MOODY'S ANALYTICS

AUGUST 2019

Analytical Solutions for Multi-Period Credit Portfolio Modelling

Prepared by

Juan Licari
Juan.Licari@moodys.com
Managing Director,
Moody's Analytics

Gustavo Ordonez-Sanz
gustavocqd@gmail.com
Head of Economic Capital
Analytics at HSBC

ABSTRACT

This papers presents a framework for credit portfolio modelling where exact analytical solutions can be obtained for key risk measures such as portfolio volatility, risk contributions to volatility, Value-at-Risk (VaR) and Expected Shortfall (ES). The framework is generic and can accommodate structural, reduced-form and macroeconomic-type models. It is also exible enough to allow for multi-period models with credit transitions and other risks such as credit spread, interest rate, FX or instrument optionalities such as loan pre-payment. Furthermore, the same mathematical results can be used to define an importance sampling algorithm that can be used to dramatically accelerate the Monte Carlo simulations that are commonly used to calculate the portfolio loss distribution. Another key result from this framework is the ability to run reverse stress testing analyses analytically. Finally, these solutions can also be readily used to obtain approximation to VaR/ES through the saddlepoint method.

Contact Us

Email
help@economy.com

U.S./Canada
+1.866.275.3266

EMEA

+44.20.7772.5454 (London)
+420.224.222.929 (Prague)

Asia/Pacific
+852.3551.3077

All Others
+1.610.235.5299

Web
www.economy.com
www.moodysanalytics.com

Analytical Solutions for Multi-Period Credit Portfolio Modelling

Juan M. Licari[1] and Gustavo Ordóñez-Sanz*[2]

[1]Managing Director at Moody's Analytics
[2]Head of Economic Capital Analytics at HSBC

August 2019

Abstract

This paper presents a framework for credit portfolio modelling where exact analytical solutions can be obtained for key risk measures such as portfolio volatility, risk contributions to volatility, Value-at-Risk (VaR) and Expected Shortfall (ES). The framework is generic and can accommodate structural, reduced-form and macroeconomic-type models. It is also flexible enough to allow for multi-period models with credit transitions and other risks such as credit spread, interest rate, FX or instrument optionalities such as loan pre-payment. Furthermore, the same mathematical results can be used to define an importance sampling algorithm that can be used to dramatically accelerate the Monte Carlo simulations that are commonly used to calculate the portfolio loss distribution. Another key result from this framework is the ability to run reverse stress testing analyses analytically. Finally, these solutions can also be readily used to obtain approximation to VaR/ES through the saddlepoint method.

1 Introduction

One of the cornerstones of credit portfolio management is the estimation of the probability distribution of losses. Due to the complexity of the inter-dependency between the different assets in the portfolio numerical simulations (Monte Carlo) are generally used to estimate the loss distribution. For an overview of these simulation techniques see [1].

One of the main drawbacks of using Monte Carlo is that it is computationally very intensive, especially when trying to calculate the value of risk measures in the tail of the distribution such as VaR/ES. The lack of convergence of Monte

*corresponding author (gustavocqd@gmail.com). Any views and opinions expressed in the work are the author and do not necessarily represent the policy or position of any HSBC Group Members

Carlo techniques, even after a large number of simulations, is especially acute when attempting to allocate risk measures (such as VaR/ES) to the different assets in the portfolio. Because of this limitation, often portfolio managers choose to rely on risk measures that can be calculated analytically such as the Expected Loss (EL), or by resorting to analytical approximation of the probability distribution of losses in the portfolio. One of the better known approximations to the portfolio loss distribution is the large homogeneous portfolio approximation (LHP) introduced in [2] and which forms the basis for Internal Rating Based (IRB) RWA formula in the Basel regulation [3]. Other analytical approximations to the loss distribution use specific probability distribution assumptions like in the case of CreditRisk+ [4], or use other numerical approximations such as the saddlepoint method [5]. The gain in computation speed and numerical tractability, however, comes at the expense of a loss in accuracy in the calculation of the different risk measures.

There are three main families of credit models used for Monte Carlo simulations (see [6] for a brief review of these different approaches). The first family is the structural models which are based on the ground-breaking work carried out by Robert C. Merton [7]. These models attempt to explain the default process assuming that the asset price of a given institution is a stochastic process, and that if the price falls below a default barrier within a given time horizon (typically one-year), the institution will default. There have been numerous extensions to the original work by Merton, refer to [8] for some examples of such extensions. Although it is possible to extend structural models to multi-period settings, typically a single-period simulation is used in order to reduce the computation time. In recent years, however, and because of the regulatory focus on stress testing and the implementation of the IFRS9 accounting rules, the need for multi-period credit models is becoming more and more apparent to portfolio managers.

The second family of credit portfolio models are known as "reduced-form models" of default. These models do not attempt to explain the underlying process of the firms' default, but instead assume that the time to default is a random variable, and that the defaults occur with some intensity (or hazard rate) which is modelled as a stochastic process. This is then calibrated to observed market data such as bond or CDS spreads term structures. This makes these types of models a natural choice when multi-period credit portfolio analysis and also make them very popular for pricing credit derivatives [9]. See [10] for an empirical comparison of these two families of models. The third family of models use an econometric approach to link macroeconomic variables to the probability of default (or credit transitions, credit spreads, pre-payment rates, etc.) of the different names/assets in the portfolio. A well known model using this approach is McKinsey's Credit Portfolio View (see [11] and [12] for details).

In the first two families of models, the systemic dependency of default (or credit transitions) is typically captured using factor models. Under this approach, once the systemic risk factors are determined the default (or transition) probability of each firm/account in the portfolio becomes independent of each other. This concept is known as "conditional independence".

In the third family of models, the systematic risk is usually assumed to be determined by the macro-economy. Instruments become independent conditional on a given macroeconomic scenario. This approach is becoming more popular as it allows to build a natural bridge between credit portfolio management and stress testing and IFRS9. Moreover, and quoting a staff the Federal Reserve of New York [13]:

"In principle VaR models can be thought of as the result of thousands of individual scenarios, weighted by their probability. In practice however the distributions are not tied to real-world variables other than the observed empirical distributions of the values of various assets."

In this paper we present an analytical framework that allows us to calculate portfolio credit risk measures analytically. The only assumption made is that of conditional independence, which is common to all three families of models covered above. We will show that thanks to the generality of this analytical framework we will be able to work beyond the single-period default/no-default setting and extend it to a multi-period credit model with credit transitions, spread-risk, pre-payment risk, etc.

We start deriving an analytical expression for the portfolio volatility using two well-known techniques in probability theory, the law of total variance; and the moment generating function (MGF) for the loss probability distribution. This will allow us to calculate the variance (and other moments) of the loss distribution analytically and with the same ease as EL.

We will also show how we can calculate the risk contributions (RC) to the portfolio loss volatility and how these results can be used to perform fast risk-based pricing and optimal portfolio allocation. We will also show how we can extend the results obtained in the derivation of the loss volatility using the MGF method to calculate risk contributions to VaR/ES analytically. Another advantage of this approach is that the reverse stress testing results follow immediately from this analysis. Where reverse stress testing is defined as the most probable combination of risk factors that determine a given loss in the tail.

In summary, this work presents a general framework for credit portfolios that would allow us to calculate analytically:

- Credit Loss Volatility and risk contributions (RC)
- Optimal portfolio allocation and fast risk-based pricing (e.g. RAROC)
- Tail-risk contributions (TRC) for the allocation of VaR/ES
- Reverse stress testing

This framework can also be used to define an "optimal" importance sampling to accelerate the convergence of Monte Carlo simulation in the calculating of VaR/ES. Lastly, it can be used to extend the solutions under the saddle-point approximation to a multi-period credit portfolio models including credit migration and other risk types.

2 Analytical Solutions in Credit Portfolio Modelling

In this section we present our analytical framework, starting by deriving an expression for the calculation of the volatility for the distribution of portfolio losses using two different methods: 1) The law of total variance, 2) The moment generation function (MGF).

We then show how the loss volatility can be allocated to each instrument in the portfolio and how these "Risk Contributions" can be used to perform fast risk-based pricing calculations and portfolio optimisation.

We then revisit the MGF method and explore how the results obtained can be interpreted as risk measures conditional on a given loss amount. These are then used to calculate the tail-risk contributions and perform reverse stress testing.

Finally, we discuss how to approximate the entire loss distribution using the saddlepoint approximation and how significantly accelerate the Monte Carlo simulation using optimal importance sampling.

2.1 Portfolio Loss Definition

We start the derivation with the formal definition of credit loss. Assuming conditional independence we can define the portfolio loss as:

$$L = \int L_z dz$$

where z represents the "state of the world" (or macroeconomic scenario). If z is a discrete random variable then:

$$L = \sum_{z=1}^{Z} w_z L_z$$

Where we have discretised the integral over the probability states into Z buckets and therefore each state of the world occurs with probability w_z. The loss L_z under scenario z can be defined as:

$$L_z = \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{s'=1}^{S} \sum_{s=1}^{S} V_{i,s' \to s,t|z} I_{i,s' \to s,t|z} \tag{1}$$

Here $I_{i,s' \to s,t|z}$ is an indicator function that, conditional on state of the world z, is equal to one if the obligor i transitions from credit state s' to credit state s at time t and zero otherwise. $V_{i,s' \to s,t|z}$ is the value of the loss due to instrument i transitioning from credit state s' to credit state s at time t conditional on state of the world z. It is important to notice that this loss is in general time dependent, capturing the effect of amortisations or utilisation schedules, pre-payment, or the exposure term structure in the case of derivative transactions. The loss also depends on credit quality which might affect the utilisation level on revolving credit lines or the risks of a firm refinancing a loan when its credit quality improves. The dependency of the loss level on state of the world aims to capture the effect of macroeconomic indicators on the value of the instrument, e.g. the level of credit spreads, interest rates or FX.

The last credit state S represent the state of default which may or may not be an absorbent state. The value of the loss when instrument i defaults at time t is:

$$V_{i,s'\to S,t|z} = LGD_{i,t} \cdot e_{i,s',t|z} \tag{2}$$

where $e_{i,s',t|z}$ is the exposure of the instrument n at time t, given credit state s' and conditional on state of the world z. In this case, the Loss Given Default (LGD) is time dependant, to capture the fact that the recovery amount could depend on the time taken for the recovery process to complete.

More generally, the loss due to a credit transition from credit state s' to credit state s at time t is:

$$V_{i,s'\to s,t|z} = e_{i,s,t|z} - e_{i,s',t|z} \tag{3}$$

The value of the instrument $e_{i,s,t|z} = f\left(n_{it}, s, t, T_i, z\right)$ is a function of $n_{i,t}$, the number of units of instrument i at held at time t. In the case of loans, this could represent the exposure amount or the Exposure at Default (EAD). The value also depends, of course, on credit state s, maturity T_i and the state of the world z (which as mentioned before could determine the level of credit spreads, interest rates, FX or pre-payment).

