NBP Narodowy Bank Polski

NBP Working Paper No. 181

Estimating the risk of joint defaults: an application to central bank collateralized lending operations

Dariusz Gatarek, Juliusz Jabłecki

NBP Working Paper No. 181

Estimating the risk of joint defaults: an application to central bank collateralized lending operations

Dariusz Gatarek, Juliusz Jabłecki

Economic Institute
Warsaw, 2014

Dariusz Gatarek – HVB Unicredit and Systems Research Institute, Polish Academy of Sciences, dariusz.gatarek@unicredit.de
Juliusz Jabłecki – Faculty of Economic Sciences Warsaw University and Economic Institute, Narodowy Bank Polski, juliusz.jablecki@nbp.pl

We thank Ulrich Bindseil and Fernando Monar Lora from the ECB for useful comments. All remaining errors are our own.

Print:
NBP Printshop

Published by:
Narodowy Bank Polski
Education & Publishing Department
ul. Świętokrzyska 11/21
00-919 Warszawa, Poland
phone +48 22 653 23 35
www.nbp.pl

ISSN 2084-624X

Contents

Abstract

Central bank lending to commercial banks is typically collateralized which reduces central bank's credit risk exposure to "double default events" when the counterparty and the issuer of the underlying collateral asset both default in a short period of time. This paper presents a simple model for correlated defaults which are the key drivers of residual credit risk in central bank's repo portfolios. In the model default times of counterparties and collateral issuers are determined by idiosyncratic and systematic factors, whereby a name defaults if it is struck by either factor for the first time. The novelty of our approach lies in representing systematic factors as increasing sequences of random variables. Such a setting allows to build a rich dependence structure that is free of the flaws inherent in the Gaussian copula-based approaches currently regarded as state of the art solutions for central banks.

Keywords: joint defaults, collateralized lending, residual credit risk

JEL classification codes: G12, G13

1 Introduction

That central bank credit should be collateralized is one of the most firmly established principles of central banking, unscathed even by the innovations introduced in response to the recent crisis. In theory one could imagine that the provision of reserves to the banking system might take the form of unsecured lending. However, Bindseil and Papadia (2009) note that there are a number of good reasons for central banks to insist on adequate collateralization, the most important of which is that they have neither the mandate nor the practical expertise to take on and manage credit risk. Obviously perfect protection against credit losses is impossible and in any case not really desired as it would most likely conflict with the pursuit of central bank's monetary policy objectives. Still, the requirement that central bank counterparties submit collateral of sufficient quality and quantity significantly mitigates the extent of credit risk taken by the central bank since, in the event of counterparty default, the central bank can still close the position and return its balance sheet to original size by selling the pledged assets. In such secured transactions – whether they have the form of collateralized lending or, more commonly, repurchase agreements (repos) – the central bank still faces residual credit risk that can materialize only in so called double default events when the counterparty who has submitted the collateral and the collateral issuer default within a short period of time. Thus, the key aspect of managing credit risk in the context of central bank's policy operations consists in modeling joint defaults of counterparties and issuers.

These issues have gained in importance during the recent crisis, following the marked lengthening of central bank balance sheets in most industrialized economies, and the associated rise in risk exposures (see e.g. the BIS study by Archer and Moser-Boehm, 2013).[1] To the extent that a central bank follows the so called "inertia principle," i.e. maintains its risk control framework at least inert in a crisis, it is likely to see its risk-taking increase as a result of an interplay of the following main factors: (i) the increase of probabilities of default of central bank counterparties and issuers of debt instruments used as collateral; (ii) increase of correlation risks between central bank counterparties and collateral credit quality; (iii) the shift of central bank lending towards stressed counterparties who tend to lose market access; (iv) the lengthening of central bank's balance sheet due to a flight of households out of bank deposits into banknotes. There are good reasons for the central bank to provide elastic credit, even though this leads to higher and more concentrated exposures. The original rationale for the inertia principle was presented already by Bagehot (1873) and illustrated more recently by Bindseil (2009). Bindseil and Jablecki (2013) go further to argue

[1] In another testimony to the relevance of these issues Buiter and Rahbari (2012) wonder whether the ECB's loss absorption capacity ("deep pockets of the ECB") is enough to counter potential losses on its collateralized operations. They conclude that it is.

that since in a systemic crisis the probabilities of default of central bank counterparties and collateral issuers are likely to be endogenously related to the extent of central bank's liquidity provision, a more forthcoming liquidity policy of the central bank and a loosening of the risk control framework can, paradoxically, reduce the central bank's ultimate risk exposure. This is not to say that central banks can be indifferent to the extent of credit risk they are taking. To the contrary, as public institutions, central banks have special duties in terms of transparency and accountability, and thus they should be able to accurately measure and account for the risk they are taking, irrespective of whether the act of taking such risk is economically justified or not.

The need to properly measure the residual central bank credit risk exposure stemming from collateralized transactions is relevant not only in a crisis but also in normal times. This is underscored by the fact that – from the central bank's perspective – the actual use of collateral, as well as the resulting concentration of counterparties and issuers in the portfolio cannot be fully anticipated. Hence, the actual risk-taking cannot be known ex ante and the central bank can only ensure that it remains within its accepted risk budget by closely monitoring and stress testing the use of collateral and the level of portfolio concentration and their impact on the relevant risk measures.[2] It is therefore of critical importance that the central bank understands and can properly measure its residual credit risk exposure and the sensitivity of that exposure to concentration and correlation of counterparty and issuer defaults.

Growing awareness of the issues related to portfolio concentration and default correlation has pushed central banks to develop and implement portfolio credit risk models making use of simulation techniques and copula dependence structures. For example, a recent study by a task force of Market Operations Committee of the European System of Central Banks found that "the CreditMetrics™ methodology (...) is used or being tested by most central banks participating in the task force, either directly, using the CreditManager® software, or through in-house systems" (ECB, 2007). The ECB itself as the first central bank has implemented a state-of the-art approach to estimating tail risk measures for a portfolio of collateralized lending operations (Heinle and Koivu, 2009). If collapsed drastically to the single issue of modeling default correlation, such models are based on the industry standard of Gaussian copula (cf. any of the classic textbooks Bluhm, Overbeck, and Wagner, 2002; Schönbucher, 2003 or Lando, 2004 for an overview of the main concepts and methods). Unfortunately,

[2] As explained by ECB risk managers Heinle and Koivu (2009): "Any efficient collateralization framework will provide some discretion to counterparties on what types of collateral to use, and to what extent. This discretion implies that the actual risk taking, for instance driven by concentration risks, cannot be anticipated. The central bank only can ensure that the outcome is actually acceptable by closely monitoring the actual use of the collateralization framework by counterparties, and establishing a sound methodology to measure residual risks. If it is not acceptable, specific changes to the framework are necessary to address the non-anticipated (concentration) risks that arose."

Gaussian copula is not suited very well for analyzing the concentration of defaults in time. As recently argued by Morini (2011) in the context of the subprime crisis, Gaussian copula allows for paradoxical and misleading results such as an inverse relation between correlation and the model probability of loss concentration (see also e.g. O'Kane, 2008).[3] As a remedy to the failings of the Gaussian copula, Morini proposes the Marshall and Olkin (1967) fatal shock model, which although appears to have some theoretical advantages over the latter, is not very useful in practice as a portfolio credit risk model (cf. in particular Andersen and Sidenius, 2004 for a discussion of problems associated with calibrating a Marshall-Olkin model).

Against this background we propose a simple generalization of the Marshall-Olkin approach based on the redefinition of a systematic factor as a sequence of positive random variables, rather than a single random variable. The model, first suggested in Gatarek and Jabłecki (2013), is naturally suited to handle dependencies between default times of the kind encountered e.g. in modeling residual credit risk exposure in a repo portfolio. The modification we prose is inspired by the problem of modeling default correlation in the context of CDOs, as explained in Gatarek (2010). Our approach has three main advantages over the Marshall-Olkin approach. First, on the intuitive level, it is more natural to think that systematic defaults occur in a sequence, not all at once as per Marshall-Olkin. Truly simultaneous defaults can be expected only for entities with strong capital links (parent-subsidiary) or in an "end of the world scenario." Second, on the practical level, our model requires calibration of much fewer parameters than the multivariate Marshall-Olkin model (for d obligors and N systematic factors we get $d \times N$ factor loadings vs. 2^d in multivariate Marshall-Olkin copula) – and in practice a single factor construction allows for sufficient flexibility in modeling default correlation. And finally, on the formal level, the model preserves the stopping time property of defaults lost both in copulas and in the Marshall-Olkin model. Thus, modeling credit risk can be handled in the familiar mathematical framework of martingale methods, which is consistent with the approaches used for other asset classes. Given the systemic importance of central banks in many industrialized countries and the recent lengthening of their balance sheets drawing attention to central banks' risk exposures, we show how the model can be applied to managing residual credit risk in a central bank's repo portfolio. Obviously, however, the issue of estimating joint default probabilities is pervasive, especially in view of the growing importance of repo markets globally. Hence, the ideas presented are

[3]Interestingly, Morini (2011) believes that the much discussed lack of tail dependence in the Gaussian copula was a lesser problem in the crisis: "the problems in the credit derivatives market and structured finance that led to the credit crunch did not come from this kind of loss concentration [in a single portfolio]. They came from actual or feared losses that were not high in a single portfolio (...) but were concentrated in a short period of time, creating panic and liquidity difficulties" (p. 139). On the problem of tail dependence see Embrechts, Lindskog, and McNeil (2003).

valid directly for any institution with a large repo portfolio, and in fact can be used also in the context of pricing and risk-managing more complex financial instruments, such as credit default swaps, default swaptions, and CDOs (cf. Gatarek, 2010; Gatarek and Jabłecki, 2013).

