Global Sensitivity Analysis for Latent Factor Credit Risk Models*

Dirk Baur†

Jessica Cariboni‡

Francesca Campolongo§

IPSC - ESAF, Joint Research Center, EU Commission

November 2004

Abstract

This paper proposes the use of a global sensitivity analysis to evaluate the risk associated with a credit portfolio model. The main features of this approach are its ability to assess sensitivities in the presence of non-linearities and to rank the input factors with respect to their relevance for the output variable. The commonly used local sensitivity analysis which is nested in the global model cannot provide this information. We analyze the static and time-varying uncertainties of three key input factors in a latent factor credit risk model, i.e. the multivariate distribution (copula), the default correlation and the default probability. Results show that the importance of the factors strongly depends on the average default probability of the portfolio and the analyzed quantiles of the default distribution. The proposed technique also provides information about trade-offs between the factors, e.g. between the default correlation and the copula.

JEL Classification: G31, C15, G11

KEYWORDS: credit risk model, latent factor model, copula, global sensitivity analysis

*We thank Wim Schoutens, Marco Ratto and two anonymous referees for helpful comments.

†corresponding author, Email: dirk.baur@jrc.it, Phone: +39 0332 785713, Fax: +39 0332 785733

‡K.U.Leuven, U.C.S., W. De Croylaan 54, B-3001 Leuven, Belgium. Email: jessica.cariboni@jrc.it

§Email: francesca.campolongo@jrc.it

1 Introduction

The decision of the Basel Committee for Banking Supervision to allow banks to use their own internal credit risk models pointed out the need of an evaluation of such models. This is true both for policy makers and for banks. A crucial point in a credit risk model is the obligors' dependence. In this work we choose to model this linkage by following the latent variables approach where dependencies of individual obligors are modelled via dependencies of underlying latent variables (e.g. see Crouhy et al., 2000 for an industry version of such a model).

Objective of our study is to investigate the sensitivity of the risk of a portfolio with respect to different input factors. Local techniques have been used by Frey et al. (2001) and Kiesel and Kleinow (2002). The main result obtained from these analyzes can be condensed to the fact that the copula is considerably contributing to the number of joint defaults for a given quantile of the default distribution. The authors of the above papers derive this conclusion by fixing all input variables and varying only the structural dependence (i.e. the copula) of the latent factors. Unfortunately, such an analysis does not account for non-linearities in the model and cannot assess the relative importance of the copula in comparison to the other factors, e.g. the default probabilities of the obligors in the portfolio or the correlation between obligors.

The aim of this paper is to evaluate a typical latent factor credit risk model more thoroughly. We do this via the application of global sensitivity analysis and show that this approach provides the risk manager with more information (factor ranking and interactions) and also more accurate information than the commonly applied local sensitivity analysis.

The paper is structured as follows: section 2 describes a typical latent factor model, section 3 introduces sensitivity analysis and the class of variance based measures . Section

4 presents the results and section 5 concludes.

2 The model

This section presents a so-called latent variables model (e.g. see Bluhm et al., 2003 and Frey et al., 2001). The dependence of m obligors is modelled via the dependence of m underlying latent variables $\mathbf{W} = (W_1, ..., W_m)'$. We assume that each W_j $(j = 1, .., m)$ is driven by one common factor Z and an idiosyncratic shock ϵ_j as follows:

$$W_j = \sqrt{a_j} Z + \sqrt{1 - a_j} \epsilon_j \qquad \text{for} \qquad j = 1, 2, ..., m \tag{1}$$

where a_j $(\in (0, 1))$ is the loading of obligor j to factor Z. We use the terms $\sqrt{a_j}$ and $\sqrt{1 - a_j}$ to guarantee a constant variance of W_j independent of the factor loading. The random variables Z and ϵ_j are assumed to be independent and identically distributed with mean zero and variance one. The shape of the default distribution, that is the distribution of the number of joint defaults, strongly depends on the correlation among the obligors and on the type of their multivariate distribution. The correlation among the obligors is driven by the choice of the factor loadings a_j; the multivariate distribution is determined by the marginals and the copula which governs the statistical dependence among the obligors and models the occurrence of extreme joint events. Since we change both the marginals and the copula we will say in the following that we change the multivariate distribution and will not explicitly refer to the copula.

The correlated variables W_j imply obligor default dependence as follows:

$$Y_j = 1 \Longleftrightarrow W_j \leq D_j \tag{2}$$

where Y_j is equal to one if the obligor defaults and zero otherwise. The cut-off point D_j for obligor j is calibrated to the probability of default π_j as follows:

$$\pi_j = P(W_j \leq D_j). \tag{3}$$

The joint default probability π_{ij} for obligors i and j can then be written as:

$$\pi_{ij} = P(W_i \leq D_i, W_j \leq D_j). \tag{4}$$

This equation shows that correlation in default probabilities is introduced by the correlation of the underlying (latent) variables W_i and W_j.

We focus on the number of joint defaults and therefore abstract from factors such as the recovery rate or the exposure at default. The default distribution is obtained by Monte Carlo simulation as described in Appendix A; the upper tail of this distribution is analyzed in order to investigate occurrences of extreme numbers of joint defaults.

The presentation of the model above displays the elements that potentially influence the number of joint defaults: the individual default probability π_j of obligor j, the type of multivariate distribution $\boldsymbol{W}$ for the underlying latent variables and the factor loadings a_j $(j = 1, 2, \ldots, m)$ that determine the degree of correlation among the obligors.

3 Sensitivity Analysis

A model represents a formal way to map some information and assumptions into inference. Thus, in order to ensure a correct use of the model and place confidence in its predictions, it is essential to investigate the uncertainty affecting the output variable. We assume the model input to be a vector $\mathrm{x} = (x_1, \ldots, x_k)$, k being the number of input variables. Inputs can be factors, parameters or triggers and the vector x can be considered as a realization of a random vector $\mathbf{X}$ to which a probability density function (continuous or discrete) is assigned. This distribution reflects our imprecise knowledge of the values of the inputs.

In the same way the output y (which here is considered a scalar quantity but could also be a vector) can be regarded as a realization of a random variable Y. The relationship between the inputs and the output (i.e. the model) is represented through a mathematical function f, which maps the k-dimensional input space into the output space:

$$Y = f(X_1, \ldots, X_k). \tag{5}$$

Sensitivity analysis (SA) is employed to assess the relative importance of the input factors in determining the uncertainty of the outcome. Very often, SA is conceived and defined as a local measure of the effect of a given input on the output (e.g. see Frey et al., 2001 and Kiesel and Kleinow, 2002) estimated by partial derivatives such as:

$$S_j = \frac{\partial f(X_1, \ldots, X_k)}{\partial X_j}. \tag{6}$$

Hence, the result of the analysis is obtained by varying one input factor at a time, while holding all the others fixed at predetermined values.