2.1.1 Portfolio Loss Risk Measures

In this section we introduce a number of commonly used risk measures and describe how these are related to the distribution of probability of losses of a portfolio. We start with the most commonly used and better know risk measure, the expected loss (EL). In order to be able to define the EL in a multi-credit state and multi-period context let's first define $p_{i,s'\to s,t|z}$ as the probability of transition from credit state s' to credit state s at time t conditional on state of the world z and starting credit state s_0. To calculate $p_{i,s'\to s,t|z}$, we first need to know what is the probability of being in credit state s' at time t conditional on initial state s_0 at time t_0. For this, let's say that $tp_{i,s'\to s,t|z}$ are the entries of the credit transition matrix at time t and conditional on z. The probability of transition between states after $t-1$ steps of the process is:

$$ctp_{i,s_0\to s,t-1|z} = \sum_{s_1=1}^{S}\sum_{s_2=1}^{S}\cdots\sum_{s_{t-1}=1}^{S} tp_{i,s_0\to s_1,t_1|z}\cdot tp_{i,s_1\to s_2,t_2|z}\cdot\ldots\cdot tp_{i,s_{t-1}\to s,t-1|z}$$

That is, the cumulative transition matrix from initial credit state s' at time t_0 to state s at period t calculated as the product of the transition matrices for each period. Now we can calculate $p_{i,s'\to s,t|z}$ as:

$$p_{i,s'\to s,t|z} = ctp_{i,s_0\to s,t-1|z} \cdot tp_{i,s'\to s,t|z} \tag{4}$$

With these definitions the EL can be calculated as the average loss:

$$EL \equiv \mathbb{E}\left[L\right] = \sum_{z=1}^{Z} w_z \sum_{i=1}^{N}\sum_{t=1}^{T}\sum_{s'=1}^{S}\sum_{s=1}^{S} V_{i,s'\to s,t|z}\, p_{i,s'\to s,t|z} \tag{5}$$

Another commonly used risk measure is the portfolio loss volatility which can be defined as the square root of the portfolio loss variance:

$$\sigma\left(L\right) \equiv \sqrt{\sigma^2\left(L\right)}$$

And the variance is defined as:

$$\sigma^2\left(L\right) \equiv \mathbb{E}\left[L^2 - EL^2\right]$$

Next we define the Value-at-Risk or VaR as:

$$VaR_\beta\left(L\right) \equiv \inf\left\{l \in \mathbb{R} : F_L\left(l\right) \le \beta\right\}$$

In other words, the VaR is the quantile of the probability of loss distribution $F_L\left(l\right)$, such that the probability of losing more than l is less than or equal to β. The VaR is then defined at a time horizon T and confidence level β.

A risk measure that is related to the VaR is the expected shortfall ES, which is sometimes referred to as conditional VaR or CVaR. The ES can be defined as:

$$ES_\beta\left(L\right) \equiv \mathbb{E}\left[L|L > VaR_\beta\left(L\right)\right]$$

In other words the ES is the expected loss conditional on losses being larger than VaR.

The loss distribution could be discrete, for example if only defaults are taking into account as a source of losses; or continuous, if risks such as credit spreads are taken into account. Even if the loss distribution is discrete a continuous portfolio loss probability density can be defined as:

$$f\left(L\right) \equiv \sum_{k=1}^{K} p_k \delta\left(L - L_k\right)$$

where l_k represent a possible portfolio loss and that has probability p_k, $k = 1, \ldots, K$. K is the total number of possible loss combinations in the portfolio and $\delta\left(x - x_0\right)$ is the Dirac delta function. So it is always possible to define loss probability $F_L\left(l\right)$ as:

$$F_L\left(l\right) \equiv \mathrm{Prob}\left(L \le l\right) = \int_{-\infty}^{l} f\left(L\right) dL$$

And the tail loss probability:

$$P_L(l) \equiv \text{Prob}(L > l) = \int_f^{\infty} f(L)\, dL$$

In general, it is not possible to obtain a closed form solution for $f(L)$ unless stringent simplifying assumptions are made such as those in [2]. Typically, credit portfolio models would hence resort to using Monte Carlo or other numerical techniques (such as the saddlepoint method [5]) in order to approximate the loss distribution. In the next section we will show how it is possible to calculate the portfolio volatility and the risk contributions to loss volatility analytically under the assumption of conditional independence. Later, we will also show that although exact closed form solutions for VaR/ES do not exist, once an estimate for this is obtained, we can calculate risk contributions to these risk measures exactly.

2.2 Portfolio Loss Volatility

The volatility of the loss distribution (sometimes also referred to as the unexpected loss) was defined in the previous section as:

$$\sigma(L) = \sqrt{\mathbb{E}\left[L^2 - EL^2\right]}$$

In this section we will show two different methods to calculate a closed form solution for this risk measure. The first method is based on the Law of Total Variance and is easier to derive, the second method is based on the moment generating function and requires a bit more effort. The increase in complexity of the second method is justified as it serves a number of purposes. It allows us to verify the solutions obtained under the first method. It can also be used to calculate closed form solutions for higher moments of the loss distribution (such as the skewness and kurtosis). It also allows us to calculate other moments of the loss distribution conditional on a given loss level. For example we show that the conditional expected loss can be used to allocate the VaR/ES to each instrument in the portfolio. Finally, the conditional moments can also be used to perform analytical calculations for reverse stress testing.

2.2.1 Method 1: Law of Total Variance

We can use the law of total variance and the assumption of conditional independence given a state of the world z to obtain a closed form solution of the portfolio loss volatility. We can write the variance of the probability of portfolio cumulative losses $\sigma^2(L)$ up to period T as:

$$\sigma^2(L) = \mathbb{E}_z\left[Var(L|z)\right] + Var_z\left(\mathbb{E}\left[L|z\right]\right) \tag{6}$$

Starting from the first term in equation (6) and using the fact that conditional on z the expectation $\mathbb{E}[L|z]$ and variance $Var(L|z)$ are additive we get:

$$\mathbb{E}_z\left[Var\left(L_z\right)\right]=\sum_{z=1}^{Z}w_zVar\left(L_z\right)=\sum_{z=1}^{Z}w_z\left(\sum_{i=1}^{N}\sum_{t=1}^{T}\mathbb{E}\left[\left(L_z^2\right)\right]-\mathbb{E}\left[\left(L_z\right)\right]^2\right)=$$

$$=\sum_{z=1}^{Z}w_z\sum_{i=1}^{N}\sum_{t=1}^{T}\sum_{s'=1}^{S}\sum_{s=1}^{S}\left[V_{i,s'\to s,t|z}^2\cdot p_{i,s'\to s,t|z}-\left(V_{i,s'\to s,t|z}\cdot p_{i,s'\to s,t|z}\right)^2\right]$$

For the second term, using the definition of variance of a stochastic variable and the fact that the mean loss is EL, we obtain:

$$Var_z\left(\mathbb{E}\left[L_z\right]\right)=\mathbb{E}_z\left[\left(\mathbb{E}\left[L_z\right]-EL\right)^2\right]=\sum_{z=1}^{Z}w_z\left(\mathbb{E}\left[L_z\right]-EL\right)^2=$$

$$=\sum_{z=1}^{Z}w_z\left(\sum_{i=1}^{N}\sum_{t=1}^{T}\sum_{s'=1}^{S}\sum_{s=1}^{S}V_{i,s'\to s,t|z}\cdot p_{i,s'\to s,t|z}-EL\right)^2$$

Now using the definitions

$$EL_{itz}\equiv\sum_{s'=1}^{S}\sum_{s=1}^{S}V_{i,s'\to s,t|z}\cdot p_{i,s'\to s,t|z}$$

and

$$EL_{iz}\equiv\sum_{t=1}^{T}EL_{itz}$$

$$EL_z\equiv\sum_{i=1}^{N}EL_{iz}$$

$$EL_t\equiv\sum_{z=1}^{Z}w_z\sum_{i=1}^{N}EL_{itz}$$

we finally obtain:

$$\sigma^2\left(L\right)=\sum_{t=1}^{T}\sum_{z=1}^{Z}w_z\left[\sum_{i=1}^{N}\left(\sum_{s'=1}^{S}\sum_{s=1}^{S}V_{i,s'\to s,t|z}^2\cdot p_{i,s'\to s,t|z}\right.\right.$$
$$\left.\left.-EL_{iz}^2+\left(EL_{tz}-EL_t\right)\left(EL_z-EL\right)\right)\right]\tag{7}$$

It is worth noting that the last term captures the temporal (time covariance) dependency in the portfolio volatility. The intra-period variance at time t is given by:

$$\sigma_t^2(L) = \sum_{z=1}^{Z} w_z \sum_{i=1}^{N} \left[\sum_{s'=1}^{S} \sum_{s=1}^{S} V_{i,s' \to s,t|z}^2 \cdot p_{i,s' \to s,t|z} \right. \\ \left. -EL_{itz}^2 + (EL_{tz} - EL_t) \left(EL_z^{(t)} - EL^{(t)} \right) \right] \tag{8}$$

where the superscript $^{(t)}$ indicates that the sum in the time index should be performed up to time t:

$$EL_z^{(t)} \equiv \sum_{i=1}^{N} \sum_{t'=1}^{t} \sum_{s'=1}^{S} \sum_{s=1}^{S} V_{i,s' \to s,t'|z} p_{i,s' \to s,t'|z}$$

And $EL^{(t)} \equiv \sum_{z=1}^{Z} w_z EL_z^{(t)}$.

2.2.2 Method 2: Moment Generating Function

Another way to obtain the variance of a probability distribution is using the moment generating function (MGF) which is defined as:

$$MGF_{f(x)}(\alpha) \equiv \mathbb{E}\left[e^{\alpha f(x)}\right]$$

In the case of the loss distribution this becomes:

$$MGF_L(\alpha) = \mathbb{E}\left[e^{\alpha L}\right]$$

Using the fact that we are working under the framework of conditional independence we can write:

$$MGF_L(\alpha) = \sum_{z=1}^{Z} w_z M_z = \sum_{z=1}^{Z} w_z \prod_{i=1}^{N} M_{i|z} \tag{9}$$

$M_{i|z}$ is, the conditional MGF for obligor i under state of the world z:

$$M_{i|z} = \mathbb{E}\left[e^{\alpha L_{i|z}}\right]$$

$$M_{i|z} = \sum_{t=1}^{T} \sum_{s'=1}^{S} \sum_{s=1}^{S} e^{\alpha V_{i,s' \to s,t|z}} p_{i,s' \to s,t|z}$$

To calculate the moments of the loss distribution it is convenient to introduce the cumulant generating function (CGF):

$$CGF_L(\alpha) = \log\left(MGF_L(\alpha)\right)$$

The different cumulants k_j with $j = 1$ being the mean of the loss distribution, $j = 2$ the variance, $j = 3$ the skewness, $j = 4$ the kurtosis, etc. can then be calculated as:

$$k_j = \left. \frac{d^j CFG_L(\alpha)}{d\alpha^j} \right|_{\alpha \to 0}$$

For example, for the first cumulant we have:

$$EL(\alpha) = \frac{1}{MGF_L(\alpha)} \frac{dMGF_L(\alpha)}{d\alpha} \tag{10}$$

or in terms of $M_{i|z}$:

$$EL(\alpha) = \sum_{z=1}^{Z} \frac{w_z}{MGF_L(\alpha)} \sum_{i=1}^{N} \frac{dM_{i|z}}{d\alpha} \prod_{j \neq i}^{N} M_{j|z}$$

or by multiplying and dividing by $M_{i|z}$:

$$EL(\alpha) = \sum_{z=1}^{Z} \frac{w_z M_z}{MGF_L(\alpha)} \sum_{i=1}^{N} \frac{1}{M_{i|z}} \frac{dM_{i|z}}{d\alpha} \tag{11}$$

And taking derivatives:

$$EL(\alpha) = \sum_{z=1}^{Z} \frac{w_z M_z}{MGF_L(\alpha)} \sum_{i=1}^{N} V_{i,s' \to s,t|z} \frac{e^{\alpha V_{i,s' \to s,t|z}} p_{i,s' \to s,t|z}}{M_{i|z}}$$

