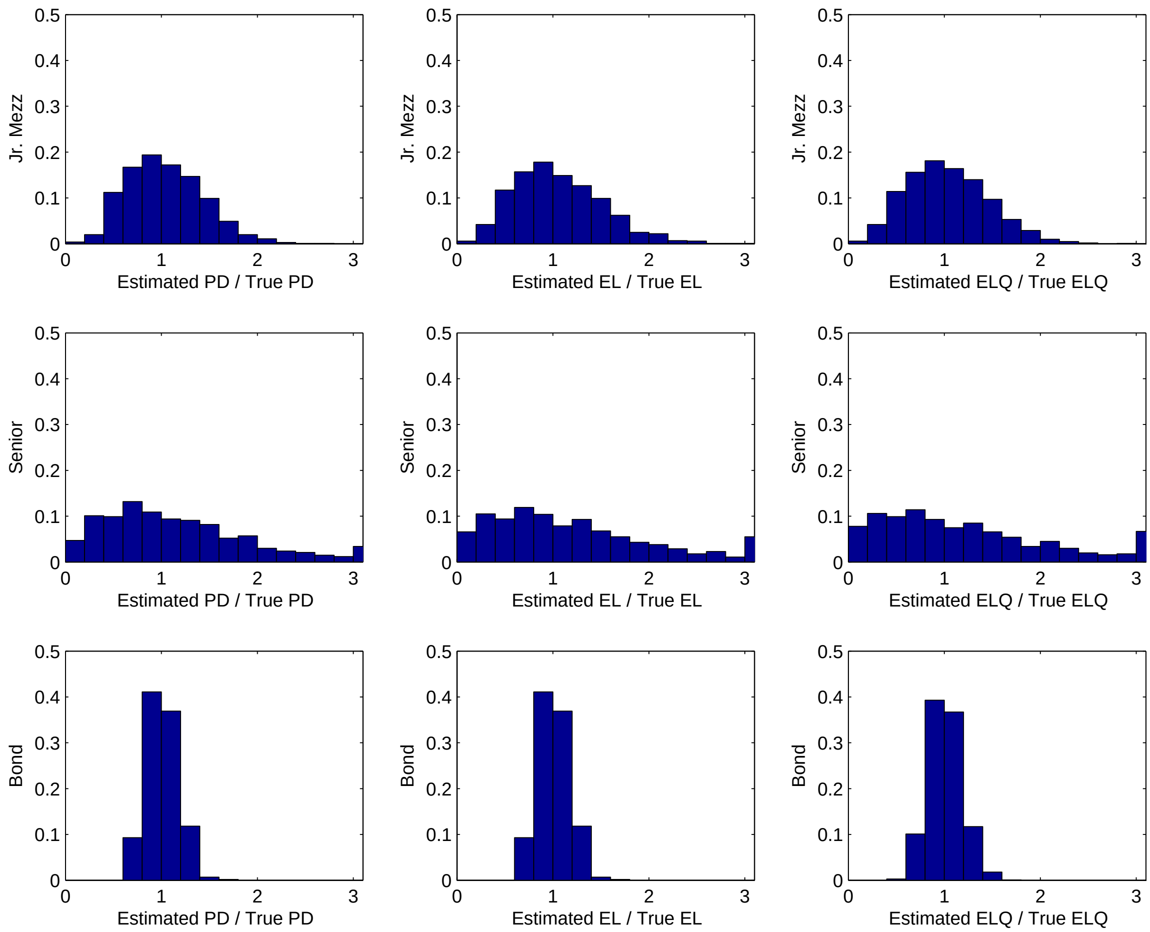

Figure 14: Distribution of estimated mezzanine CDO risk metrics given a wide data panel with unobserved credit factors ($N = 200$, $T = 10$, $\psi = 1.0$).