However, the local analysis is only reliable if the function f is linear. Since we cannot derive the form of the function f for the latent factor model due to the underlying Monte-Carlo simulation, we propose the use of a global approach. In contrast to local methods, the global approach is model-independent (reliable measures are obtained for any kind of model) and estimates the effect of a single factor while all the others are varied as well. This allows investigating potential interactions between the factors. Moreover, within the global approach it is possible to rank the factors allowing the modeler to prioritize his efforts and obtain more reliable results. This setting is known in sensitivity analysis as "factor prioritization". The question addressed is "what factor, once fixed to its true albeit unknown value, would give the greatest reduction in the uncertainty of the output?". For a review of global SA the reader can refer to Saltelli et al. (2000, 2004).

In summary, our approach includes the estimation of non-linearities of the model and of the relative importance of different factors in determining the variability of the loss distribution. To compare our results with those already published in the literature by

Frey et al. (2001) and Kiesel and Kleinow (2002), we have also performed a local sensitivity analysis.

3.1 The variance based sensitivity measures

Different global sensitivity techniques have been proposed in the literature. Among them, the class of the variance based measures play a major role because of its straightforward interpretation in terms of output variance and because of its desirable properties (Saltelli et al. (2004), chapter 2).

The variance based measures are built on the Sobol' decomposition of the function $f(\mathbf{X})$ into summands of increasing dimensionality (Sobol', 1990 and 1993):

$$Y = f(X_1, \ldots, X_k) = f_0 + \sum_i f_i(X_i) + \sum_i \sum_{j>i} f_{ij}(X_i, X_j) + \ldots + f_{12\ldots k}(X_1, \ldots, X_k) \tag{7}$$

where each term can be expressed using the mean operator E as follows:

$$f_0 = E(Y) \tag{8}$$

$$f_i(X_i) = E(Y|X_i) - f_0 \tag{9}$$

$$f_{ij}(X_i, X_j) = E(Y|X_i, X_j) - f_0 - f_i(X_i) - f_j(X_j) \tag{10}$$

and so on up to order k. Note that even to first order, the expansion (7) is not a linear superimposition, since $f_i(X_i)$ can have an arbitrary dependence on X_i. Sobol' demonstrated that for independent factors the decomposition in equation (7) is unique. It follows that the unconditional variance of the output $V(Y)$ can be written as:

$$V(Y) = \sum_i V_i + \sum_i \sum_{j>i} V_{ij} + \ldots + V_{12\ldots k} \tag{11}$$

where, for instance, $V_i = V(E(Y \mid X_i))$ and $V_{ij} = V(E(Y \mid X_i, X_j)) - V_i - V_j$. Normalizing this decomposition by the output variance we obtain:

$$1 = \sum_i S_i + \sum_i \sum_{j>i} S_{ij} + \ldots + S_{12\ldots k}. \tag{12}$$

The terms

$$S_i = \frac{V(E(Y|X_i))}{V(Y)}, i = 1, \ldots, k \tag{13}$$

are called importance measures or first order sensitivity indices and can be used to rank the factors according to the amount of output variance removed when the true value x_i^* of a given factor X_i is known (factor prioritization setting). Since we do not know what x_i^* is for each X_i, the most reliable result is achieved by considering $E(V(Y\,|\,X_i))$. The smaller is $E(V(Y\,|\,X_i))$, the more influential is the factor X_i. Since $V(Y) = E(V(Y\,|\,X_i)) + V(E(Y\,|\,X_i))$, focussing on the lowest $E(V(Y \mid X_i))$ is perfectly equivalent to focussing on the highest $V_i = V(E(Y\,|\,X_i))$ (i.e. on the highest S_i).

When $\sum_{i=1}^{k} S_i = 1$ the factors do not interact and it is said that the model is purely additive. This means that the effect of two or more factors on the output can be simply expressed by the sum of their single effects. When interactions are part of the model under investigation first order sensitivity coefficients are not capable to explain the entire variance of the output and higher order indices in equation (12) have to be taken into account. [1]

For instance the indices S_{ij} in (12) measure the second order interactions between factors i and j. These indices are related to the extra amount of variance corresponding to the factors i and j that is not explained by their individual effects. Similar interpretations can be given to all interaction terms up to order k.

The Sobol' scheme allows to define a sensitivity measure ST_i which estimates the total contribution to the variance of Y due to a certain input factor (Homma and Saltelli, 1996). This contribution includes its first order effect S_i plus all the interactions of any order with other factors. The total index ST_i is thus obtained as the sum of all the terms in equation (12) where at least one of the indices is equal to i. For instance, in the case of three input

[1] Note that there can be interaction among the factors but no correlation among them.

factors the total index ST_1 is given by:

$$ST_1 = S_1 + S_{12} + S_{13} + S_{123}. \tag{14}$$

It can be demonstrated (Saltelli 2002) that the total order sensitivity index can be estimated as:

$$ST_i = \frac{E(V(Y|X_{\sim i}))}{V(Y)}, i = 1, \ldots, k \tag{15}$$

where $\sim i$ indicates all factors but i. A comparison between the first and total sensitivity indices allows the investigation of the interactions of the model. The estimation of the first and total sensitivity indices is then enough to get a clear picture of the response of the model to different sources of uncertainty in the inputs.

Note that each operator (E or V) in the definition of the first and total sensitivity indices represent integrals in 1 or $(k-1)$ dimensions, which are usually estimated via Monte Carlo simulation.

3.2 Steps of the Analysis

The first step of the procedure is concerned with the choice of the output variables of interest since the sensitivity of the model severely depends on the choice of these variables (Saltelli et al. 2004, Chapter 2). As we are interested in investigating the highest risk associated with a specified credit portfolio, we focus on extreme quantiles of the default distribution, i.e. the distribution of the number of joint defaults. In particular we analyze the higher quantiles of the distribution of the number of joint defaults (the 95%, 99% and 99.5% quantiles). These quantiles are investigated in three scenarios, corresponding to different types of portfolios, namely high, medium and low-rated portfolios with a time horizon of 1 year. The three scenarios are characterized by different ranges of the 1-year default probabilities of the obligors, as specified below. Following Moody's classification we include ratings from Aaa to B3. We view ratings from Aaa to A3 as high-rated (scenario

1), ratings from A3 to Ba3 as medium-rated (scenario 2) and ratings from Ba3 to B3 as low-rated (scenario 3). This classification is arbitrary to some degree and for presentation purposes only.

In a second step we choose the input factors and assign the corresponding variability. In our experiment the three chosen factors are the following trigger parameters:

1. A trigger factor which defines the type of multivariate distribution **W** of the latent factors. The trigger factor may assume three possible values with equal weights corresponding to the following distributions for **W**: a Gaussian distribution, a t-distribution with 10 or 4 degrees of freedom (see Appendix B). The uncertainty we put in the type of multivariate distribution reflects our imprecise knowledge of the dependence structure among the obligors. In particular the dependence structure implicit in the multivariate t distribution (i.e. the t-copula) allows to include in the analysis the occurrence of dependent extreme values which are not captured by the Gaussian copula .