The rest of the paper proceeds as follows. In Section 2 we briefly describe current best practice solutions for risk-managing a central bank repo portfolio and discuss their main shortcomings. Section 3 presents our model of default correlation and shows how it can be applied to central bank risk management problem. Section 4 concludes.

2 Risk measurement in central bank's collateralized lending

When the central bank conducts its monetary operations through repurchase agreements or collateralized lending (which are almost indistinguishable from an economic point of view) its credit risk is limited strictly to the so called double default case when both the counterparty and the collateral issuer default in a short period of time.[4] The time frame is crucial and stems from a practical consideration that in most cases it would not be advisable to immediately sell the collateral captured after counterparty default, as such a sale could adversely impact market prices (which might already be affected by the counterparty default event itself), leading to potential losses for the central bank. Thus, strictly speaking, in the context of central bank repo operations the risk of joint defaults actually means the risk that *first* the counterparty defaults (before the maturity of the repo) and *then*, during the realization period needed to orderly liquidate the pledged collateral, the issuer defaults as well. Such probability of joint defaults depends on the following three key factors:

- the counterparty's probability of default (PD);
- the collateral issuer's PD; and
- the default correlation between the counterparty and the collateral issuer.

Hence, measuring residual credit risk exposure entails modeling the PDs of counterparties and issuers as well as their correlation. Given potentially huge number of counterparties and collateral issuers this is an extremely difficult problem in central banks, not only conceptually but also in terms of practical implementation (IT infrastructure, systems, software etc.). In what follows we focus only on the theoretical side of the problem and to motivate further discussion, we begin by presenting what seems to be the "best practice" approach to measuring residual credit risk in central banks[5], namely the Gaussian copula model, a variant of which was implemented e.g. by the ECB in 2006 (Heinle and Koivu, 2009)[6].

[4]Strictly speaking, residual credit risk is only one source of risk to which a central bank is exposed in a collateralized operation. In theory, since the central bank is not threatened by illiquidity, it could – in the event of counterparty default – hold the collateral to maturity, in which case it would only have to care about double default. In practice, however, central banks tend to liquidate collateral before maturity – probably considering this to be optimal – and thus also care about mark-to-market losses on a given asset in liquidation before maturity.

[5]Unlike private agents, central banks are not bound by any international risk management standards akin to the Basle accords, and are also less transparent about the models and techniques applied. Thus, strictly speaking, there is no such thing as a publicly available set of best practices adopted by the central banking community (a fact recently brought up by Bank of England's executive director for markets Paul Fisher).

[6]Although the ECB methodology has been updated since the publication of the Heinle and Koivu (2009) paper, the key features – especially the Gaussian copula – remain the same.

2.1 Copula-based approach to residual risk measurement

The general framework for estimating the joint default probability for a portfolio of credits (the "Black Scholes" of the correlation market) is the so called Gaussian latent factor model, inspired by the structural approach to credit risk proposed by Merton (1974), and used widely in different forms in commercial and regulatory applications.[7] In the model each obligor i is assigned a standard normal variable A_i and a time-dependent default threshold $z_i(T)$. It is assumed that default occurs before time T if the variable A_i (which itself is not observable and has no dynamics) finds itslef below the threshold $z_i(T)$. Formally,

$$\mathbb{P}(\tau_i \leq T) = \mathbb{P}(A_i \leq z_i(T)) = PD(T), \tag{1}$$

where τ_i is the default time of credit i and $PD(T)$ is the T-year probability of default. Equation (1) shows how the model can be calibrated to an obligor's credit curve. For example, given the annual PD for name i, the value of $z_i(1)$ can be easily determined, as $z_i(1) = \Phi^{-1}(PD(1))$ ($\Phi^{-1}(\cdot)$ being the inverse of the standard normal CDF), and $z_i(1)$ in turn can then be used in simulations. Typically counterparty and issuer's PDs are derived from their credit ratings, as provided either by external rating agencies, counterparties' own internal systems or in-house credit assessment models.[8]

Default correlation is introduced in the model through linear correlation of the variables A_i, so called asset correlation. Specifically, each A_i is decomposed into a systematic component Z and an idiosyncratic component Y:

$$A_i = w_i Z + \sqrt{1 - w_i^2} Y_i \tag{2}$$

where Z and Y_i are standard normal with $\text{cov}(Y_i, Y_j) = 0$ and $\text{cov}(Y_i, Z) = 0$. The asset correlation between A_i and A_j is then given by $\rho_{i,j}^{asset} = w_i w_j$.[9] Since with the size of a central bank portfolio it would be practically impossible to guarantee the technical restriction that the correlation matrix composed of all pairs $w_i w_j$ is positive definite, a common practice – at least before the crisis – was to use a fixed correlation coefficient[10]. The final key building

[7] The earliest published version of the model is a 1987 memo by Vasicek, subsequently published as Vasicek (2002). In the 1990s the model became popularized by Goupton, Finger, and Bhatia (1997) who based on it the CreditMetrics risk management framework, underlying also the ECB approach. Finally, a version of the model – called Asymptotic Single Risk Factor (Gordy, 2003) – was used for the calculation of regulatory capital under the Basel IRB framework.

[8] Sometimes central banks specify also a minimum rating/PD threshold for counterparties. For example, to limit the extent of risk taking before the euro area sovereign debt crisis, the Eurosystem had required that collateral issuer's credit must be of at least single-A quality (corresponding an annual PD of 10 bp) but the threshold has since been suspended to counteract potential collateral scarcity.

[9] Note that asset correlation is a different concept from default time correlation $\rho(\tau_i, \tau_j)$ and default indicator correlation $\rho(\mathbb{1}_{\{i\}}\mathbb{1}_{\{j\}})$. In general, conditional on default probabilities, default correlation is an increasing function of asset correlation (Hanson, Pesaran, and Schuermann, 2008).

[10] For example, as of end 2006 (the most recent date for which these calculations are publicly available),

block of the model is specification of the nature of the multivariate dependence between marginal probabilities of default. A common approach – followed i.a. by the ECB – was to use the Gaussian copula with ρ^{asset} as the correlation parameter in the copula function.

In such a setup the central bank loss function is determined by simulation, hence no additional distributional assumptions are needed. Stripped to its essentials, the simulation involves drawing pseudo-random realizations of the variable A_i and determining the number of counterparty and issuer defaults across the portfolio. Then, on the basis of the assumed recovery rates, central bank losses in each draw are calculated, allowing ultimately for the calculation of the respective tail risk measures (VaR, expected shortfall, etc.). For example, in the ECB's approach, risk is estimated over an annual horizon, so for each counterparty i it is verified whether $A_i \leq z_i(1)$ and for each issuer j whether $A_j \leq z_j(1/LH_j)$. Here, LH_j stands for the liquidation horizon of the least liquid instrument from issuer j. The scaling of issuers' PDs down to the liquidation horizon is a way to reflect in the model the temporal dimension of double defaults. This conveys the idea that while the counterparty can default at any time within the chosen horizon of 1 year, for the central bank to suffer a loss, the issuer has to default during the normally much shorter time assumed to be needed for the liquidation of the pledged assets.[11]

2.2 Some problems with copula-based estimation of double default events

Since a model is by definition only an imperfect representation of reality, no risk model is likely to be "the right one." However, as argued by Derman (2001) in one of the first essays on model risk, a useful model should provide at least a realistic or plausible description of the factors that determine its outcome. In the case of a model for estimating residual credit

the Eurosystem repo portfolio consisted of some 18,000 counterparty-issuer pairs. Thus, Heinle and Koivu (2009) report the use of a fixed correlation coefficient of 24 per cent, even though – writing in 2009 – seem somewhat dissatisfied with the low correlation input. In practical terms, what drives default correlation and the probability of joint defaults in the Eurosystem portfolio is its concentration on the level of counterparties and collateral. Such concentration can materialize along several dimensions, including concentration on the level of issuers, asset categories, industries, countries and banking groups. The Eurosystem has taken some precautions to mitigate concentration risk by prohibiting counterparties to submit collateral issued or guaranteed by themselves or entities with which they have close links. However, since the outbreak of the crisis these restrictions have been progressively relaxed to allow for own-use of covered bonds and asset backed securities as collateral. Similarly, the provisions against close links were waived for general government bonds. While these changes may have been necessary to fend off collateral scarcity and extend the ability of banks to take recourse to central bank credit operations, they may have increased residual credit risk in the Eurosystem's portfolio in a way that is not entirely clear and transparent.

[11] An important additional assumption is that collateral quality and maturity (and even prices) are constant over time, until the time of default of the counterparty, as counterparties rebalance their collateral portfolio. The ECB model assumes counterparties draw as much liquidity as they can before default so that current collateral figures determine the exposure at default.

risk of a repo portfolio, this could be interpreted as a requirement that there be a clear and intuitive relationship between the input parameters such as PDs and asset correlations and the financial risk measured. Unfortunately, the Gaussian copula does not score very well on this count. In what follows we describe the three main shortcomings of the copula-based model for estimating the joint default probability, before providing what we believe is a more accurate approach in the following section. Some of these problems have recently been identified and discussed at length by Morini (2011) in the context of credit derivatives, but as we argue below, they remain valid for central banks as risk managers of collateralized lending portfolios.