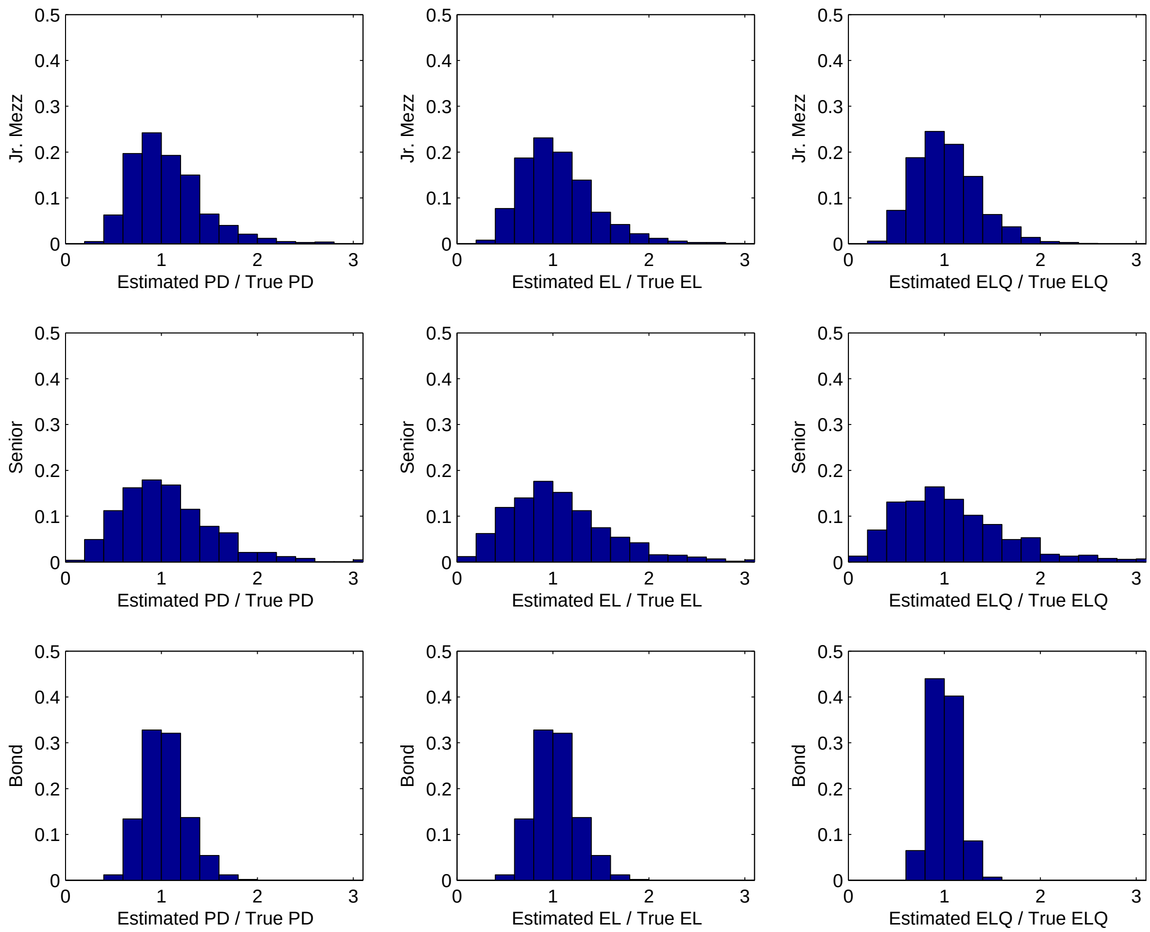

Figure 15: Distribution of estimated high-grade CDO risk metrics given a long data panel with well observed credit factors ($N = 50$, $T = 40$, $\psi = 0.2$).