2. A trigger factor, taking on three possible values with equal weights, which determines the degree of the correlation among the obligors, represented by a m-dimensional vector of factor loadings a_j. Three possible determinations of the vector of loadings are randomly generated a priori under the assumption that the a_j are uniformly distributed within the following fixed ranges:

   - low correlation i.e. $a_j \sim U[0.0, 0.4], j = 1, 2, \dots, m$
   - medium correlation i.e. $a_j \sim U[0.0, 0.6], j = 1, 2, \dots, m$
   - high correlation i.e. $a_j \sim U[0.0, 0.8], j = 1, 2, \dots, m$

   Since correlations among obligors are difficult to quantify (e.g. see Allen and Saunders, 2003 and Effenberger, 2004), we assume that neither the average correlation nor its volatility are known and that the degree of correlation does not depend on the scenarios. Note that we include relatively high correlations among the underlying latent variables since in periods of recession or financial crisis asset correlations can increase considerably. However, the average values chosen do not exceed the figures reported by Gersbach and Lipponer (2000). Note also that the (underlying) asset correlation is higher than the resulting default correlation (e.g. see Lucas, 2004).

3. A trigger factor, taking on three possible values with equal weights which determines in each scenario the m-dimensional vector π of default probabilities of the obligors. For each scenario, three possible determinations of the vector π are randomly generated a priori under the assumption that the 1-year default probabilities are uniformly distributed as specified in table 1 (see Bluhm et al., 2003).

   Also in this case we assume that neither the average of the default probabilities nor the corresponding volatility are known in each scenario. Note that we do not follow Frey et al. (2001) and Kiesel and Kleinow (2002), who assign a single value to describe the marginal default probabilities of different rating classes. On the contrary, in each scenarios we consider three possible ranges for the default probabilities. The three possible subclasses of each scenario aim to represent two types of uncertainties: (i) static errors in the default probability estimates (e.g. an obligor could be erroneously classified as A1 while it is A2) or (ii) time-varying default probabilities (e.g. an obligor initially A1-rated may after some time fall into A2).

Insert Table 1 upon here

Since in our case study only a discrete number of possible values for each input is considered, the estimates of the sensitivity indices do not require a Monte Carlo simulation. In this particular case as only 27 possible combinations of inputs are possible the integrals in the definition of the indices reduce to simple sums and the terms in the decompositions (7) and (11) can be computed analytically.

4 Results

This section is structured as follows. In the first section we present the results of the local SA experiments and compare them with the existing literature. In the second part results of the global SA are presented. In both cases the Monte Carlo sample size to build the loss distribution is $N = 10.000$ and the portfolio consists of $m = 1.000$ obligors.

4.1 Local Sensitivity Analysis

Results obtained through local sensitivity analysis are presented in Table 2. The table shows the number of joint defaults corresponding to the three fixed quantiles of the default distribution obtained by varying two factors (degree of correlation, multivariate distribution) for each determination of the vector of default probabilities and each of the three scenarios as defined above. Note that we do not estimate any derivative of the output with respect to the inputs but we just look at the variation in the number of joint defaults fixing two inputs and varying the third one. This is similar to the exercise undertaken by Frey et al. (2001) and Kiesel and Kleinow (2002).

Insert Table 2 upon here

As an example, we first focus on the 99% quantile in Table 2 and on the higher rated portfolios in each scenario (namely AAA-A1 for Scenario 1, Baa1-Ba1 for Scenario 2 and Ba3-B1 for Scenario 3). For low correlations, the distribution of the underlying latent variables seems to matter considerably (number of joint defaults varies between 4 and 2 for high-rated portfolios, between 44 and 112 for medium-rated portfolios and between 238 and 400 for low-rated portfolios). The degree of correlation of the latent variables also matters for medium and low-rated portfolios, but it seems to be almost negligible for high-rated obligors. Interestingly the relative differences are larger for Gaussian distributions ($44-81$ for medium-rated portfolios and $238-382$ for low-rated portfolios) than for t-distributions ($112-122$ for medium-rated portfolios and $400-510$ for low-rated portfolios).

If we then concentrate on the influence of changing the default probability vector we find that in all the cases the number of joint defaults increases with decreasing the rating but its effect varies considerably depending on the underlying multivariate distribution and on the scenario. For instance, in the second scenario the effect of changing the default probability vector is stronger for a Gaussian distribution ($44-102$) than for a t4-distribution ($112-243$). Note that results also differ significantly among the quantiles.

The numbers in Table 2 confirm the influence of the multivariate distribution (and consequently the copula) as found in Frey et al. (2001) and Kiesel et al. (2002). However, the outcomes also underline the influence of the other two factors, in particular the degree of correlation.

Obviously, it is difficult to assess the relative importance of the different input variables by considering the variation of the number of joint defaults produced by changing one factor at a time. For example, if we consider the medium-rated portfolios, we see that the influence of the multivariate distribution is stronger for low degrees of correlation ($44-112$ for e.g. Baa1-Ba1) than for higher degrees of correlation ($81-122$ for e.g. Baa1-Ba3) but the relative importance of the factors remains unclear.

Finally it is worth noting that a comparison between our local results and the ones present in the literature can be performed just qualitatively. In fact a one-to-one correspondence between results can not be established as the ranges of variation for the input factors differ considerably. For instance, Kiesel and Kleinow (2002) choose for each of the three rating classes a single value of the marginal default probability rather than a range.

Their marginal default probabilities are equal to the upper bounds of the ranges chosen in the present work. Furthermore Kiesel and Kleinow model the change in the correlation by adding a second common factor rather than just changing the degree of correlation value. Also Frey et al. (2001) use a slightly different approach as they consider the correlation among the obligors and the marginal default probabilities as a single factor: a couple of default probabilities and correlations is assigned to each rating class. In general, it can be stated that our analysis accounts for more credit risk scenarios than these studies.

4.2 Global Sensitivity Analysis

Results of the global sensitivity analysis are shown in Table 3 for each of the three scenarios defined above (high, medium and low-rated obligors). The first and total order sensitivity indices are listed for the three factors in correspondence to the three outputs.

Insert Table 3 upon here

Since the default distribution is obtained via Monte Carlo simulation, we have checked the reliability of our results by repeating the simulation and averaging the sensitivity indices over the simulations to compare them with the single run results. We have verified that the ranking of the factors does not change for any output variable in any scenario.

We primarily concentrate on the first order indices since they are the proper measure to rank factors in their order of importance in our setting (factor prioritization setting, see Section 3).

Results are very different depending on the scenarios determined by the rating of the obligors and also on the quantile of the default distribution. Hence, we discuss the results for each scenario starting with high-rated obligors (see Figure 1).

At the 95% quantile the type of multivariate distribution is most important ($S_{Distr.} = 0.682$) but its influence strongly decreases when more extreme quantiles are considered.

The importance of the default probability of the obligors increases with higher quantiles and becomes the most important factor in the right tail of the default distribution ($S_{Def.Prob.} = 0.595$). The degree of correlation among the obligors is the less important factor at all quantiles since it is responsible for less than 3.5% of the total output variability.

