Stress Testing in Credit Portfolio Models

Michael Kalkbrener
Deutsche Bank AG, Risk Methodology[1];

Ludger Overbeck
University of Gießen, Department of Mathematics.

January 3, 2016

Abstract

As, in light of the recent financial crises, stress tests have become an integral part of risk management and banking supervision, the analysis and understanding of risk model behaviour under stress has become ever more important. In this paper, we present a general approach to implementing stress scenarios in a multi-factor credit portfolio model and analyse asset correlations, default probabilities and default correlations under stress. We use our results to study the implications for credit reserves and capital requirements and illustrate the proposed methodology by stressing a large investment banking portfolio. Although our stress testing approach is developed in a particular credit portfolio model, the main concept - stressing risk factors through a truncation of their distributions - is independent of the model specification and can be applied to other risk types as well.

1 Introduction

Stress testing has been adopted as a generic term describing various techniques used by financial firms to analyze their potential vulnerability to extreme yet plausible events, see para 718 in Basel Committee on Banking Supervision [2006] for specific requirements on banks' stress testing programs. Stress scenarios have long been used in risk management to supplement risk measures like value-at-risk (VaR) and economic capital (EC), e.g. Kupiec [1998] and Berkowitz [2000], but stress testing has gained new prominence in the aftermath of the subprime crisis and the European sovereign debt crisis. In particular, it has become an integral part of banking supervision, which is reflected in regulatory stress testing programs such as the annual Comprehensive Capital Assessment Review (CCAR) performed by the FED since 2010 (Board of Governors of the Federal Reserve, [2012]) and the EU-Wide Stress Tests, see European Banking Authority [2011]. Principles of sound stress testing practices have been laid down by the Basel Committee on Banking Supervision [2009], analysis and surveys of macroeconomic stress testing can be found in Čihák [2007], Alfaro and Drehmann [2009], Drehmann [2009], Quagliarello [2009] and Borio et al. [2012].

An important challenge in designing effective stress tests is the selection of scenarios that are both severe and plausible. One approach frequently used by risk managers is the application of historical scenarios such as the 1987 stock market crash or the subprime crisis. By their very nature, historical scenarios are plausible and provide useful information on the sensitivity of a portfolio to specific market shocks but they restrict attention to prior stress episodes. Hypothetical scenarios, in contrast, are not constrained to replicate specific past incidents and can therefore cover a broader spectrum of potential risks. However, depending on the choice of the hypothetical scenarios, stress test results might misrepresent risks either because the most dangerous scenarios are not considered or because

[1]The views expressed in this paper are those of the author and do not necessarily reflect the position of Deutsche Bank AG.

the selected scenarios are too implausible. In order to overcome this problem, systematic approaches to scenario selection have been investigated for more than 15 years, e.g. Studer [1999]. More recent work on that subject includes Breuer et al. [2009], Breuer and Csiszár [2013], Flood and Korenko [2015] and Glasserman et al. [2015].

In this paper, we present an alternative approach to the specification of stress scenarios, which has initially been introduced in Bonti et al. [2006] for analyzing credit concentrations. Duellmann and Erdelmeier [2009] use the same methodology for stressing credit portfolios of German banks. In this approach, statistical EC or VaR models serve as quantitative framework for the specification of stress scenarios. More precisely, stress scenarios are defined through constraints on the risk factors of the model. These constraints are then used to truncate the distribution of the stressed risk factors or - in other words - restrict the state space of the model, where each state represents values of the risk factors. The response of the peripheral (or unstressed) risk factors is specified by the dependence structure of the model. As an example, consider an economic downturn in the automotive sector. In a structural credit portfolio model with industry and country factors this scenario can be implemented by truncating the systematic risk factor for the automotive industry. The severity of the downturn scenario is reflected through the truncation threshold, so that a lower threshold implies more severe stress. Since the automotive industry is positively correlated to most industry and country factors non-automotive exposures are affected as well.

The specification of stress scenarios through constraints on risk factors of VaR or EC models has a number of advantages:

1. Stress scenarios are implemented in a way that is consistent with the existing quantitative framework. This implies that the relationships between (unrestricted) risk factors remain intact and the experience gained in the day-to-day use of the model can be utilized in the interpretation of stress testing results. It has to be analyzed, however, whether historical correlation patterns, which are typically used for calibrating (unstressed) risk capital models, provide an appropriate dependence structure for stress testing, see section 4 for a sensitivity analysis of model correlations under stress.

2. In a given stress scenario, risk factors are not set to deterministic values but remain stochastic variables, i.e., stressed as well as unstressed factors follow a joint distribution conditional on the truncation thresholds that define the stress scenario. This feature distinguishes our approach from standard stress tests, which are typically based on deterministic stress scenarios. As a consequence, stressed risk measures, e.g. expected loss, value-at-risk or economic capital, can be calculated in each stress scenario.

3. The probability of each stress scenario, e.g. the probability that the risk factors satisfy all the constraints under non-stress conditions, can be easily calculated in the statistical model. This is a good indicator for the severity of a stress scenario.

Our stress testing methodology is developed in a multi-factor credit portfolio model. We provide details on the implementation of stress scenarios and discuss practical issues such as the calculation of truncation thresholds in multi-factor stress scenarios. Another objective of this paper is to review recent results on stressed asset correlations, default probabilities and default correlations presented in Kalkbrener and Packham [2015a] and Packham et al. [2014]. In these papers, the analysis is performed in a factor model that follows a normal variance mixture distribution, which covers a wide range of light-tailed to heavy-tailed distributions. Aside from analysing the behaviour under stress for given stress levels or stress probabilities, the asymptotic behaviour, that is, the behaviour under stress as the stress level becomes arbitrarily high, is investigated. Contrary to popular belief, it is shown that the impact of stress on the asymptotic behaviour is greater in light-tailed models than in heavy-tailed models. More specifically,

- asset correlations under stress are less sensitive for heavy-tailed models than light-tailed models;
- default correlations under stress converge to 0 for light-tailed models and to a number strictly greater than 0 for heavy-tailed models;
- default probabilities converge to 1 for light-tailed models and to a number strictly smaller than 1 for heavy-tailed models.

However, the asymptotic behaviour of stresses PDs is not representative for ordinary stress tests: only for rather extreme stress severities, stressed PD's become higher in light-tailed than in heavy-tailed models. Finally, these results are used to study the implications for risk measures, credit reserves and capital requirements under stress.

The paper is structured in the following way. The second section introduces the quantitative framework we will work in. The third section describes our approach to implementing stress scenarios in a multi-factor credit portfolio model. In addition, results from stressing a sample portfolio are presented. In section 4, the impact of stress on asset correlations, default probabilities and default correlations is analyzed. Section 5 concludes.

2 Quantitative framework for stress testing

The objective of this section is the introduction of a class of multi-factor credit portfolio models that serve as the formal framework for the implementation of stress scenarios.

In a typical bank, the economic as well as regulatory capital charge for credit risk far outweighs capital for any other risk class. Key drivers of credit risk capital are concentrations in a bank's credit portfolio, either caused by material concentrations of exposure to individual names or large exposures to a single sector or to several highly correlated sectors. As a consequence, the stress testing methodology for credit risk has to be implemented in a credit portfolio model that provides sufficient flexibility for modeling risk concentrations.

The IRB approach in Basel Committee on Banking Supervision [2006] does not provide an appropriate quantitative framework. It is based on a credit portfolio model that was originally designed to produce portfolio-invariant capital charges. However, it is only applicable under the assumptions that (cf. Gordy [2003])

1. bank portfolios are perfectly fine-grained and
2. there is only a single source of systematic risk.

The simplicity of the model ensures its analytical tractability. However, it makes it impossible to model risk concentrations in a reasonable way.

