Tail Behavior of Credit Loss Distributions for General Latent Factor Models*

André Lucas[‡§] **Pieter Klaassen**[†‡]

Peter Spreij[¶] **Stefan Straetmans**[‖]

This version: January 29, 2002

Abstract

Using a limiting approach to portfolio credit risk, we obtain analytic expressions for the tail behavior of credit losses. We show that in many cases of practical interest the distribution of these losses has polynomial ('fat') rather than exponential ('thin') tails. Our modeling framework encompasses the models available in the literature. Defaults are triggered by a general latent factor model involving systematic and idiosyncratic risk. If idiosyncratic risk exhibits thinner tails than systematic risk, the credit loss density actually increases towards the maximum credit loss, a feature not reported earlier in the literature. We also derive analytically the interaction between portfolio quality and credit loss tail behavior and find a striking difference between two well-known modeling frameworks for portfolio credit risk: CreditMetrics and CreditRisk$^+$. Finally, we comment on the applicability of Extreme Value Theory to describe the behavior of credit loss tails.

Key words: portfolio credit risk; extreme value theory; tail events; tail index; factor models; economic capital; portfolio quality; second-order expansions.

JEL Codes: G21; G33; G29; C19.

---

*We thank seminar participants at EURANDOM for useful comments. Correspondence to: alucasecon.vu.nl, s.straetmansberfin.unimaas.nl, pieter.klaassennl.abnamro.com, or spreijwins.uva.nl.

‡Dept. Finance and Financial Sector Management, Vrije Universiteit, De Boelelaan 1105, NL-1081HV Amsterdam, the Netherlands

§Tinbergen Institute Amsterdam, Keizersgracht 482, NL-1017EG Amsterdam, the Netherlands

†ABN AMRO Bank NV, Financial Markets Risk Management (HQ 2035), P.O.Box 283, NL-1000EA Amsterdam, the Netherlands

¶Korteweg-de Vries Institute, University of Amsterdam, Plantage Muidergracht 24, NL-1018TV Amsterdam, the Netherlands

‖Limburg Institute of Financial Economics (LIFE), Maastricht University, Tongersestraat 53, NL-6200MD Maastricht, the Netherlands

1 Introduction

Management of credit risk is a core function within banks and other lending institutions. There is an extensive literature on how to assess the credit quality of counter-parties in individual loan (or bond) transactions, see for example Altman (1983), Caouette, Altman, and Narayanan (1998), and the Journal of Banking & Finance (2001, vol. 25(1)) as starting references. Recently, we have witnessed an increased interest in the modeling and management of *portfolio* credit risk. This involves the determination of the probability distribution of potential credit losses for portfolios of loans. In this paper, we focus on the properties of this credit loss distribution in its extreme tail, i.e., where credit losses are high. Clearly, this is the part of the distribution that banks are most concerned about.

Banks usually get into trouble when in a short period of time a substantial part of the loan portfolio deteriorates significantly in quality. This can typically be traced back to some common cause, e.g., a downturn in the economy of a country or region, or problems in a particular industry sector (see also Bangia, Diebold, and Schuermann (2000)). Recent examples are the banking problems in Japan, and the Asian crisis in 1998. A bank is much less vulnerable to such systematic events when its loan portfolio is well diversified over regions, countries and industries. To evaluate and manage a bank's credit risk, it is therefore not sufficient to scrutinize individual clients to which loans are extended, but also to identify concentration of risks within the portfolio. The use of portfolio credit risk models allows banks to do that.

Banks also employ portfolio credit risk models to evaluate activities on a risk/reward basis, using measures such as risk-adjusted return on capital (RAROC) and economic-value-added (EVA), see Matten (2000). Such an evaluation can be done on the level of individual loans or clients, but also for lines of business or the bank as a whole. In addition, portfolio credit risk models can be used to evaluate the risks and merits of collateralized loan or bond obligations. A major reason for banks to enter into such structures is to obtain regulatory capital relief. In many cases, however, the majority of the economic risk of the loans involved remains with the issuing bank. A primary motivation for the current review of the 1988 Basel Accord on regulatory capital is therefore to better align regulatory capital requirements with true economic risk. In its latest proposals, the Bank for International Settlements (BIS) has in fact used a portfolio credit risk approach to set risk weights for individual counter-parties (see BIS (2001).

Several models have been put forward in the literature to capture the salient features of portfolio credit risk. The most prominent models are *CreditMetrics* of Gupton, Finger, and Bhatia (1997), *CreditRisk+* of Credit Suisse (1997), *PortfolioManager* of KMV (Kealhofer (1995)), and *CreditPortfolioView* of McKinsey (Wilson (1997a,b)). Despite the apparent differences between these approaches, they exhibit a common underlying framework, see Koyluoglu and Hickman (1998) and Gordy (2000). All models enable the computation or simulation of a probability distribution of credit losses at the portfolio level. The extreme upper quantiles of this distribution are of particular interest.

In the present paper we formulate a general modeling framework encompassing the models from the literature mentioned above. The framework enables us to derive an explicit characterization of the extreme tail behavior of credit losses in terms of underlying portfolio characteristics. Our approach extends the results in Lucas et al. (2001) and contrasts with previous studies of the behaviour of aggregate credit risk. For example, Carey (1998) uses a large database of bonds and a resampling scheme in order to investigate the tail behavior of credit loss distributions. This approach, however, does not provide an explicit relation between default correlations and credit loss tail behavior. Moreover, all results are conditional on the extent to which the database used is representative of an actual bond or loan portfolio. Alternatively, Gupton, Finger, and Bhatia (1997) uses an explicit modeling framework and a simulation set-up. The main drawback of a simulation approach is that it is difficult to obtain reliable conclusions regarding tail behavior, especially if one is concerned with extreme quantiles. Moreover, many different experiments would have to be set up in order to obtain tail properties under a variety of empirically relevant conditions. By contrast, our analytic approach allows for a direct assessment of the relation between default correlations, credit quality, distributional properties, model structure, and credit loss tail behavior.

In line with the existing literature, we decompose the risk of an individual loan into a systematic and idiosyncratic risk component. Other models have opted to fully parameterize the distribution of the risk components. For example, CreditMetrics assumes normal risk components, whereas CreditRisk$^+$ assumes Gamma distributed components. We abstain, however, from making parametric assumptions on the risk components. For our purpose, i.e. the study of the tail behaviour of portfolio credit losses, it suffices to make weak assumptions on the probability of extreme realizations (tail behaviour) of the risk factors. In addition, we allow risk factors to be related in a general, possibly non-linear way to a counter-party's creditworthiness. Using statistical Extreme Value Theory, we obtain an expansion of the tail of the credit loss distribution.

Our main contributions to the portfolio credit risk literature are the general modeling framework and the analytic results. It turns out that under quite general conditions credit losses have a polynomial (i.e., fat) rather than an exponential (thin) tail. This polynomial tail can be characterized by a single parameter, the so-called tail index. The tail index specifies the rate of decay in the tail probability. The larger the tail index, the faster the tail probability declines to zero. We show how assumptions on the extreme tail behavior of the idiosyncratic and systematic risk components determine the extreme credit loss tail behavior, i.e., the value of the tail index. In particular, we prove that thin tails for idiosyncratic risk and fat tails for systematic risk produce rather unconventional shapes of credit loss densities. These densities may be increasing near the edges of their support. To the best of our knowledge, such behavior has not been reported earlier in the literature.

We also investigate how credit quality as measured by the probability of default relates to the credit loss tail index. It turns out that credit quality affects the tail behavior of credit losses differently in the CreditMetrics framework compared to CreditRisk$^+$, which are two of the most popular portfolio models to study portfolio credit risk. More specifically, the probability of default impacts on the tail index directly in the CreditRisk$^+$ model, whereas it only has a second-order effect on the tail behaviour in the CreditMetrics framework.

The analytical framework of our paper also allows us to address questions regarding the applicability of statistical Extreme Value Theory (EVT) to empirical credit loss densities. For example, Diebold, Schuermann, and Stroughair (1998) argue that EVT is mainly suited for dealing with quantiles extremely far out in the tails. They also suggest that quantiles at 5% or 1% levels are not sufficiently extreme to warrant the use of EVT. Such percentages are typical in market risk computations, see for example Jorion (1997). In credit risk computations, typical quantiles are often much smaller than 1%, e.g. 0.1% or 0.01%. This might suggest that EVT has more potential in the credit risk context. We show for stylized portfolios and different distributions of systematic and idiosyncratic risk that one has to be very cautious in applying EVT to empirical portfolios. The point at which extreme tail behavior sets in might lie beyond empirically relevant quantiles, even in the credit risk context. The tail approximation emanating from EVT, however, might still prove a useful device. Its usefulness, however, arises more from a numerical than a theoretical point of view.

The set-up of the paper is as follows. In Section 2 we provide the basic modeling framework and derive the main results for a homogeneous bond portfolio. We also treat the CreditMetrics and CreditRisk$^+$ models as special cases. The results are generalized in Section 3 to heterogeneous portfolios. Section 4 contains a second-order approach to the tail behavior of a double Gaussian latent factor model. We highlight the differences between the CreditMetrics and the CreditRisk$^+$ approach regarding the interaction between portfolio credit quality and tail behavior. Section 5 comments on the applicability of statistical EVT for empirical credit risk management. Section 6 concludes, while the Appendix gathers all the proofs.

2 Homogeneous bond portfolios

We start our exposition with a very simple portfolio containing n bonds (or loans), each from (to) a different company.[1] The portfolio is homogeneous in the sense that all bonds have the same characteristics. This restrictive setting allows us to derive useful results on the tail behavior of portfolio credit losses. In later sections, we generalize these results to heterogeneous portfolios.