In order to simplify the notation let us introduce the following "tilted" transition probabilities ("tilted" in the sense of the Esscher transform [14]):

$$tp_{i,s' \to s,t|z}(\alpha) = tp_{i,s' \to s,t|z} e^{\alpha V_{i,s' \to s,t|z}} \frac{Y_{i,s_0,s',t-1|z}(\alpha)}{Y_{i,s_0,s',t|z}(\alpha)}$$

with $Y_{i,s_0,s',t|z}(\alpha)$ defined as:

$$Y_{i,s_0,s',t|z}(\alpha) = \sum_{s=1}^{S} \cdots \sum_{s_{t-1}=1}^{S} tp_{i,s_0 \to s_1,t|z} e^{\alpha V_{i,s_0 \to s_1,t|z}} \cdot \ldots \cdot tp_{i,s_{t-1} \to s',t|z} e^{\alpha V_{i,s_t \to s',t|z}}$$

with $Y_{i,s_0,s',T|z}(\alpha) = 1$. Note that using this definition we guarantee that sum across the rows of the tilted transition matrix adds to one, and hence these new matrices are proper transition matrices under a different probability measure. The tilted probabilities of transitions at time t, conditional on initial credit state s_0 and z, become:

$$p_{i,s' \to s,t|z}(\alpha) = ctp_{i,s_0 \to s,t-1|z}(\alpha) \cdot tp_{i,s' \to s,t|z}(\alpha) \tag{12}$$

where $ctp_{i,s_0\to s,t|z}(\alpha)$ are also defined in the new probability measure.
Now we can rewrite $EL(\alpha)$ in the more familiar form:

$$EL(\alpha) = \sum_{z=1}^{Z} w_z(\alpha) \sum_{i=1}^{N}\sum_{t=1}^{T}\sum_{s'=1}^{S}\sum_{s=1}^{S} V_{i,s'\to s,t|z} p_{i,s'\to s,t|z}(\alpha) \tag{13}$$

with

$$w_z(\alpha) \equiv \frac{w_z M_z}{MGF_L(\alpha)} \tag{14}$$

This naming convention will become clearer later in section 3.1 when we use these quantities to define the optimal change of measure that can be used to accelerate Monte Carlo simulations. Of course, we have that, $p_{i,s'\to s,t|z}(\alpha) = p_{i,s'\to s,t|z}$ and $w_z(\alpha) \to w_z$ as $\alpha \to 0$.

Now we can calculate the variance of the loss distribution in a similar way by taking the second derivative of the CGF.

$$\sigma^2(\alpha) = \frac{d^2 CGF_L(\alpha)}{d\alpha^2} =$$
$$\frac{1}{(MGF_L(\alpha))^2}\left(MGF_L(\alpha)\frac{d^2 MGF_L(\alpha)}{d\alpha^2} - \left(\frac{dMGF(\alpha)}{d\alpha}\right)^2\right) =$$
$$= \frac{1}{MGF_L(\alpha)}\frac{d^2 MGF_L(\alpha)}{d\alpha^2} - EL^2(\alpha)$$

The only term we do not yet have an expression for is the second derivative of the MGF. Using equation (11) we get:

$$\frac{d^2 MGF_L(\alpha)}{d\alpha^2} = \frac{d}{d\alpha}\left(\sum_{z=1}^{Z} w_z M_z \sum_{i=1}^{N}\frac{1}{M_{i|z}}\frac{dM_{i|z}}{d\alpha}\right) =$$
$$= \sum_{z=1}^{Z} w_z M_z \sum_{i=1}^{N}\left(\frac{1}{M_{i|z}}\frac{dM_{i|z}}{d\alpha}\right)^2 + \sum_{z=1}^{Z} w_z M_z \frac{d}{d\alpha}\left(\sum_{i=1}^{N}\frac{1}{M_{i|z}}\frac{dM_{i|z}}{d\alpha}\right) =$$
$$= \sum_{z=1}^{Z} w_z M_z \Bigg[EL_z^2 + \sum_{i=1}^{N}\sum_{t=1}^{T}\sum_{s'=1}^{S}\sum_{s=1}^{S} V_{i,s'\to s,t|z}^2 p_{i,s'\to s,t|z}(\alpha) -$$
$$- \sum_{i=1}^{N}\sum_{t=1}^{T}\sum_{s'=1}^{S}\sum_{s=1}^{S} V_{i,s'\to s,t|z}^2 p_{i,s'\to s,t|z}^2(\alpha)\Bigg]$$

After rearranging the terms in the expression above the equation for the variance becomes:

$$\sigma^2(\alpha) = \sum_{z=1}^{Z} w_z(\alpha) \left[\sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{s'=1}^{S} \sum_{s=1}^{S} \Big(V^2_{i,s' \to s,t|z} \cdot p_{i,s' \to s,t|z}(\alpha) - EL^2_{itz}(\alpha) \Big) + EL^2_z(\alpha) \right] - EL^2(\alpha)$$

Using the well known property of the variance $\sigma^2(X) = \mathbb{E}\left[(X - \mathbb{E}[X])^2\right] = \mathbb{E}\left[X^2\right] - \mathbb{E}[X]^2$ and setting $\alpha \to 0$ we recover equation (7):

$$\sigma^2(L) = \sum_{z=1}^{Z} w_z \left[\sum_{i=1}^{N} \sum_{t=1}^{T} \left(\sum_{s'=1}^{S} \sum_{s=1}^{S} V^2_{i,s' \to s,t|z} \cdot p_{i,s' \to s,t|z} - EL^2_{itz} \right) + (EL_z - EL)^2 \right]$$

Higher moments of the loss distribution can be calculated in a similar manner.

In order to illustrate the results presented in the following sections, a portfolio consisting of 50 assets with a mixture of corporate and sovereign bonds was chosen with holdings in each asset changing across 4 time periods (4 quarters in this example). Figure 1 shows the holdings in each asset (as percentage of total investment) across the 4 quarters and figure 2 shows these holdings split by credit rating. For simplicity, the probabilities of states of the world are represented by a single latent variable (systemic factor) that affect the credit quality of the issuer and which it is a assumed to be normally distributed. In a more general case actual macroeconomic scenarios could be used instead as in [15].

2.3 Risk Contributions to Portfolio Loss Volatility

Now that we have derived a closed form expression for the portfolio volatility we want to be able to allocated the result down to each instrument in the portfolio. In other words, we want to be able to calculate the risk contributions to the portfolio loss volatility. As already indicated in Section 2.1, the value of the loss $V_{i,s' \to s,t|z}$ is a function of the $n_{i,t}$, which can represent the number of units of instrument i at time t:

$$V_{i,s' \to s,t|z} = n_{it} \cdot V^*_{i,s' \to s,t|z} \tag{15}$$

where $V^*_{i,s' \to s,t|z}$ is the loss per unit of instrument (or per unit of exposure). With this, we can now define the risk contribution (RC) to portfolio loss volatility as:

$$RC_i(L) \equiv n_{it} \frac{\partial \sigma(L)}{\partial n_{it}}$$

And taking the derivatives (see Appendix A), we get:

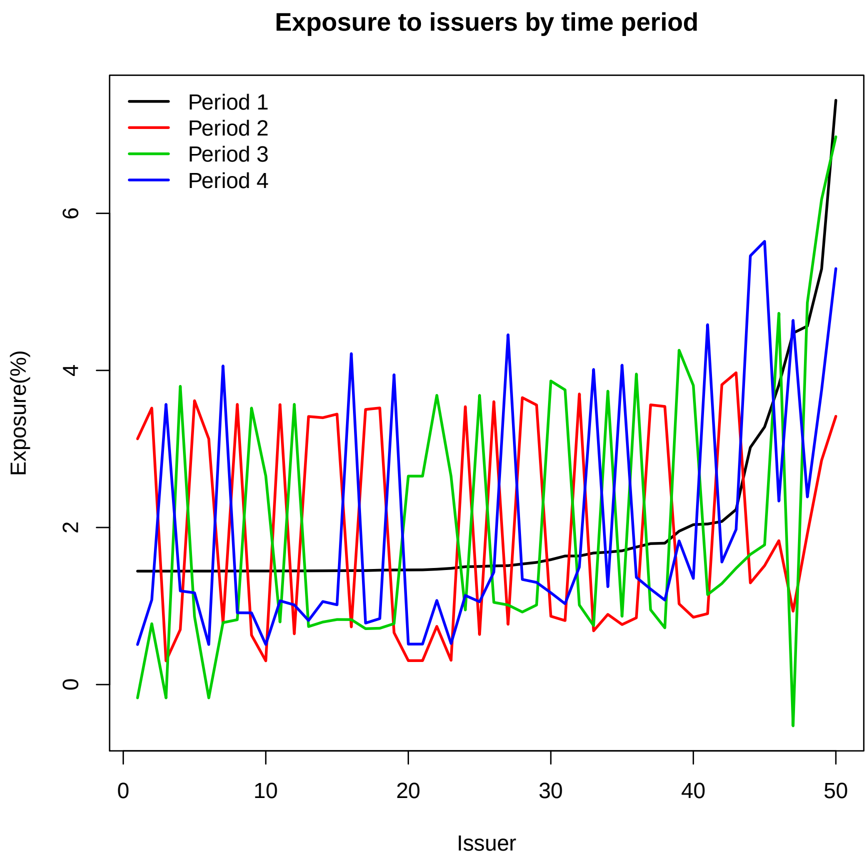

Figure 1: Exposures in percentages on an investment portfolio of 50 assets with a mixture of corporate and sovereign bonds across different time periods

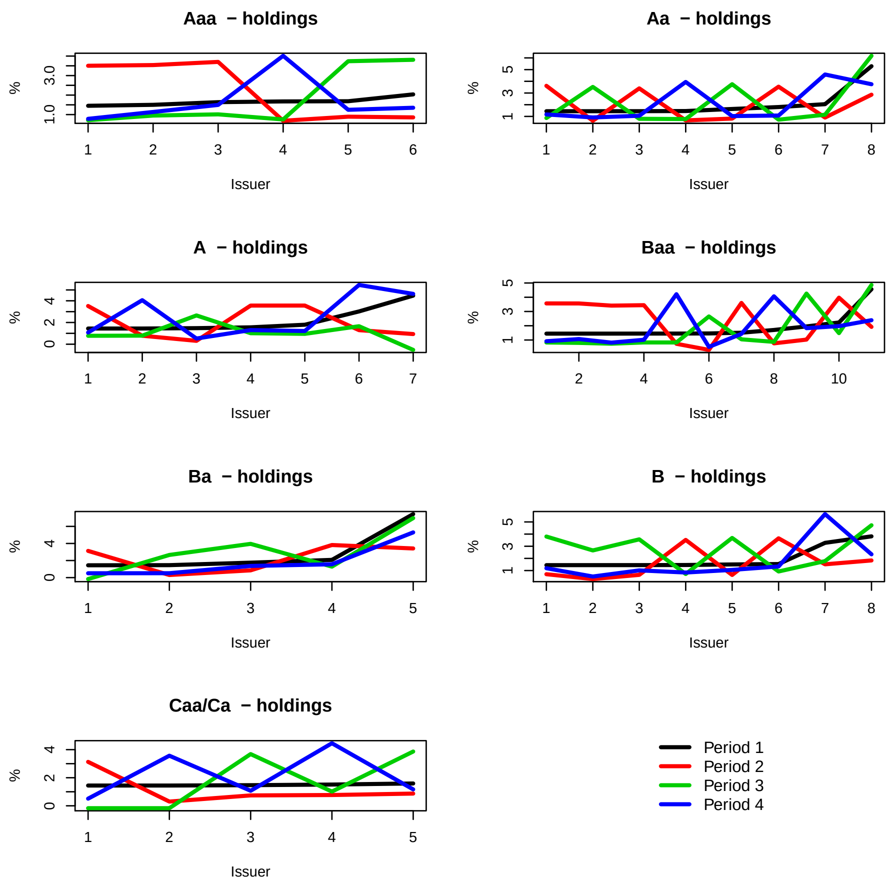

Figure 2: Holdings in percentages for the sample portfolio by credit rating across different time periods

$$RC_i\left(L\right)=\frac{1}{\sigma\left(L\right)}\sum_{t=1}^{T}\left[\sum_{z=1}^{Z}w_z\left(\sum_{s'=1}^{S}\sum_{s=1}^{S}V_{i,s'\rightarrow s,t|z}^{2}\cdot p_{i,s'\rightarrow s,t|z}\right)\right.$$
$$\left.-EL_{itz}^{2}+\left(EL_{itz}-EL_{it}\right)\left(EL_z-EL\right)\right] \quad (16)$$

where

$$EL_{i,z}\equiv\sum_{t=1}^{T}\sum_{s'=1}^{S}\sum_{s=1}^{S}V_{i,s'\rightarrow s,t|z}p_{i,s'\rightarrow s,t|z} \quad (17)$$

and

$$EL_i\equiv\sum_{z=1}^{Z}w_zEL_{i,z} \quad (18)$$

Figure 3 shows the risk contributions to the cumulative portfolio loss volatility after 4 quarters and figure 4 show the cumulative RCs split by rating. Figure 5 present the intra-period RCs across each of the the four quarters and figure 6 the same RCs split by rating.