In order to develop meaningful stress tests, we need to generalize the IRB approach to a multi-factor credit portfolio model that takes into account individual exposures and has a richer correlation structure. In this paper, we use a structural model (Merton [1974]), which links the default of a firm to the relationship between its assets and the liabilities that it faces at the end of a given time period $[0, T]$.[2]

More generally, in a structural credit portfolio model the j-th obligor defaults if its ability-to-pay variable A_j falls below a default threshold c_j: the default event at time T is defined as $\{A_j \leq c_j\} \subseteq \Omega$,

[2]A survey on credit portfolio modeling can be found in Bluhm et al. [2002] and McNeil et al. [2005].

where A_j is a real-valued random variable on the probability space $(\Omega, \mathcal{A}, \mathbb{P})$ and $c_j \in \mathbb{R}$. We denote the default indicator $\mathbf{1}_{\{A_j \leq c_j\}}$ of the j-th obligor and its default probability $\mathbb{P}(\{A_j \leq c_j\})$ by I_j and p_j respectively. The portfolio loss variable is defined by

$$L := \sum_{j=1}^{n} l_j \cdot I_j, \tag{1}$$

where n denotes the number of obligors and l_j is the loss-at-default of the j-th obligor. In order to reflect risk concentrations, a joint distribution of the A_j has to be specified that captures the dependence between defaults of different obligors. This is done via the introduction of a factor model consisting of systematic and idiosyncratic factors. More precisely, each ability-to-pay variable A_j is decomposed into a sum of systematic factors $\Psi_1, \ldots, \Psi_m$ and an idiosyncratic [or specific] factor ε_j, that is

$$A_j = \sqrt{R_j^2} \sum_{i=1}^{m} w_{ji} \Psi_i + \sqrt{1 - R_j^2} \varepsilon_j. \tag{2}$$

It is usually assumed that the vector of systematic factors $\Psi = (\Psi_1, \ldots, \Psi_m)$ follows an m-dimensional normal distribution with mean $\mathbf{0} = (0, \ldots, 0)$ and covariance matrix $\Sigma = (\Sigma_{kl})$. The systematic weights $w_{j1}, \ldots, w_{jm} \in \mathbb{R}$ determine the impact of each systematic factor on the ability-to-pay variable A_j. The systematic weights are scaled such that the systematic component

$$\phi_j := \sum_{i=1}^{m} w_{ji} \Psi_i \tag{3}$$

is a standardized normally distributed variable, i.e., ϕ_j has mean 0 and variance 1. The idiosyncratic factors $\varepsilon_1, \ldots, \varepsilon_n$ are standardized normally distributed variables, they are independent of each other as well as independent of the systematic factors. Each R_j^2 is an element of the unit interval $[0, 1]$. It determines the impact of the systematic component on A_j and therefore the correlation between A_j and ϕ_j: it immediately follows from (2) that

$$R_j^2 = \mathrm{Corr}(A_j, \phi_j)^2. \tag{4}$$

In order to quantify portfolio risk, measures of risk are applied to the portfolio loss distribution (1). The most widely used risk measures in banking are value-at-risk and expected shortfall: value-at-risk $\mathrm{VaR}_\alpha(L)$ of L at level $\alpha \in (0, 1)$ is simply an α-quantile of L whereas expected shortfall of L at level α is defined by

$$\mathrm{ES}_\alpha(L) := (1 - \alpha)^{-1} \int_\alpha^1 \mathrm{VaR}_u(L) du.$$

For most practical applications the average of all losses above the α-quantile is a good approximation of $\mathrm{ES}_\alpha(L)$: for $c := \mathrm{VaR}_\alpha(L)$ we have

$$\mathrm{ES}_\alpha(L) \approx \mathrm{E}(L | L > c) = (1 - \alpha)^{-1} \int L \cdot \mathbf{1}_{\{L > c\}} \, d\mathbb{P}.$$

These risk measures are used to determine the economic capital, which is designed to state with a high degree of certainty the amount of capital needed to absorb unexpected losses. Economic capital $\mathrm{EC}(L)$ is usually defined as value-at-risk $\mathrm{VaR}_\alpha(L)$ at a high level α, e.g., $\alpha = 0.9998$, minus the expected loss $\mathrm{E}(L)$ of L:

$$\mathrm{EC}(L) := \mathrm{VaR}_\alpha(L) - \mathrm{E}(L),$$

where the subtraction of the expected loss reflects the fact that only unexpected losses are covered by economic capital.

2.1 Definition of asset and default correlations

The critical quantities entering the risk measures defined above are the default probabilities and the risk concentrations of the default indicators I_j, either specified by default or asset correlations. In this subsection, we provide a formal definition of these quantities, an analysis of default or asset correlations under stress is performed in section 4.

The default or event correlation ρ_{ij}^D of obligors i and j, with $i \neq j$, is defined as the correlation $\mathrm{Corr}(I_i, I_j)$ of the corresponding default indicators. Because

$$\mathrm{Var}(I_j) = \mathrm{E}(I_j^2) - p_j^2 = p_j - p_j^2,$$

the default correlation equals

$$\rho_{ij}^D = \mathrm{Corr}(I_i, I_j) = \frac{\mathrm{E}(I_i I_j) - p_i p_j}{\sqrt{(p_i - p_i^2)(p_j - p_j^2)}}. \tag{5}$$

The indicator variables I_j are defined in terms of ability-to-pay variables A_j, which are typically interpreted as log-returns of asset value processes. The correlation $\mathrm{Corr}(A_i, A_j)$ is therefore called the asset correlation ρ_{ij}^A of obligors $i \neq j$. As an immediate consequence of (2), the correlation as well as the covariance of the ability-to-pay variables of the counterparties i and j are given by

$$\mathrm{Corr}(A_i, A_j) = \mathrm{Cov}(A_i, A_j) = \sqrt{R_i^2}\sqrt{R_j^2}\sum_{k,l=1}^{m} \omega_{ik}\omega_{jl}\mathrm{Cov}(\psi_k, \psi_l). \tag{6}$$

There exists an obvious link between default and asset correlations. For given default probabilities, the default correlation ρ_{ij}^D is determined by $\mathrm{E}(I_i I_j)$ according to (5), and

$$\mathrm{E}(I_i I_j) = \mathbb{P}(A_i \leq c_i, A_j \leq c_j) = \int_{-\infty}^{c_i}\int_{-\infty}^{c_j} f_{ij}(u, v)\,du\,dv,$$

where $f_{ij}(u, v)$ is the 2-dimensional joint density function of A_i and A_j. Hence, default correlations depend on the joint distribution of A_i and A_j. If (A_i, A_j) is bivariate normal the correlation of A_i and A_j determines the copula of their joint distribution and hence the default correlation:

$$\mathrm{E}(I_i I_j) = \frac{1}{2\pi\sqrt{1 - {\rho_{ij}^A}^2}}\int_{-\infty}^{c_i}\int_{-\infty}^{c_j} \exp\!\left(-\frac{1}{2(1 - {\rho_{ij}^A}^2)}(u^2 - 2\rho_{ij}^A uv + v^2)\right)du\,dv. \tag{7}$$

Note, however, that for general ability-to-pay variables outside the multivariate normal class, the asset correlations do not fully determine the default correlations.

3 Factor Stress Methodology

In this section, we describe each of the steps of the stress testing process:

1. Specification of an economic stress scenario or scenario based on the characteristics of the portfolio

2. Translation of the scenario into constraints on the systematic factors of the credit portfolio model

3. Quantification of the impact of the stress scenario by calculating the conditional expected loss and other statistics of the portfolio

3.1 Specification of stress scenarios

The following classification should serve as a rough guide and distinguish different types of stress scenarios.

1. *Macroeconomic scenarios.* A macroeconomic scenario usually requires the use of a macroeconomic model. It specifies an exogenous shock to the whole economy that is propagated over time and may impact the banking system in various ways. This type of stress scenario is used by financial regulators or central banks in order to gain an understanding of the resilience of financial markets or the banking system as a whole.

2. *Market shocks.* These scenarios specify shocks to financial markets. This category also includes certain shocks of a "systemic" nature affecting credit risk (such as a sudden flight to liquidity), or sectoral shocks, for instance the deterioration in credit spreads in the TMT (Technology Media-Telecommunications) sector. Historical scenarios are frequently used for this type of shocks in order to increase the plausibility of these stress scenarios.