Each bond in the portfolio specifies a future pay-off stream of coupons and/or principal. The value of this stream depends on the creditworthiness of the company issuing the bond. The value of an identical stream of future cash flows will be lower if the company is more likely to default, i.e., has a lower creditworthiness. In our benchmark setting, each company j, where $j = 1, \ldots, n$, is characterized by a two-dimensional vector

$$(S_j, s^*). \tag{1}$$

Here, S_j is a latent variable that triggers a company's default. A prime candidate for S_j is the company's 'surplus' or equity value, i.e., the difference in market value of assets and liabilities. If this surplus falls below the threshold s^*, default occurs. As our focus in the present paper is on extreme tail behavior of credit losses, we concentrate on defaults only and abstract from credit losses due to credit rating migrations, see Gupton, Finger, and Bhatia (1997). Further, for simplicity we set the recovery rate to 0, implying that the loss given default is 100%. This means that in case of default, the complete amount invested is lost. Alternatively, one can use more realistic values like historical averages of recovery rates. This, however, does not affect the rate of tail decay of portfolio credit losses as derived later on. We assume that the initial value of each bond is unity (i.e., each bond values to par at the start). The credit loss on an individual bond j is now given by the random variable

$$1_{\{S_j < s^*\}}, \tag{2}$$

where 1_A is the indicator function of the set A.

[1] We focus on bonds and loans for expositional purposes, but the basic modeling framework remains applicable in case of alternative credit risky securities.

We assume that S_j obeys the general factor model

$$S_j = g(f, \varepsilon_j), \tag{3}$$

where f is a common factor, ε_j is a firm-specific risk factor, and $g(\cdot,\cdot)$ defines the functional form of the factor model. In this section, we restrict the factor model to be the same for each firm j. This assumption is relaxed in the next section. The formulation in (3) comprises the well-known factor models from the literature. For example, if we set $g(f, \varepsilon_j) = \beta f + \varepsilon_j$ for some factor loading $\beta \in \mathbb{R}$ with Gaussian f and ε_j, we obtain (a restricted version of) the CreditMetrics model introduced by Gupton, Finger, and Bhatia (1997). In our present static context, this also coincides with the formulation of CreditPortfolioView of McKinsey, see Wilson (1997a,b). Alternatively, if $g(f, \varepsilon_j) = \varepsilon_j/(\beta f)$ with $\beta > 0$ and ε_j and f exponentially and Gamma distributed, respectively, we obtain the Creditrisk$^+$ specification of Credit Suisse, see Gordy (2000) and Credit Suisse (1997).

For sake of simplicity we consider a one-factor version of (3) only. Some results for linear multi-factor models are given in Lucas et al. (2001). The key ingredient in (3) is the common risk factor f. Due to the functional dependence of all S_j on this common f, we are able to accommodate default correlations. For example, average default rates can be much higher during recessions than during booms of the economy, which can be captured by an adequate choice of f.

Given the formulation of the individual credit loss in (2), the credit loss for a portfolio of n loans expressed as a fraction of the amount invested is now given by

$$C_n = n^{-1} \sum_{j=1}^{n} 1_{\{S_j < s^*\}}, \tag{4}$$

i.e., the sum of the individual losses divided by the total initial value of the bonds. Looking at the extreme tail behavior of C_n is rather trivial as the support of C_n is discrete. We obtain a continuous credit loss distribution only if we let the number of loans n go to infinity, as in Lucas et al. (2001). We will follow this approach as it allows us to establish explicit links between default correlations and credit loss tail thickness.[2] Define

$$C = \lim_{n\to\infty} C_n, \tag{5}$$

[2]The introduction of stochastic recovery rates can also make the credit loss distribution continuous for finite n, but this would not provide the desired insight into the relation between default correlations and credit loss tail behavior. Indeed, the extreme tail behavior of portfolio credit losses would be directly equal to the assumed tail behavior for the recovery rates.

where the limit is assumed to exist almost surely. As is shown in Theorem 1, C only depends on the systematic risk factor f and not on the idiosyncratic risk factors ε_j. Using the formulation in (5) rather than (4), we therefore limit the number of stochastic components considerably. This facilitates the study of the tail behavior of credit losses. As was shown in Lucas et al. (2001), empirically relevant quantiles of C_n, e.g., 99% or 99.9%, can be approximated well by quantiles of C, provided the credit portfolio contains at least a few hundred exposures. These values of n are quite small given the usually large numbers of exposures in typical bank portfolios, such that this requirement is usually met.

We now introduce our key assumptions on the factor model $g(\cdot,\cdot)$ and the risk factors f and ε_j. Purely for expositional purposes, we again use more restrictive assumptions than necessary. In the discussion of the assumptions, we point out which conditions can be relaxed. Some of these relaxations are worked out in later sections. We use the notation $\bar{F}(x) = 1 - F(x)$ where $F(\cdot)$ is a distribution function.

Assumption 1 *(i) $\{\varepsilon_j\}_{j=1}^{\infty}$ is an i.i.d. sequence that is independent of f.*
(ii) $g(f, \varepsilon_j)$ is monotonically increasing in f and ε_j for all j. Moreover, let $\mathcal{S}$, $\mathcal{F}$, and $\mathcal{E}$ denote the supports of S_j, f, and ε_j, respectively. Then for all $s \in \mathcal{S}$ and $f \in \mathcal{F}$ an inverse function $\varepsilon(f, s) \in \mathcal{E}$ exists, and for all $s \in \mathcal{S}$ and $\varepsilon \in \mathcal{E}$ an inverse function $f(s, \varepsilon) \in \mathcal{F}$ exists, so that

$$s = g(f(s, \varepsilon), \varepsilon) = g(f, \varepsilon(f, s)). \tag{6}$$

(iii) The supports $\mathcal{E}$ of the ε_j and $\mathcal{F}$ of f are unbounded to the right and left, respectively.

Part (i) of the assumption is standard. The identically distributed requirement is less crucial and will be relaxed in the next section. Part (ii) of Assumption 1 requires the factor model to be increasing in the risk factors. The focus on *increasing* $g(\cdot)$ is not very restrictive per se. For example, the specification of CreditRisk$^+$ ($S_j = \varepsilon_j/(\beta f)$) does not satisfy the assumption directly, as it is decreasing in f. If this is the case, however, we can usually easily transform variables and consider $g(\tilde{f}, \varepsilon_j)$ with $\tilde{f} = -f$, which is increasing in $\tilde{f}$. The additional condition in part (ii) requires invertibility of the factor model $g(\cdot)$. The inverses must be well defined and lie in the appropriate supports of the original risk factors. In particular, we only consider factor models from which we can always uniquely retrieve an element from the vector (S_j, f, ε_j) given the other two elements. Note that both the linear CreditMetrics model ($S_j = \beta f + \varepsilon_j$ with S_j, f, and ε_j in $\mathbb{R}$) and the multiplicative CreditRisk$^+$ model ($S_j = \varepsilon_j/(\beta f)$ with S_j, f, and ε_j in $\mathbb{R}^+$) satisfy this criterion. Note that (ii) also implies that any realization of f can be offset by a suitable realization of ε_j to produce the same value of S_j. This excludes for example 'unbalanced' factor models as $S_j = \exp(f) + \varepsilon_j$ with $S_j, f, \varepsilon_j \in \mathbb{R}$. Here, $g(\cdot)$ is increasing in both its arguments, but for $s < \varepsilon$ the inverse $f(s, \varepsilon)$ is not properly defined and (6) cannot be met. Part (iii) states that the supports of f and the ε_j are unbounded from below and above, respectively. This assumption is not crucial, but greatly simplifies subsequent notation. Again, there is no loss in generality as situations with a bounded support can be accommodated by an appropriate change of variables. Combining parts (ii) and (iii), it follows that if ε_j is pushed to the upper end of its support ($+\infty$), f has to be pushed to its lower end ($-\infty$) in order to keep the value of the surplus variable S_j constant. The intuition in economic terms is as follows. Consider the borderline case where a firm j is almost pushed into bankruptcy. If common risk factors (f), e.g., the state of the business cycle, are extremely adverse, then firm specific conditions (ε_j) have to be extremely favorable to prevent the firm from going bankrupt. We thus exclude bankruptcies that are solely induced by adverse values of f regardless of firm specific risk ε_j (or vice versa).

A second set of assumptions constrains the different types of tail behaviour for the risk factors f and ε_j whose impact on the aggregate credit risk tail will be investigated. In studying the tail behaviour of aggregate credit losses we will either start from polynomially declining tails for the underlying risk components (Assumption 2A) or exponentially declining risk component tails (Assumption 2B). Let $F(\cdot)$ and $G(\cdot)$ denote the (almost everywhere continuously differentiable) distribution functions of the ε_j and f, respectively.

Assumption 2A *(i) $F(\cdot)$ has a right-hand tail expansion of the form*

$$\bar{F}(\varepsilon) = \varepsilon^{-\nu_1} \cdot L_1(\varepsilon), \tag{7}$$

where $L_1(\varepsilon)$ is a slowly varying function at infinity, meaning that

$$\lim_{\varepsilon \uparrow \infty} L_1(t\varepsilon)/L_1(\varepsilon) = 1 \quad \forall\ t > 0.$$

In addition, $G(\cdot)$ has a left-hand tail expansion of the form

$$G(f) = (-f)^{-\mu_1} \cdot L_2(-f), \tag{8}$$

with $L_2(\cdot)$ a slowly varying function for $f \to -\infty$.
(ii) $\lim_{\varepsilon \uparrow \infty} \ln |f(s^*, \varepsilon)| / \ln |\varepsilon| =: \zeta_1$, *with* $0 < |\zeta_1| < \infty$.

Assumption 2B *(i) As opposed to Assumption 2A, the right-hand tail expansion of $F(\cdot)$ has the form*

$$\bar{F}(\varepsilon) = \varepsilon^{-\nu_1} \cdot \exp\left(\nu_3 \varepsilon^{\nu_2}(1 + o(1))\right) \cdot L_1(\varepsilon), \tag{9}$$

with $\nu_3 \neq 0$ and $\nu_2 \neq 0$. Similarly, the left-hand tail expansion of $G(\cdot)$ has the form

$$G(f) = (-f)^{-\mu_1} \cdot \exp\left(\mu_3 (-f)^{\mu_2}(1 + o(1))\right) \cdot L_2(-f), \tag{10}$$

with $\mu_3 \neq 0$ and $\mu_2 \neq 0$.
(ii) $\lim_{\varepsilon \uparrow \infty} (-f(s^*, \varepsilon))^{\mu_2} / \varepsilon^{\nu_2} =: \zeta_2$, *with* $0 < |\zeta_2| < \infty$.

