Credit contagion and credit risk

JPL Hatchett† **and R Kühn**§

† Hymans Robertson LLP, One London Wall, London EC2Y 5EA, UK

§ Department of Mathematics, King's College London, The Strand, London WC2R 2LS, United Kingdom

Abstract. We study a simple, solvable model that allows us to investigate effects of credit contagion on the default probability of individual firms, in both portfolios of firms and on an economy wide scale. While the effect of interactions may be small in typical (most probable) scenarios they are magnified, due to feedback, by situations of economic stress, which in turn leads to fatter tails in loss distributions of large loan portfolios.

PACS numbers: 02.50.-r, 05.40.-a, 89.65.Gh, 89.75.Da

E-mail: `jon.hatchett@hymans.co.uk` `reimer.kuehn@klc.ac.uk`

1. Introduction

Modelling credit risk in a coherent yet applicable manner is an important yet challenging problem. The difficulties arise from the combination of a large, and co-dependent set of risk parameters such as default rates, recovery rates, or exposures, which are correlated and non-stationary in time. An additional issue is that of credit contagion [1, 2, 3, 4, 5, 6, 7, 8, 9], which examines the role of counter-party risk in credit risk modelling. If a firm is in economic distress, or defaults, this will implications for any firm which is economically influenced by this given firm, for example, a service provider to it, purchaser of its goods or a bank with a credit line to the firm. The direct correlations between firms caused by credit contagion lead to further complications in modelling the overall, either portfolio or economy wide, level of risk. Davis and Lo [1] considered a model in which defaults occur either directly, or through infection by another defaulted firm, with probabilities for direct default or infection taken *uniform* throughout the system (or throughout sectors, assuming independence across sectors). Defaults occurring due to both, endogenous or exogenous causes were not considered in their set-up. Jarrow and Yu [2] introduced a framework of primary and secondary firms, the former would default depending on some background stochastic process while the latter were affected by a stochastic process and the performance of the primary firms. They argued that this was a reasonable level of detail for their purposes and it also simplifies matters as there are no feedback loops in the system. Secondary firms depend only on primary firms whose performance is independent of the secondary firms. Rogge and Schönbucher [3] use copula functions to quantify correlations in default dynamics, and in particular to determine the impacts a defaulting obligor will have on the hazard rates of other obligors in a portfolio — conditioned on a specification of the set of countdown levels of surviving obligors and on the set of defaults that have already occurred at the given time. While the expression of the conditional impact parameters in terms of a covariance of macro-economic factors is intriguingly simple at a formal level, the evaluation can, as pointed out in [3], be cumbersome in practice once a sizeable number of defaults has occurred. Also, actually *solving* the dynamics of hazard rates using the impact parameters thus computed is an entirely non-trivial affair. This is hardly addressed in [3] even for moderately sized portfolios. Another approach for modelling credit contagion dynamics was provided by Giesecke and Weber [4] who used the well known voter process [10], from the theory of interacting particle systems, to model interactions between firms. They assumed a regular structure for their firms (a regular infinite hyper-cubic lattice) and focussed on the equilibrium properties of the model. The model is highly idealised. It would seem that both the regularity (and the symmetry) of the interaction pattern would have to be abandoned, if one were to calibrate a model of this kind to represent realistic patterns of mutual dependencies. Egloff et al. [5] model contagion using a linear coupling of asset returns between business counterparts to describe the micro-structure of mutual dependencies. This leads to a self-consistent description of mutual dependencies in equilibrium (though an autoregressive mechanism is mentioned to capture non-equilibrium situations), which allows analytic solutions even for the case of asymmetric and heterogeneous impacts. This feature would seem to open the way for a proper calibration of their model, though it has been argued [11] that feedback mechanisms via rating events rather than (unobservable) asset returns would be preferable in realistic models of contagion. Frey and Backhaus [6] and Kraft and Steffensen [9] use continuous time Markov models to describe the dynamics of transitions of the indicator variables describing rating classes of the obligors in a portfolio. The major problem here is that the state space of the system grows exponentially in portfolio-size. Frey and Backhaus circumvent this problem by using a mean-field approximation for *large* portfolios, assuming that these portfolios contain only a small number of different sectors, and that contagion effects are homogeneous within sectors, whereas Kraft and Steffensen [9] concentrate on *small* portfolios (involving 2 or 3 firms), and so-called n-to default baskets with small n chosen such that the dimension of the state space remains small, allowing them, among other things, to derive explicit results for loss-distributions, and also to address pricing issues in some detail.