3. *Portfolio specific worst case scenarios.* The objective of this worst case analysis is to identify scenarios that are most adverse for a given portfolio. The specification of worst case scenarios can either be based on expert judgement or quantitative techniques.

These scenario types serve different purposes. Economic stress scenarios and market shocks are usually specified by risk management. The objective is to quantify the impact of a plausible economic downturn or a market shock on a credit portfolio.

The aggregated loss of portfolio specific worst case scenarios, on the other hand, serves more as a benchmark to create some awareness of the current market situation. The construction of these scenarios is driven by portfolio characteristics instead of economic considerations.

Regardless of the motivation for considering a particular scenario, there exist a number of criteria that characterize useful stress scenarios:

1. *Plausible.* Stress scenarios must be realistic, e.g. have a certain probability of actually occurring. Risk management will not take any actions based on scenarios that are regarded as implausible.

2. *Consistent.* One objective is to implement stress scenarios in a way that is consistent with the existing quantitative framework. This has the advantage that the relationships between risk factors remain intact and the experience gained in the day-to-day use of the model can be utilized in the interpretation of stress testing results.

3. *Adapted.* Stress tests should include scenarios that are specifically designed for the portfolio at hand. They should reflect certain portfolio characteristics and particular concerns in order to give a complete picture of the risks inherent in the portfolio.

4. *Reportable.* Stress scenarios should provide useful information for risk management purposes, which can be translated into concrete actions. For reporting purposes, it is crucial that the stress scenario is characterized by a clearly identifiable set of stressed risk factors, sometimes called the "core" factors. The remaining "peripheral" factors should then move in a consistent way with those "core" factors.

When designing specific stress scenarios, we usually focus on a small number of directly stressed factors, e.g. those factors that correspond to the sectors of interest. In addition, a small number of stressed factors makes it easier to transform the stress results into concrete management actions.

The response of the other risk factors is specified by the dependence structure of the model. This approach is also a superior way to identify risk concentrations compared to just aggregating exposures per sector, because there it can happen that concentrations in distinct but highly correlated sectors remain undetected.

3.2 Implementation of stress scenarios in credit portfolio models

In order to translate a given stress scenario into model constraints, a precise meaning has to be given to the systematic factors of the portfolio model. Recall that each ability-to-pay variable

$$A_j = \sqrt{R_j^2}\sum_{i=1}^{m} w_{ji}\Psi_i + \sqrt{1-R_j^2}\varepsilon_j$$

is a weighted sum of m systematic factors $\Psi_1, \ldots, \Psi_m$ and one specific factor ε_j. The systematic factors often correspond either to countries (or geographic regions) and industries. Equity data is frequently used to construct time-series for the systematic factors. Statistical techniques are then applied to these time-series to derive the joint distribution of the systematic factors. The systematic weights w_{ji} are chosen according to the relative importance of the corresponding factors for the given counterparty. They are either based on economic information or calculated via statistical techniques such as linear regression.

The economic interpretation of the systematic factors is essential for implementing stress scenarios in the model. The actual translation of a scenario into model constraints is done in two steps:

1. Identification of the appropriate risk factors based on their economic interpretation
2. Truncation of their distributions by specifying upper bounds that determine the severity of the stress scenario

Using the credit portfolio model introduced in section 2 as quantitative framework, the specification of the model constraints is formalized as follows. A subset $S \subseteq \{1, \ldots, m\}$ is defined, which identifies the stressed factors Ψ_i, $i \in S$. For each of these factors a cap $C_i \in \mathbb{R}$ is specified. The purpose of the thresholds C_i, $i \in S$, is to restrict the sample space of the model. More formally, the restricted sample space $\bar{\Omega} \subseteq \Omega$ is defined by

$$\bar{\Omega} := \{\omega \in \Omega \mid \Psi_i(\omega) \leq C_i \text{ for all } i \in S\}. \tag{8}$$

In other words, $\omega \in \Omega$ is an element of the restricted sample space $\bar{\Omega}$ if none of the stressed factors exceeds its threshold in the event ω. Note that the probability $\mathbb{P}(\bar{\Omega})$ of the restricted sample space $\bar{\Omega}$ under the original probability measure $\mathbb{P}$ provides information on the likelihood of the stress scenario.

Although the formal framework for implementing stress scenarios is simple the actual translation of scenarios into model constraints can be rather complex depending on the specification of the scenario. If a scenario is defined in terms of constraints on the existing systematic country and industry factors the implementation is straightforward. However, even the identification of systematic risk factors is a difficult problem if the given scenario specification involves economic variables that cannot easily be mapped to the country and industry classification used in the model, e.g. the implementation of a drop in US house prices would require an analysis of the potential impact on different countries and industries before the scenario can be translated into model constraints. A more transparent approach, however, is

1. to add a US house price index to the set of systematic factors,
2. to extend the joint distribution of systematic factors in order to capture the dependence between US house prices and the country and industry factors of the model and

3. to implement this stress scenario through a constraint on the new factor.

It is important to note that the new macroeconomic factor - in the present example the US house price index - is not included in the decomposition of the ability-to-pay variable in (2), i.e. the US house price index has a weight of zero in all ability-to-pay variables. As a consequence, the behaviour of the unstressed model is not affected.

However, the dependence between new macroeconomic factors, denoted by $\Xi_1, .., \Xi_k$, and the industry and country factors $\Psi_1, ..\Psi_m$ is captured in the extended covariance matrix of the larger factor model $(\Psi_1, ..\Psi_m, \Xi_1, .., \Xi_k)$. In a stress scenario, the conditional distribution $\mathcal{L}((\Psi_1, ..\Psi_m)|\Xi_1 \leq C_1, .., \Xi_k \leq C_k)$ of the country and industry factors given the constraint on the macroeconomic factors is used in (2) to obtain the stressed ability-to-pay variables. Therefore, the constraints on macroeconomic factors have an impact on the distribution of the country and industry factors in a stress scenario and, consequently, also on the ability-to-pay variables of all counterparties.

The above example illustrates that, in principle, the initial set of country and industry factors can be extended by a large number of macroeconomic and market factors in order to provide a comprehensive model for stress testing. However, the specification of the joint distribution of these different factors $(\Psi_1, ..\Psi_m, \Xi_1, .., \Xi_k)$ is a challenging problem due to differences in the data frequency, e.g. quarterly GDP data versus daily market data, potential time lags between market and macroeconomic variables, etc.

Stress tests are frequently specified by setting the respective risk factors to specific values, e.g. a 10% drop in US house prices in a stress scenario compared to a 2% increase in the baseline scenario. In order to implement this scenario in our model the 10% drop has to be translated into a truncation threshold:

1. using historic house price volatility together with the baseline scenario we calibrate a distribution of US house prices changes and

2. based on that distribution, we specify the truncation threshold C such that the conditional mean, i.e., the average of US house price changes below C, equals the 10% drop.

This technique can be generalized to a multi-factor stress scenario. However, if a stress scenario is not consistent with the correlation structure of the model, e.g. if two factors behave differently in the stress scenario although they are almost perfectly correlated in the underlying model, it will not be possible to precisely replicate the specified stress values through multi-dimensional thresholds. In this case, an optimization problem has to be solved instead that results in thresholds that provide the best possible replication but not a perfect match.

Restricting the state space through constraints on systematic factors is a flexible technique to incorporate stress scenarios into the portfolio model. So far, we have only considered stress scenarios that are defined by truncating factor distributions. Alternatively, stress scenarios could be defined via defining more complex constraints than simple caps on individual factors. One possibility is to restrict the state space of the model in such a way that the dependence of particular risk factors is increased. This technique provides an interesting alternative to simply changing correlation parameters of the model. By keeping the original model parameters intact, consistency problems are avoided such as maintaining the positive semi-definiteness of the correlation matrix of the systematic factors.