Assumptions 2A and 2B place further restrictions on the stochastic behavior of f and ε_j and on the factor model $g(\cdot, \cdot)$. Assumption 2A states that f and ε_j have polynomial left-hand and right-hand tails, respectively. Note that we only make assumptions about the extreme tail behavior of these random variables. By contrast, both CreditMetrics and CreditRisk$^+$ make much more restrictive assumptions on the complete stochastic behavior of the risk factors. Using part (i) of Assumption 2A, we allow for any tail shape that lies in the domain of attraction of a Fréchet (or a Weibull) law, see Embrechts et al. (1997). An example of this is a distribution with polynomial tails, e.g., the Student t distribution. Part (ii) of Assumption 2A further limits the number of allowed factor model specifications. For example the specification $g(f, \varepsilon_j) = \varepsilon_j \exp(f)$ is not allowed as it is not balanced in f and ε_j. Again, such unbalancedness can usually be resolved by an appropriate change of variables.

Assumption 2B resembles Assumption 2A except for the fact that we now have exponential rather than polynomial tails. Though our formulation is not as general as that in Theorem 3.3.26 of Embrechts et al. (1997), we still cover a wide range of distributions that are commonly used in empirical exercises, e.g., the normal and the Gamma distributions of CreditMetrics and CreditRisk$^+$, respectively. Part (ii) of Assumption 2B is a modified balancedness condition, similar to part (ii) of Assumption 2A. This is seen more directly if we write down the implication of this condition as $\lim_{\varepsilon \uparrow \infty} \ln |f(s^*, \varepsilon)| / \ln(\varepsilon) = \nu_2 / \mu_2$. Whereas Assumption 2A requires this limit to exist and be unequal to zero, Assumption 2B only allows one particular value for this limit, namely ν_2 / μ_2. Assumptions 2A and 2B are easily applied to the standard credit risk models as well as to straightforward extensions of these. We do this later in the paper when we give explicit examples.

The following theorem now follows directly from Williams (1991), Theorem 12.13.

Theorem 1 *Given Assumption 1, the random variable C is well defined in (5) as an a.s. limit and satisfies*

$$C = P\left[S_j < s^* | f \right]. \tag{11}$$

Note that C is still stochastic due to its dependence on f. We now study the extreme tail behavior of portfolio credit losses C. The following theorems are proved in the Appendix.

Theorem 2 *Let H be the distribution function of C. Given Assumptions 1 and 2A, H lies in the maximum domain of attraction of the Weibull distribution with tail index*

$$\alpha = \zeta_1 \mu_1 / \nu_1,$$

which means that we can write $H(c)$ with a slowly varying function at infinity L as

$$H(c) = (1-c)^{\alpha} \cdot L(1/(1-c)), \tag{12}$$

as c is tending to the maximum credit loss 1.

Theorem 2 directly reveals the extreme tail behavior of credit losses. In particular, the fact that C lies in the domain of attraction of the Weibull distribution implies that the distribution $H(\cdot)$ of C has the form given in (12). The theorem further reveals how the tail index of the credit loss distribution (α) depends on the tail indices of the latent factors (f and ε_j) and on the factor model $g(\cdot)$. The dependence on the factor model enters through ζ_1, which is controlled by the balancedness condition (ii) in Assumption 2A. If the tails of f and ε_j are of the Fréchet or Weibull type, see Embrechts et al. (1997), the theorem shows that the tail index of the credit loss distribution is directly proportional to the ratio of the tail index of f to that of ε_j. We illustrate this further below using Student t distributions, which lie in the maximum domain of attraction of the Fréchet distribution (with the tail index equal to the degrees of freedom parameter). The tail index of C can thus be very low if ν_1 is much higher than μ_1. Put differently, the tails of the credit loss distribution may be very fat if the idiosyncratic risk factor is much lighter tailed than the systematic risk factor. This makes economic sense. If f has fatter tails than ε_j, extreme realizations of S_j occur more often due to bad realizations of f than of ε_j. Consequently, it is more likely that large portions of the portfolio default simultaneously (due to common risk) rather than individually (due to firm-specific risk). Because of this clustering effect, extreme realizations of portfolio credit losses also become more likely, resulting in a slower rate of tail decay (i.e., fatter tails).

We obtain a similar theorem for the case of exponential tails.

Theorem 3 *Given Assumptions 1 and 2B, H lies in the maximum domain of attraction of the Weibull distribution with tail index*

$$\alpha = \zeta_2 \mu_3 / \nu_3.$$

An interesting implication of this theorem is that the tail index of credit losses can be finite even if the underlying risk factors f and ε_j are both thin-tailed, see also Lucas et al. (2001) and Figure 1 below. We consider two examples: the CreditMetrics model of Gupton, Finger, and Bhatia (1997), and the CreditRisk$^+$ model of CreditSuisse. First, consider the linear factor model of CreditMetrics, $S_j = \beta f + \varepsilon_j$, with f and ε_j both standard normally distributed and $\beta > 0$. As f and ε_j are normal, we have $\nu_1 = \mu_1 = 1$, $\nu_2 = \mu_2 = 2$, and $\nu_3 = \mu_3 = -1/2$. Also note that $f(s, \varepsilon) = (s - \varepsilon)/\beta$ and consequently that

$$\zeta_2 = \lim_{\varepsilon \uparrow \infty} \frac{(\varepsilon - s^*)^2/\beta^2}{\varepsilon^2} = \beta^{-2},$$

and $\alpha = \beta^{-2}$. This confirms the results in Lucas et al. (2001). A higher systematic risk component (i.e., higher β) transforms into a lower tail index of C. This is intuitively clear: more systematic risk results in fatter tails for portfolio credit losses. For CreditRisk$^+$ we have the specification $S_j = \varepsilon_j/(-\beta f)$, where ε_j is standard exponentially distributed, and $(-f)$ has a Gamma distribution with parameters γ_1 and γ_2. We have $\nu_1 = 0$, $\nu_2 = 1$, $\nu_3 = -1$ for the exponential, and $\mu_1 = 1 - \gamma_1$, $\mu_2 = 1$, $\mu_3 = -1/\gamma_2$ for the Gamma, see Abramowitz and Stegun (1970) equation 6.5.32. Furthermore, we have $f(s, \varepsilon) = \varepsilon/(-\beta s)$, such that

$$\zeta_2 = \lim_{\varepsilon \uparrow \infty} -f(s^*, \varepsilon)/\varepsilon = (\beta s^*)^{-1}.$$

Therefore, following Theorem 3 the tail index of portfolio credit losses is given by $\alpha = (\beta s^* \gamma_2)^{-1}$. Again we see that a more dominant common risk component (higher β) results in a lower rate of tail decline. In contrast to the CreditMetrics model, however, we also see that the portfolio quality enters the tail index. This quality is measured by the magnitude of the default threshold s^*. Portfolios with a higher quality level will have a lower value for s^*, and thus a higher tail index. In Section 4 we prove that also the CreditMetrics model is affected by portfolio quality. In contrast to the CreditRisk$^+$ specification where the effect is of first-order, portfolio quality only has a second-order effect in the CreditMetrics specification (i.e., only directly affects the slowly varying function and not the tail index). This provides yet another difference between the two modeling frameworks, see also Gordy (2000).

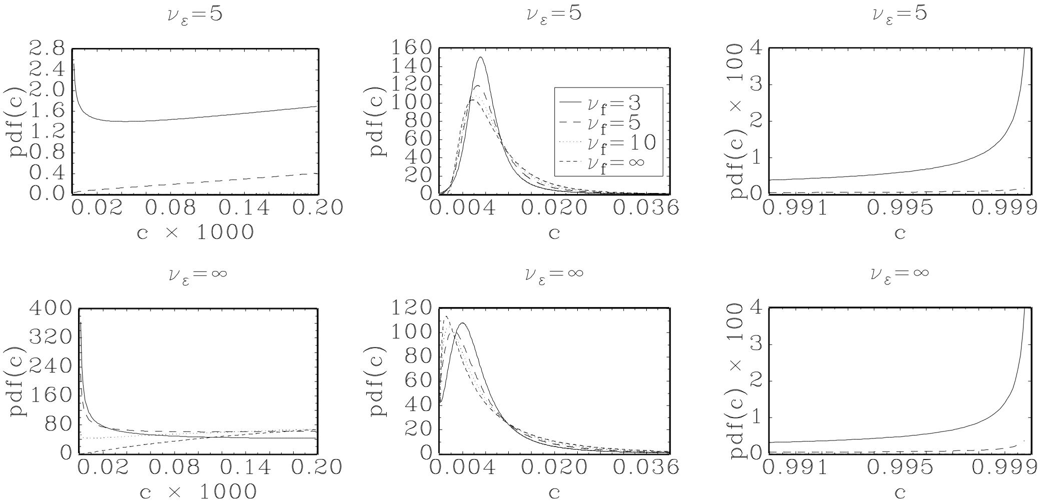

Figure 1: Credit loss distributions with different tail indices

The figure contains the credit loss densities for a homogeneous portfolio. The underlying factor model is linear, $g(f, \varepsilon_j) = a \cdot f + b \cdot \varepsilon_j$, with $a = \rho(1 - 2/\mu_1)^{1/2}$, $b = (1 - \rho^2)^{1/2}(1 - 2/\nu_1)^{1/2}$, and $\rho = 0.15$. All ε_j are identically distributed. The risk factors f and ε_j both follow a standardized Student t distribution with μ_1 and ν_1 degrees of freedom, respectively. The default probability is 1%. The left-hand plots display the credit portfolio loss density's behavior in the extreme left-hand tail. The middle plots display the behavior in the middle of the support, and the right-hand plots give the extreme right-hand tail behavior. Note the different scaling of the axes, especially the horizontal axis in the left-hand plots and the vertical axis in the right-hand plots.