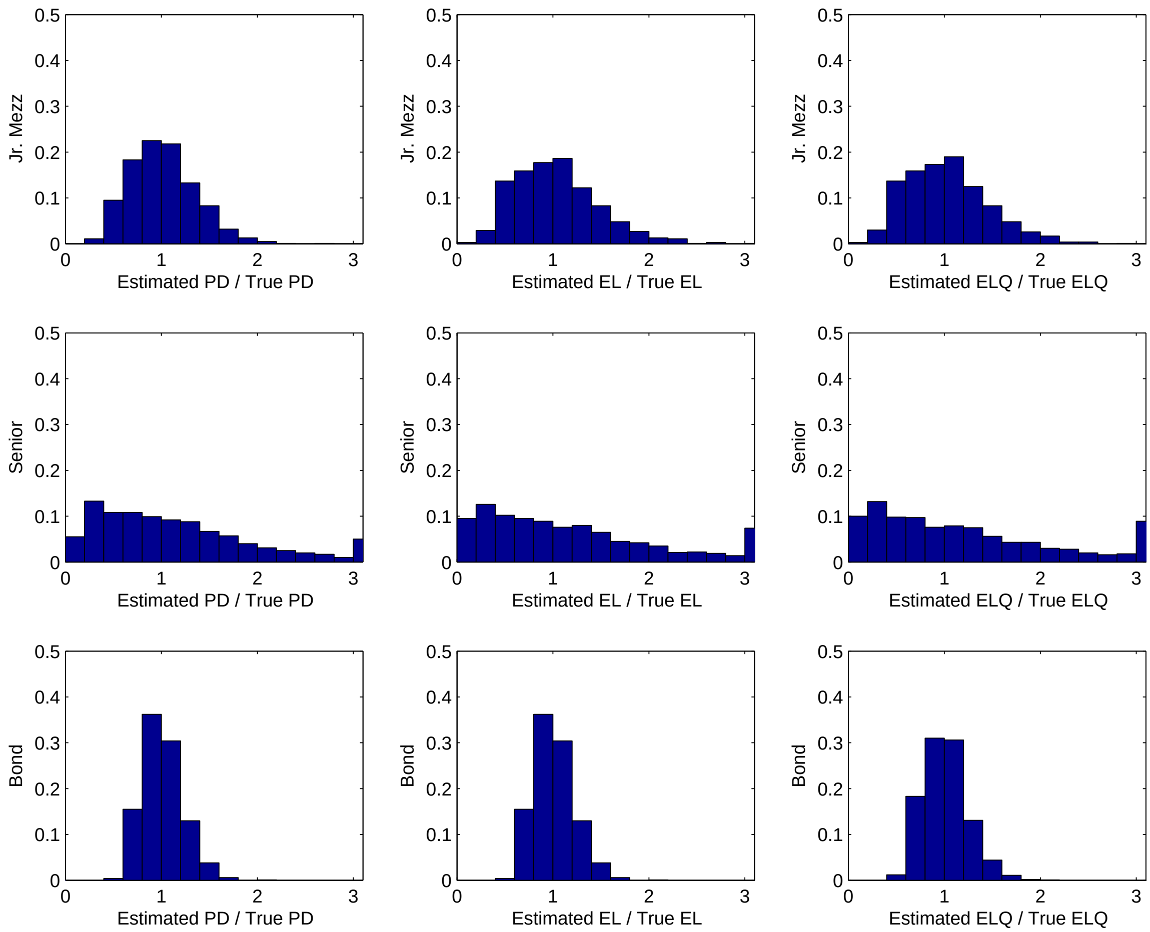

Figure 16: Distribution of estimated high-grade CDO risk metrics given a wide data panel with well observed credit factors ($N = 200$, $T = 10$, $\psi = 0.2$).