3.3 Calculation of stressed risk capital

The actual calculation of the stressed loss distribution of the portfolio is done through Monte Carlo simulation on the restricted model space $(\bar{\Omega}, \bar{\mathcal{A}}, \bar{\mathbb{P}})$, see (8). It is therefore straightforward to calculate risk measures like expected loss, value-at-risk or expected shortfall for the loss distribution under stress and to use statistical techniques such as QQ-plots to study its behavior.

It depends on the particular purpose of a stress test which of those risk measures is used to quantify the impact of a stress test on the credit portfolio. One possibility is to analyze whether current capital requirements cover realized losses in stress scenarios and to use stress tests for the calculation of the conditional expected loss. Another application of stress tests is the analysis of future capital requirements, e.g. the bank wishes to satisfy its EC constraint one year into the future. If the stress event arrives within the one year horizon, then the bank will need capital sufficient to meet its EC requirement conditional on that stress event. This type of analysis requires the calculation of the VaR of the stressed portfolio. Finally, the future regulatory capital requirements in stress scenarios can be assessed by recalculating the Basel II formula with the stressed PDs from the multi-factor model. Since regulatory capital requirements are essential for capital management and strategic planning we regard this impact analysis as an important component of the stress testing methodology in a financial institution.

In the following, we will describe our approach by means of a specific scenario. As an example, consider a downturn scenario for the automotive industry. The simplest implementation in the portfolio model is the following restriction of the state space of the model: only those samples are considered in the Monte Carlo simulation where the automotive industry factor decreases by a certain percentage, say at least 2%. In other words, the distribution of the automotive industry factor is truncated from above at -2%. More precisely, the steps in the calculation of stressed EL and EC are:

- simulate risk factors under their original (non-stress) joint distribution,
- dismiss any simulation not satisfying the scenario constraints,
- derive EL, EC and other statistics from the loss distribution specified by the MC scenarios that satisfy the constraints.

Note that the automotive downturn scenario does not only have an impact on the automotive industry factor: because of correlations, other country factors as well as industry factors are also affected. Figure 1 shows the stressed distribution of the automotive industry factor (left) and the impact on the factor for the chemical industry (right): the distribution of the automobile factor has been truncated, while the distribution of the chemical industry factor is no longer centered but has moved to the left.[3]

3.4 Case study

We consider the following downturn scenario for the automotive industry: the industry production is forecast to drop by 8% during next year. Using the methodology presented in section 3.2 this forecast is translated into a cap on the distribution of the automobile factor.

In this case study, the stress is applied to a sample investment banking portfolio, which consists of 25000 loans with an inhomogeneous exposure and default probability distribution. Its total exposure is 1000 mn EUR, average exposure size is 0.004% of the total exposure and the standard deviation of the exposure size is 0.026%. Default probabilities vary between 0.02% and 27%. Figure 2 exhibits the portfolio's exposure by rating class both for automotive companies and all other borrowers.

Application of the downturn scenario yields the following risk estimates:

[3]The distributions in figure 1 can be represented in a simple way: if $F_{auto}(x)$ denotes the (Gaussian) distribution of the automobile factor, its truncated distribution is given by $F_{auto}(x)/F_{auto}(-2\%)$ for $x \leq -2\%$. The factor for the chemical industry is called an incidentally truncated variable. Its marginal distribution is given by $F_{auto,chem}(-2\%, y)/F_{auto}(-2\%)$, where $F_{auto,chem}$ denotes the joint distribution of the two industry factors.

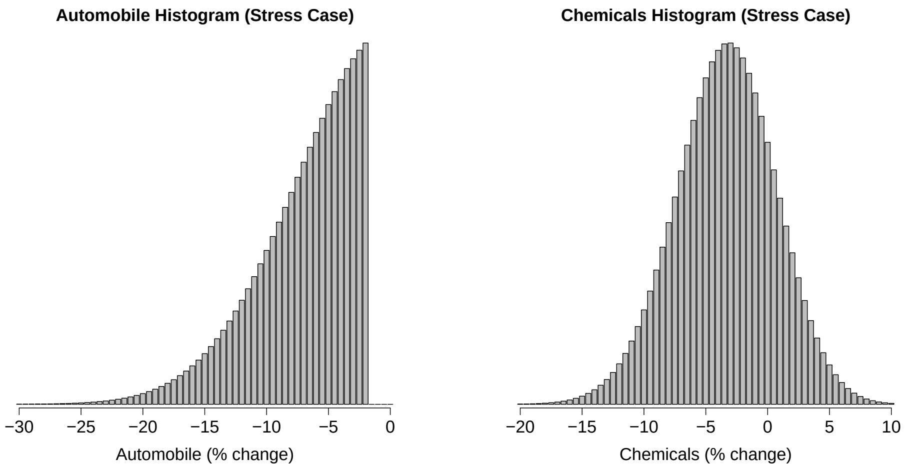

Figure 1: Histogram of simulated factor changes (stress case)

	Non-stress	Stress	% chg.
Expected loss	7.03	10.94	55.6
99.98% VaR	103.23	122.80	19.0
Expected shortfall at 99.98%	119.68	145.45	21.5
Economic capital	96.20	111.86	16.3

Table 1: Portfolio risk estimates

These key statistics provide important information on the impact of the stress scenario. The 99.98% confidence interval has been chosen because we use the corresponding value-at-risk for the EC calculation. Note that the relative EL increase of 55.6% is significantly higher than the 19% increase of the 99.98% VaR. This results in a 16.3% increase of economic capital defined as 99.98% VaR minus EL.

Figure 2 exhibits the portfolio's exposure by rating class both in the non-stress and stress case. The analysis is done separately for automotive companies and all other borrowers. Figure 2 clearly shows that exposure is shifted from investments grades (BBB or above) to non-investment grades. As expected, the deterioration of ratings is more pronounced for the automotive industry. Note, however, that due to the dependence structure of the portfolio this stress scenario also has a significant impact on other borrowers.

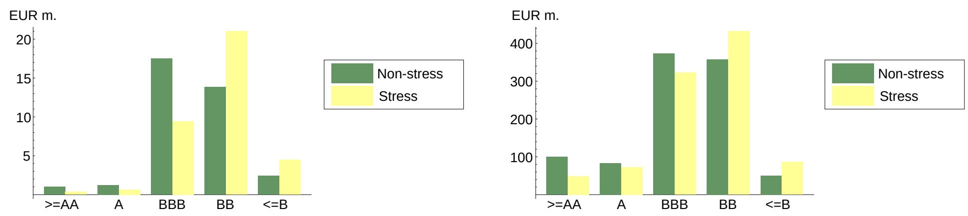

Figure 2: Exposure by rating class for automotive companies (left) and all other borrowers (right)

Rather than just looking at certain quantiles or other summary statistics, we can get a better understanding of the impact of a stress scenario by studying the whole loss distribution before and after the stress. In order to see the effect of the automotive stress scenario on the portfolio loss, the left graph of figure 3 shows the original (circles) and the stressed (triangles) loss densities, together with fitted Vasicek distributions (curves). The corresponding QQ-plot, i.e., the quantiles of the two distributions plotted against each other, is shown in the right graph.

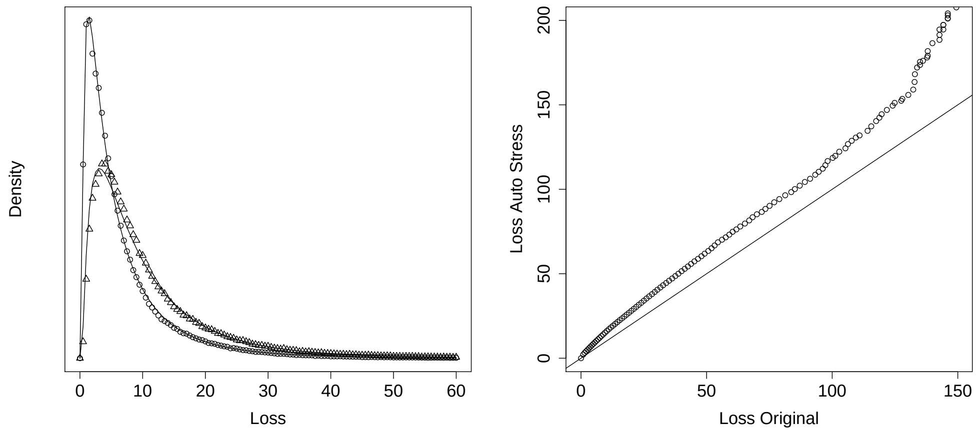

Figure 3: Left graph: Density plots of original (circles) and stressed (triangles) loss distributions, together with fitted Vasicek curves. Right graph: QQ plot of original against stressed loss distribution.