To fully grasp the analytic results in Theorems 2 and 3, we present some credit loss densities for risk factor distributions with different tail indices like in the theorems above. We consider the linear factor model of CreditMetrics, slightly reparameterized as

$$S_j = \rho(1 - 2/\mu_1)^{1/2} f + [(1 - \rho^2)(1 - 2/\nu_1)]^{1/2} \varepsilon_j, \tag{13}$$

with $\rho = 0.15$. Results are similar for other values of ρ between 0 and 1. We further assume that f and ε_j follow a Student t distribution with degrees of freedom μ_1 and ν_1, respectively. Note that $(1-2/\mu_1)^{1/2} f$ and $(1-2/\nu_1)^{1/2}\varepsilon_j$ now both have zero mean and unit variance. We set the probability of default to 1%. The resulting credit loss densities are given in Figure 1 over various relevant regions of the domain $C \in [0, 1]$. If $\nu_1, \mu_1 < \infty$, Theorem 2 applies, such that the tail index of C is given by $\alpha = \mu_1/\nu_1$. If $\nu_1, \mu_1 \uparrow \infty$, we obtain normally distributed risk factors and the tail index of C is given by $\alpha = (1 - \rho^2)/\rho^2$ as derived earlier.

The first thing to note in Figure 1 are the middle plots. These reveal the typical shape of credit loss distributions known in the literature. Due to the common dependence on f, defaults are correlated. This in turn gives rise to a portfolio credit loss density that is right-skewed and has a fat right-hand tail. More peculiar are the steeply decreasing and increasing shapes of the density in the extreme left-hand (see left-hand plots) and right-hand tail (see right-hand plots), respectively. These characteristics only show up in the plots if either the density of f or ε_j has polynomial rather than exponential tails. This is due to the specific value of ρ chosen. If $\rho^2 > 0.5$, similar patters can show up if both tails are of the exponential type, e.g., normal. As the assumption of thin tails for f and ε_j has been predominant in the literature, it is not surprising that these unconventional shapes of the credit loss density have not been observed earlier. The intuition for this behaviour of the densities is as follows. Situations in which all firms default or no firm defaults will almost always correspond to extremely negative and positive realisations, respectively, of the systematic factor f. These situations are more likely to occur if the distribution of f has fat tails, and if it is less likely that the realisation of the idiosyncratic risk factor ε_j offsets the realisation of f, i.e., if ε_j is thinner tailed.

The phenomena displayed in Figure 1 can also be illustrated using the analytical expression of the credit loss density. From the proof of Theorem 2 in the Appendix, it follows that for a linear factor model $S_j = af + b\varepsilon_j$, this density $H'(c)$ has the form

$$H'(c) = \frac{b}{a} \cdot \frac{G'\left(\frac{s^*}{a} - \frac{b}{a}F^{-1}(c)\right)}{F'\left(F^{-1}(c)\right)}, \tag{14}$$

where F', G', and H' are the derivatives of the distribution functions F, G, and H, respectively. If the tails of f are thinner than those of ε_j, the numerator tends faster to zero for c tending to either 0 or 1 (and thus $F^{-1}(\cdot)$ tending to $-\infty$ or $+\infty$). By contrast, if the tails of ε_j are thinner, the denominator tends to zero at a faster rate. As a consequence, the density diverges to ∞ for both $c \downarrow 0$ and $c \uparrow 1$. If both tails are equally thin, e.g., normal, (14) shows that what matters at the extremes of the support is the size of b/a. For example, for polynomial tails of f and ε_j that are equally fat, it follows from (14) that the density tends to a non-zero limit at the edge of its support if $|b| < |a|$.

The results so far also have a practical edge for credit risk management. The likelihood of extreme credit losses is increased if the common risk factor has fatter tails than the idiosyncratic risk factor. As it can be difficult to reliably estimate the tail-fatness of f and ε_j from the empirical data that is typically available, a more conservative approach than that based on normally distributed risk factors can be warranted for prudent risk management. Especially in the upper quantiles of the credit loss distribution, more probability mass might be concentrated than suggested by the normality assumption for common and idiosyncratic risk (see also the numerical results in Lucas et al. (2001)).

3 Heterogeneous bond portfolios

So far we have concentrated on a homogeneous portfolio and a one-factor model. We now extend the results to a heterogeneous portfolio. For simplicity, we focus on a portfolio consisting of m homogeneous groups. We use i as the index of group i, $i = 1, \ldots, m$. Each group consists of $n_i = n_i(n)$ companies with $\sum_{i=1}^{m} n_i = n$. Notice also that for each company j there exists exactly one $i = i_j$ such that this company belongs to group i. We now have a company/group specific factor model, such that for all $j = 1, \ldots, n$ it holds that

$$S_j = g_i(f, \varepsilon_j),$$

for some $i = 1, \ldots, m$. We modify the assumptions from Section 2 accordingly.

Assumption 1′ *The same as Assumption 1, except for the following modifications:*
(i) The ε_j are still independent and are within each group identically distributed. The common distribution function in group i is denoted by F_i.
(ii) The factor models $g_i(f, \varepsilon)$ are increasing in both arguments and the inverse functions $f_i(s, \varepsilon)$ and $\varepsilon_i(f, s)$ exist and are well defined for all s, ε, f in their relevant supports.
(iii) Unchanged.

Assumption 2A′ *Similar to Assumption 2A, except:*
(i) Each F_i has a right-hand tail expansion as in (7), but with parameter ν_{1i}.
(ii) $\lim_{\varepsilon \uparrow \infty} \ln |f_\iota(s_\iota^*, \varepsilon)| / \ln |\varepsilon| =: \zeta_{1\iota}$, with $0 < |\zeta_{1\iota}| < \infty$, and ι as defined in Assumption 3 further below.

Assumption 2B′ *Similar to Assumption 2B, except:*
(i) Each F_i has a right-hand tail expansion as in (7), but with parameters ν_{1i}, ν_{2i}, and ν_{3i}.
(ii) $\lim_{\varepsilon \uparrow \infty} (-f_\iota(s_\iota^*, \varepsilon))^{\mu_2} / \varepsilon^{\nu_{2\iota}} =: \zeta_{2\iota}$, with $0 < |\zeta_{2\iota}| < \infty$, and ι as defined in Assumption 3 further below.

The main relaxations with respect to the previous set of assumptions concern the group-specific factor models and distributions of the idiosyncratic risk components. Also note that the credit quality as measured by s_i^* may differ accross groups. In order to avoid uninteresting pathological situations in the present context we also make the assumption that the relative sizes of the groups $\lambda_i(n) = \frac{n_i(n)}{n}$ eventually stabilize. That is $\lambda_i = \lim_{\to\infty} \lambda_i(n)$ is assumed to exist for all i. The assertion of Theorem 1 now takes a different form, which can be proved similarly, with only a slightly more delicate way of reasoning. Different from (11) we have the following formulation of portfolio credit losses:

$$C = \sum_{i=1}^{m} \lambda_i \cdot P[g_i(f, \varepsilon_i) < s_i^* | f] = \sum_{i=1}^{m} \lambda_i \cdot F_i(\varepsilon_i(f, s_i^*)), \tag{15}$$

where we have replaced the firm index j of ε by the group index i. For each firm in group i, ε_i follows the distribution F_i, and the ε_i are independent. The constants s_i^* determine the default probability in group i. As said before, the constants λ_i denote the (asymptotic) relative size of group i. Alternatively, one can allow for different loan sizes or recovery rates between groups and incorporate these in λ_i. This does not affect the rate of tail decay, but may impact the upper endpoint of the support of C. For simplicity, we do not consider this case here.

Assumption $1'$ on the factor models all being increasing in f is more restrictive for heterogeneous portfolios than for homogeneous portfolios. In particular, it is no longer always possible to meet this assumption by an appropriate change of variables. As an example, consider two groups where one has a factor model that is increasing in f, while the other factor model is decreasing in f. By changing variables from f to $\tilde{f}$ to make the latter model increasing in $\tilde{f}$, one makes the former model decreasing in $\tilde{f}$. Such situations are, however, of limited practical interest as they imply both positive and negative correlation between companies' surplus variables and macroeconomic conditions for significant parts of the portfolio.

As can be seen from (15), only the groups with a positive λ_i contribute to the asymptotic credit loss. We now discard all bonds in group i' for which $\lambda_{i'} = 0$. The resulting portfolio now contains $n' = n - n_{i'}(n)$ bonds and the relative sizes of the groups become $\lambda_{i'}(n') = \frac{n_i(n')}{n'}$. It is however fairly easy to see that still $\lim_{n'\to\infty} \lambda_{i'}(n') = \lambda_i$. Therefore equation (15) is still valid for the smaller portfolio, since for the original portfolio the i'-th group contributed nothing to the asymptotic credit loss. Henceforth we assume a portfolio for which all λ_is are strictly positive.

Parts (ii) of Assumptions 2A$'$ and 2B$'$ mention an index ι, which is defined in the following assumption.

Assumption 3 *There exists an index* $\iota \in \{1, \ldots, m\}$ *and a constant* K *such that*

$$\frac{1 - F_i(\varepsilon_i(f, s_i^*))}{1 - F_\iota(\varepsilon_\iota(f, s_\iota^*))} \leq K \tag{16}$$

for all $i = 1, \ldots, m$ *and all* f *sufficiently large and negative.*

This assumption requires that for extreme common risk factor realizations f one of the idiosyncratic tails dominates the other tails. The tail behavior is not checked explicitly for f, but by feeding the inverse function $\varepsilon_i(f, s_i^*)$ through the idiosyncratic distribution F_i for group i. Parts (ii) of Assumptions 2A$'$ and 2B$'$ now only need to be satisfied for group ι rather than for every group $i = 1, \ldots, m$.

We have the following theorem on the tail index of credit losses for heterogeneous portfolios. The theorem is proved in the Appendix.

Theorem 4 *Let Assumptions 1$'$, 2A$'$, and 3 be satisfied, then* H *(the cdf of* C*) lies in the maximum domain of attraction of a Weibull distribution with tail index*

$$\alpha = \zeta_{1\iota}\mu_1/\nu_{1\iota}.$$

We obtain a similar theorem for exponential tails.