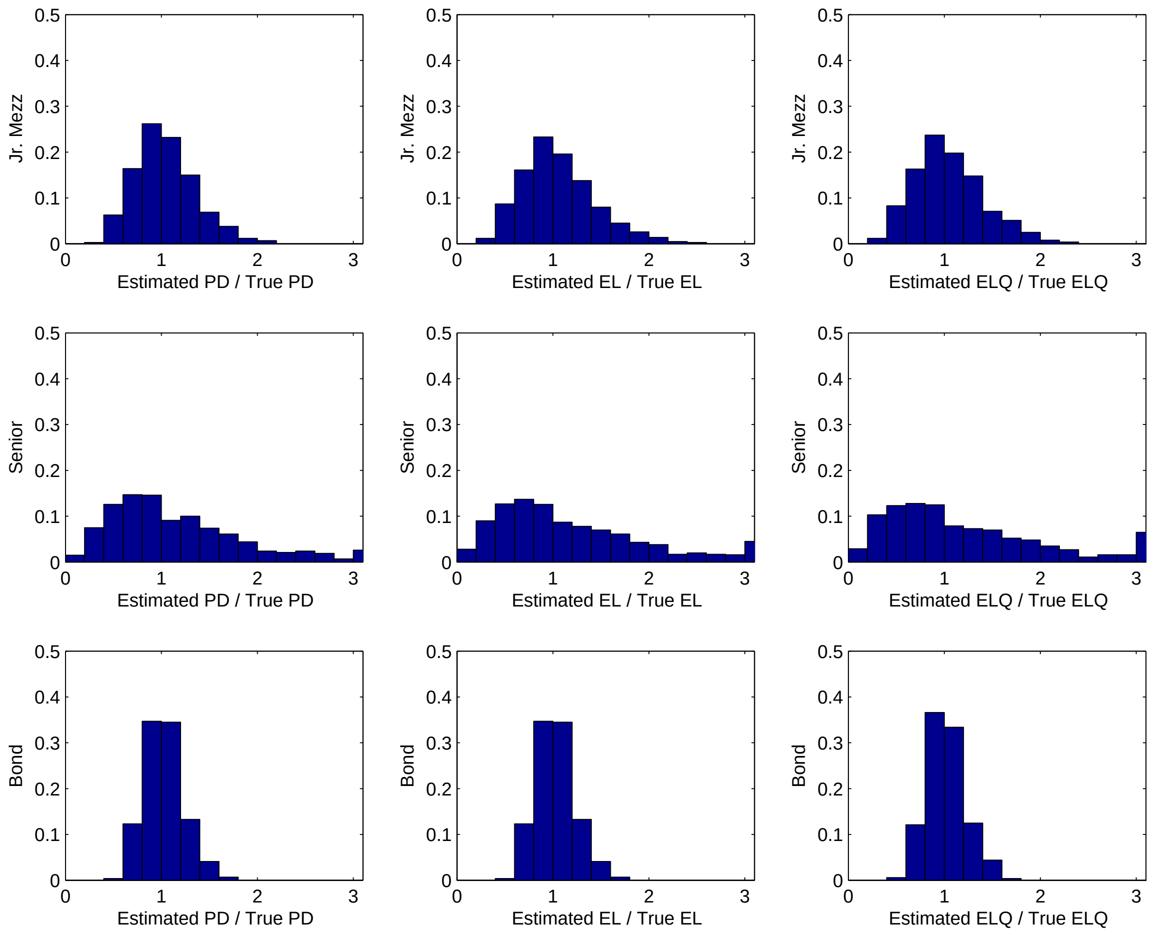

Figure 17: Distribution of estimated high-grade CDO risk metrics given a long data panel with unobserved credit factors ($N = 50$, $T = 40$, $\psi = 1.0$).