The final step in this case study is the calculation of the regulatory capital requirements conditional on the stress event: recalculating the Basel II formula with the stressed PDs increases the regulatory capital from 131.41mn to 156.48mn. In this example, the increase of 19% is in line with the increase of the 99.98% quantile (see table 1).

4 Stressed correlations and default probabilities

In the above case study, the expected loss of the portfolio is increased by more than 50% under stress whereas the proportional EC increase is significantly lower. In order to better understand the high sensitivity of the expected loss we analyse the behaviour of default probabilities in stress scenarios, see section 4.3. Whereas default probabilities are the only relevant component for the EL, stressed EC also depends on the correlations in the stressed model. Section 4.2 deals with stressed asset correlations, an analysis of stressed default correlations is part of section 4.3. Our presentation follows Kalkbrener and Packham (2015b).

It is not surprising that the joint distribution of risk factors has a significant impact on the behaviour of default probabilities and correlations under stress. In order to cover a wide range of light-tailed to heavy-tailed distributions we perform our analysis in factor models that follow a normal variance mixture distribution, which is introduced in section 4.1.

4.1 Distribution of model variables

The standard approach in credit risk management is to model the risk factors and ability-to-pay variables through a joint multi-variate normal (aka Gaussian) distribution. In order to specify a more flexible dependence structure we introduce an additional random variable W, the so-called *mixing variable*, which is strictly positive and independent of the systematic and idiosyncratic factors. The definition of the ability-to-pay variables is generalized to

$$A_j = W(\sqrt{R_j^2}\sum_{i=1}^{m} w_{ji}\Psi_i + \sqrt{1-R_j^2}\varepsilon_j) \tag{9}$$

$$= \sqrt{R_j^2}\sum_{i=1}^{m} w_{ji}W\Psi_i + \sqrt{1-R_j^2}W\varepsilon_j \tag{10}$$

and the systematic and idiosyncratic risk factors now have the form $W\Psi_i$ and $W\varepsilon_j$ respectively. The ability-to-pay variables and risk factors specified in this way follow a so-called multivariate *normal variance mixture (NVM) distribution.* The most important distribution classes covered in this general model are the multivariate normal distribution, in which case the variable W equals 1, and the multivariate Student-t distribution, where W^2 follows an inverse gamma distribution. The Student-t distribution allows for more extreme events than the normal distribution and is therefore a commonly used alternative in financial modelling. Compared to the normal distribution, it takes one additional parameter, the so-called *degrees of freedom*, denoted by ν, that controls the heaviness of the tails. For more details we refer to McNeil et al. (2005).

In general, the tail behaviour of the risk factor W determines the so-called *heaviness* of the tails of the A_j: If the tail function $\mathbb{P}(W \geq x)$ follows a power law, e.g. $\mathbb{P}(W \geq x) \approx x^{-\nu}$ for a $\nu > 0$ and large x, then the ability-to-pay variables are said to have heavy tails. If W is bounded or its tail function decays exponentially, e.g. $\mathbb{P}(W \geq x) \approx e^{-x}$ for large x, then $A_1, \ldots, A_n$ are light-tailed.[4] The normally distributed model and the Student-t distributed model are examples of light-tailed and heavy-tailed models, respectively.

For the sake of simplicity, it will always be assumed that the first risk factor $W\Psi_1$ is truncated. We denote this factor by $V := W\Psi_1$.

4.2 Asset correlations under stress

For ability-to-pay variables (or asset returns) A_i and A_j we denote their (unconditional) correlation by ρ_{ij}, the correlation of A_i with risk factor V will be denoted by ρ_i.

It turns out that asset correlations are less sensitive to stress in heavy-tailed models than in light-tailed models. For illustration, we assume that A_1 and A_2 are normally distributed and set $\rho_{12} = 0.4$. Figure 4 shows the impact on asset correlations when risk factor V is truncated: The left plot shows a scatter plot of 5000 simulated samples of A_1 and A_2. All simulated scenarios are relevant in the unstressed model. In the right plot only those scenarios are shown where the stressed risk factor V does not exceed a threshold C, where C is chosen such that the stress probability $\mathbb{P}(V \leq C)$ equals 10%. As a consequence, only approximately 500 of the 5000 scenarios are considered under stress. Since the A_i and V have a positive correlation of 0.6 the average value of A_i in the stressed model is negative, which results in a higher number of defaults. It can also be observed that the asset correlation of 0.4 is significantly reduced under stress, i.e., the correlation of A_1 and A_2 drops to 0.1.

For comparison, we now repeat the calculation for heavy-tailed t-distributed A_i using the same correlation assumptions as in Figure 4. The left graph of Figure 5 shows stressed asset correlations, where instead of the stress level C, stress is expressed by stress probabilities, which are just the probabilities associated with the stress event, $\mathbb{P}(V \leq C)$. For instance, values at 10^{-1} correspond to a stress scenario with probability 10%. Stressed asset correlations are shown for normally distributed and t-distributed assets with degrees of freedom $\nu = 10$ and $\nu = 4$.

[4] The precise definition is based on the theory of regular variations, see McNeil et al. (2005). Heavy-tailed models correspond to a regularly varying tail function of W, whereas a model is light-tailed if W is bounded or its tail function is rapidly varying.

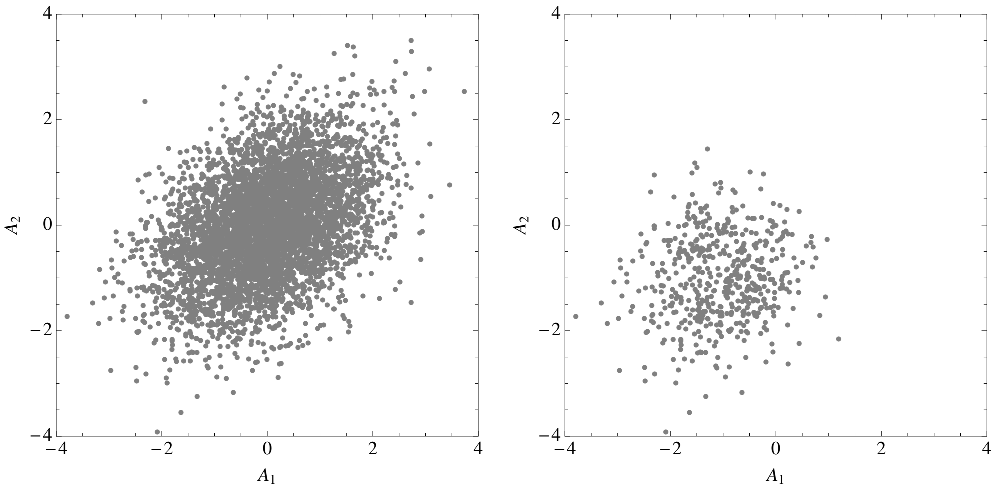

Figure 4: Left: Simulated normally distributed asset returns A_1 and A_2 with correlation 0.4; A_1 and A_2 are correlated to the joint driving risk factor V with correlation 0.6. Right: Samples conditional on $V \leq -1.28$ which corresponds to a stress event with probability 10%; the correlation of the sample is 0.1, which is far smaller than the original correlation of 0.4.