Theorem 5 *Let Assumptions 1$'$, 2B$'$, and 3 be satisfied, then* H *lies in the maximum domain of attraction of a Weibull distribution with tail index*

$$\alpha = \zeta_{2\iota}\mu_3/\nu_{3\iota}.$$

It is easy to see that Theorems 4 and 5 generalize Theorems 2 and 3. An important implication of Theorems 4 and 5 is that in order to characterize the extreme tail behavior of portfolio credit losses, we do not have to take the complete portfolio into account. Only segment ι is important to compute the tail index. In fact, the tail index is the same for a heterogeneous portfolio compared to a homogeneous portfolio of the same size consisting of loans to group ι only. This also follows from the fact that the size of the investment in group ι (λ_ι), does not enter the expression for the tail index. To provide some further insight, we focus on the definition of ι. For concreteness, assume a factor model that is identical accross groups, $g_i(f, \varepsilon) \equiv g(f, \varepsilon)$, whereas the distributions of the idiosyncratic risk factors are still allowed to follow different distributions accross groups. According to (16), ι then characterizes the group that has the thickest right-hand tails for the idiosyncratic risk component. So in the context of a heterogeneous portfolio, the group with the most heavy-tailed idiosyncratic risk factor dictates portfolio credit loss tail behavior. In particular, the heavier this tail compared to the tail of f, the lighter the tail of portfolio credit losses C. The intuition for this result follows from the limiting approach taken. Idiosyncratic risk is diversifiable and therefore not incorporated in C, which only depends on common risk f. If a part of the portfolio has a strong idiosyncratic risk component, this part of the portfolio is less likely to be pushed into default by movements in common risk only. In the extreme right-hand tail of credit losses, all bonds in the portfolio have to default due to adverse common risk realizations only. As argued, the most problematic cases in this respect are precisely the bonds in group ι, which are more easily pushed into default by idiosyncratic risk compared to common risk. Therefore, this group determines the tail behavior near the maximum credit loss.

4 Second order tail expansion

In the previous sections, we showed that portfolio quality only determined the tail index of credit losses in the CreditRisk$^+$, and not in the CreditMetrics framework. We also showed that heterogeneous portfolios show the same tail index as a homogeneous portfolio consisting only of loans to a particular segment from the heterogeneous portfolio. In the present section, we again focus on a homogeneous portfolio and the linear factor model $S_j = \rho f + (1-\rho^2)^{1/2}\varepsilon_j$ with Gaussian risk factors. We aim to show that portfolio quality also affects the tail shape in the CreditMetrics framework, but that this is a second-order effect. By contrast, we proved in Section 2 that in the CreditRisk$^+$ specification portfolio quality has a first-order impact.

In order to study the second-order tail expansion, we derive an expression for the slowly varying function $L(\cdot)$ in (12) that is correct up to first order. In the Appendix, we prove the following theorem.

Theorem 6 *Given the homogeneous Gaussian linear factor model setting*

$$S_j = \rho f + \sqrt{1-\rho^2}\,\varepsilon_j,$$

for $\rho \in [-1,1]$, the distribution of C has a tail expansion for $c \uparrow 1$ of the form

$$P[C > c] = (1-c)^{(1-\rho^2)/\rho^2} \cdot L(1/(1-c)), \tag{17}$$

where L is a function that is slowly varying at infinity and that satifies

$$L(x) = \frac{\rho(\ln(x^2))^{\frac{1-3\rho^2}{2\rho^2}}}{\sqrt{1-\rho^2}} \exp\left[-\frac{(s^*)^2}{2\rho^2} + \frac{s^*\sqrt{\ln(x^2)}\sqrt{1-\rho^2}}{2\rho^2}\right] \cdot (1+o(1)). \tag{18}$$

The theorem gives a more explicit form of the slowly varying function $L(\cdot)$ in the tail expansion. Gathering the components of $L(x)$ that depend on x, we have

$$L(x) \propto \exp\left[\frac{1-3\rho^2}{2\rho^2}\ln(\ln(x^2)) + \frac{s^*\sqrt{\ln(x^2)}\sqrt{1-\rho^2}}{2\rho^2}\right]. \tag{19}$$

The dominant term in $L(x)$ as a function of $x \uparrow \infty$ is therefore

$$\exp\left[\frac{s^*\sqrt{\ln(x^2)}\sqrt{1-\rho^2}}{2\rho^2}\right]. \tag{20}$$

First note that $s^* = \Phi^{-1}(p)$ for a default probability p. For p less than 50%, the default threshold s^* will be negative. Moreover, s^* is increasing in p. If $s^* < 0$, (20) is decreasing in x, because $\rho^2 \leq 1$. The smaller the default probability p, the faster the rate of decline of (20) in x. A higher level of portfolio quality, i.e., a lower p and more negative s^*, increases the rate of tail decline for credit losses. Therefore, less far out in the credit loss tail, tails may appear thinner than suggested by the result in Theorem 3. This effect, however, is only of second order. In the extreme tail, the slowly varying function is again dominated by the factor $(1-c)^{(1-\rho^2)/\rho^2}$ in (17). This contrasts with the finding for the CreditRisk$^+$ model in Section 2, where s^* entered the tail index of credit losses directly.

5 Applying EVT to credit loss tails

So far, we have concentrated on deriving the extreme rate of tail decay theoretically. It is interesting, however, to see whether these results can be applied to empirically relevant credit loss distributions. Following suggestions by Diebold et al. (1998) we might expect some merit of the use of Extreme Value Theory (EVT) for credit losses, because the quantiles used for credit loss management typically correspond to very low confidence levels (1% or lower),.

We consider the linear factor model (13) and vary the degrees of freedom parameters of the Student t distributions for f and ε_j. Instead of picking the parameter ρ and the default threshold c directly, we select the expectation and standard deviation of credit losses. Tuning the model on these parameters rather than on ρ and c directly seems more natural, see also Gordy (2000). This means that we choose ρ and c such that $E(C) = q$ and $\text{Var}(C) = \sigma_c$, with q and σ_c prespecified constants. Given ρ and c, we can compute the Economic Capital (EC) corresponding to a confidence level p. The EC is defined as the pth quantile of the credit loss distribution minus the expected loss. Note that the expected loss is given by q. Table 1 presents the values of EC for various parameter configurations.

It is clear from Table 1 that the tail shape of the systematic and idiosyncratic risk components has a major impact on the level of EC. For example, if the expected credit loss (q) is 1% and we consider the $p = 99\%$ confidence level, EC may increase from 1.21 to 6.30 for low credit loss variability ($\sigma_c = 1.5\%$) or from 2.66 to 13.60 for high variability ($\sigma_c = 3\%$), both increases by more than 400%. Similarly for high confidence quantiles $p = 99.99\%$ may increase dramatically from Gaussian risk factors to Student $t(3)$ distributed factors. Given the imprecise data generally available for determining the tail behavior of the latent risk factors, the sensitivity displayed in Table 1 is worrying for empirical risk management point of view. There appear to be few apparent systematic structures in the levels of EC. If the quality of the portfolio is good (i.e., low expected losses or $q = 0.1\%$) and one is interested in not too extreme confidence levels ($p = 99\%, 99.9\%$), then the Gaussian model appears best for prudent risk management: it provides an upper bound for the EC level for alternative values of μ_1 and ν_1. A similar result holds for portfolios with $q = 1\%$ and moderate confidence levels ($p = 99\%$). For other parameter configurations, the patterns is less clear. In some cases ($q = 1\%$ and $p = 99.99\%$), low values of μ_1 and ν_1 lead to the largest values of EC. In other cases ($q = 0.1\%$, $p = 99.99\%$ and $\sigma_c = 0.15\%$), thin-tailed f (high values of μ_1) and fat-tailed ε_j (low values of ν_1) lead to the highest levels of EC.

Though the level of EC plays an important role in contemporary credit risk management and performance evaluation, see Bessis (1998), it mainly reveals information on the incidence of loss. If one holds the expected loss and EC as a capital buffer, one is able to adverse economic conditions with a probability of p. In the remaining $(1 - p)$ cases, the capital buffer is inadequate to cover all losses. One can then use the tail shape beyond the capital buffer to determine the magnitudes and probabilities of extreme losses. These can be used to construct alternative risk measures, like the expected shortfall given buffer depletion or variants thereof. We now investigate whether statistical EVT can be used to approximate the tail shape beyond the capital buffer reliably for various levels of confidence. The setting is the same as for Table 1. Given the pth quantile c_p^* of the credit loss distribution $H(c)$, we approximate the conditional credit loss distribution $H(c|c \geq c_p^*)$ for $c \in [c_p^*, 1]$ by

$$\tilde{H}(c|c \geq c_p^*) = 1 - \frac{(1-c)^\alpha}{(1-c_p^*)^\alpha}, \qquad (21)$$

with conditional density

$$\tilde{h}(c|c \geq c_p^*) = \frac{\alpha}{(1-c_p^*)^\alpha}(1-c)^{\alpha-1}. \tag{22}$$

The form of this approximation follows directly from Section 2. We use two different approaches to determine the index α. First, we set α to its theoretical value following from the EVT results in Section 2. This should give a good approximation of the credit loss distribution sufficiently far out in the tails, i.e., for sufficiently high values of p. Our computations in this section serve to show whether 'sufficiently high values of p' lie at empirically relevant quantiles of the credit loss distribution or beyond. If the extreme tail behavior as derived in previous sections using EVT only sets in at too distant quantiles from an empirical point of view, one has to conclude that EVT has to be applied with caution in the credit risk context. Our second approach to determine α is by conditional maximum likelihood (ML). In particular, we solve

$$\max_{\alpha \geq 0} \int_{c_p^*}^{1} \log(\tilde{h}(c|c \geq c_p^*) h(c|c \geq c_p^*) dc. \tag{23}$$

By straightforward manipulation, one can solve (23) analytically as

$$\hat{\alpha}_{ML} = \max\left(0, -\left[\int_{-\infty}^{G^{-1}(1-p)} \frac{\log\left(1 - F\left(\frac{c-a\cdot f}{b}\right)\right) g(f)}{1-p} df - \log(1-c_p^*)\right]^{-1}\right), \tag{24}$$

with $a = \rho(1-2/\mu_1)^{1/2}$ and $b = ((1-\rho^2)(1-2/\nu_1))^{1/2}$. Note that (24) can be approximated accurately by means of numerical integration. The intuition underlying this approach is as follows. One uses the general form of the tail approximation following from EVT, but rather than focusing on the extreme tail only, one sets α such that (22) fits adequately on average in both the extreme and less extreme tail area. The results of $\hat{\alpha}_{ML}$ are presented in Table 2.