Insert Figure 1 upon here

For portfolios comprising only medium-rated obligors (see Figure 2), the default probability explains almost all the output variance at the 95% quantile ($S_{Def.Prob.} = 0.920$) but its influence decreases with higher quantiles. The degree of correlation and the type of multivariate distribution increase their relevance with higher quantiles. In particular, the importance of the multivariate distribution increases more than that of the degree of correlation, and it becomes the driving factor when extreme events are considered ($S_{Distr.} = 0.435$).

Insert Figure 2 upon here

For low-rated portfolios (see Figure 3) the default probability is the most important factor ($S_{Def.Prob.} = 0.628$), the correlation degree is second important and the distribution is the least important factor at the 95% quantile. The correlation strongly increases its relevance with higher quantiles, and it accounts for more than 44% of the total variance at quantile 99.5%. The default probability becomes the less important factor while the type of multivariate distribution increases up to $S_{Distr.} = 0.374$.

Insert Figure 3 upon here

The importance of each factor with respect to each scenario and quantile can also be seen in Figure 4. Note that fixing one scenario and one output the relative importance of the three factors is clearly quantified.

Insert Figure 4 upon here

Although our analysis is based on a model that does not explicitly account for rating changes, our findings help answering the question whether rating migrations should be included in the model or whether it is sufficient to focus on default or non-default states. The importance of the rating changes depends on the type of the portfolio and the chosen quantile: the changes of the default probability are most important at lower quantiles for medium and low-rated portfolios and at higher quantiles for high rated obligors.

Table 3 shows that the differences between ST_i and S_i are pronounced only for high-rated obligors and increase at higher quantiles. The difference ($ST_i - S_i$) explains the sum of all interactions of factor i with the others. For high rated portfolios interactions account for a maximum of 19% at the 99.5% quantile. For medium and low-rated portfolios interactions are almost negligible since for all the quantiles they explain a maximum 5% of the total variability in the number of joint defaults correspondent to the selected quantiles.

Following Section 3.1 we have also decomposed each output variable into summands of increasing dimensionality to investigate potential non-linearities or non-monotonicities of the model (for an example of the same type of analysis see e.g. Welch et al. (1992)). Figure 5 plots the first order functions in decomposition (7), which give the effect associated to a fixed input acting independently upon the output. Each plot refers to a specified input within a fixed scenario and is labelled as $f^i(X_j)$ where $i = High, Medium, Low$ indicates the scenario and X_j specifies the factor. Within each plot a single line represents the behavior of the first order function of a fixed quantile varying the considered input. Note that the lines are simply a linear interpolation between the dots. In most of the cases the behavior of the functions is monotone.

Insert Figure 5 upon here

Only high-rated portfolios ($f^{High}(Deg.Corr.)$ and $f^{High}(Distr.)$) present non-monotonic effects, which become more pronounced at higher quantiles.

Each line in a graph can be used to quantify the average variation in the number of joint defaults changing the selected input. For instance, if we choose the second scenario (second line of plots in figure 5) and the 99.5% quantile (circled dots) , the average change in the number of joint defaults obtained by increasing the default probability vector (e.g. from Baa1-Ba1 to Baa1-Ba2) is about 80 (first plot). For the correlation degree this change is close to 50 joint defaults (e.g. from low to medium) (second plot) while a variation of the multivariate distribution has also an effect of about 80 (e.g. from t_4 to t_{10}) (third plot).

To obtain a complete decomposition of the outputs, we have also computed the values of the second order functions of each couple of factors in correspondence to each quantile and within each scenario. Since interactions do not play a major role in the present model, a discussion of these results is beyond our purpose. Appendix C briefly summarizes these results.

5 Conclusions

This paper has introduced the concept of global sensitivity analysis to evaluate credit risk models. We have considered a factor model to simulate the loss distribution, i.e. the distribution of the number of joint defaults, of a credit portfolio. We have analyzed the influence of three factors, namely the type of multivariate distribution of the latent factor, the degree of correlation among the obligors and the default probability of the obligors, on the volatility of some specified quantile in the right tail of the loss distribution.

In contrast to a local analysis, we have showed that a global approach allows to investigate non-linearities of the model and to provide information on the relative importance of the three factors on the output.

Our results have demonstrated that the model presents non-linearities only when high-rated portfolios and high quantiles of the loss distribution are considered. Moreover, we have found that the influence of the three factors on the output can be rather different. This depends both on the portfolio rating and on the risk the manager is considering, that is on the quantile of the default distribution. In particular, for the 99.5% quantile, the degree of correlation is the most important factor for low-rated portfolios, the type of multivariate distribution is most relevant for medium-rated obligors while the default probability explains most of the variance for high-rated portfolios. If the focus is on less extreme events (e.g. the 95% quantile), the type of multivariate distribution plays a major role for high-rated portfolios while the variance of medium and low-rated portfolios is mostly driven by the default probability.

Future applications of the global sensitivity analysis could include the recovery rate as additional factors and evaluate other and richer credit risk models. For example, the method could be used to analyze the relative importance of the model assumptions in a collateralized debt obligation (CDO).

6 References

Allen, L., Saunders, A., 2003, A survey of cyclical effects in credit risk measurement models, BIS Working Paper No. 126

Bluhm, Ch., Overbeck, L. and Ch. Wagner, 2003, An Introduction to Credit Risk Modeling, Chapman & Hall/ CRC

Crouhy, M., Galai, D. and R. Mark, 2000, A comparative analysis of current credit risk models, Journal of Banking and Finance 24(1-2), 59-117.

Effenberger, D., 2004, Credit Derivatives: effects on the stability of financial markets, Deutsche Bank Research, Current Issues

Embrechts, P., Lindskog, F. and A. McNeil, 2001, Modelling dependence with copulas and applications to risk management, Handbook of Heavy Tailed Distributions in Finance, ed. S. Rachev, Elsevier, Chapter 8, 329-384.

Frey, R., McNeil, A.J. and M. Nyfeler, 2001, Copulas and credit models. RISK, October, 111-114.

Gersbach H., Lipponer, A., 2000, The Correlation Effect, Working Paper University of Heidelberg

Kiesel, R. and T. Kleinow, 2002, Sensitivity analysis of credit portfolio models in Applied Quantitative Finance, eds: W. Härdle, T. Kleinow, G. Stahl., 140-152.

Homma, T. and A. Saltelli, 1996, Importance measures in global sensitivity analysis of model output. Reliability Engineering and System Safety 52(1), 1-17.

Lucas, D., 2004, Default Correlation: From Definition to Proposed Solutions, UBS Working Paper

Saltelli, A., 1999, Sensitivity analysis. Could better methods be used?, Journal of Geophysical Research, 104(D3), 3789-3793.

Saltelli, A., Chan, K. and E. M. Scott, 2000, Sensitivity Analysis, John Wiley & Sons publishers, Probability and Statistics series.