2.3.1 Analytical Portfolio Optimisation

One of the key advantages of having a closed form solution for the risk contribution to the portfolio volatility is that we can perform fast portfolio optimisation. This is, in general, a very difficult problem to solve and, if a simulation approach is used, is also very resource intensive and time consuming. Under this framework, however, the optimisation can also be done analytically. For example, assuming that the return of instrument i is a linear function of the number of units held n_{it} at time t:

$$\mu_{i,s'\rightarrow s,t|z}=n_{it}\cdot\mu_{i,s'\rightarrow s,t|z}^{*}$$

where $\mu_{i,s'\rightarrow s,t|z}$ is the return for instrument i at time t given a credit transition $\mu_{i,s'\rightarrow s,t|z}^{*}$ and conditional on state of the world z. The expected return of the portfolio can be written as:

$$R\left(n_{it}\right)=\sum_{z=1}^{Z}w_z\sum_{i=1}^{N}\sum_{t=1}^{T}\sum_{s'=1}^{S}\sum_{s=1}^{S}\mu_{i,s'\rightarrow s,t|z}p_{i,s'\rightarrow s,t|z}$$

And the average value of the portfolio at time $t=1$ as:

$$P\left(n_{it}\right)=\sum_{z=1}^{Z}w_z\sum_{i=1}^{N}e_{i,s=s_0,t=1|z}$$

where s_0 is the initial credit state. What we want to do is to find the portfolio with the largest expected return $\bar{R}$ and smallest possible return volatility subject

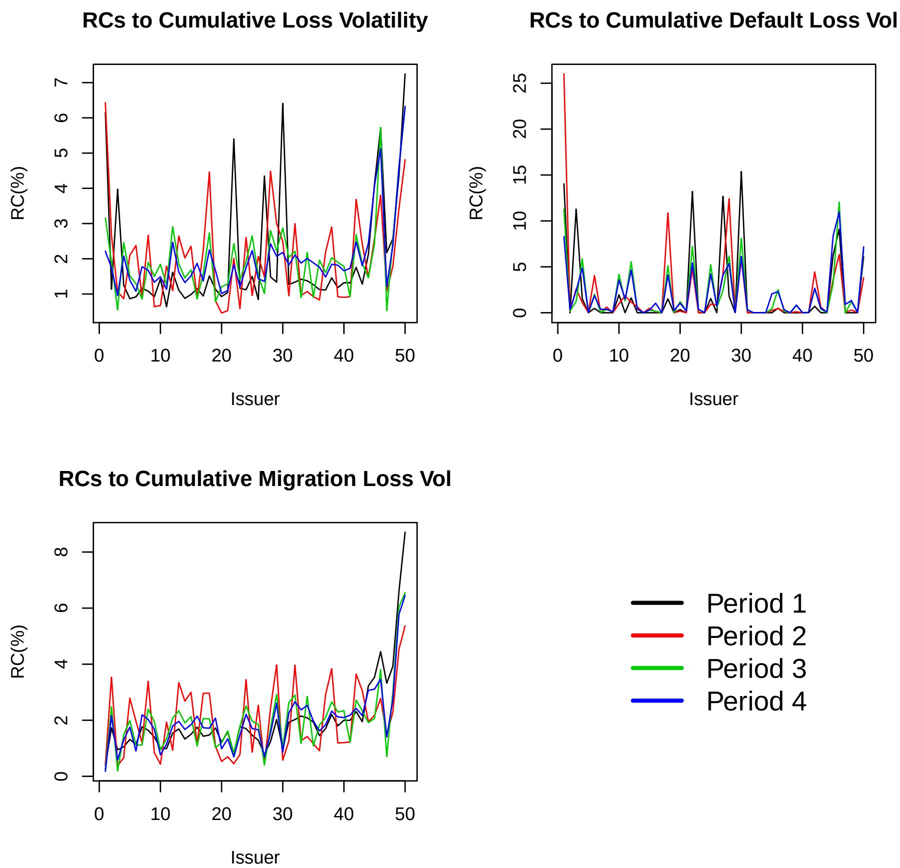

Figure 3: Risk contributions to portfolio cumulative loss volatility across different time periods for total loss (top left), default loss (top right) and migration loss (bottom left)

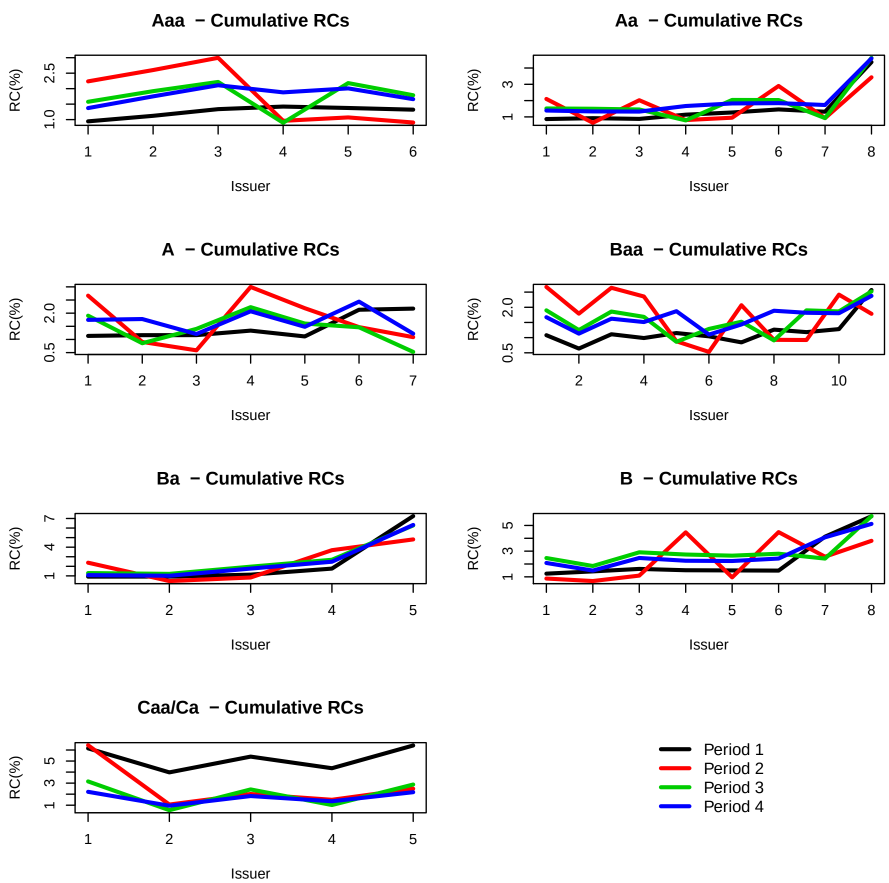

Figure 4: Risk Contributions to portfolio cumulative loss volatility across different time periods by credit rating

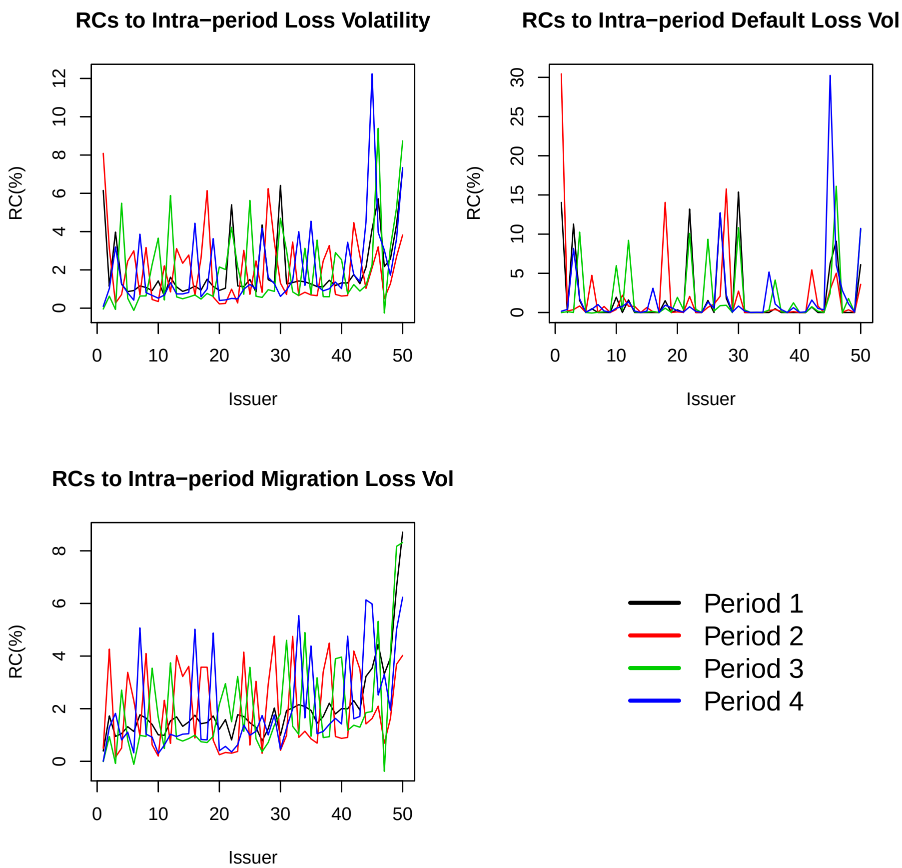

Figure 5: Risk Contributions to portfolio intra-period loss volatility across different time periods for total loss (top left), default loss (top right) and migration loss (bottom left)