Problem 1. No meaningful default clustering under perfect asset correlation

This problem is well known to practitioners and relates to the inability of the Gaussian latent factor model to produce concentration of defaults in time if obligors in the portfolio have different conditional default probabilities (see e.g. O'Kane 2008 in the context of valuing credit derivatives). To understand this phenomenon better, consider two credits with ratings AAA and BBB- and corresponding annual probabilities of default $PD_1(1) = 0.01\%$ and $PD_2(1) = 0.40\%$. Assuming conditional default probabilities of the two obligors are constant, we can easily obtain the respective hazard rates $\lambda_1 \approx 0.0001$ and $\lambda_2 \approx 0.0040$ and then also the two credit curves $z_1(t) = \Phi^{-1}(1 - \exp(-\lambda_1 t))$ and $z_2(t) = \Phi^{-1}(1 - \exp(-\lambda_2 t))$. In the Gaussian latent factor model default times are simply the values of t for which $A_1 = z_1(t)$ and $A_2 = z_2(t)$. Hence, realizations of correlated default times of both credits can be simulated using the formulas:

$$\begin{aligned} -\frac{\ln\left(1 - \Phi(A_1)\right)}{\lambda_1} &= \tau_1 \\ -\frac{\ln\left(1 - \Phi(A_2)\right)}{\lambda_2} &= \tau_2 \end{aligned} \tag{3}$$

with A_1 and A_2 being correlated standard normal random variables. Perfect asset correlation of the two names means of course that $A_1 = A_2$ which implies a deterministic and fixed relationship between the two defaults: $\tau_1 = \tau_2(\lambda_2/\lambda_1)$, i.e. in this case $\tau_1 = 40\tau_2$. In other words, despite imposing perfect asset correlation on the two obligors – which is the key measure of dependence in the Gaussian latent factor model – we were unable to produce in the model meaningful clustering of defaults in time, which would additionally require that $\lambda_1 \approx \lambda_2$ and $PD_1(1) \approx PD_2(1)$. Whether this is a problem in practice depends on the specific application of the model and its ultimate use. However, it seems that failure to produce defaults near in time for credits with different probabilities of default is a shortcoming not to be taken lightly by central banks. After all, the key factors in estimating residual credit risk

of a repo portfolio are the probability that defaults of the counterparty and the issuer are near in time – e.g. within several weeks needed to sell the collateral – and the dependence of that probability on portfolio concentration and correlation. If a model cannot produce defaults that are close in time for highly dependent obligors than it may lead to erroneous conclusions about the central bank's credit risk exposure.

Problem 2. Unstable relation between joint default probability and asset correlation

The second problem relates to the fact that the probability of joint defaults can be either an increasing or decreasing function of asset correlation depending on the time frame under consideration and obligors' PDs. Under normal circumstances, the central bank is only interested in a spot-starting probability of joint defaults, i.e. the probability of double default events over the next year starting from now. However, in some cases it might be necessary to calculate the relevant risk measures over a future horizon, e.g. over a year starting from some specified future date. This might be relevant in stress testing exercises or when the central bank is particularly concerned about the materialization of adverse scenarios over a certain period of time in the future.[12] As a simple illustration, suppose that counterparty and issuer hazard rates are $\lambda_1 = 0.01$ and $\lambda_2 = 0.20$ respectively and consider the probability that both default during a 3-year horizon starting now, $\mathbb{P}(\tau_1 \leq 3, \tau_2 \leq 3)$, and two years from now, $\mathbb{P}(2 \leq \tau_1 \leq 5, 2 \leq \tau_2 \leq 5)$ (for simplicity we abstract from the liquidation horizon). Figure 1 shows that for the spot-evaluated case, the probability curve has an increasing shape, in line with expectations. However, the forward-starting probability of joint defaults is monotonously decreasing in asset correlation. Indeed, we already know from earlier discussion (see eq. (3)) that with $\lambda_1 = 0.01$, $\lambda_2 = 0.20$ and perfect asset correlation, default times will be deterministically related, $\tau_1 = 20\tau_2$, and so intervals $2 \leq 20\tau_2 \leq 5$ and $2 \leq \tau_2 \leq 5$ never intersect. This reasoning also shows that by changing λ_1 and λ_2 we can easily guarantee that the intervals for default times do intersect and thus the probability of joint defaults is nonzero under perfect asset correlation. For example, setting $\lambda_1 = \lambda_2 = 0.01$ we obtain exponentially increasing probability of joint defaults as a function of correlation (Figure 1, black dashed line). Although it is the spot-starting probability of joint defaults that matters most in practical applications, the instability of the relation between the key model parameter – asset correlation – and financial risk detected in forward-

[12]Suppose, for example, that the ECB lends money to a counterparty in a 3Y LTRO collateralized by a sovereign bond from the eurozone periphery. Suppose also that the ECB fears that the fiscal situation of the sovereign issuer is likely to deteriorate over time while in addition new strict capital rules, penalizing banks for holding government bonds, will become binding in a year from now. In such a case the ECB might want to estimate the probability of joint default at a forward date, e.g. the probability that both the counterparty and the issuer default between 1 and 3 years from now.

Figure 1: Spot- and forward-starting probability of joint defaults as a function of asset correlation

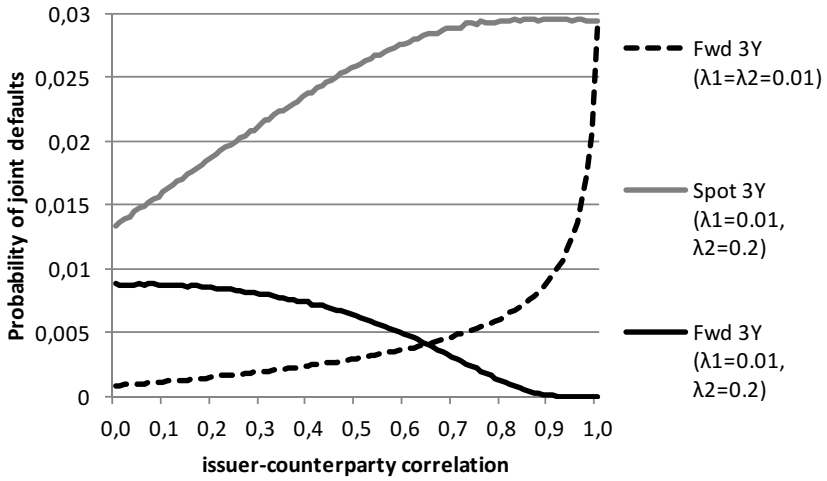

"Spot 3Y" is the probability of joint defaults avaluated on a spot basis: $\mathbb{P}(\tau_1 \leq 3, \tau_2 \leq 3)$; "Fwd" is the probability that both names default jointly over a 3-year horizon in the future $\mathbb{P}(2 \leq \tau_1 \leq 5, 2 \leq \tau_2 \leq 5)$.

starting calculations is also very important. At minimum, it indicates that the model used to estimate risk can in some cases produce an unrealistic and implausible relation between input parameters and output, and hence is associated with nontrivial model risk.

Problem 3. Probability of defaults near in time is non-monotonic with respect to asset correlation

The third key problem with using Gaussian copula to model residual credit risk is that the probability of defaults happening near in time and one after another – as in a repo double default scenario – can for given conditional hazard rates be a non-monotonic, hump shaped, function of asset correlation. Consider the follwoing example. The central bank provides cash to a counterparty for the maturity of 12 months. The counterparty with annual PD equal 0.05 (equivalent to a constant hazard rate $\lambda_1 = -\ln(0.95)$) submits as collateral securities issued by an obligor with annual PD 0.2 (think of a bank with decent credit quality submitting poor quality ABS). Assume further that the collateral has a liquidation horizon of 1 month. In such a simplified yet realistic example the central bank could suffer losses if the following two events materialize: (i) the counterparty defaults at time τ_1 within the next 12 months; (ii) the issuer defaults after τ_1, but not later than $\tau_1 + 1\text{M}$. Formally, the estimation of residual credit risk for this portfolio entails quantifying the following probability:

$$\mathbb{P}_{\rho_{1,2}}(\tau_1 \leq 1 \wedge \tau_1 \leq \tau_2 \wedge (\tau_2 - \tau_1) \leq 0.08) \tag{4}$$

Figure 2: Probability that the defaults of the counterparty and issuer are within 1M as a function of Gaussian copula correlation parameter.

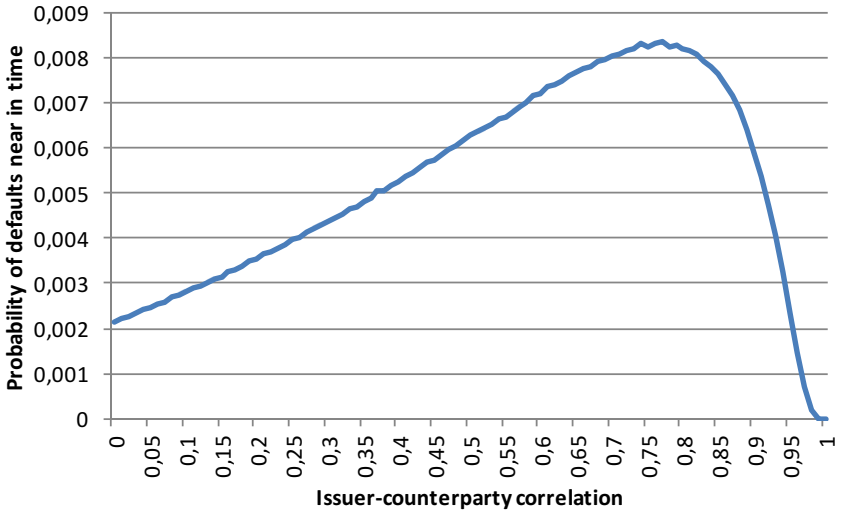

Note: $\lambda_1 = 0.05$, $\lambda_2 = 0.2$; probability is evaluated based on formula (4).

as a function of issuer-counterparty correlation $\rho_{1,2}$. As shown in Figure 2, the probability of joint defaults – and thus central bank's residual credit risk – first increases as a function of $\rho_{1,2}$, assumes its maximum value of 84 bp for $\rho_{1,2} = 0.77$ and then falls, reaching a minimum for $\rho_{1,2} = 1$. The problem with such a non-monotonic pattern is not only that it is difficult to understand from an economic point of view, but also that it makes it in general impossible to know a priori whether – with a given set of PDs – an increase in correlation will increase or decrease central bank risk exposure. To determine that, one would need a full simulation of the relevant risk measures for different correlation assumptions, which would obviously be very computationally intensive.