The general pattern emerging from Table 2 is that the discrepancies between the analytic value of α based on EVT and the ML estimates are quite large and vary considerably across risk factor distributions, portfolio quality, and confidence levels. Discrepancies appear to diminish if one gets further out into the tails, i.e., if p is increased. Hovwever, even for very low crash levels $1-p$ of 1 basis point ($p = 99.99\%$), the mismatch between α and $\hat{\alpha}_{ML}$ is quite often considerable. To obtain further insight into the causes of these differences and the (in)applicability of EVT for credit loss distributions, we highlight 8 cases from Table 1. The cases are displayed in Figure 2 and represent severaly prototypical shapes of the credit loss distribution and its approximations.

For low confidence levels ($p = 99\%$), tail approximations based on EVT or ML are markedly different. For example, for $\mu_1 = 5$ and $\nu_1 = 3$, the EVT approximation is concave, whereas the ML approximation is convex. For $\mu_1 = 3$, $\nu_1 = 5$ both approximations are convex, but EVT gives an increasing tail, whereas ML gives a decreasing one. Generally speaking, the ML approach results in a better overall fit to the true density. This is not surprising, as the ML estimate of α is set to match the true density as close as possible over the entire range $c_p^* \dots 1$. The EVT approach on the other hand is only valid in the *extreme* tail. It is clear from the figure that $p = 99\%$ is not extreme enough in general for EVT to be applicable.

For more extreme quantiles ($p = 99.99\%$) with $\mu_1 = 3$ and $\nu_1 = 5$ or for $\mu_1 = 10$ and $\nu_1 = 5$, the EVT and ML approach are very similar. This also follows from the values of the tail indices in Table 2. This results in two conclusions. First, the EVT approach can provide a useful characterization of the tail shape at empirically relevant quantiles. Second, the ML estimate of α may very considerably with the range over which the tail is approximated. For example, consider the case $\mu_1 = 3$ and $\nu_1 = 5$. For $p = 99\%$, the ML estimate of α exceeds 1, whereas for $p = 99.99\%$ we have $\hat{\alpha}_{ML} < 1$. This results for the former in a decreasing tail, and for the latter in an increasing tail. The tail approximation thus depends very much on the chosen (tail) area of interest. By contrast, the rate of tail decline based on the EVT result is independent of p.

For $\mu_1 = 5$ and $\nu_1 = 3$, we see a peculiar feature of the credit loss distribution. The tail of the conditional pdf is convexo-concave for $p = 99.99\%$. Over the dominant part of the support, the shape is convex. As a result, the ML fit results in a convex and diverging tail approximation ($\hat{\alpha}_{ML} < 1$). By contrast, the EVT tail approximation is concave. Though the concavity of the true pdf only sets in far beyond the 99.99th percentile, it is clear from the right-hand panel in the second row of Figure 2 that the EVT approximation is the correct one in the *extreme* tail. In this case, however, extremity is defined as much more extreme than 99.99%, which by itself is already very extreme for empirical purposes. So though the EVT approximation is ultimately the correct one, its applicability is beyond the scope of direct empirical interest.

For the double Gaussian factor model $\mu_1 = \nu_1 = \infty$, we see that the ML approximation is very good. By contrast, the EVT approximation is less suitable. There appears to be a minor improvement from $p = 99\%$ to $p = 99.99\%$, but the match is still far from perfect. The potential cause of this persistent mismatch for the double Gaussian model lies in the form of the slowly varying function, as explained in Section 4. So the EVT result appears less applicable empirically for this model. The algebraic tail approximation, however, is very good, but the tail index has to be estimated by ML rather than by analytic EVT methods.

The important conclusion from this section is that EVT can provide a useful approximation to the tail shape of credit loss distributions beyond conventional levels of expected losses and economic capital. If the confidence levels are set too low, however, the approximation is better when one uses ML estimates of polynomial tail approximations rather than tail approximations based on analytic EVT results. The ML approximation, however, is only valid locally and has to be adapted when the confidence level p of economic capital is changed. By contrast, the EVT approximations depend much less on p, but their applicability may lie outside the range of empirical interest. In short, there is some scope for a cautios application of EVT results, but the potential is more promising for stress testing than for regular determination of Economic Capital. This is true regardless of the fact that one usually employs much higher confidence levels p for credit risk management than for market risk management, compare Diebold et al. (1998) and Basel Committee on Banking Supervision (2001).

6 Concluding remarks

In this paper, we followed a limiting approach to determining the distribution of aggregate portfolio credit risk. Using a general (nonlinear) latent factor model, we decomposed credit risk into a systematic and an idiosyncratic risk factor. We allowed for different tail behavior of both risk components. With these ingredients, we proved that under a wide variety of circumstances, the distribution of portfolio credit losses exhibits heavy tails.

For a homogeneous portfolio and a single factor, we obtained explicit expressions linking the tail index of portfolio credit losses directly to the factor model structure and the tail indices of systematic and idiosyncratic risk. The results were illustrated by computing the tail decay rate of aggregate credit losses for two of the most common credit risk portfolio models available in the literature. This revealed a striking difference: the portfolio quality has a first-order effect on the rate of tail decline under the specification of CreditRisk$^+$, but not under that of CreditMetrics.

Using the explicit expressions for the tail index of credit losses, we showed that if the idiosyncratic risk has much thinner tails than the systematic risk, the tail index of portfolio credit losses is very small. In particular, the density of credit losses may then be *increasing* towards the edges of its support. The increasing part of the density may already start before quantiles of empirical interest, e.g., 99%. This means that extreme credit losses may show up with a much larger probability than suspected on the basis of a factor model with both Gaussian systematic and idiosyncratic risk.

We generalized our findings to the case of heterogeneous portfolios allowing for different distributions of company specific risk factors and different default probabilities and loan exposures. The results turned out to be very similar. Credit loss tails are dictated by that part of the portfolio that has the thickest idiosyncratic tails. The thicker these tails, the thinner the tail of portfolio credit losses. In particular, the credit loss tail shape of a heterogeneous portfolio is the same as that of a homogeneous portfolio consisting solely of the bonds with the thickest idiosyncratic tails.

We also derived a second order result on tail behavior for a doubly Gaussian factor model, which is the specification used by CreditMetrics. We showed that portfolio quality has an effect on the rate of tail decline of credit losses, but that this is only a second-order effect. The effect may be important, though, if one is interested in quantiles less far out in the tails.

Finally, we checked the applicability of statistical Extreme Value Theory (EVT) for credit loss tails. We showed that the tail approximation emanating from EVT provides a reasonable approximation to the credit loss tail shape if the tail decay parameter is estimated by Maximum Likelihood. Using the true parameter value of extreme tail decay only provides a good approximation sufficiently far out in the tails. In many cases, 'sufficiently far out in the tails' implies that the focus is on quantiles beyond those of empirical relevance. One thus has to be very cautious in applying the EVT results on tail decay for empirical credit loss distributions.

Appendix: Proofs

Proof of Theorem 2: First note that under Assumption 1, $\varepsilon(f,s)$ is decreasing in f. Moreover, the inverse of $\varepsilon(f,s)$ with respect to f is given by $f(s,\varepsilon)$. From (11), we have

$$C = P\left[\varepsilon_j < \varepsilon(f,s^*)|f\right] = F\left[\varepsilon(f,s^*)\right], \tag{A1}$$

where the inequality holds because of Assumption 1(ii). Therefore,

$$\begin{aligned} P[C > c] &= P[\varepsilon(f,s^*) > F^{-1}(c)] \\ &= P[f < f(s^*, F^{-1}(c))] \\ &= G[f(s^*, F^{-1}(c))], \end{aligned} \tag{A2}$$

where the inequality is reversed because $\varepsilon(f,s^*)$ is decreasing in f. Using the substitution $\varepsilon = F^{-1}(c)$, we obtain

$$\lim_{c\uparrow 1} \frac{\ln P[C > c]}{\ln[1-c]} = \lim_{\varepsilon\uparrow\infty} \frac{\ln G[f(s^*,\varepsilon)]}{\ln[1-F(\varepsilon)]}. \tag{A3}$$

The importance of this result lies in the fact that it allows us to compute the tail index of the distribution of C. From Corollary 3.3.13 of Embrechts, Klüppelberg, and Mikosch (1997) and de l'Hôpital's rule, it follows that if (A3) equals $\alpha \neq 0$, then C lies in the maximal domain of attraction of a Weibull law with tail index α.