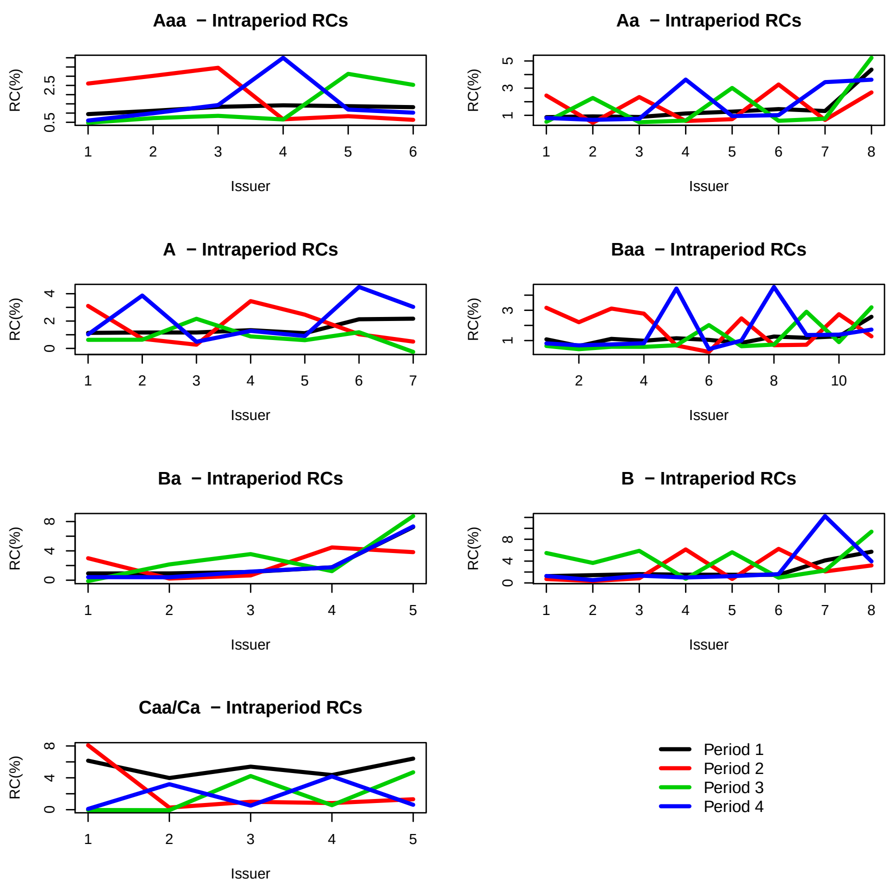

Figure 6: Risk contributions to portfolio intra-period loss volatility across different time periods by credit rating

to budget constrain $\bar{P}$. We solve this optimisation problem using the Lagrange multipliers method:

$$\Lambda\left(n_{it}\right) = \sigma\left(L; n_{it}\right) - \lambda_1\left(P\left(n_{it}\right) - \bar{P}\right) - \lambda_2\left(R\left(n_{it}\right) - \bar{R}\right)$$

Taking derivatives we have:

$$\begin{aligned}
\frac{\partial\Lambda\left(n_{it}\right)}{\partial n_{it}} = \frac{\partial\sigma\left(L, n_{it}\right)}{\partial n_{it}} - \lambda_1 P_i^* - \lambda_2 R_i^* \\
\frac{\partial\Lambda\left(n_{it}\right)}{\partial\lambda_1} = \bar{P} - P\left(n_{it}\right) \\
\frac{\partial\Lambda\left(n_{it}\right)}{\partial\lambda_2} = \bar{R} - R\left(n_{it}\right)
\end{aligned}$$

All that is left to do, is to find the solution to the following system of equations:

$$\begin{aligned}
n_{it}RC_i - \lambda_1 P_i - \lambda_2 R_i = 0 \\
\sum_{i=1}^{N} n_{it}P_i^* - \bar{P} = 0 \\
\sum_{i=1}^{N} n_{it}R_i^* - \bar{R} = 0
\end{aligned}$$

2.3.2 Fast New Deal Analysis

When making an investment decision or when deciding whether to grant a new credit facility to a client it is important to understand, in a timely manner, what would be the return that this new investment will produce given the increased risk in the context of the current portfolio. A common measure of the risk-weighted profitability of an investment is based on the ratio of expected return μ_i over the risk contribution to portfolio volatility RC_i (this measure is also known as the Sharpe ratio). The main advantage of having a closed form solution for the risk contributions to portfolio volatility is that profitability calculations can be performed quickly and without compromise on accuracy need to rely on slow and computationally intensive Monte Carlo simulations.

Theoretically, in order to improve the profitability of the portfolio, the new Sharpe ratio after on-boarding the new deal in the portfolio needs to be larger than the original Sharpe ratio:

$$\frac{\mu_P + \mu_i}{\sigma\left(L\right) + RC_i} > \frac{\mu_P}{\sigma\left(L\right)}$$

where μ_P is the expected return of the portfolio before the new deal is added. More generally, in order to take into account costs and profitability targets, the new Sharpe ratio is usually compared against a minimum hurdle rate H:

$$\frac{\mu_P + \mu_i}{\sigma(L) + RC_i} > H$$

2.4 Conditional Portfolio Statistics

In section 2.2.2 we showed how to calculate the different moments of the loss distribution using the MGF and taking the limit $\alpha \to 0$. In this section, we take this a step further and show that, for any other value of α, the solutions represent moments of the loss distribution conditional on a given level of loss $\bar{L}$. We also show that the expected loss for instrument i conditional on a given loss level equal its tail-risk contributions (TRCs). In addition, the same analytical calculation provides can be used to perform reverse stress testing analytically.

2.4.1 Tail Risk Contributions

Similar to the risk contributions to portfolio loss volatility, the tail risk contribution (TRC) can be defined as the contribution to the overall loss level $\bar{L}$ of instrument i,

$$TRC_i \equiv n_{it} \frac{\partial \bar{L}}{\partial n_{it}} \tag{19}$$

where n_{it} is the holding amount of the instrument as defined in equation (15). Under certain conditions (see for example [16]), it can be shown that this is equivalent to:

$$TRC_i \equiv \mathbb{E}\left[L_i | L = \bar{L}\right] \tag{20}$$

To calculate the TRCs for a portfolio using Monte Carlo simulation, we need to define a small loss region (ϵ) around the target loss level $\bar{L}$ and select only simulations that produce losses within that interval (i.e. a simple form of rejection sampling). An approximated TRC for instrument i can the be calculated as:

$$TRC_i^* = \mathbb{E}\left[L_i | L - \epsilon < \bar{L} < L + \epsilon\right] \tag{21}$$

Note that TRC_i and TRC_i^* should convergence as $\epsilon \to 0$. Hence, the smaller the chosen interval the lesser the bias in the calculation of the TRCs. On the other hand, choosing a value of ϵ that is too small will cause too many of the Monte Carlo simulations to be rejected and its convergence will be slow. It is important to note that this definition is not without issues (see section 3.1 of [17] for an illuminating example). Luckily, with the techniques that we have already developed, we can avoid these types of issues as well as the slow convergence of Monte Carlo simulations for calculating tail risk measures. We can calculate the TRCs as:

$$TRC_i = EL(\alpha) \tag{22}$$

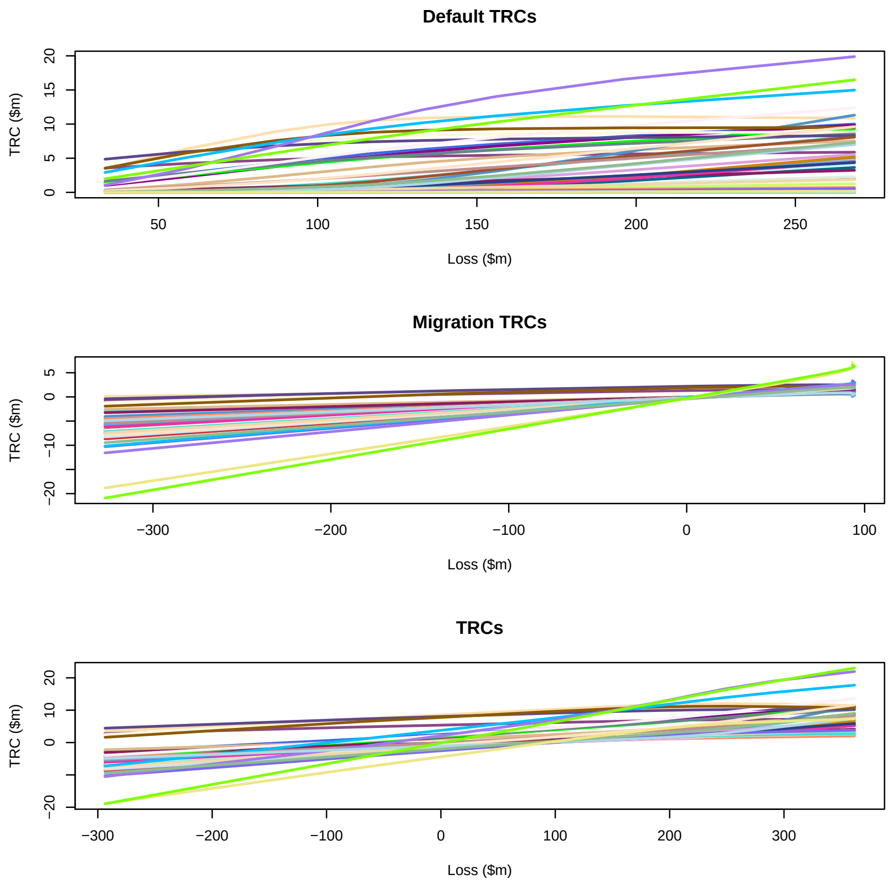

Figure 7: Tail Risk Contributions for different profit and loss (P&L) levels for default loss only (top), migration P&L (middle) and total P&L

where α is obtained by solving for α the following equation:

$$EL(\alpha) = \bar{L} \tag{23}$$

In other words, portfolio statistics conditional on loss level $\bar{L}$ can be simply calculated by solving equation (23) for α. Obviously, choosing $\bar{L} = EL$ implies that $\alpha = 0$ and we recover the expression for the EL. Refer to Appendix B for a proof of this result.

Figure 7 shows how the Tail Risk Contributions change with the chosen level of loss in the portfolio and figure 8 shows how α (Alpha) and portfolio loss level relate to each other.

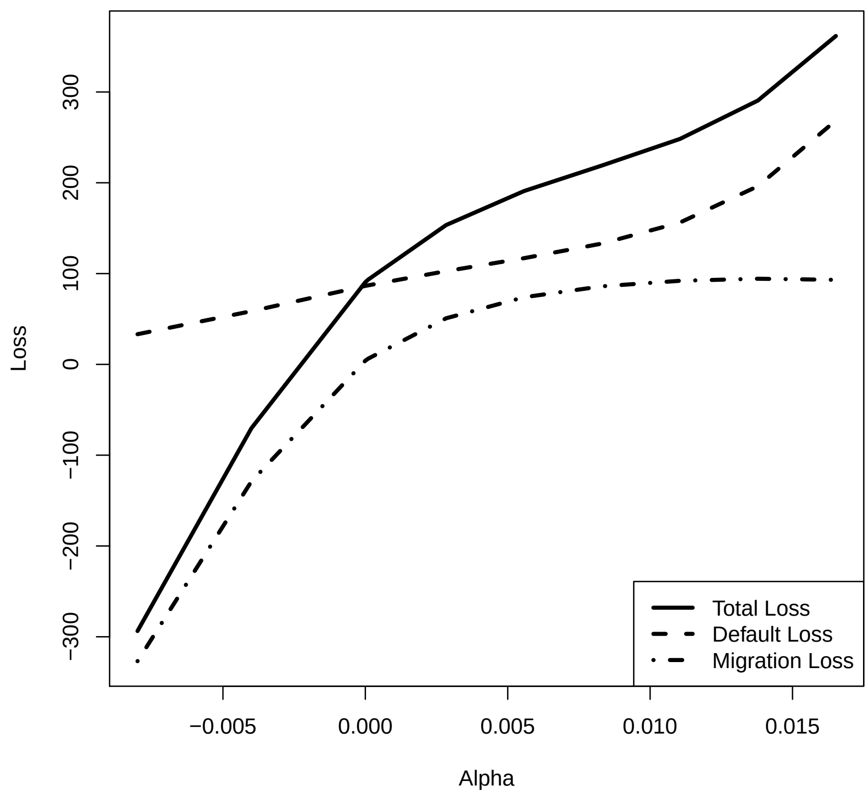

Figure 8: This figure shows how the portfolio loss level varies with the value of Alpha (α), it is interesting to note that for negative values of Alpha (corresponding to profit in the P&L distribution), the P&L is dominated by migrations, however as Alpha increases so does the contribution to default losses to the total P&L

2.4.2 Reverse Stress Testing

In the previous section we showed that there is a one-to-one correspondence between the level of loss and the parameter α of the moment generating function, linked by the equation (23). Hence, for any loss level there is one value of α that we can use to calculate conditional portfolio moments. For the case of the expected loss, conditioning on a given loss level means shifting the probabilities of transition from $p_{i,s' \to s,t|z}$ to $p_{i,s' \to s,t|z}(\alpha)$ and the probability of the states of the world from w_z to $w_z(\alpha)$. In other words, by conditioning on that loss level we are making certain transition probabilities more likely, while at the same time, we are changing the distribution of probabilities of states of the world. So for example, conditioning on a large loss would have the effect of making "bad states of the world", together with credit transitions that would result losses, more likely. Therefore, selecting the most probable scenarios z conditional on loss level is equivalent to perform a reverse stress testing analysis.