This unwelcome result of non-monotonic relation between probability of joint defaults and correlation can be avoided if the timing of defaults is reflected by the scaling down of issuers' PDs to the liquidation period while considering counterparties' PDs in annual terms (instead of modeling directly the clustering of defaults in time as in (4)). In other words, default times are evaluated on a spot-starting basis, albeit with a different horizon: annual for counterparties and weekly, bi-weekly, monthly etc. (the equivalent of liquidation periods) for issuers. In practice, a joint default event that is critical for a collateralized loan has a very clear forward-starting nature, as the default of the issuer is only dangerous if it occurs *after* the default of the counterparty *and* in the period needed to sell the collateral. It is natural for models to simplify an otherwise complex reality, but in case of risk management models such simplifications should in principle err on the side of caution, i.e. overstate, not understate the risks taken. Moreover, even if the linear scaling of annual PDs and default thresholds is a conservative approach with respect to "true" short-term PDs, it is not

consistently conservative in modeling the clustering of defaults in time. To see this, consider once again the previous example. In the benchmark case the hazard rate of the counterparty is as before $\lambda_1 = 0.05$ and that of the issuer $\lambda_2 = 0.2$. But we consider also a reversed situation where the counterparty has lower credit quality than the collateral it pledges, i.e. $PD_1(1) = 0.2$ and $PD_2(1) = 0.05$. The remaining assumptions on the maturity of the collateralized loan and the liquidation period are unchanged. In both cases we estimate the probability of joint defaults using two methods – first as in (4), and then using the approach of scaling the default probabilities down to the liquidation horizon:

$$\mathbb{P}^{scaled}_{\rho_{1,2}}\left(A_1 \leq \Phi^{-1}(PD_1(1))), A_2 \leq \Phi^{-1}(PD_2(1) \times LH_2)\right) \tag{5}$$

where A_1, A_2 are correlated standard normal variables and LH_2 is the liquidation horizon corresponding to 1 month (22 business days). Figure 3 shows that in the benchmark case the model based on scaling the default threshold down to the liquidation horizon gives indeed a more conservative assessment of the risk that both the counterparty and issuer default at the same time. However, when the counterparty is less risky than the issuer, the method of directly modeling default clustering in time yields a conservative assessments for higher values of counterparty-issuer correlation. Unfortunately, even in the highly simplified case of just one issuer-counterparty pair it is impossible to know ex ante which method of estimating double defaults yields a more conservative assessment. Although in some cases it would be more prudent to estimate residual risk based on (4) rather than (5), the former approach is, as we have already explained, non-monotonic with respect to correlation and hence also not very practical.

Figure 3: Probabilities of joint default as functions of copula correlation parameter: a comparison of calculation methods

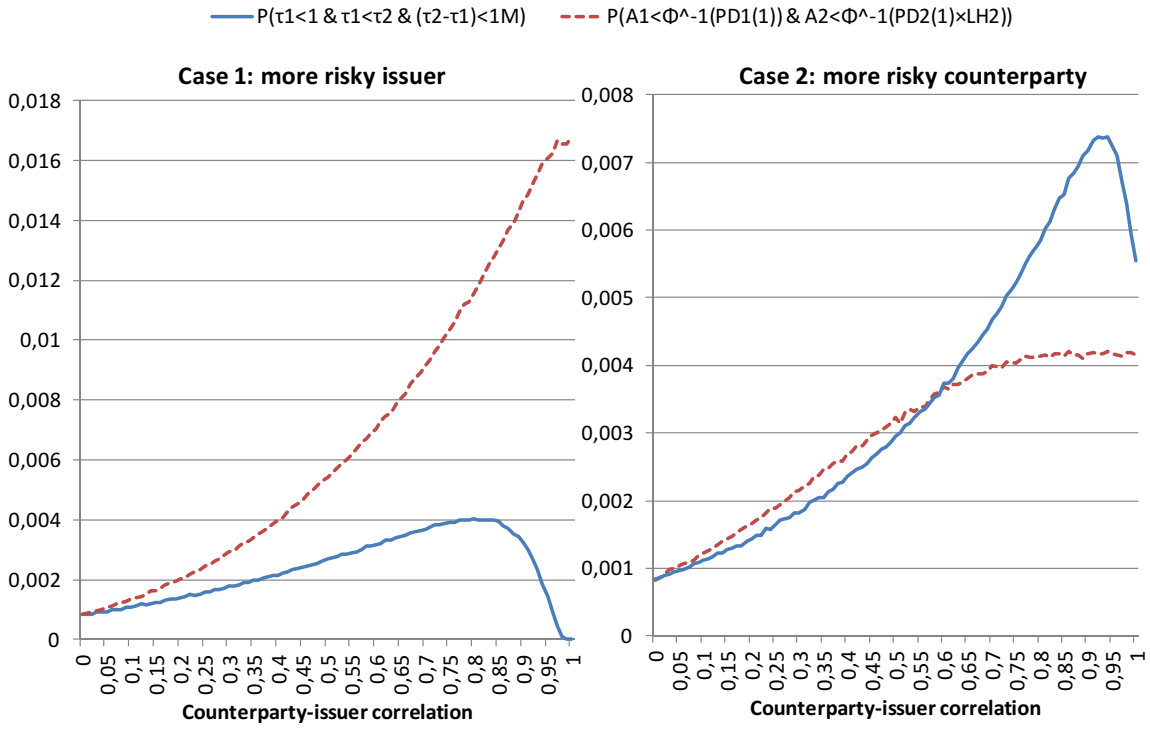

Note: Parameterization for Case 1 assumes $PD_1(1) = 0.05$ and $PD_2(1) = 0.2$ and vice versa for Case 2.

3 A simple factor model of joint defaults

In this section we propose our own model for correlated defaults that is free from the shortcomings inherent in the use of Gaussian copula approach. To facilitate exposition, we introduce the model in the simplest single factor and constant hazard rate setting and discuss the similarities of our approach to the Marshall-Olkin model as advocated recently e.g. by Morini (2011), and earlier also by Elouerkhaoui (2006) and Giesecke (2003). We also discuss a method of calibrating the model to empirical data and finally also explain how the model can be extended to a multifactor stochastic hazard rate setting.

3.1 Single factor case

Consider d obligors with default times $\tau_1, ..., \tau_d$. Assume for now (we relax this assumption in section 3.3) that default times are exponentially distributed with parameters $\lambda_1, ..., \lambda_d$, which admit natural interpretation as hazard rates or conditional default probabilities. As before, we introduce dependence between default times by stating that each default can result from the materialization of either an idiosyncratic factor or a systematic factor – whichever hits sooner. Being hit by either factor has the mathematical interpretation of the first jump of a specific Poisson process. Hence, for each obligor i the time until the arrival of the idiosyncratic factor is represented simply by an exponential variable Y_i with parameter λ_i^{idio}. Where we differ from previous approaches is that, unlike in (2) where the systematic factor was a single random variable, we think of a systematic factor as an increasing sequence of exponential variables $Z_1 \leq ... \leq Z_d$ with parameters $\lambda_1^{sys}, ..., \lambda_d^{sys}$. This should be intuitive as the most natural interpretation of dependence for random variables expressing time is an ordering relation.[13] Under such assumptions, individual obligors' default times can be represented as:

$$\tau_i = \min\{Y_i, Z_i\}, \tag{6}$$

where $Y_1, ..., Y_d$ and $Z_1, ..., Z_d$ are independent exponential variables. Obviously, default times of all obligors, τ_i, are also exponentially distributed with parameters $\lambda_i = \lambda_i^{idio} + \lambda_i^{sys}$ and survival probabilities

$$\mathbb{P}(\tau_i > T) = e^{-\lambda_i T}. \tag{7}$$

Consider now two useful properties of the proposed model.

[13] We show below that defining a systematic factor as an increasing family of random variables rather than a single random variable allows to preserve a useful formal property that default times are also stopping times. Such redefinition also formalizes an important practical intuition that a systematic factor need not cause the default of the whole economy at once. It is much more natural to expect that some credits will default sooner and some later, depending on their sensitivity to the given systematic factor.

Remark 1. Idiosyncratic defaults tend to be more frequent than systematic defaults.

Denote the first default of an idiosyncratic type by $Y_{first} = \min\{Y_i : 1 \leq i \leq d\}$ and the first default of a systematic type by Z_1. Then

$$\mathbb{P}\left(Y_{first} \geq t\right) = \mathbb{P}\left(\min\{Y_i : 1 \leq i \leq d\} \geq t\right) = \exp\left(-t\sum_{i=1}^{d}\lambda_i^{idio}\right) \tag{8}$$

Since systematic defaults are by definition ordered, we know which obligor defaults first (although we do not now exactly when):

$$\mathbb{P}\left(Z_1 \geq t\right) = \exp\left(-t\lambda_1^{sys}\right). \tag{9}$$

Hence,

$$\frac{\mathbb{P}(t \leq Z_t < t + dt | Z_1 > t)}{\mathbb{P}(t \leq Y_{first} < t + dt | Y_{first} > t)} = \frac{\lambda_1^{sys}}{\sum_{i=1}^{d}\lambda_i^{idio}} \tag{10}$$

and – under normal conditions – the chance that the first default is of an idiosyncratic character should be considerably greater.

Remark 2. The definition of a systematic factor as an increasing sequence of random variables allows to capture the phenomenon of default clustering.