There are a variety of techniques for modelling the correlations between firms' default behaviour, which is a major complication in credit risk modelling. The binomial expansion technique assumes independence between firms so that the number of defaults in a portfolio is described by a binomial distribution. In order to capture the effects of correlations a binomial distribution with an "effective" number of firms is assumed which is *smaller* than the actual number in the portfolio, but the weight given to each firm scaled so as to keep the mean number of defaults constant, while the variance of the overall number of defaults is increased. The relationship between the true number of firms and the effective reduced number is a modelling choice that depends on the diversity of the firms in terms of sectors, geographic locations or any other identifiable trait that would lead to strong correlations in default behaviour. JP Morgans' CreditMetrics approach [12] and Credit Suisse First Financial Products CreditRisk$^+$ [13] (see [14] for a detailed comparison between the two) uses the correlations in equity values as a surrogate for the correlations in credit quality. The structural modelling approach goes back a long way to work by Merton [15] which directly models the dynamics of a firm's assets, with default being triggered by the asset value hitting some predetermined value (which henceforth we take without loss of generality to be zero). Correlations between firms are due to correlations in the dynamics of different firms' assets. This approach is very general, as it is relatively transparent to identify different driving forces of asset levels and straightforward to include them in the model (though the resulting model itself will be non-trivial). However, it suffers from the fact that the asset level is not an observable quantity [11]. On the other hand, the reduced form approach gives default rates for a given firm without modelling the underlying default process. Correlations are then directly introduced between the default rates. There was some discussion in the literature about whether the reduced form model could describe the true level of default correlations seen empirically. Yu [16] seems to have answered this question in the affirmative if a suitable structure between the default rates is taken into account, while the results of Das et al. [17] seem to imply that the reduced form model is insufficient to fully account for observed default correlations and direct contagion would indeed be required for a full explanation.

The approach we take here is a discrete time Markov process (at the microeconomic level) where the probability of a default of a given firm in a particular time step depends materially on the state of its economic partners at the start of that time step, as well as on macro-economic influences. The model we look at was proposed by Neu and Kühn [7], and initially studied using Monte-Carlo simulations. The model improves upon [6] by introducing a fully heterogeneous specification, with model parameters given in terms of unconditional and conditional default probabilities, or alternatively as a variant of [5] in which a full dynamic description is maintained, and contagion is via rating events (defaults) rather than via unobservable asset returns. Using techniques developed in the statistical mechanics of disordered systems [18], and recently applied to this specific model in [8], we are able to solve the dynamics of our model exactly, and given our assumptions that we describe shortly, this solution takes a particularly simple form despite the fact that in principle we have feedback correlations, non-equilibrium dynamics and in principle non-Markovian behaviour at the macroscopic (economy/portfolio wide) level. We note that it is possible to frame our model in either the structural approach or the reduced form approach, depending on requirements and taste, although the interpretation of the variables in the two approaches will of course be different. We find that the correlations introduced through credit contagion lead to large increases in default rates in times of economic stress, above and beyond those introduced by simple macro-economic dependencies. This has strong implications for portfolio risk management.

2. The microeconomic framework

We will analyse an economy of N firms in the large N limit. Generally, we focus on the characteristic changes in the economy due to interactions between firms, which will be described in a probabilistic manner.

As mentioned in the introduction we take a discrete time approach. For clarity we restrict our discussion to a one year time frame split into twelve steps; this is not essential, but parameters may need rescaling depending on the set-up. We use a binary indicator variable $n_{i,t}$ to denote whether firm i is solvent at time t ($n_{i,t} = 0$) or has defaulted ($n_{i,t} = 1$). The default process is a function of an underlying stochastic process for each firm in terms of a "wealth" variable $W_{i,t}$, where we assume default if the wealth drops below zero. We shall assume that recovery from default over the time horizon of a year is not possible, so that the defaulted state is absorbing. As a function of the wealth, therefore, the indicator variables evolve according to

$$n_{i,t+1} = n_{i,t} + (1 - n_{i,t})\Theta(-W_{i,t}) \ , \tag{1}$$

where $\Theta(\ldots)$ is the Heavyside function.