Figure 9 shows how conditioning on a given loss level changes the shape of the probability distribution of "states of the world". If discrete macroeconomic scenarios were used instead, for example by using procedures such as the ones described in [15], conditioning on a given level of loss would make some of the scenarios more likely, thus facilitating the selection of candidate scenarios to perform reverse stress testing.

2.5 Analytical Solutions for Portfolios with Pooled Exposures

For larger portfolios with many relatively small exposures, it is often desirable to pool similar exposures into cohorts using some common characteristic (e.g. origination date, credit quality, LTV, etc.). If these pools are homogeneous enough, it is possible to assume that all the instruments in the pool are equal to each other and have the same sensitivity to the state of the world z. Hence if a given cohort c is formed of $N_{c,t|z}$ instruments at time t, conditional on state of the world z, the value of the loss for the cohort will be:

$$V_{c,s' \to s,t|z} = N_{c,t|z} \cdot v_{c,s' \to s,t|z} \tag{24}$$

where $v_{c,s' \to s,t|z}$ represents the average value of the loss for a given instrument in the cohort. Using these pools it is possible to simplify the calculation of the analytical solutions for large portfolios. For cohorted exposures the analytical formula for the portfolio variance becomes:

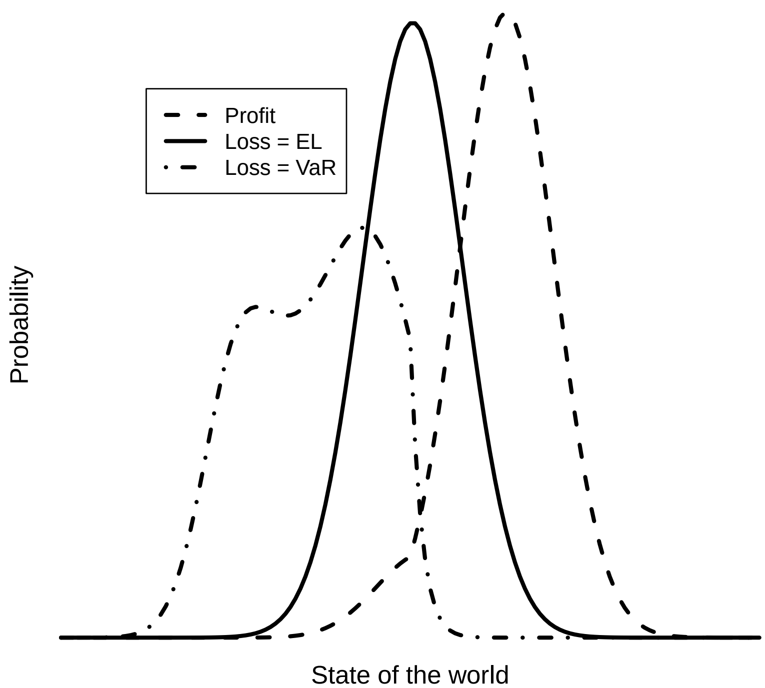

Figure 9: Distribution for the probability of latent variable representing the "state of the world" for three cases: when the the conditional loss is set equal to the expected loss (solid line) the unconditional distribution is recovered. When the conditional loss is set to a profit (dash-dot line), the conditional probability shifts towards good "states of the world". Likewise, when set to a large loss like the 99.9% VaR, the probability shifts towards bad "states of the world"

$$
\begin{aligned}
\sigma^2(L) &= \sum_{t=1}^{T}\sum_{z=1}^{Z} w_z \left[\sum_{i=1}^{N}\left(\sum_{s'=1}^{S}\sum_{s=1}^{S} v_{i,s'\to s,t|z}^2 \cdot p_{i,s'\to s,t|z}\right)\right. \\
&\qquad\qquad \left. -EL_{iz}^2 + (EL_{tz} - EL_t)(EL_z - EL)\right] \\
&= \sum_{t=1}^{T}\sum_{z=1}^{Z} w_z \left[\sum_{c=1}^{C}\frac{1}{N_{c|z}}\left(\sum_{s'=1}^{S}\sum_{s=1}^{S} V_{c,s'\to s,t|z}^2 \cdot p_{c,s'\to s,t|z}\right)\right. \\
&\qquad\qquad \left. -EL_{iz}^2 + (EL_{tz} - EL_t)(EL_z - EL)\right]
\end{aligned}
$$

where we have used the fact that $um_{i=i}^{N} N_{i|z} = \sum_{c=1}^{C} N_{c|z}$ and equation (24) together with the fact that two instruments in the same cohort share the same values of $v_{i=1,s'\to s,t|z}, \cdots, v_{i=N,s'\to s,t|z} = v_{c,s'\to s,t|z}$ and $p_{i=1,s'\to s,t|z}, \cdots, p_{i=N,s'\to s,t|z} = p_{c,s'\to s,t|z}$. The equations for risk contributions for pooled exposures to portfolio volatility follow trivially from this.

The expressions for the conditional loss statistics are also easy to find. Noting that, equation (9) becomes:

$$
MGF_L(\alpha) = \sum_{z=1}^{Z} w_z M_z = \sum_{z=1}^{Z} w_z \prod_{i=1}^{N} M_{i|z} = \sum_{z=1}^{Z} w_z \prod_{c=1}^{C} \left(M_{c|z}\right)^{N_{c|z}}
$$

And from this is straight forward to show that:

$$
EL(\alpha) = \sum_{z=1}^{Z} w_z(\alpha) \sum_{c=1}^{C}\sum_{t=1}^{T}\sum_{s'=1}^{S}\sum_{s=1}^{S} V_{c,s'\to s,t|z}\, p_{c,s'\to s,t|z}(\alpha)
$$

2.5.1 Accounting for Heterogeneity in Portfolios with Pooled Exposures

It is possible to reduce the error introduced when creating homogeneous pools by allowing certain degree of heterogeneity in the cohorts. For example, imagine that the instrument in cohort c is such that the probability of transition take value $tp_{c,s'\to s,t|z}(x)$ according to a density distribution function $f(x)$, e.g.

$$
tp_{c,s'\to s,t|z} = \int tp_{c,s'\to s,t|z}(x)\, f(x)\, dx
$$

subject to the condition:

$$
\sum_{s'=1}^{S} tp_{c,s'\to s,t|z}(x) = 1
$$

If we assume that the deviations $tp_{c,s'\to s,t|z}(x)$ from the mean of the cohort $tp_{c,s'\to s,t|z}$ are proportional shocks $h(x)$ of the form:

$$tp_{c,s'\to s,t|z} = \int tp_{c,s'\to s,t|z} \cdot h(x) f(x)\, dx$$

it follows that

$$\int h(x) f(x)\, dx = 1$$

and since $f(x)$ is a density function we also have that:

$$\int f(x)\, dx = 1$$

Using these properties is easy to see that the expression for the loss variance for a cohorted portfolio with heterogeneous pools is:

$$\sigma^2(L) = \sum_{t=1}^{T}\sum_{z=1}^{Z} w_z \left[\sum_{c=1}^{C} \frac{1}{N_{ct|z}} \left(\sum_{s'=1}^{S}\sum_{s=1}^{S} V^2_{c,s'\to s,t|z} \cdot p_{c,s'\to s,t|z} \right) \right.$$
$$\left. - EL^2_{iz} \int h^2(x) f(x)\, dx + (EL_{tz} - EL_t)(EL_z - EL) \right]$$

And the $EL(\alpha)$ becomes:

$$EL(\alpha) = \int f(x) \left(\sum_{z=1}^{Z} w^*_z(\alpha) \sum_{c=1}^{C}\sum_{t=1}^{T}\sum_{s'=1}^{S}\sum_{s=1}^{S} V_{c,s'\to s,t|z} p^*_{c,s'\to s,t|z}(\alpha)\, h(x) \right) dx$$

with $p^*_{c,s'\to s,t|z}$ and $w^*_z(\alpha)$ having the same expression as equations 12 and 14 but with all the $p_{c,s'\to s,t|z}$ replaced by the product $h(x) \cdot p_{c,s'\to s,t|z}$.

3 Analytical Approximations & Optimal Simulation in Credit Portfolio Modelling

In this section we present two approaches that can be used to estimate the VaR/ES measures for a portfolio, first we show how Monte Carlo convergence can be accelerated by using an optimal importance sampling technique. Then we show the VaR and ES can be estimated using the saddlepoint approximation.

3.1 Optimal Importance Sampling

In Section 2.4.2 we showed how reverse stress testing, i.e. finding the states of the world that would more likely result in a given portfolio loss, could be calculated analytically. In this section we are going to formalise this result and show how the same idea can be used to define a change of measure that results in an "optimal" importance sampling method. For example, to accelerate the convergence of the estimation of tail risk measures, it would be ideal if we could somehow select the area of interest around a section of the tail of the distribution, and concentrate the simlated losses around that area. Say that we are trying to estimate the tail probability $P(L > l)$, the importance sampling estimator will be:

$$\hat{P}(L > l) = \mathbb{E}_g \left[\mathbb{I}_{L>l} \frac{f(L)}{g(L)} \right]$$

What we would like to do is to choose the new importance measure $g(x)$ such as to minimise the variance of the estimator:

$$min_q \left[Var_q \left(\mathbb{E}_g \left[\mathbb{I}_{L>l} \frac{f(L)}{g(L)} \right] \right) \right] \tag{25}$$

In theory, the optimal importance sampling density function would be:

$$g_{opt}(L) = \mathbb{I}_{L>l} \frac{f(L)}{P(L > l)} \tag{26}$$

which has zero variance, however, implies that we know the very same parameter we are trying to estimate, $P(L > l)$. Instead, we look for the "closest" density function to g_{opt} that can be calculated easily. For this we use a variant of the cross-entropy method [18]. If we define "closest" in terms of the the Kullback-Leibler (K-L) divergence:

$$D_{KL}(g||h) \equiv \mathbb{E}_g \left[\log \left(\frac{g(L)}{h(L)} \right) \right]$$

We can define the importance sampling density $g(L)$ that minimises:

$$D_{KL}(g||f) = \int g(L) \log \left(\frac{g(L)}{f(L)} \right) dL$$

subject to the constrain

$$\int L g(L)\, dL = l$$

This ensures that using importance sampling the simulations will be concentrated in the tail of the distribution. And, at the same time, the importance sampling density $g(L)$ will be "as close as possible", in the K-L sense, to the real loss distribution $f(L)$ in the tail. We can carry out the minimisation using the Lagrange multipliers method:

$$\begin{aligned} \Lambda(g(L), \lambda) = & \int g(L)\, log \frac{g(L)}{f(L)} dL - \lambda \left(\int Lg(L)dL - l \right) \\ & \int \left(log \frac{g(L)}{f(L)} - \lambda L \right) g(L)\, dL - \lambda l \end{aligned}$$

Minimising this Lagrangian we get:

$$\delta\Lambda\left(g\left(L\right),\lambda\right)=\int\left(log\frac{g\left(L\right)}{f\left(L\right)}-\lambda L\right)\delta g\left(L\right)=0$$

and thus

$$g\left(L\right)\propto f\left(L\right)e^{\lambda L}$$

In order to ensure that $g\left(L\right)$ is a proper probability density we need to impose the normalisation condition:

$$\int g\left(L\right)dL=1$$

which yields

$$g\left(L\right)=f\left(L\right)\frac{e^{\lambda L}}{\mathbb{E}_f\left[e^{\lambda L}\right]} \tag{27}$$

Now, we have already seen that the denominator is actually the $MGF_L\left(\alpha\right)$ with $\lambda\to\alpha$ and we already have the expression for $EL\left(\alpha\right)=l$ from equation (10) . The constrain can then be expressed as:

$$\int Lg\left(L\right)dL=EL\left(\alpha\right)=\sum_{z=1}^{Z}w_z\left(\alpha\right)p_{i,s'\to s,t|z}\left(\alpha\right)=l$$

From which we can be readily implemented without needing an expression for $f\left(L\right)$, let alone $g\left(L\right)$. The algorithm to estimate $g\left(L\right)$ would be as follows: for each Monte Carlo run, draw the state of the world z with probability $w\left(\alpha\right)$ and then, for each instrument in the portfolio, determine the credit state at time t according to the probability of transition $p_{i,s'\to s,t|z}\left(\alpha\right)$. Finally, we can recover the $f\left(L\right)$ from $g\left(L\right)$ inverting equation (27).