In fact in our model only systematic defaults can be multiple. To see this define the following point processes counting the defaults triggered by the respective factors: $N(t) = \text{card}(\tau_i < t)$, $M(t) = \text{card}(Z_i < t)$ and $N_j(t) = \text{card}(i > j : Y_i < t)$. Using the property that Z_i are ordered we easily get

$$\mathbb{P}(M(t) = j) = \mathbb{P}(Z_j < t < Z_{j+1}) = \left(1 - e^{-t\lambda_j}\right) - \left(1 - e^{-t\lambda_{j+1}}\right) = e^{-t\lambda_{j+1}} - e^{-t\lambda_j}. \tag{11}$$

Consequently,

$$\mathbb{P}(N(t) = m) = \sum_{j=0}^{m}\mathbb{P}(M(t) = j)\mathbb{P}(N_j(t) = m - j) = \sum_{j=0}^{m}\left(e^{-t\lambda_{j+1}} - e^{-t\lambda_j}\right)\mathbb{P}(N_j(t) = m - j). \tag{12}$$

If we assume that defaults of individual obligors can repeat themselves,[14] then the point process $N_j(t)$ is a Poisson process with intensity $\sum_{i=j+1}^{d}\lambda_i^{idio}$. Hence,

[14]Since in practice defaults of individual obligors cannot repeat themselves – which is equivalent to randomly drawing defaulting names *without* replacement – this assumption slightly overstates the total number of defaults. A more rigorous approach to estimating such probabilities is based on the Bernoulli triangle as suggested originally by Hull and White (2004).

Figure 4: Default distribution for a portfolio of 100 credits assuming: (i) idiosyncratic; and (ii) systematic defaults.

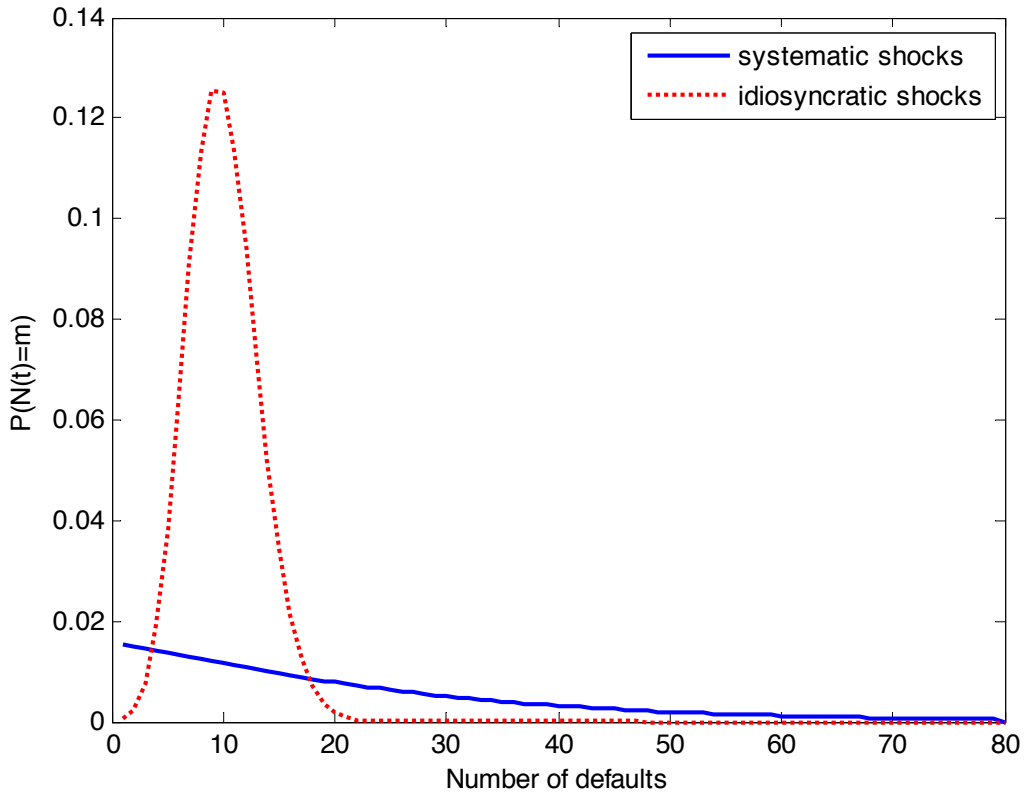

Note: The same set of 100 hazard rates is used to generate both plots, but in one case the shocks are treated as idiosyncratic and in the other as systematic.

$$\mathbb{P}(N_j(t) = m) \approx \frac{1}{m!} \left(t \sum_{i=j+1}^{d} \lambda_i^{idio} \right)^m \exp\left(-t \sum_{i=j+1}^{d} \lambda_i^{idio} \right). \tag{13}$$

Figure 4 shows how the ordering of random variables comprising a systematic factor affects the aggregate default distribution. The same set of 100 hazard rates (ranging from 10% to 0.6%) and equations (11)-(13) are used to produce default distributions in one case assuming default times are idiosyncratic and in the other that they are ordered. The ordering of default times clearly increases the chance of having multiple defaults within a given time horizon.

Example 1. Marshall-Olkin fatal shock model

Our model can be considered a generalization of the approach originated by Marshall and Olkin (1967) and initially used in reliability theory to model the failure of multi-component systems. In the model each "system component" (i.e. each obligor) is subject to a shock Y that is fatal to itself (i.e. idiosyncratic) as well as a common shock Z affecting also all other components (i.e. systematic). Thus, similarly as above, the time until failure (default) of each obligor is given by

$$\tau_i = \min\{Y_i, Z\}. \tag{14}$$

Figure 5 presents a stylized decomposition of hazard rates of 10 obligors into idiosyncratic and systematic components. Because Marshall and Olkin devised their model largely with engineering applications in mind, a common shock in their setup is represented by a single random variable. Such an approach seems natural in situations when common cause failures occur at the same time (or very near in time) – as e.g. due to some mechanical or electrical malfunction.[15] There is, however, much less automatism in economic phenomena which calls for a more nuanced treatment of a systematic shock. Moreover, the Marshall-Olkin model has a number of shortcomings that limit its practical applications in modeling portfolio credit riks. First, note that by construction, the systematic part of the credit spread cannot exceed the lowest spread itself, i.e.

[15]For a typical engineering application of the Marshall-Olkin approach see e.g. Vesely (1977) who uses it to test the reliability of a nuclear reactor emergency shutdown system, so called scram. Scram reliability depends on timely insertion of control rods into the core of a reactor which increases neutron absorption and decreases reactor power. Testing scram reliability consists i.a. in evaluating the risk of a failure to insert control rods due to some common cause, typically mechanical in nature, such as filter plugging, rod drive blockage etc. For example, a Dresden 2 reactor malfunction, revealed in June 1975, had the following description "rods inserted [too] slowly due to high regulated pressure in the scram valve air header." Since mechanical failures tend to cause the malfunction of some or all rods virtually simultaneously, it is natural to model the arrival time of such a common shock using a single random variable.

Figure 5: Marshall-Olkin model: a decomposition of hazard rates into idiosyncratic and systematic components

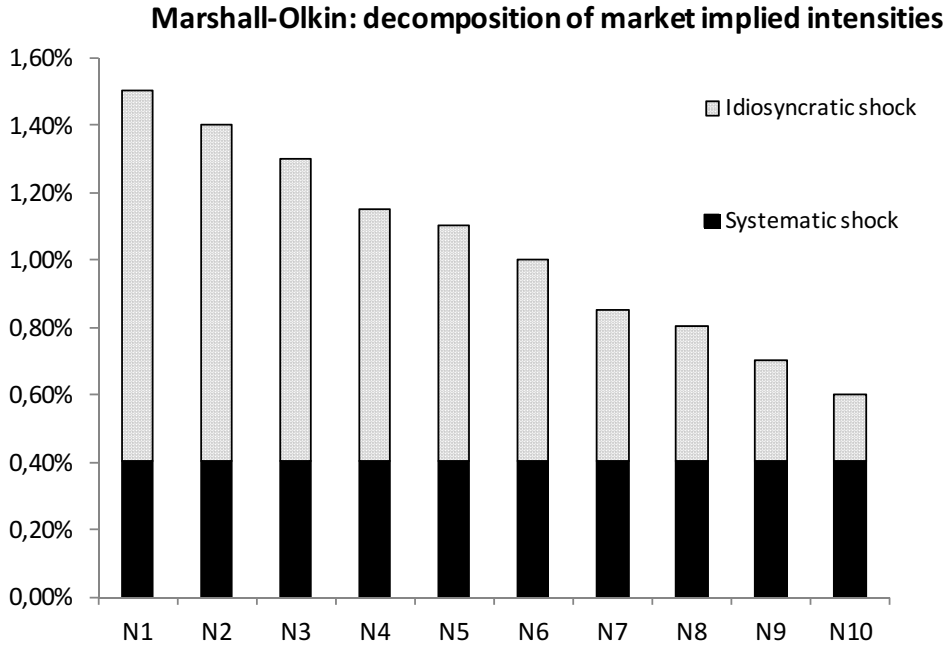

Note: Here $\lambda^{sys} = 0.40\%$ while $\lambda_i^{idio} = \lambda_i - \lambda^{sys}$.

$$\forall i \ \lambda^{sys} \leq \min\{\lambda_i : i = 1, ...d\} \tag{15}$$

Equally important, in the Marshall-Olkin model more risky obligors must always be more idiosyncratic than the less risky ones, which does not seem to be an accurate description of reality and might prove to be a limitation in practical applications. Yet another problem relates to the fact that, except in the trivial case when all the spreads are equal, it is in general impossible to have a common meaningful measure of credit correlation for the whole portfolio – a useful property that has cemented the widespread use of the Gaussian copula correlation parameter. Finally, the multivariate extension of the one factor case is computationally challenging, since if each nonempty subset of obligors $\{1, ..., d\}$ is assigned a shock which is fatal to all names in that subset, we get effectively $2^d - 1$ factor loadings – considerably more than the $N \times d$ parameters needed to calibrate an N-factor Gaussian copula model.[16]

[16]Thus, Andersen and Sidenius (2004) note: "(...) exact parameterization of the [multivariate] Marshall-Olkin copula is a rather formidable problem, given the abstract nature (and sheer number) of its parameters. We note that to make the model consistent across different CDOs. one really must calibrate a single matrix and a single set of intensities for all credits in the universe of traded credit default swaps. Such a calibration would likely be difficult to make robust, and strong assumptions will be needed to make it feasible. We also point out that even if a calibration method could be constructed. the model remains quite unwieldy and involve a number of non-trivial operational and computational issues...".