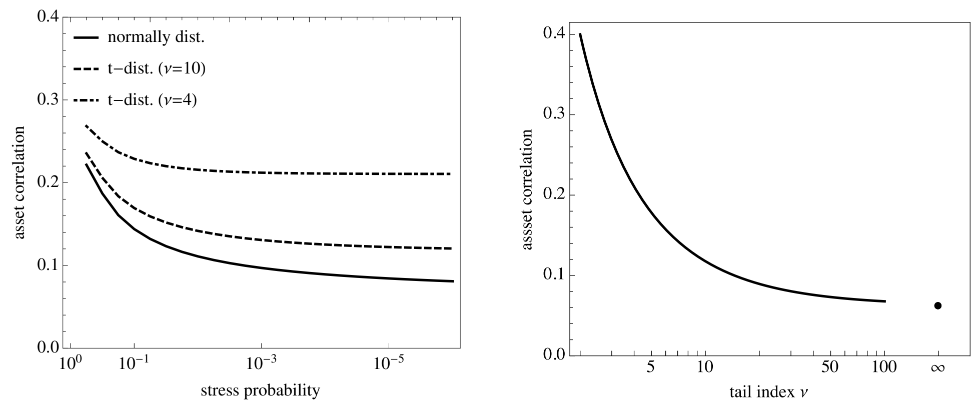

Figure 5: Left: Stressed asset correlation for different distribution assumptions as a function of the stress probability. Right: Stressed asset correlation as a function of the tail index when the stress event is taken to the limit $-\infty$. Correlations are as in Figure 4.

Stressed asset correlations may be either greater or smaller than the unconditional asset correlation depending largely on the correlations between the risk factor and the respective asset returns. As illustrated in Figure 5, when the assets in question are sufficiently correlated with the risk factor, the stressed correlation is typically smaller than the unstressed correlation. Loosely speaking, in such a case systematic risk is reduced by conditioning on the risk factor, whereas unsystematic risk remains.

The stressed correlations in the left graph of Figure 5 are calculated with analytic formulas derived in Kalkbrener and Packham (2015a). For normally distributed asset returns A_i, A_j their asset correlations conditional on stress level C are given by

$$\mathrm{Corr}^C(A_i, A_j) = \frac{\rho_i\,\rho_j\,\mathrm{Var}^C(V) + \rho_{ij} - \rho_i\,\rho_j}{\sqrt{(\rho_i^2\,\mathrm{Var}^C(V) + 1 - \rho_i^2)(\rho_j^2\,\mathrm{Var}^C(V) + 1 - \rho_j^2)}}, \tag{11}$$

with

$$\mathrm{Var}^C(V) = 1 - \frac{C\,\phi(C)}{\mathrm{N}(C)} - \frac{(\phi(C))^2}{(\mathrm{N}(C))^2},$$

where ϕ denotes the standard normal density function and N denotes the standard normal distribution function. A corresponding, but more involved, formula is also derived for the Student t-distribution.

The severity of the stress is increased by setting the stress level C to higher negative values, or equivalently, reducing the probability of the stress scenario specified by $\mathbb{P}(V \leq C)$. By letting $\mathbb{P}(V \leq C)$ converge to 0, e.g. by moving to the right in the left graph of Figure 5, we arrive at the asymptotic limit, which is of particular importance for understanding the model behaviour under stress. The right-hand side graph of Figure 5 shows the asymptotic limit of stressed asset correlations for t-distributed assets with different values for ν, where $\nu = \infty$ corresponds to the normally distributed case. The asymptotic analysis confirms the higher sensitivity of light-tailed asset variables under stress.

We have also derived concrete formulas for the asymptotic case, see Kalkbrener and Packham (2015a). These formulas hold in the more general setup of normal variance mixture models. For heavy-tailed NVM models the asymptotic limit of the stressed correlation of A_i and A_j equals

$$\frac{\rho_i\,\rho_j + (\rho_{ij} - \rho_i\,\rho_j)\,(\nu - 1)}{\sqrt{(\rho_i^2 + (1-\rho_i^2)\,(\alpha - 1))\,(\rho_j^2 + (1-\rho_j^2)\,(\nu - 1))}}, \quad \nu > 2, \tag{12}$$

if the risk factor is stressed asymptotically, i.e., if V is truncated at a threshold C, and C converges to $-\infty$. The parameter ν specifies the tail index of the asset returns and the risk factor in the heavy-tailed case and corresponds just to the degrees of freedom defined for t-distributions. The case when the variables are light-tailed corresponds to the limit as $\nu \to \infty$, in which case the asymptotic limit of the conditional correlation between A_i and A_j is

$$\frac{\rho_{ij} - \rho_i\,\rho_j}{\sqrt{(1-\rho_i^2)(1-\rho_j^2)}}. \tag{13}$$

Finally, note that the analysis in this section is not restricted to credit portfolio models but holds for any portfolio model with asset variables and risk factors that follow a normal variance mixture distribution.

4.3 Default probabilities and default correlations under stress

The credit-specific quantities entering credit portfolio models are the default probabilities and default correlations. Just as for asset correlations, their asymptotic behaviour depends on whether the credit portfolio model follows a light- or heavy-tailed NVM distribution. In the light-tailed case, default probabilities converge to 1 under extreme stress and default correlations converge to 0.[5] In other words, default of the entire portfolio becomes a sure event under extreme stress and correlations between default indicators become irrelevant.

In contrast, asymptotic default probabilities and asymptotic default correlations are in $(0, 1)$ in the heavy-tailed case. Both quantities depend on the tail index ν and can be expressed in terms of the Student t-distribution function. More specifically, the asymptotic default probability under stress for a model with tail index ν is given by Abdous et al. (2005) and Packham et al. (2014):

$$\lim_{C\to-\infty} \mathbb{P}(A_1 \leq D_1 | V \leq C) = t_{\nu+1}\left(\frac{\sqrt{\nu+1}\,\rho_1}{\sqrt{1-\rho_1^2}}\right) \in [1/2, 1), \tag{14}$$

where t_ν is the distribution function of the Student-t distribution with parameter ν. A formula for bivariate default probabilities – albeit more involved – and an integral representation for multivariate default probabilities that can be calculated numerically, are derived in Packham et al. (2014).

[5] In this subsection, we assume that the unconditional correlations between asset returns $A_1, \ldots, A_n$ and the risk factor V are positive and less than 1, i.e., $\rho_i, \rho_{ij} \in (0, 1)$ for $i, j \in \{1, \ldots, n\}$.

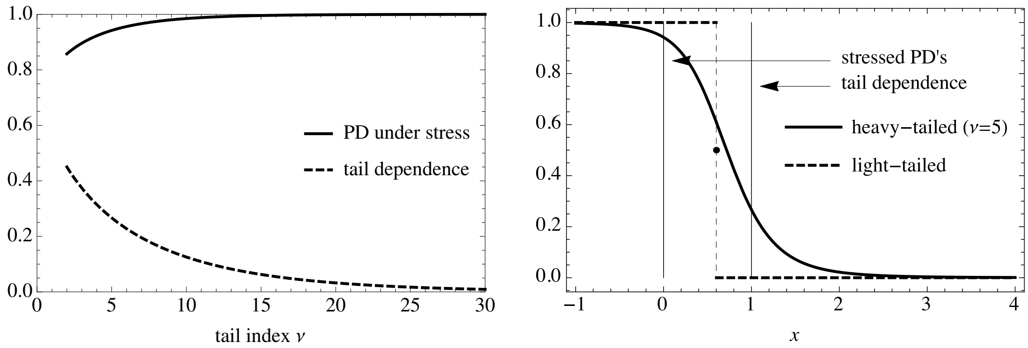

Figure 6: Left: Tail dependence coefficient and asymptotic PD under stress as a function of the tail index ν. Right: Tail dependence function $\lambda(V, A_1, x)$ for light- and heavy-tailed variables; special cases arise at $x = 0$ (stressed PD's) at $x = 1$ (tail dependence). The initial correlation between the ability-to-pay variable and the risk factor is 0.6 in both cases.

In all models – whether heavy-tailed or light-tailed – the asymptotic limit of stressed default probabilities and default correlations does not depend on the unstressed default probabilities. For the heavy-tailed case, the tail index and the unstressed correlations enter the asymptotic results.