The intuition behind this, is that in order to attain losses in the tail of the distribution two things are needed. Firstly, a "bad" state of the world needs to be drawn (as large losses are more likely for "bad" z's). Secondly, the simulation needs to produce a state of the portfolio where large number of credit downgrades and defaults occur (note that the probability of downgrade/default $p_{i,s'\to s,t|z}\left(\alpha\right)$ increases with increasing values of α).

Figures 10 and 11 show Monte Carlo simulation results for the reference portfolio with an without importance sampling.

3.2 Saddlepoint Approximation for Credit Portfolios in a Multiperiod and Credit Migration Setting

In this section, we present a powerful analytical approximations to these risk measures which are derived using the saddlepoint approximation. For this, we follow a similar approach as in [19], however using the findings obtained

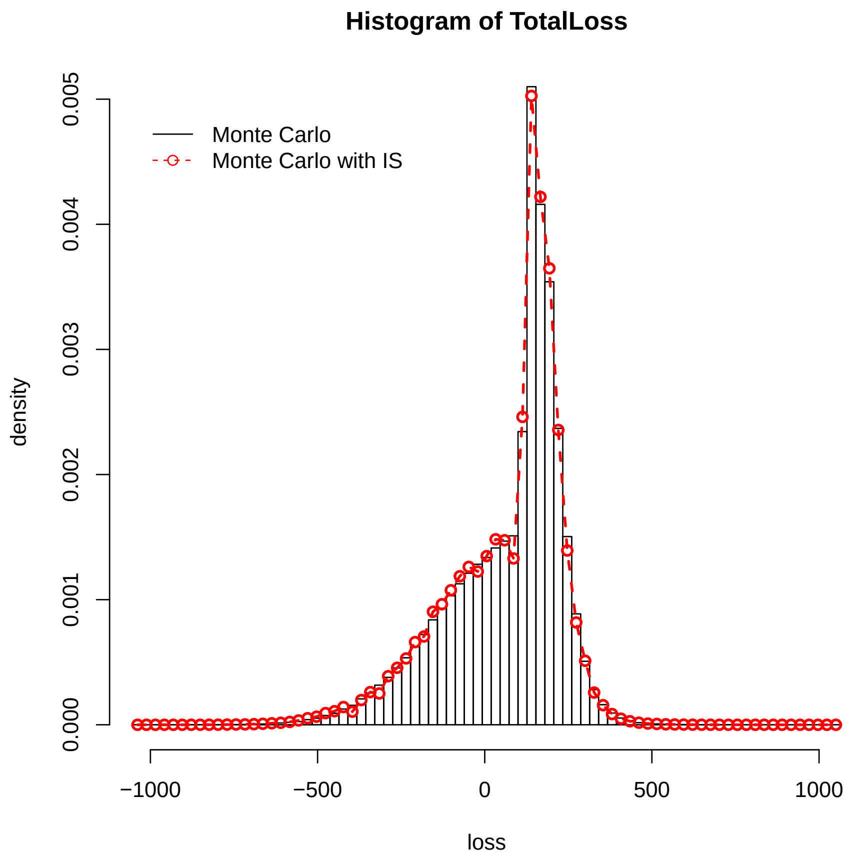

Figure 10: Monte Carlo simulation with (red dots and dash line) and without (bars) importance sampling. For the simulation with importance sampling 3 runs with 10,000 Monte Carlo samples where performed, for the first run alpha was chosen to concentrated to the simulations in the left tail of the distribution. In the second run, alpha was set to zero, thus corresponding to a simulation without importance sampling to ensure that the body of the distribution was estimated properly. In the last run, a positive value of alpha was chosen to simulated the right tail of the loss distribution. For the Monte Carlo simulation without importance sampling 1 million samples where taken

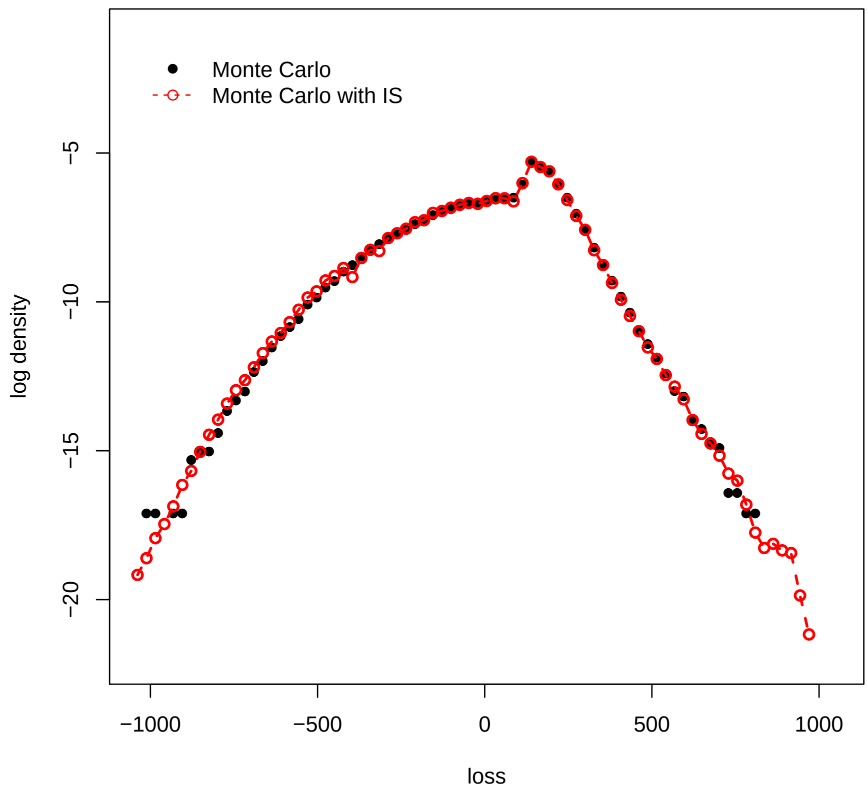

Figure 11: The same Monte Carlo results are shown in logarithmic scale so that the full power of the optimised importance sampling approach can be better appreciated. Note that to achieve the same level of accuracy without importance sampling many millions of simulations would be required therefore, the importance sampling introduced here can potentially speed up convergence by several orders of magnitude

in previous sections, we can extend the solutions from a single period default no-default setting into the generic framework covered here.

In the following we will use some of the definitions introduce in Appendix B. Please note that as in [19] we are going to assume independence of the credit losses. Generic results under conditional independence can be obtained by calculating weighted averages (using the weights w_z) of the solutions presented below.

The probability density function can be defined in terms of the characteristic function as:

$$f(L) \equiv \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{-i\omega L}\phi_L(\omega)\,d\omega$$

or where it exists (as is our case here), in terms of the moment generating function, we can write it as:

$$f(L) \equiv \frac{1}{2\pi i}\int_{c-i\infty}^{c+i\infty} e^{-\alpha L} MGF_L(\alpha)\,d\alpha$$

We can also write the above equation in terms of the cumulant generating function, which from here on we denote $K_L(\alpha)$ to simplify the notation. With this we get:

$$f(L) = \frac{1}{2\pi i}\int_{c-i\infty}^{c+i\infty} e^{K_L(\alpha)-\alpha L} d\alpha$$

The saddlepoint method consists in deforming the path of integration in this contour integral so that it lies along the path of steepest descent. This path is determined by:

$$\frac{\partial\left(K_L(\alpha)-\alpha L\right)}{\partial\alpha} = 0 = \frac{\partial K_L(\alpha)}{\partial\alpha} - L$$

which is equivalent to $EL(\alpha) = L$. So, in this case, the path of steepest descent goes parallel to the imaginary line and passes through $\hat{\alpha}$:

$$K_L'(\hat{\alpha}) = \left.\frac{\partial K_L(\alpha)}{\partial\alpha}\right|_{\alpha=\hat{\alpha}}$$

We can approximate the exponent $K_L(\alpha)-\alpha L$ using its Taylor expansion around $\hat{\alpha}$:

$$\begin{aligned} K_L(\alpha)-\alpha L &\approx K_L(\hat{\alpha}) - \alpha L + (\alpha-\hat{\alpha})K_L'(\hat{\alpha}) + \frac{1}{2}(\alpha-\hat{\alpha})^2 K_L''(\hat{\alpha}) \\ &= K_L(\hat{\alpha}) - \hat{\alpha}L + \frac{1}{2}(\alpha-\hat{\alpha})^2 K_L''(\hat{\alpha}) \end{aligned}$$

Noting that $K_L''(\hat{\alpha}) = \sigma_L^2(\hat{\alpha})$, we get:

$$f(L) \approx \frac{e^{K_L(\hat{\alpha})-\hat{\alpha}L}}{2\pi i} \int_{\hat{\alpha}-i\infty}^{\hat{\alpha}+i\infty} e^{\frac{1}{2}(\alpha-\hat{\alpha})^2\sigma_L^2(\hat{\alpha})} d\alpha$$

which is a Gaussian integral,

$$f(L) \approx \frac{e^{K_L(\hat{\alpha})-\hat{\alpha}L}}{\sqrt{2\pi} \cdot \sigma_L(\hat{\alpha})}$$

If further terms are taken in the Taylor expansion, the expression can be expanded to:

$$f(L) \approx \frac{e^{K_L(\hat{\alpha})-\hat{\alpha}L}}{\sqrt{2\pi K_L''(\hat{\alpha})}} \left[1 + \left(\frac{K_L'''(\hat{\alpha})}{8K_L''(\hat{\alpha})^2} - \frac{5K_L'''(\hat{\alpha})^2}{24K_L''(\hat{\alpha})^3} + \ldots\right)\right]$$

To calculate the VaR, we can use the expression for the tail loss probability:

$$P(l) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-i\omega L} \phi_L(\omega) \frac{d\omega}{\omega}$$

which, in terms of the cumulant generating function becomes:

$$P(l) = \frac{1}{2\pi i} \int_{c-i\infty}^{c+i\infty} e^{K_L(\alpha)-\alpha l} \frac{d\alpha}{\alpha}$$

Proceeding as above, it is easy to show that for $l > EL$ we get:

$$P(l) \approx e^{K_L(\hat{\alpha})-\hat{\alpha}l+\frac{1}{2}\hat{\alpha}^2\sigma_L^2(\hat{\alpha})} N\left(-\sqrt{2\pi} \cdot \sigma_L(\hat{\alpha})\right)$$

where $N(x)$ is the cumulative Normal distribution. Higher-order expansions for the tail risk contribution can be found in [20]. For another version of the saddlepoint approximation of the tail probability formula please refer to the Lugannani-Rice formula [21].