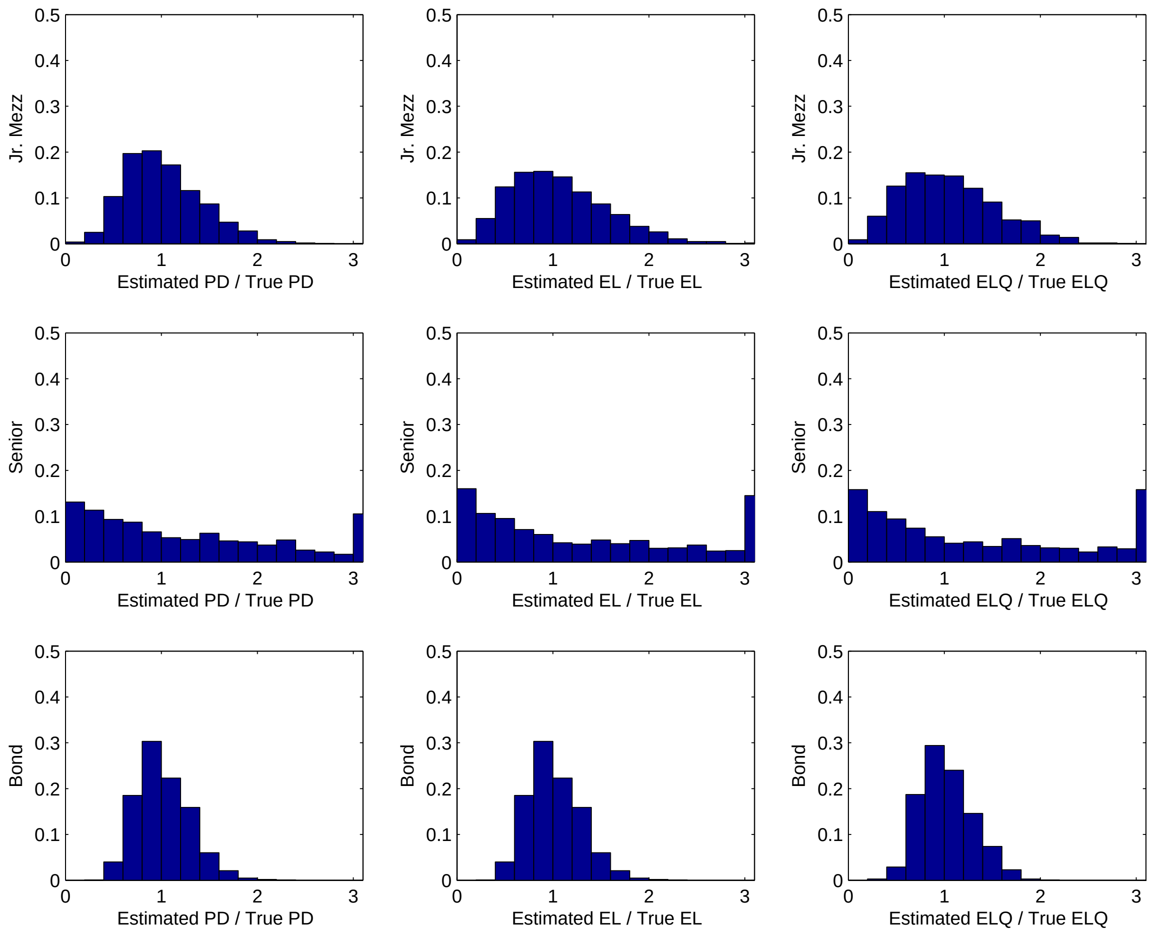

Figure 18: Distribution of estimated high-grade CDO risk metrics given a wide data panel with unobserved credit factors ($N = 200$, $T = 10$, $\psi = 1.0$).