Saltelli A. Tarantola S., Campolongo, F. and Ratto, M., 2004, Sensitivity Analysis in Practice. A Guide to Assessing Scientific Models, John Wiley & Sons publishers, Probability and Statistics series

Saltelli, A., 2002, Making best use of model valuations to compute sensitivity indices, Computer Physics Communications, 145, 280-297.

Saltelli, A. and S. Tarantola, 2002, On the relative importance of input factors in mathematical models: safety assessment for nuclear waste disposal, Journal of the American Statistical Association, 97(459), 702-709.

Sobol' I.M., 1990 (in Russian), 1993 (in English), Sensitivity Estimates for Nonlinear Mathematical Models, Mathematical Modelling and Computational Experiments, 1(4) 407-414.

Welch W.J., Buck R.J., Sacks J., Wynn H.P., Mitchell T.J. and Morris M.D., 1992, Screening, Predicting and Computer Experiments, Technometrics, 34(1), 15-25.

Appendix A

The default distribution is obtained by Monte Carlo simulation as follows:

1. a Monte Carlo loop is performed to determine the distribution W_j for each obligor. Each distribution is obtained by randomly generating T values of the common factor Z and of the shocks ϵ_j and deriving the values of W_j from equation (1) as a function of the assigned factor loading (degree of correlation).

2. from the W_j distributions and the input default probabilities π_j the cutoff points of the obligors $\{D_j, j = 1, 2, \ldots, m\}$ are derived by inverting equation (3).

3. The distribution of the number of joint defaults is obtained comparing the simulated values of the latent variables W_j with the cutoff points $\{D_j, j = 1, 2, \ldots, m\}$.

Appendix B

In order to obtain a multivariate Student t distribution from a gaussian we consider the multivariate normal mixture models.

A member of the family of variance mixtures of normal distribution is equal in distribution to (see Embrechts et al., 2001):

$$\boldsymbol{W}^{\star} = M \cdot \boldsymbol{W}$$

where $\boldsymbol{W}$ is a multivariate normal random variable with mean zero and covariance matrix Σ and M is a scalar, positive and independent of $\boldsymbol{S}$ with finite second moment.

It can be shown that $Corr(W_i^{\star}, W_j^{\star}) = Corr(W_i, W_j)$.

$\boldsymbol{W}^{\star}$ has a m-dimensional Student t distribution with ν degrees of freedom ($\boldsymbol{W}^{\star} \sim t_m(\nu, \mathbf{0}, \Sigma)$) if $M = \sqrt{\dfrac{\nu}{T}}$ where T has a chi-square distribution with ν degrees of freedom.

Appendix C

Second order functions are related to the second order terms in equation (7), i.e. to the cooperative effects of each couple of inputs on the output. Note that in the definition of the second order functions, the first order effects are subtracted. Figures 6 show some examples of the obtained surfaces. Each plot refers to a fixed quantile within a specific scenario. The surfaces are labelled as $f^i(X_j, X_k)$ where i labels the scenario and j, k labels the fixed factors, as specified in Section 4.2. Note that the surfaces are obtained by interpolating the values of the functions at fixed points.

Since interactions are pronounced only in the first scenario (high-rated obligors, see Table 3), the reported examples refers to these cases. Explicitly we have reported for the 99.5% quantile $f^{High}(Deg.Corr., Distr.)$ for the 99% quantile and $f^{High}(Def.Prob., Distr.)$, $f^{High}(Deg.Corr., Def.Prob.)$.

In most of the cases the shape of the surfaces resemble saddles (e.g. first graph in Figure 6), but some other more peculiar structures are found (e.g. second and third graphs in Figure 6). Except for high-rated portfolios the values of the second order functions are lower than those of the first order functions, since most of the outputs variance is explained by first order sensitivity indices.

Scenario 1	$\pi(\mathbf{Aaa-A1})$	$\pi(\mathbf{Aaa-A2})$	$\pi(\mathbf{Aaa-A3})$
high-rated portfolio	U[0, 0.05%]	U[0, 0.075%]	U[0, 0.10%]
Scenario 2	$\pi(\mathbf{Baa1-Ba1})$	$\pi(\mathbf{Baa1-Ba2})$	$\pi(\mathbf{Baa1-Ba3})$
medium-rated portfolio	U[0.10%, 1.80%]	U[0.10%, 2.80%]	U[0.10%, 3.80%]
Scenario 3	$\pi(\mathbf{Ba3-B1})$	$\pi(\mathbf{Ba3-B2})$	$\pi(\mathbf{Ba3-B3})$
low-rated portfolio	U[3.00%, 6.00%]	U[3.00%, 10.00%]	U[3.00%, 14.00%]

Table 1: Uniform distributions chosen to generate the vectors of default probabilities in each scenario.

		Quantile 95%			Quantile 99%			Quantile 99.5%		
	Correlation	$\mathbf{W} \sim \mathbf{t_4}$	$\mathbf{W} \sim \mathbf{t_{10}}$	$\mathbf{W} \sim \mathcal{G}$	$\mathbf{W} \sim \mathbf{t_4}$	$\mathbf{W} \sim \mathbf{t_{10}}$	$\mathbf{W} \sim \mathcal{G}$	$\mathbf{W} \sim \mathbf{t_4}$	$\mathbf{W} \sim \mathbf{t_{10}}$	$\mathbf{W} \sim \mathcal{G}$
Scenario 1	Low	0	1	1	2	5	4	8	10	5
Aaa-A1	Medium	0	0	1	2	4	4	6	10	7
	High	0	0	1	2	4	4	6	8	8
Scenario 1	Low	0	1	2	5	7	5	17	14	8
Aaa-A2	Medium	0	1	2	5	8	6	13	16	9
	High	0	1	1	3	7	7	9	18	10
Scenario 1	Low	0	2	2	7	11	6	21	22	9
Aaa-A3	Medium	0	1	2	9	11	8	23	21	13
	High	0	1	2	6	10	8	15	18	15
Scenario 2	Low	24	24	19	112	75	44	167	106	56
Baa1-Ba1	Medium	20	23	20	125	83	58	180	126	86
	High	16	18	19	122	96	81	212	171	124
Scenario 2	Low	52	48	36	178	125	75	264	160	93
Baa1-Ba2	Medium	48	48	41	208	140	105	289	213	137
	High	46	43	43	231	174	129	309	244	172
Scenario 2	Low	80	70	52	243	156	102	333	199	129
Baa1-Ba3	Medium	81	70	61	266	199	144	337	268	190
	High	80	70	67	292	244	176	397	341	232
Scenario 3	Low	211	167	140	400	299	238	481	360	283
Ba3-B1	Medium	222	196	166	446	384	321	531	445	388
	High	240	216	193	510	447	382	622	527	472
Scenario 3	Low	265	218	189	443	380	296	504	434	345
Ba3-B2	Medium	291	251	223	523	436	370	596	497	422
	High	308	278	258	582	520	485	670	580	548
Scenario 3	Low	302	270	235	490	410	356	556	475	400
Ba3-B3	Medium	343	305	281	548	482	440	634	547	501
	High	379	357	326	618	562	533	696	614	585

Table 2: Results of local sensitivity analysis: number of joint defaults correspondent to quantiles 95%, 99% and 99.5% for a portfolio of 1.000 obligors. $\mathcal{G}$ indicates the multivariate gaussian distribution.