The approximation for the expected shortfall can be written in terms of the results above as:

$$ES(l) \equiv \mathbb{E}[L|L > l] \approx EL \cdot P(l) + \frac{l - EL}{\hat{\alpha}} f(l)$$

To show this we start from:

$$\mathbb{E}[L|L > l] = \int_{-l}^{\infty} Lf(L)\,dL = \frac{1}{2\pi i} \int_{-l}^{\infty} \int_{c-i\infty}^{c+i\infty} Le^{K_L(\alpha)-\alpha L} d\alpha dL$$

using the fact that

$$\frac{\partial\left(e^{K_L(\alpha)-\alpha L}\right)}{\partial\alpha} = \left(K_L'(\alpha) - L\right) e^{K_L(\alpha)-\alpha L}$$

we can rewrite the expression above for the ES as:

$$\mathbb{E}\left[L|L>l\right]=\int_{-l}^{\infty}Lf\left(L\right)dL=$$

$$\frac{1}{2\pi i}\int_{-l}^{\infty}\int_{c-i\infty}^{c+i\infty}\left[K_L'\left(\alpha\right)-e^{-K_L(\alpha)+\alpha L}\frac{\partial\left(e^{K_L(\alpha)-\alpha L}\right)}{\partial\alpha}\right]e^{K_L(\alpha)-\alpha L}d\alpha dL$$

$$=\frac{1}{2\pi i}\int_{c-i\infty}^{c+i\infty}K_L'\left(\alpha\right)e^{K_L(\alpha)-\alpha l}\frac{d\alpha}{\alpha}-\frac{1}{2\pi i}\int_{-l}^{\infty}\int_{c-i\infty}^{c+i\infty}\frac{\partial\left(e^{K_L(\alpha)-\alpha L}\right)}{\partial\alpha}d\alpha$$

And as the second integral vanishes we are left with:

$$ES\left(l\right)=\frac{1}{2\pi i}\int_{c-i\infty}^{c+i\infty}K_L'\left(\alpha\right)e^{K_L(\alpha)-\alpha l}\frac{d\alpha}{\alpha}$$

A number of different expressions exist for the saddlepoint approximations of the ES, [22] is a good reference. Also, in [23] one can find expressions for the saddlepoint approximation of more general conditional expectations.

Appendix A Risk Contributions to Porfolio Loss Volatility

In this appendix we derive equation (16) starting from $RC_i(L) = n_{it}\frac{\partial\sigma(L)}{\partial n_{it}}$. we know that:

$$\frac{\partial\sigma(L)}{\partial n_{it}} = \frac{1}{2\sigma(L)}\frac{\partial\sigma^2(L)}{\partial n_{it}}$$

Re-arranging equation (7) by bringing the sums in i and t to the front, we get:

$$\sigma^2(L) = \sum_{i=1}^{N}\sum_{t=1}^{T}\left[\sum_{z=1}^{Z} w_z\left(\sum_{s'=1}^{S}\sum_{s=1}^{S} V^2_{i,s'\to s,t|z}\cdot p_{i,s'\to s,t|z}\right)\right.$$
$$\left. -EL^2_{itz} + (EL_{itz} - EL_{it})(EL_z - EL)\right]$$

with

$$EL_{itz} \equiv \sum_{s'=1}^{S}\sum_{s=1}^{S} V_{i,s'\to s,t|z}\, p_{i,s'\to s,t|z}$$

and

$$EL_{it} \equiv \sum_{z=1}^{Z} w_z \sum_{t=1}^{T} EL_{itz}$$

Now noting that from equation (15) we have:

$$n_{it}\frac{\partial V_{i,s'\to s,t|z}}{\partial n_{it}} = n_{it}V^*_{i,s'\to s,t|z} = V_{i,s'\to s,t|z}$$

And $\sigma^2(L)$ is quadratic in $V_{i,s'\to s,t|z}$ we obtain:

$$n_{it}\frac{\partial\sigma^2(L)}{\partial n_{it}} = 2\sigma^2(L)$$

recovering equation (16):

$$RC_i(L) = \frac{1}{\sigma(L)}\sum_{t=1}^{T}\left[\sum_{z=1}^{Z} w_z\left(\sum_{s'=1}^{S}\sum_{s=1}^{S} V^2_{i,s'\to s,t|z}\cdot p_{i,s'\to s,t|z}\right)\right.$$
$$\left. -EL^2_{itz} + (EL_{itz} - EL_{it})(EL_z - EL)\right]$$

Appendix B Analytical Solution to the Tail Risk Contributions

In this appendix we will show the equivalence of equations (20) and (22). Let's start by going through some definitions.

The **characteristic function** of the portfolio loss distribution can be defined as:

$$\phi_L(\omega) \equiv \mathbb{E}\left[e^{i\omega L}\right]$$

the characteristic function is closely related to the **moment generation function** of the portfolio loss distribution (MGF_L). This can be defined as:

$$MGF_L(\alpha) \equiv \mathbb{E}\left[e^{\alpha L}\right]$$

The characteristic function of a probability distribution is always well defined. On the other hand, the moment generating function does not necessarily exist for all probability distributions. However, in our case, the loss distribution is bounded (i.e. the maximum loss is not infinite), and all the moments of the loss distribution can be calculated. This means that the MGF_L exists and the transformation ($\alpha \to i\omega$) is well defined. So we can write:

$$\phi_L(\omega) = MGF_L(i\omega)$$

Using the characteristic function, the **portfolio loss probability density function** can be defined as:

$$f(L) \equiv \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{-i\omega L}\phi_L(\omega)\,d\omega$$

The **loss cumulative probability function** can be defined as:

$$F(l) \equiv \int_{-\infty}^{l} f(L)\,dL$$

And from this, the **tail loss** can be defined as:

$$P(l) \equiv 1 - F(l) = Pr(L > l) = \int_{l}^{\infty} f(L)\,dL$$

which in terms of the characteristic function becomes:

$$P(l) = \int_{l}^{\infty} f(L)\,dL = \frac{1}{2\pi}\int_{l}^{\infty}\int_{-\infty}^{\infty} e^{-i\omega L}\phi_L(\omega)\,d\omega dL$$

Taking the integral over L and using the fact that the probability vanishes as $L \to \infty$, we get:

$$P(l) = \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{-i\omega L}\phi_L(\omega)\,\frac{d\omega}{\omega}$$

To calculate the tail risk contributions, we want to vary n_{it} while keeping the tail loss probability constant and fixed at $L = \bar{L}$ (note that the derivative of the probability is zero since it is fixed for a particular confidence level and does not depend on n_{it}). In other words,

$$0 = \left.\frac{\partial P(l)}{\partial n_{it}}\right|_{L=\bar{L}} = \frac{1}{2\pi}\int_{-\infty}^{\infty}\left[\frac{\partial \phi_L(\omega)}{\partial n_{it}} - \frac{\partial \bar{L}}{\partial n_{it}} i\omega\phi_L(\omega)\right] e^{-i\omega\bar{L}}\phi_L(\omega)\frac{d\omega}{\omega}$$

using the fact that the integrand must vanish we get:

$$\frac{\partial \bar{L}}{\partial n_{it}} = \left.\frac{1}{i\omega\phi_L(\omega)}\frac{\partial \phi_L(\omega)}{\partial n_{it}}\right|_{L=\bar{L}}$$

And multiplying by n_{it} we recover the definition of the tail risk contributions:

$$TRC_i = n_{it}\frac{\partial \bar{L}}{\partial n_{it}} = \left.\frac{n_{it}}{i\omega\phi_L(\omega)}\frac{\partial \phi_L(\omega)}{\partial n_{it}}\right|_{L=\bar{L}}$$

Now using the transformation $i\omega \to \alpha$ we get:

$$TRC_i = n_{it}\frac{\partial \bar{L}}{\partial n_{it}} = \left.\frac{n_{it}}{\alpha MGF_L(\alpha)}\frac{\partial MGF_L(\alpha)}{\partial n_{it}}\right|_{L=\bar{L}}$$

From equations (9) and (15) is easy to see that:

$$n_{it}\frac{\partial MGF_L(\alpha)}{\partial n_{it}} = \alpha\frac{\partial MGF_L(\alpha)}{\partial \alpha}$$

And therefore:

$$TRC_i = n_{it}\frac{\partial \bar{L}}{\partial n_{it}} = \left.\frac{1}{MGF_L(\alpha)}\frac{\partial MGF_L(\alpha)}{\partial \alpha}\right|_{L=\bar{L}}$$

which from equation (10), can be solved for α using

$$EL(\alpha) = \bar{L}.$$

References

[1] Brereton T. J., Kroese D. P. & Chan J. C. *Monte Carlo Methods for Portfolio Credit Risk, in Credit Securitizations and Derivatives: Challenges for the Global Markets*. John Wiley & Sons Ltd, 2013.

[2] Vasicek O. Loan portfolio value. *Risk*, 2002.

[3] Basel Committee on Banking Supervision. International convergence of capital measures and capital. *Bank of Interntional Settlement*, 2005.

[4] Creditrisk+: A credit risk management framework. *Credit Suisse*, 1997.

[5] Martin R., Thompson K. & Browne C. Taking to the saddle. *Risk*, 1997.

[6] Derbali A. & Hallara S. The current models of credit portfolio management: A comparative theoretical analysis. *Int. J. Manag. Bus. Res., 2 (4), 271-292*, 2012.

[7] Merton R. C. On the option pricing of corporate debt: The risk structure of interest rates. *Journal of Finance, Vol. 29, No. 2*, 1974.

[8] Tarashev N. A. An empirical evaluation of structural credit risk models. *Bank of Interntional Settlement*, 2005.

[9] Schönbucher P. J. *Credit Derivatives Pricing Models: Models, Pricing and Implementation*. John Wiley & Sons Ltd, 2003.

[10] Arora N., Bohn J. R. & Zhu F. Reduced form vs. structural models of credit risk: A case study of three models. *Moody's KMV*, 2005.

[11] Wilson T. Credit portfolio risk: Part i. *Risk*, 1997.

[12] Wilson T. Credit portfolio risk: Part ii. *Risk*, 1997.

[13] Hirtle B. & Lehnert A. Supervisory stress testing. *Federal Reserve Bank of New York Staff Reports*, 2014.

[14] Brereton T. J., Kroese D. P. & Chan J. C. *Stochastic Simulation*. Springer. p. 330, 2007.

[15] deRittis C., Licari J. M., Ordóñez-Sanz G. Macroeconomic forecasting and scenario design for IFRS9 and CECL. *The new Impairment Model under IFRS9 and CECL, Risk books*, 2018.

[16] Tasche D. Risk contributions and performance measurement. *Working paper, Technische Universität München*, 2009.

[17] Martin R. Saddlepoint methods in portfolio theory. *arXiv*, 2012.

[18] Rubinstein R. Optimization of computer simulation models with rare events. *European Journal of Operational Research*, 1997.

[19] Martin R. The saddlepoint method and portfolio optionalities. *Risk*, 2006.

[20] Huang X., Oosterlee C. & van der Weide J. Higher order saddlepoint approximations in the vasicek portfolio credit loss model. *Delft University of Technology*, 2007.

[21] Lugannani R. , Rice S. Saddle point approximation for the distribution of the sum of independent random variables. *Advances in Applied Probability*, 1980.

[22] Broda S. A. & Paolella M. S. Expected shortfall for distributions in finance. *Statistical Tools for Finance and Insurance*, 2011.

[23] Sojung K. & Kyoung-Kukm K. Saddlepoint methods for conditional expectations with applications to risk management. *arXiv*, 2015.