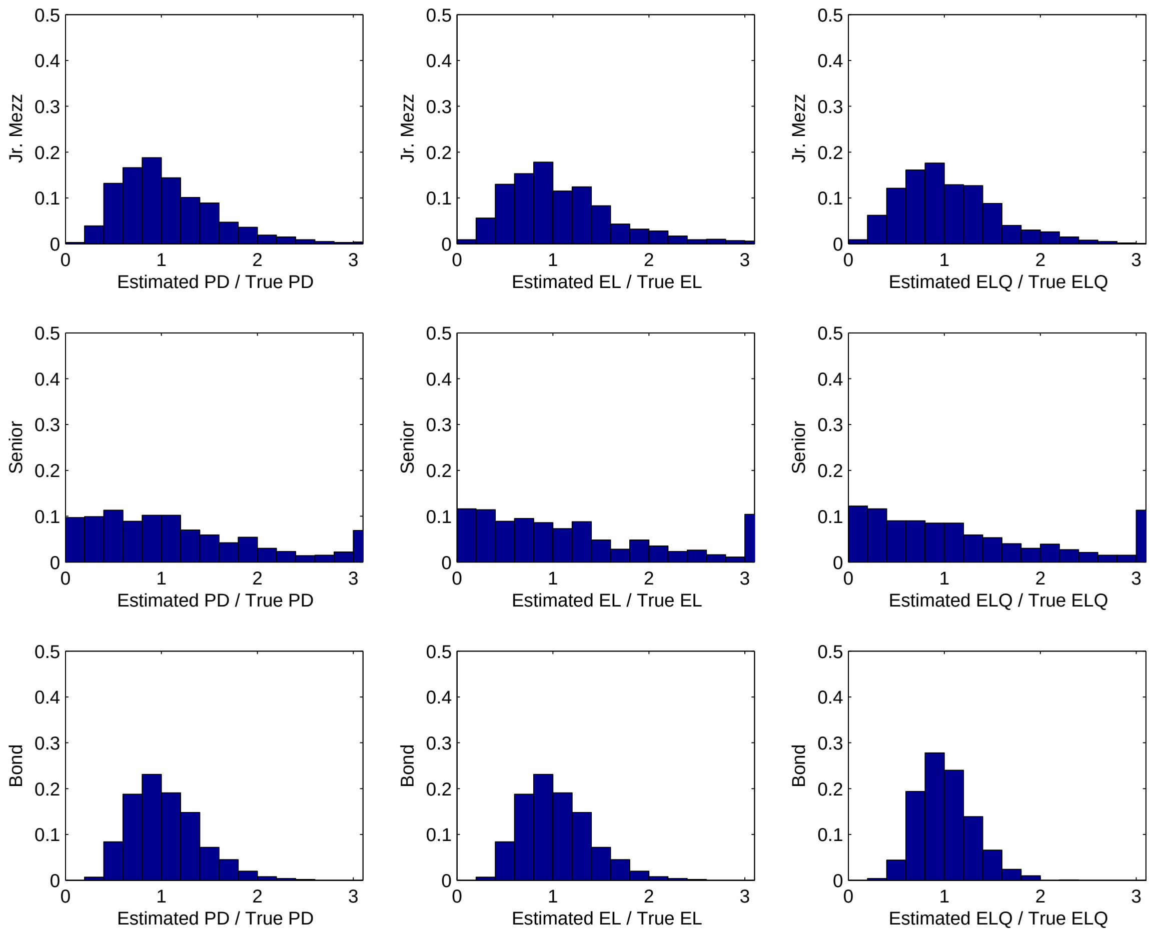

APPENDIX: Derivation of DGP likelihood

Conditional on X, Y^* and Y have the joint distribution,

$$\begin{bmatrix} Y^* \\ Y \end{bmatrix} \mid X \sim \mathcal{N}\left(\begin{bmatrix} \sqrt{\omega}\sqrt{1-\psi}\sqrt{\rho}X \\ \sqrt{\rho}X \end{bmatrix}, \begin{bmatrix} \omega(1-\psi)(1-\rho)+\psi & \sqrt{\omega}(1-\rho)\sqrt{1-\psi} \\ \sqrt{\omega}(1-\rho)\sqrt{1-\psi} & 1-\rho \end{bmatrix}\right).$$

This implies that conditional on Y^* and X, Y is distributed

$$Y \mid Y^*, X \sim \mathcal{N}\left(\frac{\psi\sqrt{\rho}X + (1-\rho)\sqrt{1-\psi}\,(Y^*/\sqrt{\omega})}{(1-\psi)(1-\rho)+\psi}, \frac{\psi(1-\rho)}{(1-\psi)(1-\rho)+\psi}\right).$$

Thus, we can write the joint distribution of Y^* and Y conditional on X as

$$\begin{aligned} f\left(y^*, y|X\right) = \phi & \left(\frac{(y^*/\sqrt{\omega}) - \sqrt{1-\psi}\sqrt{\rho}X}{\sqrt{(1-\psi)(1-\rho)+\psi}}\right) \\ & \times \phi\left(\frac{((1-\psi)(1-\rho)+\psi)\,y - \psi\sqrt{\rho}X - (1-\rho)\sqrt{1-\psi}\,(y^*\sqrt{\omega})}{\sqrt{(1-\psi)(1-\rho)+\psi}\sqrt{\psi(1-\rho)}}\right). \end{aligned} \tag{6}$$

The first right-hand term is the PDF of Y^* given X. The second term is the PDF of Y given Y^* and X.

Recall that D is an indicator variable that is equal to one if $Y \leq \Psi^{-1}(\pi)$ and zero otherwise. Y^* is only observable when $D = 0$. If $D = 0$ the joint likelihood of Y^* and D given X is

$$\begin{aligned} f(y^*, D=1|X) = \int_{\Phi^{-1}(\pi)}^{\infty} f(y^*, y|X)dy = \phi & \left(\frac{y^* - \sqrt{1-\psi}\sqrt{\rho}X}{\sqrt{S}}\right) \times \\ & \left(1 - \Phi\left(\frac{S\Phi^{-1}(\pi) - \psi\sqrt{\rho}X - (1-\rho)\sqrt{1-\psi}\,(y^*/\sqrt{\omega})}{\sqrt{S}\sqrt{\psi(1-\rho)}}\right)\right) \end{aligned} \tag{7}$$

where $S = (1-\psi)(1-\rho)+\psi$. The second equality follows directly from (6). If $D = 1$, Y^* is not observable, and the likelihood is

$$f(D=0|X) = \int_{-\infty}^{\Phi^{-1}(\pi)} f(y|X)dy = \Phi\left(\frac{\Phi^{-1}(\pi) - \sqrt{\rho}X}{\sqrt{1-\rho}}\right). \tag{8}$$

Combining (7) and (8) yields (5).