		Quantile 95%		Quantile 99%		Quantile 99.5%	
	Factors	**S**	**ST**	**S**	**ST**	**S**	**ST**
Scenario 1	Corr. Degree	0.032	0.124	0.006	0.089	0.013	0.144
High-rated	Distribution	0.682	0.829	0.214	0.327	0.201	0.384
portfolio	Def. Prob.	0.129	0.249	0.659	0.735	0.595	0.700
Scenario 2	Corr. Degree	0.000	0.019	0.090	0.111	0.131	0.153
Medium-rated	Distribution	0.036	0.074	0.356	0.391	0.435	0.471
portfolio	Def. Prob.	0.920	0.952	0.503	0.551	0.389	0.424
Scenario 3	Corr. Degree	0.179	0.198	0.422	0.435	0.445	0.457
Low-rated	Distribution	0.174	0.179	0.313	0.327	0.374	0.388
portfolio	Def. Prob.	0.628	0.644	0.249	0.259	0.166	0.174

Table 3: Global sensitivity analysis results for the selected quantiles of the distribution of number of joint defaults: first order indices S and total indices ST for the three considered input factors in the three scenarios.

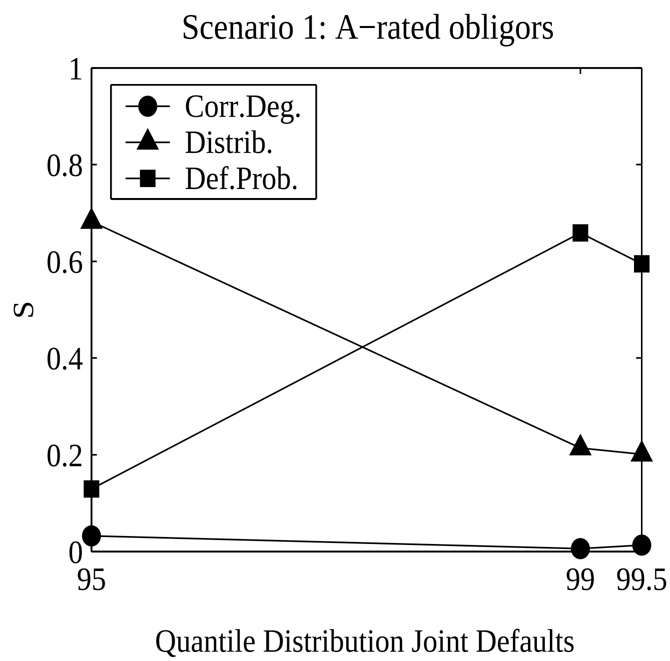

Figure 1: Evolution of first order sensitivity index as a function of the quantile of the loss distribution for the high-rated portfolio.

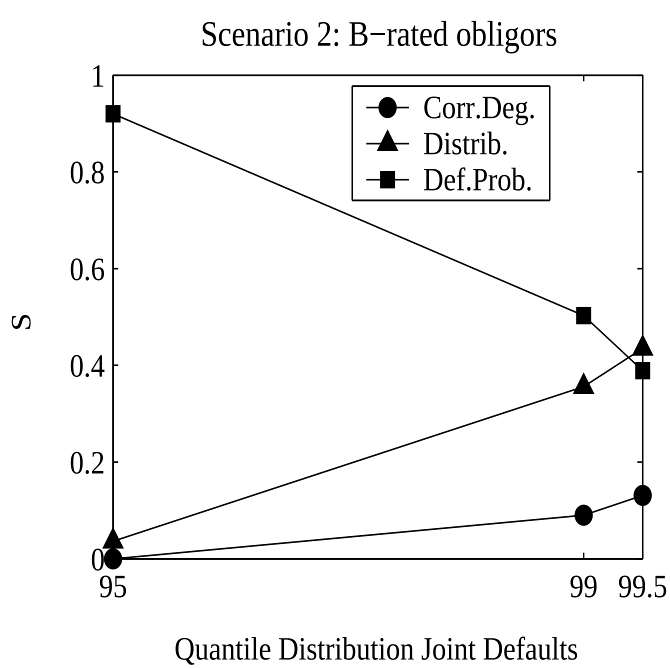

Figure 2: Evolution of first order sensitivity index as a function of the quantile of the loss distribution for the medium-rated portfolio.

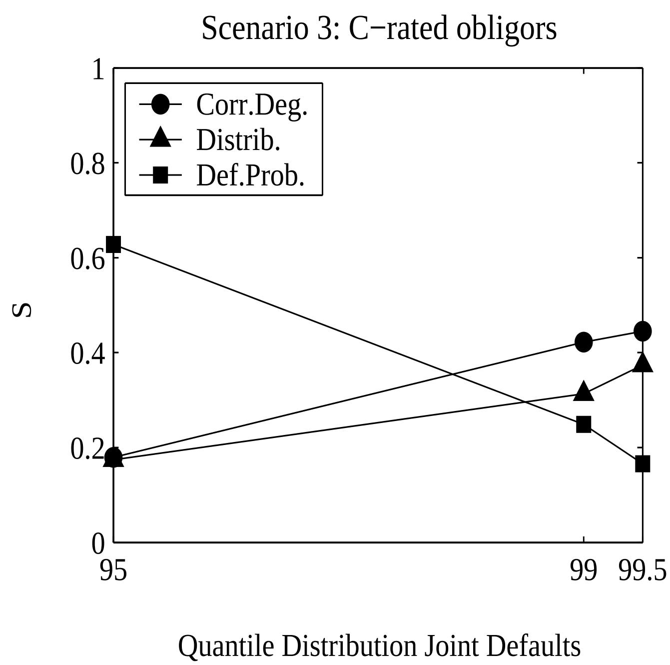

Figure 3: Evolution of first order sensitivity index as a function of the quantile of the loss distribution for the low-rated portfolio.