Using the tail conditions in Assumption 2A and the result in (A3), we obtain

$$\begin{aligned}\lim_{c \uparrow 1} \frac{\ln[P(C > c)]}{\ln[1-c]} &= \lim_{\varepsilon \uparrow \infty} \frac{\ln\left[(-f(s^*,\varepsilon))^{-\mu_1} \cdot L_2(-f(s^*,\varepsilon))\right]}{\ln\left[\varepsilon^{-\nu_1} \cdot L_1(\varepsilon)\right]} \\ &= \zeta_1 \mu_1 / \nu_1,\end{aligned}$$

where ζ_1 was defined in Assumption 2A(ii). As mentioned, the result now follows from Corollary 3.3.13 of Embrechts, Klüppelberg, and Mikosch (1997) by applying the rule of de l'Hôpital. ■

Proof of Theorem 3: Using (A3) and Assumption 2B, we have

$$\lim_{c \uparrow 1} \frac{\ln[P(C > c)]}{\ln[1-c]} = \lim_{\varepsilon \uparrow \infty} \frac{\mu_3 (-f(s^*,\varepsilon))^{\mu_2}}{\nu_3 \varepsilon^{\nu_2}} = \frac{\mu_3 \zeta_2}{\nu_3},$$

with ζ_2 as defined in Assumption 2B(ii). The result again follows from Corollary 3.3.13 of Embrechts, Klüppelberg, and Mikosch (1997). ■

Proof of Theorem 4: Note that

$$\begin{aligned}P[C > c] &= P\left[\sum_{i=1}^{m} \lambda_i [1 - F_i(\varepsilon_i(f, s_i^*))] < 1 - c\right] \\ &= P\left[[1 - F_\iota(\varepsilon_\iota(f, s_\iota^*))] \left(\sum_{i=1}^{m} \lambda_i \frac{[1 - F_i(\varepsilon(f, s_i^*))]}{[1 - F_\iota(\varepsilon_\iota(f, s_\iota^*))]}\right) < 1 - c\right]. \qquad \text{(A4)}\end{aligned}$$

Also note that using Assumption 3,

$$\begin{aligned}&\lim_{c \uparrow 1} \frac{\ln P\left[(1 - F_\iota(\varepsilon_\iota(f, s_\iota^*)))K \sum_{i=1}^{m} \lambda_i < 1 - c\right]}{\ln[1-c]} \leq \\ &\lim_{c \uparrow 1} \frac{\ln P\left[[1 - F_\iota(\varepsilon_\iota(f, s_\iota^*))] \left(\sum_{i=1}^{m} \lambda_i \frac{[1 - F_i(\varepsilon(f, s_i^*))]}{[1 - F_\iota(\varepsilon_\iota(f, s_\iota^*))]}\right) < 1 - c\right]}{\ln[1-c]} \leq \\ &\lim_{c \uparrow 1} \frac{\ln P\left[(1 - F_\iota(\varepsilon_\iota(f, s_\iota^*)))\lambda_\iota < 1 - c\right]}{\ln[1-c]}. \qquad \text{(A5)}\end{aligned}$$

Combining this with (A4) and the fact that $\ln[\lambda_\iota \cdot (1-c)]/\ln[1-c]$ tends to 1 for $\lambda_\iota > 0$ and $c \uparrow 1$, we have

$$\lim_{c \uparrow 1} \frac{\ln P[C > c]}{\ln[1-c]} = \lim_{c \uparrow 1} \frac{\ln P\left[F_\iota(\varepsilon_\iota(f, s_\iota^*)) > c\right]}{\ln[1-c]}.$$

The proof now follows along the lines of that of Theorem 2 for a homogeneous portfolio. ■

Proof of Theorem 5: Similar to the proof of Theorem 4. ■

Proof of Theorem 6: Using the fact that for $x \downarrow -\infty$ we have $\Phi(x) = \phi(x)/|x|(1 + O(|x|^{-2}))$, we obtain for $\xi \downarrow 0$ that

$$
\begin{aligned}
P[C > 1 - \xi] &= \Phi\left(\frac{s + \Phi^{-1}(\xi)\sqrt{1-\rho^2}}{\rho}\right) \\
&\sim \frac{\phi\left(\frac{s+\Phi^{-1}(\xi)\sqrt{1-\rho^2}}{\rho}\right)}{\frac{|s+\Phi^{-1}(\xi)\sqrt{1-\rho^2}|}{\rho}} \\
&= \exp\left(-\frac{s^2}{2\rho^2} - \frac{s\Phi^{-1}(\xi)\sqrt{1-\rho^2}}{2\rho^2}\right)\left[\frac{\phi\left(\Phi^{-1}(\xi)\right)}{|\Phi^{-1}(\xi)|}\right]^{\frac{1-\rho^2}{\rho^2}} \frac{|\Phi^{-1}(\xi)|^{(1-\rho^2)/\rho^2}}{\left|\frac{s+\Phi^{-1}(\xi)\sqrt{1-\rho^2}}{\rho}\right|} \\
&\sim \exp\left(-\frac{s^2}{2\rho^2} - \frac{s\Phi^{-1}(\xi)\sqrt{1-\rho^2}}{2\rho^2}\right)[\xi]^{\frac{1-\rho^2}{\rho^2}} \frac{|\Phi^{-1}(\xi)|^{(1-\rho^2)/\rho^2}}{\frac{|s+\Phi^{-1}(\xi)\sqrt{1-\rho^2}|}{\rho}}. \qquad \text{(A6)}
\end{aligned}
$$

Let $\hat{\Phi}(x) = \phi(x)/|x|$, then

$$\hat{\Phi}^{-1}(\xi) = \frac{-\exp[-\frac{1}{2}\ell(1/(2\pi\xi^2))]}{\sqrt{2\pi\xi^2}},$$

with $\ell(\cdot)$ the Lambert-W function, i.e., the solution to

$$\ell(x) \cdot \exp[\ell(x)] = x.$$

For large positive x, we have asymptotically that

$$\ell(x) = \ln(x) - \ln(\ln(x)) + o(\ln(\ln(x))),$$

such that

$$\hat{\Phi}^{-1}(\xi) \overset{\xi\downarrow 0}{=} -\sqrt{-\ln(2\pi\xi^2)}. \qquad \text{(A7)}$$

Substituting $\Phi^{-1}(\xi)$ in (A6) by (A7), we obtain the desired result. ■

References

Abramowitz, M. and I. Stegun (1970). *Handbook of Mathematical Functions.* New York: Dover.

Altman, E. (1983). *Corporate financial distress. A complete guide to predicting, avoiding, and dealing with bankruptcy.* New York: Wiley.

Bangia, A., F. Diebold, and T. Schuermann (2000). Ratings migration and the business cycle, with application to credit portfolio stress testing. Technical report, Stern School, New York University. http://www.stern.nyu.edu/~fdiebold.

Basel Committee on Banking Supervision (2001, January). The new Basel capital accord. Technical report, Bank of International Settlements.

Bessis, J. (1998). *Risk Management in Banking*. New York: Wiley.

Caouette, J., E. Altman, and P. Narayanan (1998). *Managing credit risk. The next great financial challenge*. New York: Wiley.

Carey, M. (1998). Credit risk in private debt portfolios. *Journal of Finance 53*(4), 1363–1387.

Credit Suisse (1997). *CreditRisk*$^{+}$. Downloadable: http://www.csfb.com/creditrisk.

Diebold, F., T. Schuermann, and J. Stroughair (1998). Pitfalls and opportunities in the use of extreme value theory in risk management. In A.-P. N. Refenes, A. Burgess, and J. Moody (Eds.), *Decision Technologies for Computational Finance*, pp. 3–12. Deventer: Kluwer.

Embrechts, P., C. Klüppelberg, and T. Mikosch (1997). *Modeling Extremal Events*, Volume 33 of *Applications of Mathematics; Stochastic Modelling and Applied Probability*. Heidelberg: Springer Verlag.

Gordy, M. (2000). A comparative anatomy of credit risk models. *Journal of Banking and Finance 24*(1–2), 119–149.

Gupton, G., C. Finger, and M. Bhatia (1997). *CreditMetrics — Technical Document* (1st ed. ed.). http://www.riskmetrics.com.

Jorion, P. (1997). *Value at Risk — the New Benchmark for Controlling Risk*. New York: McGraw-Hill.

Kealhofer, S. (1995). Managing default risk in derivative portfolios. In *Derivative Credit Risk: Advances in Measurement and Management*. London: Risk Publications.

Koyluoglu, H. and A. Hickman (1998, October). Reconcilable differences. *Risk*, 56–62.

Lucas, A., P. Klaassen, P. Spreij, and S. Straetmans (2001). An analytic approach to credit risk of large corporate bond and loan portfolios. *Journal of Banking and Finance 25*, 1635–1664.

Matten, C. (2000). *Managing bank capital* (2nd ed.). New York: Wiley.

Williams, D. (1991). *Probability with Martingales*. Cambridge: Cambridge University Press.

Wilson, T. (1997a, September). Portfolio credit risk, part I. *Risk*, 111–117.

Wilson, T. (1997b, October). Portfolio credit risk, part II. *Risk*, 56–61.

Table 1: Economic Capital for Linear Factor Models with Student t Risk Factors

μ_1	$q=1\%$				$q=0.1\%$			
	ν_1				ν_1			
	3	5	10	∞	3	5	10	∞
	$p=99\%$, $\sigma_c=1.5\%$				$p=99\%$, $\sigma_c=0.15\%$			
3	1.21	1.72	2.23	2.82	0.03	0.04	0.06	0.11
5	2.34	3.19	3.91	4.55	0.07	0.11	0.17	0.26
10	3.70	4.69	5.32	5.74	0.17	0.27	0.39	0.50
∞	5.25	6.00	6.26	6.30	0.38	0.52	0.59	0.62
	$p=99.9\%$, $\sigma_c=1.5\%$				$p=99.9\%$, $\sigma_c=0.15\%$			
3	8.63	11.67	13.61	14.93	0.07	0.13	0.22	0.39
5	13.19	15.46	16.23	16.23	0.17	0.30	0.52	0.85
10	16.37	16.93	16.38	15.30	0.45	0.84	1.21	1.43
∞	17.50	16.08	14.54	12.98	1.20	1.54	1.53	1.39
	$p=99.99\%$, $\sigma_c=1.5\%$				$p=99.99\%$, $\sigma_c=0.15\%$			
3	91.02	85.98	80.14	72.90	0.30	0.64	1.27	2.36
5	78.76	67.90	58.76	50.00	0.53	1.17	2.12	3.22
10	61.94	49.18	40.53	33.27	1.53	2.98	3.83	3.76
∞	42.46	31.85	25.64	20.92	3.79	3.91	3.18	2.47
	$p=99\%$, $\sigma_c=3\%$				$p=99\%$, $\sigma_c=0.3\%$			
3	2.66	3.98	5.25	6.61	0.04	0.06	0.09	0.16
5	4.63	6.58	8.14	9.57	0.10	0.16	0.26	0.43
10	6.80	9.06	10.60	11.80	0.22	0.38	0.58	0.83
∞	9.36	11.58	12.81	13.60	0.48	0.76	0.99	1.19
	$p=99.9\%$, $\sigma_c=3\%$				$p=99.9\%$, $\sigma_c=0.3\%$			
3	55.13	51.87	49.22	46.50	0.11	0.21	0.38	0.73
5	52.21	48.17	45.02	41.93	0.29	0.58	1.05	1.78
10	49.15	44.65	41.28	38.09	0.75	1.51	2.36	3.06
∞	45.48	40.68	37.20	34.05	1.94	3.07	3.55	3.55
	$p=99.99\%$, $\sigma_c=3\%$				$p=99.99\%$, $\sigma_c=0.3\%$			
3	98.59	98.68	98.74	98.77	0.66	1.65	3.39	5.96
5	97.09	96.20	94.57	91.31	1.49	3.66	6.28	8.49
10	92.83	88.11	82.07	74.43	3.97	7.67	9.54	9.56
∞	82.23	72.18	63.60	55.65	8.84	10.29	9.18	7.49

This table contains the values of Economic Capital (EC) defined as the pth quantile of the credit loss distribution minus the expected loss. The EC is expressed as a percentage of the invested notional amount. The expected credit loss (or default frequency) equals q, whereas the credit loss standard deviation is σ_c. The factor model is given by (13), with f and ε_j Student t distributed with degrees of freedom parameters μ_1 and ν_1, respectively, zero means, and unit variances.