3.2 Calibration and application to residual credit risk estimation

Calibration of the proposed model entails distributing the market-implied (or otherwise obtained) hazard rates of individual obligors across the idiosyncratic and systematic shock components. The first step in the process is the construction of a systematic factor which can be done e.g. in the following way. Start by sorting obligors in the order reflecting their overall exposure to the systematic factor in question. Suppose, for example, that credit "1" is the most sensitive to the systematic factor, credit "2" is less so but still highly sensitive, and so on, while credit d is the least exposed. Hence, the systematic factor should first trigger the default of name "1", then "2" etc. before ultimately hitting d. To reflect this, assign to each name i a Poisson process $\widetilde{Z}_i$, with intensity λ_i^{sys}, whose arrival triggers the default of credit i, but also – due to the ordering relation – also the default of all the more systematically risky names $i-1$, $i-2$, ..., 1 (we assume that Poisson processes $\widetilde{Z}_i$ are independent). Thus, the systematic intensity of each obligor i will be a sum of its own intensity λ_i^{sys} and the intensities of the Poisson processes triggering defaults of more senior names, i.e. $\sum_{j=i}^{d} \lambda_j^{sys}$. This can be formalized by setting

$$Z_i = \min\left\{\widetilde{Z}_i : i \geq j\right\}, \tag{16}$$

where Z_i is the Poisson process representing the total systematic exposure of obligor i. Note that $Z_i \leq Z_{i+1}$ for $i = 1, ..., d-1$, so indeed the family $Z_1, ..., Z_d$ is a systematic factor. Since each obligor is also affected by an idiosyncratic shock Y_i with intensity λ_i^{idio}, default times τ_i are exponentially distributed with parameters $\lambda_i = \lambda_i^{idio} + \sum_{j=i}^{d} \lambda_j^{sys}$ and survival probabilities

$$\begin{aligned}\mathbb{P}(\tau_i > T) &= \mathbb{P}\left(\min\left\{Y_i, \min\left\{\widetilde{Z}_i, \widetilde{Z}_{i+1}, ..., \widetilde{Z}_d\right\}\right\} > T\right) = \mathbb{P}\left(\min\left\{Y_i, \widetilde{Z}_i, \widetilde{Z}_{i+1}, ..., \widetilde{Z}_d\right\} > T\right) = \\ &= \mathbb{P}(Y_i > T)\mathbb{P}(\widetilde{Z}_i > T)\mathbb{P}(\widetilde{Z}_{i+1} > T)...\mathbb{P}(\widetilde{Z}_d > T) = e^{-\lambda_i T}.\end{aligned} \tag{17}$$

The second step of the calibration process consists in allocating the total hazard rate of each obligor (obtained e.g. from rating agencies, market spreads or in house credit models) to the idiosyncratic and the systematic component. The approach typically followed with respect to the Marshall-Olkin approach is to assume that factor loadings are set as a fixed percentage of the total intensity of a given name (see e.g. Lindskog and McNeil, 2003; Duffie and Pan, 2001). We modify this convention slightly and introduce a parameter $\rho \in [0, 1]$ which determines the extent to which default times in the economy are independent or ordered – i.e. triggered by idiosyncratic or systematic factors. Specifically, given a portfolio of d obligors with conditional default probabilities $\lambda_1, ..., \lambda_d$ and a parameter ρ, we set the systematic exposure of each credit as follows:

$$
\begin{aligned}
\lambda_d^{sys} :=& \rho\lambda_d \\
\lambda_{d-1}^{sys} :=& \rho\lambda_{d-1} - \lambda_d^{sys} = \rho(\lambda_{d-1} - \lambda_d) \\
\lambda_{d-2}^{sys} :=& \rho\lambda_{d-2} - \lambda_{d-1}^{sys} - \lambda_d^{sys} = \rho(\lambda_{d-2} - \lambda_{d-1}) \\
&\vdots \quad \vdots \\
\lambda_1^{sys} :=& \rho\lambda_1 - \lambda_2^{sys} - ... - \lambda_d^{sys} = \rho(\lambda_1 - \lambda_2)
\end{aligned} \tag{18}
$$

The set of identities (18) determines the individual systematic exposure of each name. The corresponding total systematic exposure, taking into account the cascading nature of the systematic shock, is determined according to (16). Thus, obligor d has the lowest systematic riskiness and systematic shock intensity given simply by $\rho\lambda_d$.[17] Obligor $d-1$ is slightly more risky with a total intensity comprising both the individual systematic sensitivity λ_{d-1}^{sys} and the default intensity of the less risky obligor λ_d^{sys}. Finally, the most systematically risky obligor has total intensity equal to the sum of its individual systematic sensitivity and the intensities of all better credits, which reduces to $\rho(\lambda_1 - \lambda_2)$. On top of that, each obligor is also hit by an independent idiosyncratic shock whose intensity is determined by subtracting the total systematic intensity from that name's hazard rate, i.e.:

$$
\lambda_i^{idio} := \lambda_i - \sum_{j=i}^{d} \lambda_j^{sys} = (1-\rho)\lambda_i \tag{19}
$$

The proposed decomposition (18) allows then to characterize and analyze a portfolio of credits in terms of a uniform measure of co-dependency ρ which determines the extent to which default times of obligors in that portfolio are systematic, i.e. ordered. Notice that although parameter ρ can be interpreted as a natural counterpart of the correlation parameter in the Gaussian copula (as it characterizes dependence within the portfolio), the logic in our method is reversed as true default correlation will be an output rather than an input of the modeling procedure.

The final step of the calibration process entails determining the parameter ρ, and hence by (18), also the breakdown of hazard rates into systematic and idiosyncratic components. The true ρ is obviously not observable, but a proxy can be established for example by running a principal components analysis on CDS spreads of a representative group of obligors, e.g. names included in some credit derivatives index (such as Itraxx or CDX).[18] Once principal components have been determined, the signs of the respective factor loadings can be used

[17] For ease of presentation we omit time subscripts and assume constant and deterministic hazard rates. In practice, a non-trivial term structure of hazard rates should be calibrated.

[18] For practical considerations, the dimension of the problem might have to be reduced. This can be done for example by breaking down the portfolio of obligors $\{1, ..., d\}$ into rating classes $C_1, C_2, , ..., C_k \subseteq \{1, 2, ..., d\}$ and assuming that each rating class C_j is characterized by its own hazard rate λ_j.

to distinguish between idiosyncratic and systematic factors, since loadings on the systematic factor – as opposed to idiosyncratic factors – should generally have the same sign. The approach is explained in greater detail in the example below.

Example 2. Calibrating ρ to the Itraxx Europe index

The Itraxx Europe index (compiled by Markit) comprises 125 equally weighted credit default swaps on investment grade European corporate entities. To calibrate ρ and thus determine the extent to which credits spreads (hazard rates) are driven by idiosyncratic or systematic factors we run a princiapl components analysis on the 125 spreads over the period 2007-2014. The estimation reveals that the first principal component accounts for roughly 64% of total variation in spreads, the second for roughly 32% of total variation in spreads, and the first three components together account for over 97% of total variance. A look at the signs of the coefficients reveals that loadings on the second factor are more consistently positive so it seems reasonable to interpret the second principal component as systematic. Since our systematic factor explains about 30% of the total index variance then this can suggest that the systematic component of hazard rates amounts to roughly 30%. Hence, given $\rho = 0.3$ we use (18) and set:

$$\begin{aligned}
\lambda_{125}^{sys} &:= 0.2\lambda_{125} \\
\lambda_{124}^{sys} &:= 0.2\lambda_{124} - \lambda_{125}^{sys} = 0.2(\lambda_{124} - \lambda_{125}) \\
&\vdots \quad \vdots \\
\lambda_{1}^{sys} &:= 0.2\lambda_{1} - \lambda_{2}^{sys} - ... - \lambda_{125}^{sys} = 0.2(\lambda_{1} - \lambda_{2})
\end{aligned} \tag{20}$$

whereby the names $1, 2, ..., 125$ have been ordered with respect to their exposure to the systematic factor. Given the λ_i and the breakdown into idiosyncratic and systematic components, one can simulate default times and default times correlation. Note that in this approach correlation will be model output rather than input.

Having explained the key ideas behind calibration, we now present a simple example that shows how the model can be applied to estimate residual credit risk exposure on a collateralized lending portfolio. The analysis focuses in particular on how the probability of double defaults changes with parameter ρ.

Example 3. Estimating probability of credit losses on a collateralized lending portfolio

Assume the central bank offers its 5 counterparties $N_1, ..., N_5$ a 3-year refinancing operation, allowing them to submit as collateral securities issued by 5 issuers $N_6, ..., N_{10}$. As a result, the

Figure 6: The probability of joint defaults in central bank's repo portfolio as a function of portfolio correlation

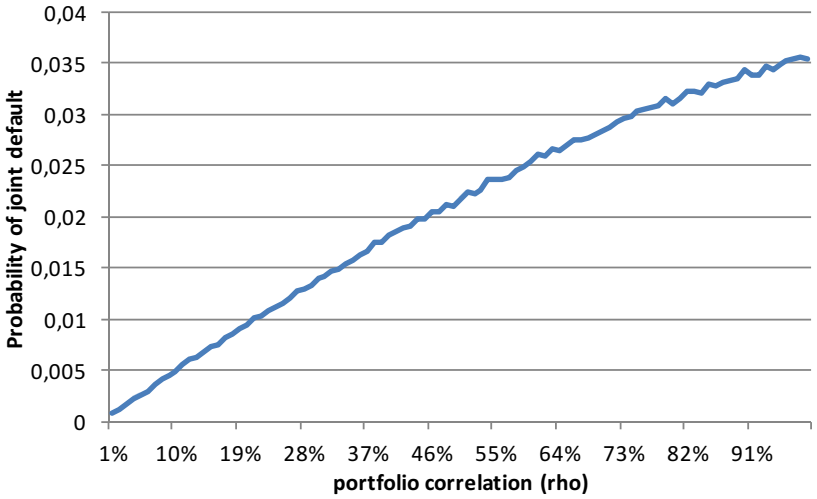

central bank's repo portfolio is made up of 5 counterparty-issuer pairs of the form (N_i, N_{i+5}), for $i = 1, ..., 5$, with credit spreads (hazard rates) of the respective obligors given in the table below. Asssume the liquidation time for all assets is 1 month.