In summary, the impact of stress on the asymptotic limit of default probabilities and correlations is greater in light-tailed models than in heavy-tailed models. This is a remarkable observation since light-tailed models, in particular normally distributed models, are usually considered less sensitive to extreme stress than heavy-tailed models: a popular measure in finance to assess the ability of a bivariate distribution to generate joint extreme events – the tail dependence – is zero in light-tailed models, whereas it is a positive number in heavy-tailed models. In order to better understand this phenomenon we now compare the behaviour of limiting default probabilities to tail dependence.

The tail dependence, or more precisely, the *coefficient of (lower) tail dependence* of the identically distributed variables V and A_1 is defined as[6]

$$\lambda_l(V, A_1) := \lim_{C \to -\infty} \mathbb{P}(A_1 \leq C | V \leq C). \tag{15}$$

Hence, the tail dependence of V and A_1 measures the probability $\mathbb{P}(A_1 \leq C)$ conditional on the event $\{V \leq C\}$ for stress levels C converging to $-\infty$. If the NVM distributed random variables V, A_1 are heavy-tailed with tail index ν, the tail dependence coefficient is given by

$$\lambda_l(V, A_1) = 2t_{\nu+1}\left(-\sqrt{\frac{(\nu+1)(1-\rho_1)}{1+\rho_1}}\right),$$

see McNeil et al. (2005). It follows that the tail dependence is strictly positive for heavy-tailed models, provided that $\rho_1 > -1$. For light-tailed NVM distributions, the tail dependence is zero. This includes, of course, the normal distribution, which is still the de-facto standard for modelling risk factors and asset log-returns in structural credit portfolio models, such as CreditMetrics™ (Gupton et al., 1997) and Moody's KMV Portfolio Manager™ (Crosbie and Bohn, 2002).

The zero tail dependence is in contrast to the asymptotic default probability in the light-tailed case, where default is a sure event. Similarly, tail dependence and asymptotically stressed default probabilities disagree in the heavy-tailed case. The left graph of Figure 6 illustrates the difference between tail dependence and asymptotic stressed PD's as a function of the tail index ν.

[6] In the general case, when V and A_1 are not identically distributed, the tail dependence coefficient is defined via quantiles, see McNeil et al. (2005).

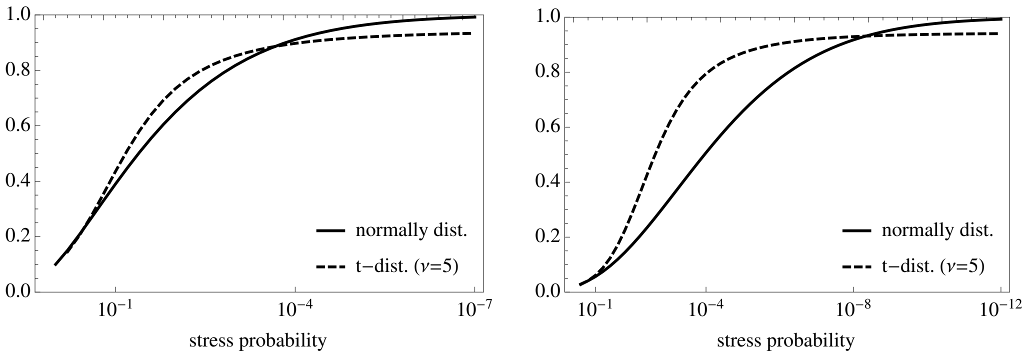

Figure 7: PD's under stress as a function of the stress probability. Models considered are the normal distribution and the t-distribution with parameter $\nu = 5$. Correlations are 0.6. Left: unconditional PD is 0.1. Right: unconditional PD is 0.01.

To make the relation between tail dependence and asymptotic stressed PD's more precise, we introduce an additional parameter $x \in \mathbb{R}$ and measure the probability $\mathbb{P}(A_1 \le x \cdot C)$ conditional on the event $\{V \le C\}$ for stress levels C converging to $-\infty$. More formally, we consider the function

$$\lambda(V, A_1, x) := \lim_{C \to -\infty} \mathbb{P}(A_1 \le x \cdot C | V \le C), \quad x \in \mathbb{R},$$

which provides an elegant generalization of both concepts: the tail dependence coefficient of V and A_1 equals $\lambda(V, A_1, 1)$, whereas the asymptotic stressed PD corresponds to $\lambda(V, A_1, 0)$. A closed-form expression for $\lambda(V, A_1, x)$ can be obtained via elementary transformations from Abdous et al. (2005), see also Packham et al. (2014). The tail dependence function is illustrated in the right graph of Figure 6.

The analysis of the function $\lambda(V, A_1, x)$ illustrates the fundamentally different behaviour of the tail dependence coefficient and asymptotic stressed PD's in light-tailed and heavy-tailed credit portfolio models. In the light-tailed case, the asset variable A_1 converges to $-\infty$, more specifically, it is concentrated at $\rho_1 \cdot C$ when $V \le C$ and $C \to -\infty$. In the heavy-tailed case, however, A_1 does not show the same uniform asymptotic behaviour: $0 < \lambda(V, A_1, x) < 1$ holds for all $x \in \mathbb{R}$ and, in particular, tail dependence as well as stressed default probabilities are in $(0, 1)$.

In summary, this analysis clearly shows that the tail dependence coefficient only provides partial information on a model's ability to produce extreme (joint) events. A more comprehensive picture is given by function $\lambda(V, A_1, x)$, which also explains the observed differences between tail dependence and asymptotic stressed PD's.

So far, our analysis has focused on asymptotic stressed default probabilities. For practical purposes, the model behaviour at smaller and therefore more realistic stress levels is even more important. Hence, we now take a closer look at PD's under stress for various stress levels C and compare them in light- and heavy-tailed models. Figure 7 shows PD's under stress for both normally distributed and t-distributed ($\nu = 3$) models as a function of the stress probabilities. The unconditional correlation between the ability-to-pay variable and the risk factor is 0.6. Despite converging to a value smaller than 1, PD's under stress in the t-distributed model dominate the normally distributed case unless the stress probability is very small: If the unconditional PD is 10%, then for stress probabilities greater than approximately $10^{-3.5}$, the PD under stress in the t-distributed model is greater than the respective PD in the normal model. If the unconditional PD is 1%, then the threshold lies beyond 10^{-8}.

This example shows that for realistic stress tests the impact on PD's is usually greater in heavy-tailed models. Only for rather extreme stress severities, stressed PD's become higher in light-tailed

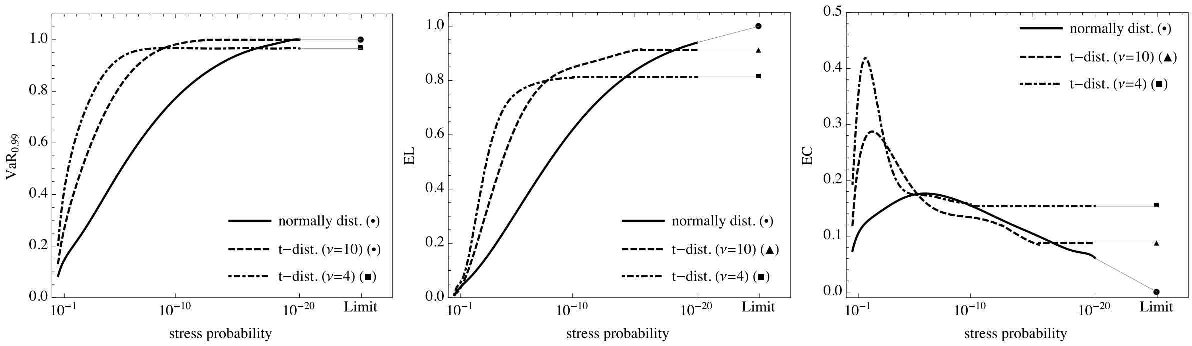

Figure 8: Risk measures for portfolio consisting of 60 homogeneous counterparties, each with a PD of 1%. Left: Value-at-risk at 99% confidence level; middle: Expected loss; right: Economic capital.

models and eventually converge to 1.