References

Andersen, L. and J. Sidenius, "CDO pricing with factor models: survey and comments," *Journal of Credit Risk*, 2005, *1* (3), 71–88.

—— and —— , "Extensions to the Gaussian copula: random recovery and random factor loadings," *Journal of Credit Risk*, 2005, *1* (1), 29–70.

Basel Committee on Banking Supervision, *International Convergence of Capital Measurement and Capital Standards: A Revised Framework*, Basel: Bank for International Settlements, June 2004.

Burtschell, X., J. Gregory, and J.-P. Laurent, "Beyond the Gaussian copula: stochastic and local correlaiton," *Journal of Credit Risk*, 2007, *3* (1).

Frey, R., A. McNeil, and M. Nyfeler, "Copulas and Credit Models," *Risk*, October 2001, pp. 111–114.

Frey, Rudiger and Alexander J. McNeil, "Dependent Defaults in Models of Portfolio Credit Risk," *Journal of Risk*, 2003, *6*, 59–92.

Gilkes, Kai, "CDO Evaluator Version 3.0: Technical Document," Technical Report, Standard and Poor's December 2004. Structured Finance Criteria Report.

Gordy, Michael and Erik Heitfield, "Estimating Default Correlations from Short Panels of Credit Rating Performacne Data," January 2002. Working Paper.

Gordy, Michael B., "A Risk Factor Model Foundation for Ratings-Based Bank Capital Rules," *Journal of Financial Intermediation*, 2003, *12*, 199–232.

Gupton, Greg, Christopher Finger, and Mickey Bhatia, "CreditMetrics – Technical Document," Technical Report, The RiskMatrics Group April 1997.

Hamerle, Alfred and Daniel Rosch, "Misspecified copulas in credit risk models: how good is gaussian?," *Journal of Risk*, 2005, *8* (1).

Heitfield, Erik, Steve Burton, and Souphalla Chomsisengphet, "Systematic and Idiosyncratic Risk in Syndicated Loan Portfolios," *Journal of Credit Risk*, 2006, *2* (3), 3–31.

Jolivet, Guillaume, Laurent Lassalvy, Michael Mueller-Heumann, and Julien Sieler, "CDOROM 2.3 Users Guide," Technical Report, Moody's Investors Services June 2006.

Kealhofer, Stephen, "Portfolio Management of Default Risk," Technical Report, KMV Corporation February 1998.

Koo, Mia, Jeffrey Cromartie, and Paul Vedova, "Global Rating Criteria for Collateralised Debt Obligations," Technical Report, DerivativeFitch October 2006. Credit Products Criteria Report.

Li, David X., "On Default Correlations: A Copula Function Approach," *Journal of Fixed Income*, March 2000, *9*, 43–54.

Loffler, Gunter, "The effects of estimation error on measures of portfolio credit risk," *Journal of Banking and Finance*, October 2001, *27*, 1427–1453.

Mason, Joseph R. and Josh Rosner, "Where did the risk go? How misapplied bond ratings cause mortgage backed securities and collateralized debt obligation market disruptions," May 2007. Working Paper (SSRN 1027475).

Merton, Robert C., "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance*, 1974, *29*, 449–470.

Pykhtin, Michael, "Multi-factor Adjustment," *Risk*, 2003, *17* (3), 88–90.

___ **and Ashish Dev**, "Credit Risk in Asset Securitisations: an Analytical Model," *Risk*, May 2002, pp. S16–S20.

Tarashev, Nikola and Haibin Zhu, "Modelling and Calibration Errors in Measures of Portfolio Credit Risk," June 2007. Working Paper.