Table 2: Tail Index Estimates of Polynomial Tail Approximations to Credit Loss Distributions

μ_1	$q=1\%$, $\sigma_c=1.5\%$				$q=1\%$, $\sigma_c=3\%$				$q=0.1\%$, $\sigma_c=0.15\%$				$q=0.1\%$, $\sigma_c=0.3\%$			
	ν_1				ν_1				ν_1				ν_1			
	3	5	10	∞	3	5	10	∞	3	5	10	∞	3	5	10	∞
	EVT α				EVT α				EVT α				EVT α			
3	1.00	0.60	0.30	0	1.00	0.60	0.30	0	1.00	0.60	0.30	0	1.00	0.60	0.30	0
5	1.67	1.00	0.50	0	1.67	1.00	0.50	0	1.67	1.00	0.50	0	1.67	1.00	0.50	0
10	3.33	2.00	1.00	0	3.33	2.00	1.00	0	3.33	2.00	1.00	0	3.33	2.00	1.00	0
∞	—	—	—	4.21	—	—	—	1.30	—	—	—	7.77	—	—	—	3.07
	$\hat{\alpha}_{ML}$, $p=99\%$				$\hat{\alpha}_{ML}$, $p=99\%$				$\hat{\alpha}_{ML}$, $p=99\%$				$\hat{\alpha}_{ML}$, $p=99\%$			
3	11.1	11.0	11.0	11.0	2.93	2.97	2.98	2.72	1778	1144	762	498	332	310	263	202
5	12.3	12.9	13.9	15.3	3.49	3.80	4.14	4.50	752	573	419	309	228	193	155	122
10	14.0	15.8	18.1	21.4	4.13	4.78	5.52	6.46	496	322	253	239	191	138	109	97.0
∞	16.8	20.8	25.4	31.5	5.00	6.11	7.29	8.73	253	223	248	304	122	93.7	90.1	98.6
	$\hat{\alpha}_{ML}$, $p=99.9\%$				$\hat{\alpha}_{ML}$, $p=99.9\%$				$\hat{\alpha}_{ML}$, $p=99.9\%$				$\hat{\alpha}_{ML}$, $p=99.9\%$			
3	1.40	1.55	1.68	1.76	0.52	0.54	0.53	0.44	162	133	104	80.4	30.8	31.3	30.1	27.8
5	1.99	2.54	3.16	3.99	0.77	0.89	1.02	1.10	99.2	89.1	77.9	71.5	26.4	25.4	24.2	24.1
10	3.25	4.82	6.77	9.61	1.18	1.56	1.99	2.58	92.2	77.1	77.0	94.2	27.6	25.9	27.0	32.3
∞	6.70	11.3	16.7	24.1	2.09	3.10	4.27	5.83	76.9	95.6	139	212	26.6	29.9	39.4	56.8
	$\hat{\alpha}_{ML}$, $p=99.99\%$				$\hat{\alpha}_{ML}$, $p=99.99\%$				$\hat{\alpha}_{ML}$, $p=99.99\%$				$\hat{\alpha}_{ML}$, $p=99.99\%$			
3	0.59	0.49	0.42	0.33	0.80	0.50	0.29	0.11	10.2	10.6	10.7	11.1	2.61	2.83	3.09	3.59
5	0.77	0.81	0.93	1.08	0.87	0.63	0.47	0.32	11.1	11.4	12.1	14.5	2.87	3.18	3.70	4.67
10	1.36	1.95	2.88	4.49	1.20	1.05	1.06	1.18	13.1	15.4	21.6	36.4	3.61	4.79	7.00	11.5
∞	3.81	7.43	12.42	19.88	2.08	2.50	3.28	4.54	21.0	43.4	85.4	160	5.92	11.3	20.8	38.2

This table contains the estimates of α of the tail expansion of the credit loss distribution, where the approximation has the form $\tilde{H}(c|c > c_p^*) = 1 - ((1-c)/(1-c_p^*))^\alpha$, with c_p^* the pth quantile of the credit loss distribution $H(c)$. The parameter α is determined by extreme value theory (upper panels, see Section 2), or by conditional Maximum Likelihood (ML) over the range $c_p^* \ldots 1$. The factor model is linear with Student t distributed risk factors, see (13). The parameters μ_1 and ν_1 denote the degrees of freedom parameters of the systematic risk factor f and the idiosyncratic risk factor ε_j, respectively. The expected credit loss and its standard deviation are denoted by q and σ_c, respectively.

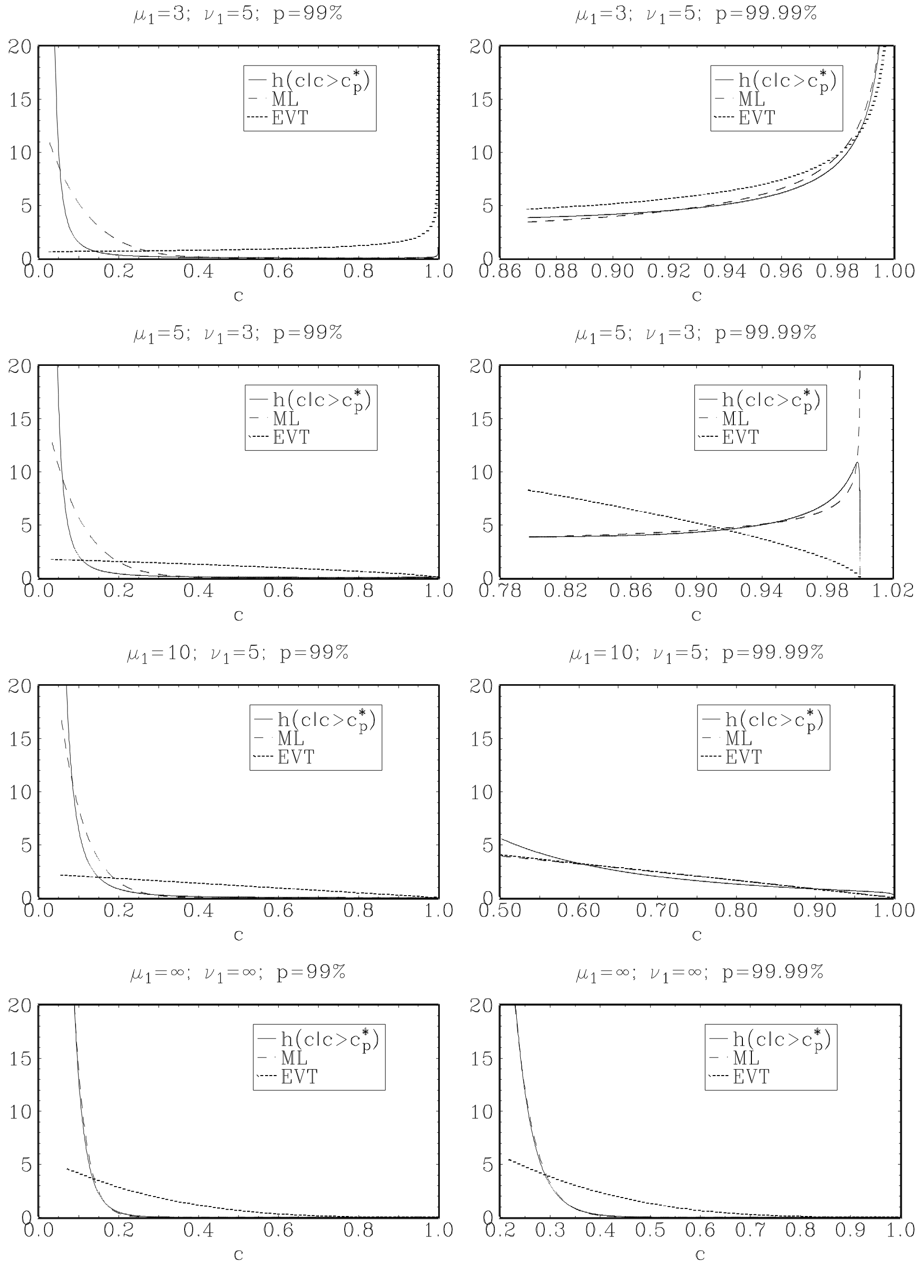

Figure 2:
The figure contains the tail shape of credit loss distributios for the linear factor model (13) with Student t distributed risk factors. The degrees of freedom parameters for systematic risk f and idiosyncratic risk ε_j are μ_1 and ν_1, respectively. The expected loss is 1%, and its standard deviation 1.5%. Each panel depicts the true conditional pdf of credit losses $h(c|c > c_p^*)$, its approximation based on EVT, and its approximation based on an ML estimate over the range $c_p^* \dots 1$, with c_p^* the pth quantile of the credit loss distribution. We consider $p = 99\%$ and $p = 99.99\%$.