5 Risk measures

The different behaviour of light-tailed and heavy-tailed models has implications on the credit reserves and capital requirements in stress scenarios, as demonstrated by the following stylized example. Consider a homogeneous portfolio consisting of 60 counterparties. Each counterparty has notional and loss-at-default of 1/60 and defaults with a probability of 1%. The asset variables of the counterparties are correlated through one risk factor, with $\rho = 0.4$ the correlation between any one counterparty and the risk factor. This implies that the counterparties are correlated with $\rho^2 = 0.16$.

Figure 8 shows the value-at-risk, the expected loss and the resulting economic capital for the portfolio under different distribution assumptions, i.e., under a normal distribution and t-distributions with $\nu = 4$ and $\nu = 10$, and for different stress levels. As before, the stress level C is translated into a stress probability, which denotes the probability that a certain stress event occurs. The left graph shows the 99%-value-at-risk of the portfolio. Despite being lower under moderate stress, the VaR in a normally distributed model converges to 1, whereas the VaR for a very heavy-tailed model (t-distribution with $\nu = 4$) converges to a number strictly smaller than 1, see Packham et al. (2014) for the calculation of the asymptotic results. When comparing the two t-distributed models, the more heavy-tailed model with $\nu = 4$ has higher risk for moderate stress levels, but lower risk for less probable stress events.

Similar observations hold for the expected loss (middle graph). The expected loss under stress corresponds just to the probability of default under stress, since the recovery rate is 0 in this example. The asymptotic results, Equation (14), confirm that the EL converges to 1 in the light-tailed case, whereas it converges to a number strictly smaller than 1 in the heavy-tailed cases. Finally, economic capital converges to zero for normally distributed models and to a number strictly greater than zero for heavy-tailed models (a confidence level of 99% for economic capital may not be realistic in practice, but serves well to illustrate some key characteristics of the stressed portfolios). Because stress has different impact on value-at-risk and expected loss, economic capital is not monotone, but increases under moderate stress and decreases for greater stress levels.

To conclude, in light-tailed models, extreme stress scenarios tend to heavily increase the credit reserves specified by the expected loss whereas economic capital, which defines capital requirements, converges to 0. The impact of extreme stress on expected loss and economic capital is more balanced in heavy-tailed models, whose asymptotic limit retains a richer dependence structure.

6 Conclusion

In this paper, we have presented a general approach to implementing stress scenarios in a multi-factor credit portfolio model. The general philosophy behind this type of stress test is that stress scenarios are implemented through a restriction of the probability space of the model or, in other words, certain future scenarios are no longer considered possible. The calculation of the stressed portfolio loss distribution is done under a probability measure that contains additional information. The scenarios are then implemented in a way that is consistent with the quantitative framework, i.e., without destroying the dependence structure of risk factors in the model. This is achieved by translating the economic stress scenarios into constraints on the systematic factors. The main prerequisite here is that the systematic factors of the credit portfolio model can be linked to economic variables.

Although the methodology has been developed in a particular factor model, the main concept - implementing stress scenarios through a truncation of the distribution of the risk factors - is completely independent of the model specification and the way that default dependencies are parameterized, e.g. whether asset or default correlations are used. In fact, it can be applied to factor models for market and operational risk as well. However, the model choice has significant implications for the behavior of correlations under stress. In ordinary stress tests, stressed PD's are usually higher in heayy-tailed models. Contrary to popular belief, however, the impact of stress on the asymptotic behaviour is greater in light-tailed models than in heavy-tailed models.

References

Abdous, B., Fougères, A.-L., and Ghoudi, K. (2005). Extreme behaviour for bivariate elliptical distributions. *Canadian Journal of Statistics* **33**(3), 317-334.

Alfaro, R., and Drehmann, M. (2009). Macro Stress Tests and Crises: What Can We Learn? *BIS Quarterly Review* (December 2009), 2941.

Basel Committee on Banking Supervision (2006). International Convergence of Capital Measurement and Capital Standards: A Revised Framework, Comprehensive Version. Technical report, Bank for International Settlements.

Basel Committee on Banking Supervision (2009). Principles for sound stress testing practices and supervision. Technical report, Bank for International Settlements.

Berkowitz, J. (2000). A Coherent Framework for Stress-Testing. *Journal of Risk* **2**(2), 5-15.

Bluhm, C., Overbeck, L., and Wagner, C. (2002). *An Introduction to Credit Risk Modeling.* CRC Press/Chapman & Hall.

Board of Governors of the Federal Reserve (2012). Comprehensive Capital Analysis and Review 2012: Methodology and Results for Stress Scenario Projections. Technical Report, March 2012.

Bonti, G., Kalkbrener, M., Lotz, C., and Stahl, G. (2006). Credit Risk Concentrations under Stress. *Journal of Credit Risk* **2**(3), 115-136.

Borio, C., Drehmann, M., and Tsatsaronis, K. (2012). Stress Testing Macro Stress Testing: Does It Live up to Expectations? BIS Working Papers no. 369, Bank for International Settlements.

Breuer, T. and Csiszár, I. (2013). Systematic stress tests with entropic plausibility constraints. *Journal of Banking and Finance* **37**(5), 1552-1559.

Breuer, T., Jandačka, M., Rheinberger, K., and Summer, M. (2009). How to find plausible, severe, and useful stress scenarios. *International Journal of Central Banking*, **5**(3), 205-224.

Čihák, M. (2007). Introduction to Applied Stress Testing. IMF Working Paper No 07/59, International Monetary Fund, Washington DC.

Crosbie, B. and Bohn, J. (2002). Modelling default risk. Working paper, KMV Corporation. Available from www.kmv.com.

Drehmann, M. (2009). Macroeconomic stress testing banks: A survey of methodologies. In Quagliariello, M. (ed.): *Stress testing the banking system: Methodologies and applications*, Cambridge University Press, 37-67.

Duellmann, K. and Erdelmeier, M. (2009). Crash testing German banks. *International Journal of Central Banking*, **5**(3), 139-175.

European Banking Authority (2011). 2011 EU-Wide Stress Test Aggregate Report. Technical report, London, July 2011.

Flood, M. and Korenko, G. (2015). Systematic scenario selection: stress testing and the nature of uncertainty. *Quantitative Finance* **15**(1), 43-59.

Glasserman, P., Kang, C., and Kang, W. (2015). Stress scenario selection by empirical likelihood. *Quantitative Finance* **15**(1), 25-41.

Gordy, M.B. (2003). A Risk-Factor Model Foundation for Ratings-Based Bank Capital Rules. *Journal of Financial Intermediation* **12**(3), 199-232.

Gupton, G., Finger, C., and Bhatia, M. (1997). CreditMetrics: Technical Document. Technical report, JP Morgan & Co..

Kalkbrener, M. and Packham, N. (2015a). Correlation under stress in normal variance mixture models. *Mathematical Finance* **25**(2), 426-456.

Kalkbrener, M. and Packham, N. (2015b). Stress testing of credit portfolios in light- and heavy-tailed models. *Journal of Risk Management and Financial Institutions*, **8**(1), 34-44.

Kupiec, P. (1998). Stress Testing in a Value at Risk Framework. *Journal of Derivatives* **24**, 7-24.

McNeil, A.J, Frey, R., and Embrechts, P. (2005). *Quantitative Risk Management.* Princeton University Press, Princeton.

Merton, R. (1974). On the Pricing of Corporate Debt: the Risk Structure of Interest Rates. *Journal of Finance* **29**, 449-470.

Packham, N., Kalkbrener, M., and Overbeck, L. (2014). Multivariate default probabilities and default correlations under stress. *Journal of Applied Probability*, forthcoming.

Quagliariello, M. (2009). Macroeconomic stress-testing: definitions and main components. In Quagliariello, M. (ed.): *Stress testing the banking system: Methodologies and applications*, Cambridge University Press, 18-36.

Studer, G. (1999). Market risk computation for nonlinear portfolios. *Journal of Risk* **1**(4), 33-53.