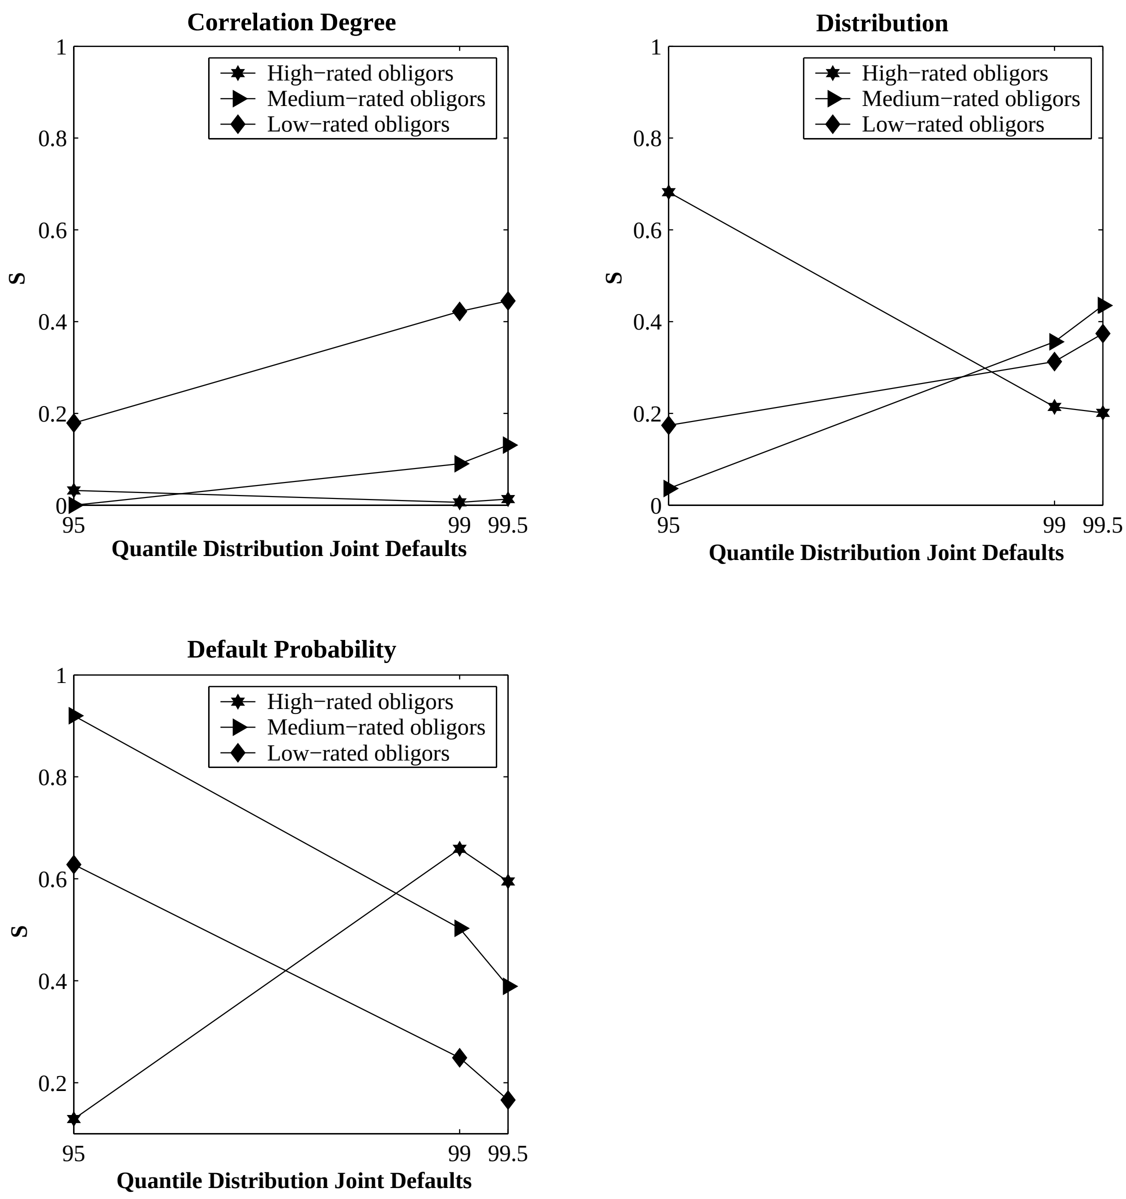

Figure 4: Evolution of first order sensitivity indices in the three scenarios for each fixed factor.