N_1	N_2	N_3	N_4	N_5	N_6	N_7	N_8	N_9	N_{10}
2.10%	1.90%	1.80%	1.50%	1.42%	1.22%	1.10%	0.80%	0.55%	0.45%

The central bank's residual credit risk is driven by the probability that the counterparties default within the next three years and that – when they do so – the respective collateral issuers default within the next 1 month of each default:

$$\mathbb{P}\left(\{\tau_1 \leq 3 \wedge |\tau_1 - \tau_6| \leq 0.08\} \vee ... \vee \{\tau_5 \leq 3 \wedge |\tau_5 - \tau_{10}| \leq 0.08\}\right) \tag{21}$$

Such probability, in turn, will be driven by the level of portfolio correlation, or in other words by the degree to which portfolio defaults are determined by systematic and idiosyncratic factors. Given ρ we can use (18) to decompose the hazard rates into systematic and idiosyncratic components and then evaluate (21) using Monte Carlo simulation. Figure 6 shows that the probability of joint defaults in central bank's repo portfolio is an increasing function of the level of portfolio correlation. Figure 7 shows in addition a decomposition of hazard rates and simulated default times in one pseudo random scenario for two levels of portfolio correlation: 0.1 and 0.9.

Figure 7: A decomposition of hazard rates and simulated default times in one pseudo random scenario for two levels of portfolio correlation: 0.01 and 0.99

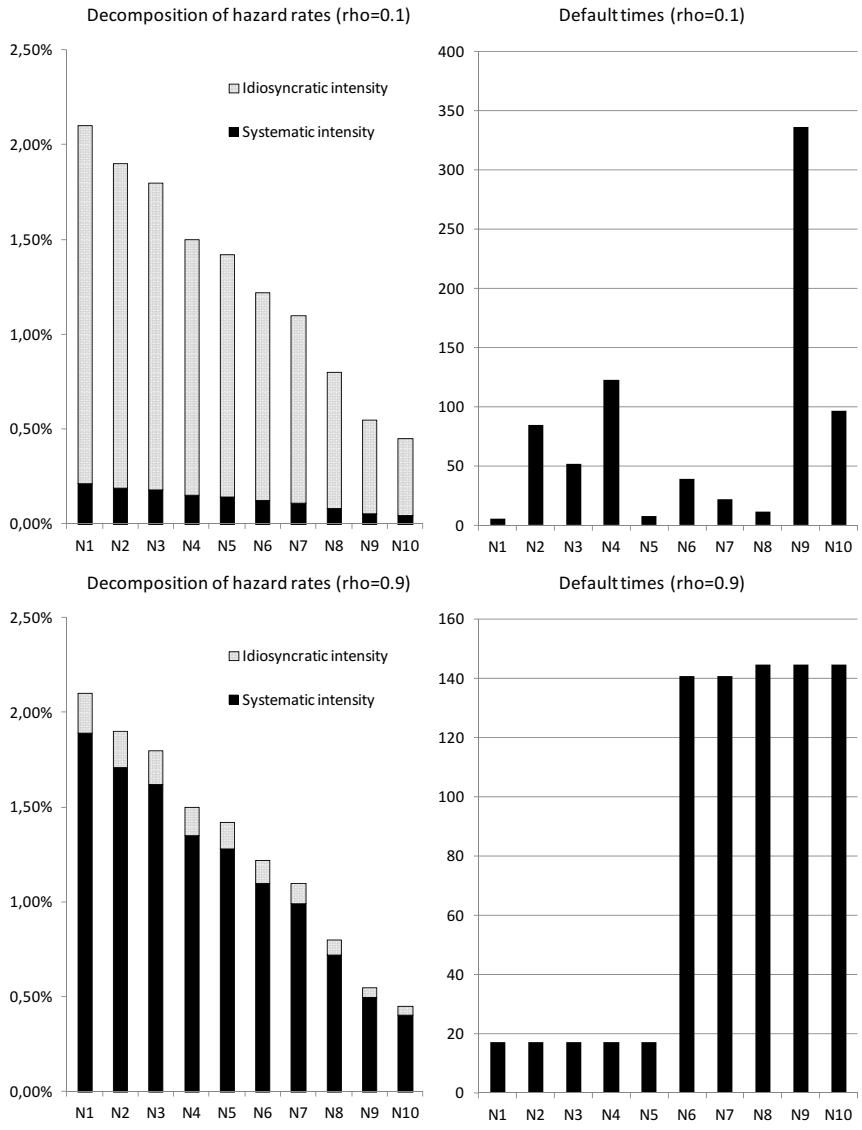

3.3 A general model for dependent defaults

In this section we show how to generalize the simple model proposed above to cater for a greater number of factors, larger class of factor distributions and stochastic hazard rates. For modeling purposes, we assume that all processes and variables are defined on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P})$, under usual conditions, with $(\mathcal{F}_t)_{t\geq 0}$ modeling the information flow and $\mathbb{P}$ being the risk-neutral (martingale) measure relative to which all security prices discounted by the risk-free interest rate are martingales. A stopping, or default, time with respect to $\mathcal{F}_t$ is a random variable τ such that $\{\tau \leq t\} \in \mathcal{F}_t$ for all $t \geq 0$. We say that a non-negative, $\mathcal{F}_t$-predictable process $\lambda(t)$ is the hazard rate, equivalently default intensity, of τ if $\mathbb{1}_{\{\tau>t\}} + \int_0^t \lambda(s)ds$ is a martingale (see Brémaud, 1981 or Schönbucher, 2003 for details).

We begin with a general definition of a factor.

Definition 1. By a systematic factor we mean a pair $\left(\{Z_i\}_{i=1}^d, \Phi\right)$, where $\{Z_i\}_{i=1}^d$ is a family of positive random variables with given distributions, and Φ is a permutation of the set $\{1, 2, \ldots, d\}$ such that the sequence $Z_{\Phi(i)}$ is increasing i.e. $Z_{\Phi(i)} \leq Z_{\Phi(i+1)}$ for $i = 1, ..., d$.

Definition 1 extends the concept of a factor as an ordered family of random variables to a multifactor setting where each factor is characterized by the order in which it triggers the default of a sequence of risky names. The ordering of defaults – reflecting different sensitivities to different factors – is modeled by associating each factor (i.e. ordering) with a permutation of the set of names $\{1, 2, ..., d\}$. Since each factor is given by a different permutation of $\{1, \ldots, d\}$, there are $d!$ possible systematic factors for a given set of obligors. The intuition for why the Marshall-Olkin "fatal shocks" approach has important practical limitations has already been given above. Here we stress the mathematical advantage of our definition which is that it allows different obligors to have different correlation with the systematic factors (i.e. different factor loadings), while preserving the "stopping time" character of credit events.[19] The latter, in turn, ensures the mathematical flexibility of martingale pricing and consistency with models developed for other asset classes, thus in principle allowing the central bank to use a consistent model for both policy and investment operations.

The construction of systematic factors in a multifactor setting is analogous to the construction presented in section 3.2. Specifically, let $\mathcal{Z} = \{\widetilde{Z}_i : i = 1, ..., d\}$ be a set of $\mathcal{F}_t$ measurable stopping times with conditionally independent respective hazard rates $\lambda(t)$, i.e. such that

$$\mathbb{1}_{\{\min_{\xi \in V} \xi > t\}} + \int_0^t \mathbb{1}_{\{\min_{\xi \in V} \xi > s\}} \sum_{\xi \in V} \lambda_\xi(s) ds \tag{22}$$

is a martingale for any $V \subseteq \mathcal{Z}$. Let $\Phi(\cdot)$ be a permutation of the set $\{1, ..., d\}$. Then, the family $Z_i = \min\{\widetilde{Z}_j : \Phi^{-1}(i) \leq j \leq d\}$ is a systematic factor. The sequence $Z_{\Phi(i)} = \min\{\widetilde{Z}_j : i \leq j \leq d\}$ is increasing and by assumption (22)

$$N_i(t) + \int_0^t N_i(s) \sum_{j=\Phi^{-1}(i)}^{d} \lambda_j(s) ds \tag{23}$$

is a martingale, and $\sum_{j=\Phi^{-1}(i)}^d \lambda_j(t)$ is the hazard rate of Z_i. To construct another factor

[19] Note that this is not true in the fatal shock Marshall-Olkin model, i.e. all factors have by definition the same sensitivity to the systematic factor. Our construction is the only way to ensure stopping time property while allowing different systematic factor loadings. To see this, suppose *a contrario* that we represent default times in the following way $\tau_i = \min\{Y_i, k_i Z\}$, where the variables k_i represent the sensitivity of name i to factor Z, which is here a standard exponential variable. Then if $k_i < 1$, we have that $t/k_i > t$ and $\{k_i Z \leq t\} = \{Z \leq t/k_i\} \notin \mathcal{F}_t$. Hence, $k_i Z$ is not a stopping time, and by implication $1/k_i$ is not a hazard rate.

we start again with a sequence of d positive random variables, choose a new permutation Ψ of the set $\{1, ..., d\}$, such that $\Psi \neq \Phi$, and proceed with the construction exactly as before. Once again, note that such representation ensures in a clear and mathematically tractable way that each of the $\{1, \ldots, d\}$ names in the economy can have a different dependence on one of the $d!$ systematic factors. This allows for an almost arbitrarily rich correlation structure, certainly beyond what used to be implied out of market data before the crisis.

We can now present the definition of correlated default times.

Definition 2. Let $U = \left\{Y_i, Z_i^1, \ldots, Z_i^N : 1 \leq i \leq d\right\}$ be the set of default times of all risk factors, both idiosyncratic, Y_i, and systematic $Z_i^1, \ldots, Z_i^N$, conditionally independent of their hazard rates. We write $U(i) = \{Y_i, Z_k^j : \Phi_j^{-1}(i) \leq k \leq d, 1 \leq j \leq N\} \subseteq U$, where $\Phi_j(\cdot)$ are the permutations ordering $(Z_{i.}^j)_{i=1}^d$, $j = 1, ..., N$ and we define dependent default times as

$$\tau_i = \min\left\{Y_i, Z_i^1, \ldots, Z_i^N\right\} = \min_{\xi \in U(i)} \xi. \tag{24}$$

The assumptions now guarantee that for all $t < T$ conditional survival probabilities under the natural filtration $\mathcal{F}_t$ are given by

$$\mathbb{P}(\xi > T \mid \mathcal{F}_t) = \mathbb{E}\left\{\exp\left(-\int_t^T \lambda_\xi(s)ds\right) \mid \mathcal{F}_t\right\} \mathbb{1}_{\{\xi > t\}} \tag{25}$$

and

$$\mathbb{P}(\tau_i > T \mid \mathcal{F}_t) = \mathbb{E}\left\{\exp\left(-\int_t^T \sum_{\xi \in U(i)} \lambda_\xi(s)ds\right) \mid \mathcal{F}_t\right\} \mathbb{1}_{\{\tau_i > t\}}. \tag{26}$$

4 Conclusions

The unprecedented lengthening of global central bank balance sheets during the recent crisis, accompanied by the deterioration of average credit quality of collateral issuers and central bank counterparties, has raised the issue of central bank finances, and in particular the extent of credit risk accumulated on those balance sheets. Central banks are not normally subject to capital requirements – and rightly so – as their losses, at least up to a point, do not compromise their ability to deliver on their obligations and to conduct monetary policy. Nevertheless, since the taking of credit risk necessarily entails an allocation of capital, it falls more appropriately under the scope of fiscal policy executed with a democratic mandate. This implies that in cases when a central bank must stand ready to take on credit exposure – e.g. to contain a financial crisis – the natue and scope of this exposure should be well understood, properly measured and managed. Over and above the requirements of transparency and democratic accountability, central banks have a highly developed sensitivity for reputation risk which is crucial for their ability to achieve policy goals and to "preach" credibly to the rest of the world what is right and wrong. Mismanaging credit risk and suffering losses – even small relative to balance sheet size – could be perceived as irresponsible or amateurish and would thus weaken credibility.

In this paper we have argued that the conceptual framework underlying the risk management solutions adopted (explicitly or implicitly) by central banks – i.e. the Gaussian latent factor model – is inadequate to the study of default clustering in time. This is unfortunate, since to the extent that a central bank conducts its policy operations through collateralized lending, its residual credit risk is limited exactly to default clustering: when default of the counterparty and the collateral issuer are very near in time (within collateral liquidation horizon). We show that using the Gaussian latent factor model in the time domain to study default times often produces unintuitive results. Moreover, in many practically relevant cases there is no clear and stable link between the main model parameter – i.e. asset correlation – and model outcomes (risk parameters), which is the main source of model risk. Obviously, no model is ever correct in the sense of being a perfect description of reality – and nor should it be – but given the special role played by central banks and their sensitivity to reputation risk, it appears that exposure to credit risk should be well understood and exposure to model risk minimized. We believe the Gaussian latent factor model, as applied to the measurement of residual credit risk, is not compatible with these goals and propose a more intuitive solution.

Our model can be considered an extension of the Marshall and Olkin (1967) approach in that it assumes that the default of each obligor in the portfolio can be triggered either by an idiosyncratic factor or by any of the systematic factors – whichever hits first. To intuitively capture dependence in the time domain, we redefine a systematic factor as a

sequence of increasing random variables, rather than a single random variable as per Marshall and Olkin. The redefinition allows to capture the practically relevent feature of obligors' different dependence on systematic factors without sacrificing the stopping time property of default times (which in turn is a mathematical condition allowing us to use the powerhorse of martingale methods). The model produces default clustering in time and features an intuitive relation between a portfolio correlation parameter and the probability of joint defaults near in time. The model is also straightforward to calibrate and implement. Finally, since the model is cast in the martingale framework, it is consistent with approaches used for other asset classes, and hence can potentially be used for integrated credit and market risk management – both by central banks and more broadly within the private sector.

References

Andersen, L., and J. Sidenius (2004): "Extensions to the Gaussian copula: Random recovery and random factor loadings," *Journal of Credit Risk Volume*, 1(1), 05.

Archer, D., and P. Moser-Boehm (2013): "Central bank finances," BIS Papers 71, Bank for International Settlements.

Bagehot, W. (1873): *Lombard Street: A description of the Money Market*. H. S. King.

Bindseil, U. (2009): "Central bank financial crisis management from a risk management perspective," in *Risk Management for Central Banks and Other Public Investors*, ed. by U. Bindseil, F. Gonzalez, and E. Tabakis, pp. 394–440. Cambridge University Press.

Bindseil, U., and J. Jablecki (2013): "Central bank liquidity provision, risk-taking and economic efficiency," ECB Working Paper 1542, ECB.

Bindseil, U., and F. Papadia (2009): "Risk management and market impact of central bank credit operations," in *Risk Management for Central Banks and Other Public Investors*, ed. by U. Bindseil, F. Gonzalez, and E. Tabakis. Cambridge University Press.

Bluhm, C., L. Overbeck, and C. Wagner (2002): *An introduction to credit risk modeling*. CRC Press.

Brémaud, P. (1981): *Point Processes and Queues: Martingale Dynamics: Martingale Dynamics*. Springer.

Buiter, W., and E. Rahbari (2012): "Looking into the Deep Pockets of the ECB," Global economics view, Citigroup Global Markets.

Derman, E. (2001): "The Principles and Practice Of Verifying Derivatives Prices," Discussion paper, Firmwide Risk, Goldman Sachs & Co.

Duffie, D., and J. Pan (2001): "Analytical value-at-risk with jumps and credit risk," *Finance and Stochastics*, 5(2), 155–180.

ECB (2007): "The use of portfolio credit risk models in central banks," ECB Occasional Paper 64, Task Force of the Market Operations Committee of the European System of Central Banks.

Elouerkhaoui, Y. (2006): "Etude des Problemes de Corrélation et d'Incomplétude dans les Marchés de Crédit," Ph.D. thesis, Universite Paris IX - Dauphine.

Embrechts, P., F. Lindskog, and A. McNeil (2003): "Modelling dependence with copulas and applications to risk management," *Handbook of heavy tailed distributions in finance*, 8(1), 329–384.

Gatarek, D. (2010): "Individual names in top-down CDO pricing models," *Risk*, March, 80–84.

Gatarek, D., and J. Jabłecki (2013): "Systematic risk factors redefined," *Risk*, November, 70–74.

Giesecke, K. (2003): "A simple exponential model for dependent defaults," *Journal of Fixed Income*, 13(3), 74–83.

Gordy, M. B. (2003): "A risk-factor model foundation for ratings-based bank capital rules," *Journal of financial intermediation*, 12(3), 199–232.

Goupton, G., C. Finger, and M. Bhatia (1997): "Creditmetrics," *Technical Document, Morgan Guaranty Trust Company*.

Hanson, S. G., M. H. Pesaran, and T. Schuermann (2008): "Firm heterogeneity and credit risk diversification," *Journal of Empirical Finance*, 15(4), 583–612.

Heinle, E., and M. Koivu (2009): "Risk measurement for a repo portfolio – an application to the Eurosystem's collateralized lending operations," in *Risk Management for Central Banks and Other Public Investors*, ed. by U. Bindseil, F. Gonzalez, and E. Tabakis, pp. 359–393. Cambridge University Press.

Hull, J. C., and A. D. White (2004): "Valuation of a CDO and an n-th to default CDS without Monte Carlo simulation," *The Journal of Derivatives*, 12(2), 8–23.

Lando, D. (2004): *Credit Risk Modeling: Theory and Applications*. Princeton Series in Finance.

Lindskog, F., and A. J. McNeil (2003): "Common Poisson shock models: applications to insurance and credit risk modelling," *Astin Bulletin*, 33(2), 209–238.

Marshall, A. W., and I. Olkin (1967): "A multivariate exponential distribution," *Journal of the American Statistical Association*, 62(317), 30–44.

Merton, R. C. (1974): "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance*, 29(2), 449–70.

Morini, M. (2011): *Understanding and Managing Model Risk: A practical guide for quants, traders and validators*. Wiley.

O'Kane, D. (2008): *Modelling Single-name and Multi-name Credit Derivatives*. John Wiley and Sons.

Schönbucher, P. (2003): *Credit Derivatives Pricing Models*. Wiley.

Vasicek, O. (2002): "The distribution of loan portfolio value," *Risk*, 15(12), 160–162.

Vesely, W. (1977): "Estimating common cause failure probabilities in reliability and risk analyses: Marshall-Olkin specializations," *Nuclear systems reliability engineering and risk assessment*, pp. 314–341.

www.nbp.pl