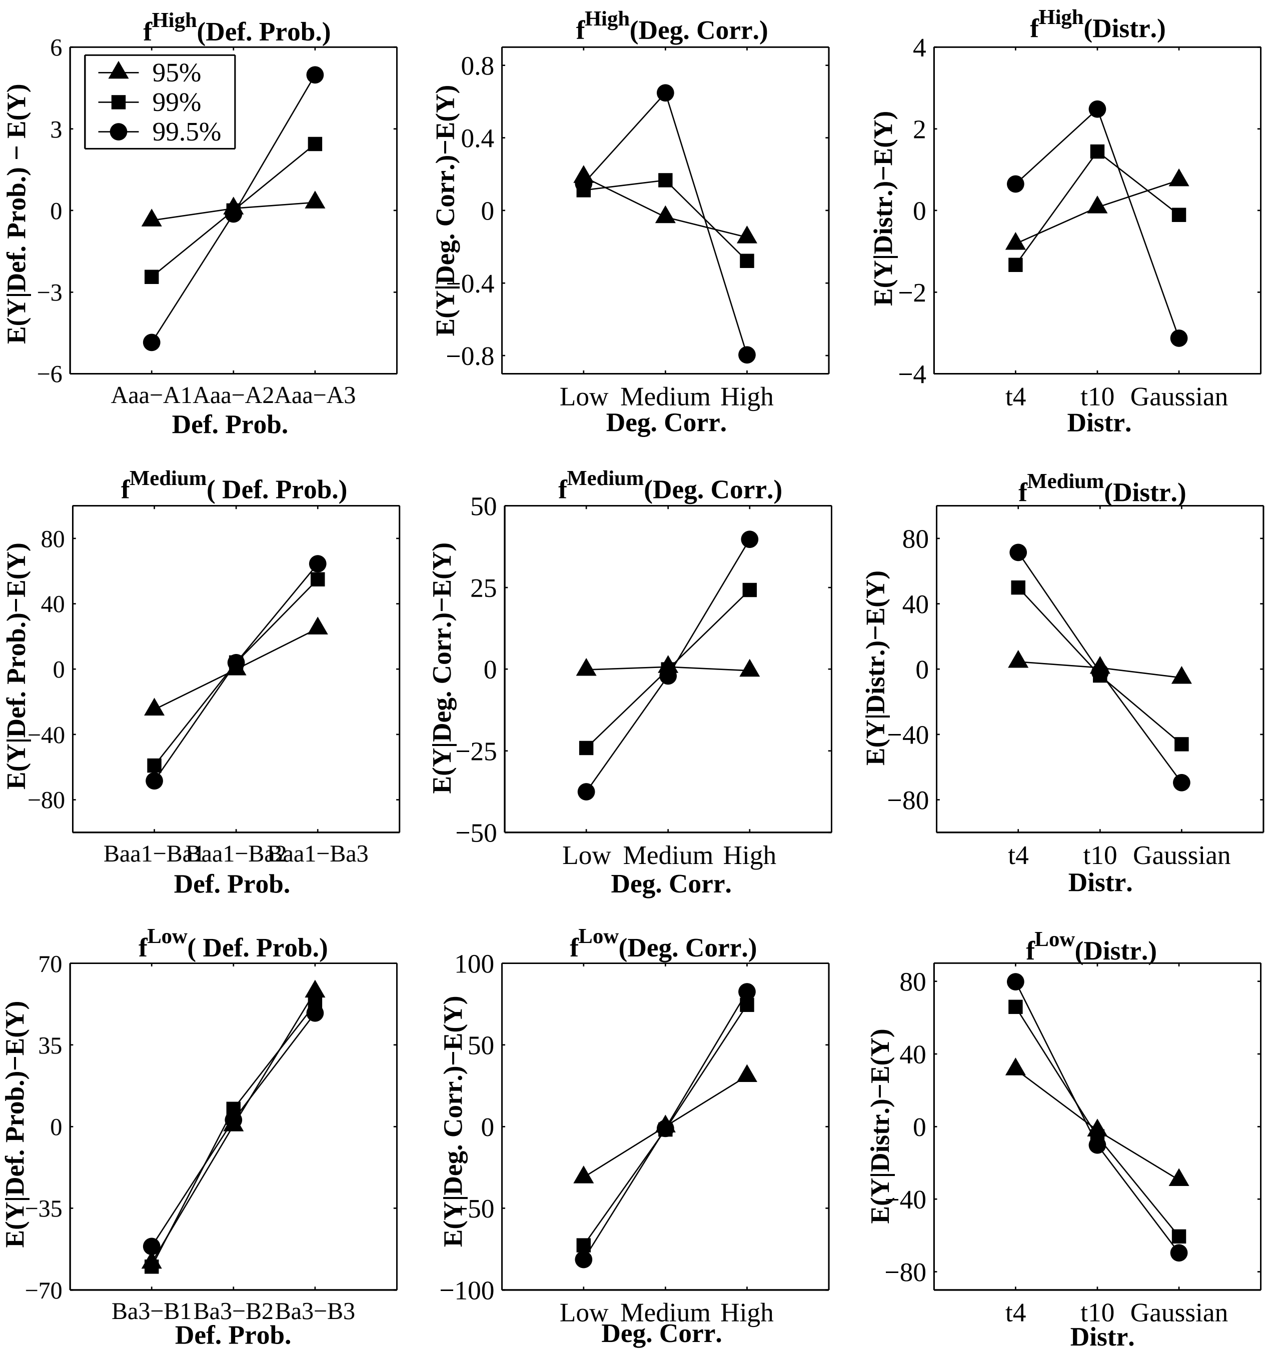

Figure 5: Evolution of first order functions in the three scenarios for each fixed factor and each quantile.

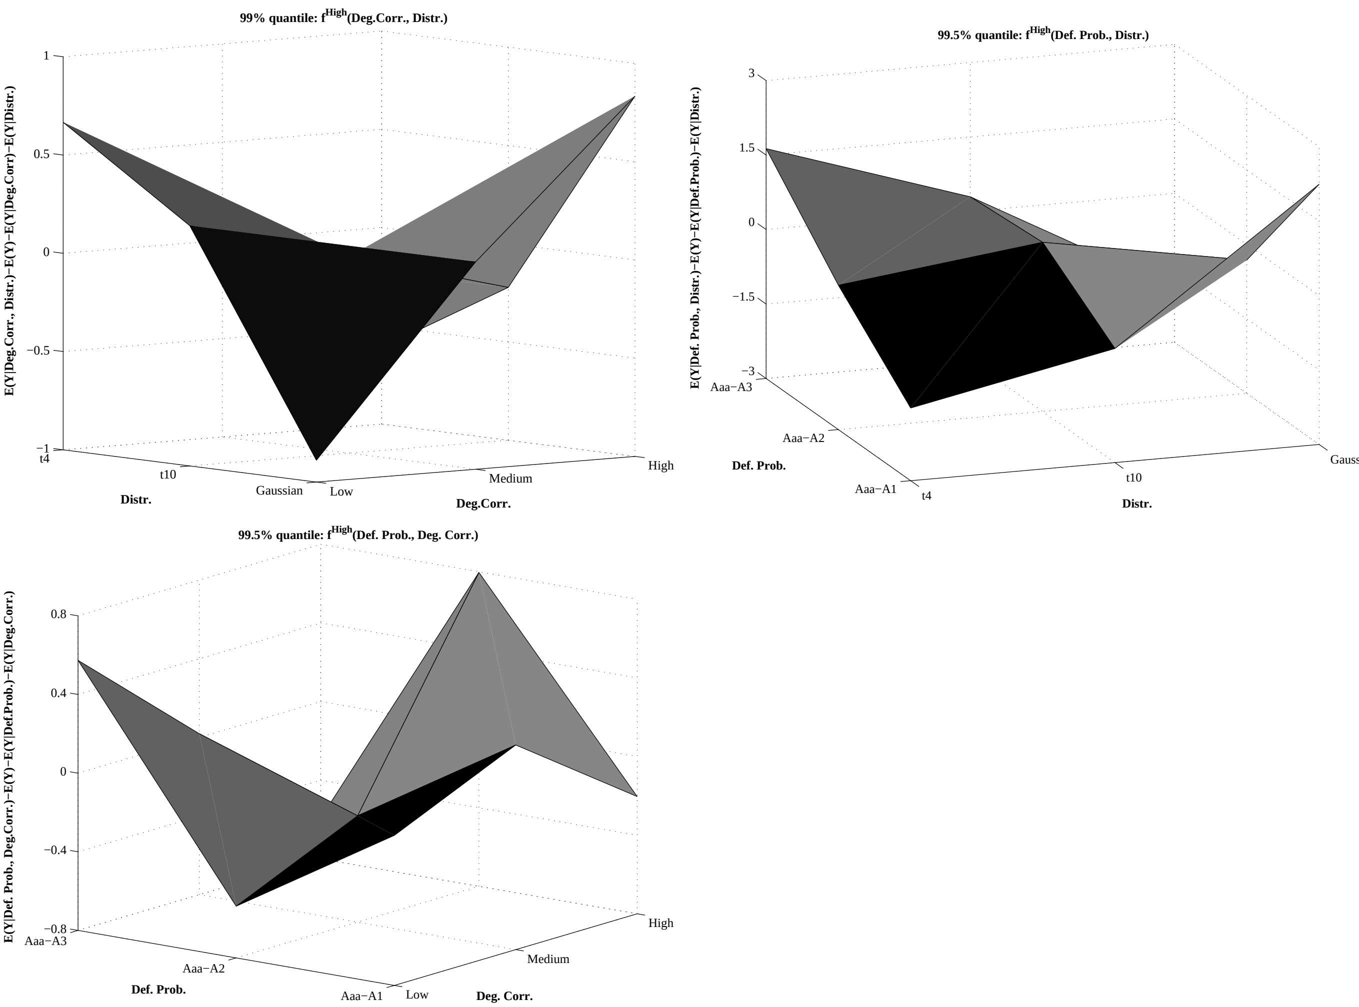

Figure 6: Evolution of three selected second order functions for high-rated portfolios.