A *dynamic model* for the indicator variables is obtained from (1) by specifying the underlying stochastic process for the wealth variables $W_{i,t}$. We shall take it to be of the form

$$W_{i,t} = \vartheta_i - \sum_{j=1}^{N} J_{ij} n_{j,t} - \eta_{i,t} \ . \tag{2}$$

Here ϑ_i denotes an *initial wealth* of firm i at the beginning of the risk horizon, and J_{ij} quantifies the material impact on the wealth of firm i that would be caused by a default of firm j. This may or may not be a reduction in wealth, depending on whether j has a cooperative ($J_{ij} > 0$) or a competitive ($J_{ij} < 0$) economic relation with i.

We shall assume that the fluctuating contributions $\eta_{i,t}$ to (2) are zero-mean Gaussians. There is still some degree of flexibility concerning the decomposition of the $\eta_{i,t}$ into contributions that are *intrinsic* to the firm and *extrinsic* contributions. The latter describe the influence of economy-wide fluctuations or fluctuations pertaining to different economic sectors, depending on the level of detail required. We restrict ourselves to a minimal model containing a *single* macro-economic factor (assumed to be constant over a risk horizon of a year), and individual fluctuations for each firm,

$$\eta_{i,t} = \sigma_i \left(\sqrt{\rho_i}\eta_0 + \sqrt{1 - \rho_i}\ \xi_{i,t} \right) \ , \tag{3}$$

where σ_i sets the scale of the individual fluctuations, and the $\{\xi_{i,t}\}$ are taken to be *independent* $\mathcal{N}(0, 1)$ Gaussians; finally, the parameters ρ_i quantify the correlations of the $\eta_{i,t}$ created via the coupling to economy-wide fluctuations η_0, also taken to be $\mathcal{N}(0, 1)$.

It is not necessary to take the $\xi_{i,t}$ independent, but it allows us to interpret the wealth of the firm as being given by $\vartheta_i - \sum_{j\neq i} J_{ij} n_j$ which describes the state of the firm (which evolves as a jump process each time step), and then the fluctuations about that state give a certain *intensity* of defaults per time step (given the wealth).

Up to this point the wealth dynamics does not contain an endogenous drift. If predictions are required over longer time periods then it may also be pertinent to introduce such a drift, e.g. by using a time-dependent ϑ_i for example, $\vartheta_{i,t} = \vartheta_i(0)\mathrm{e}^{z_i t}$, where z_i denotes an intrinsic growth rate of the average wealth of firm i (with $z_i > 0$ for a firm making profits and $z_i < 0$ for a firm making losses). However, for the current purposes of examining default rates over the medium term and especially focussing on the behaviour on the tails, this adjustment does not lead to significant changes in our overall conclusion.

The model, as formulated above, clearly takes a *structural point of view* on the problem of credit contagion. However, we note that the dynamics (1) of the indicator variables is clearly independent of the *scale* of the wealth variables $W_{i,t}$. By appropriately rescaling the initial wealths ϑ_i and the impact parameters J_{ij} we can thus assume a unit-scale $\sigma_i \equiv 1$ for the noise variables (3). Interestingly, this simple rescaling, which leaves the dynamics of the system unaffected, amounts to changing to a *reduced-form* interpretation of the dynamics.

To see this, note from (2) that the event $W_{i,t} < 0$ is equivalent to $\eta_{i,t} > \vartheta_i - \sum_{j=1}^{N} J_{ij} n_{j,t}$. With $\sigma_i \equiv 1$, we see that this occurs with probability $\Phi(\sum_j J_{ij} n_{j,t} - \vartheta_i)$ where $\Phi(\cdot)$ is the cumulative normal distribution. From a reduced form point of view this is just the intensity of default of firm i at time step t (in a given economic environment specified by the set of firms defaulted at time t). This allows us to re-interpret the (rescaled) initial wealth and impact variables ϑ_i and J_{ij} in terms of the bare default probabilities [7, 8, 20]. In other words, if company i has an expected default probability of p_i in a given time unit (e.g. one month in the present set-up) as predicted from tables from ratings agencies, then $\vartheta_i = -\Phi^{-1}(p_i)$. Similarly, the expected default probability $p_{i|j}$ of firm i, given that only firm j has defaulted leads to the value $J_{ij} = \Phi^{-1}(p_{i|j}) - \Phi^{-1}(p_i)$.

In determining the model parameters by the method suggested above we are splitting our default probability into terms that come from credit contagion and other terms such as the bare default probability that come from historical data. It could fairly be argued that the historical data already incorporate the credit contagion terms and thus we are double counting. As we will see later in numerical simulations, the *credit contagion terms make very little difference to average behaviour* and thus making estimates based on average historical data is still a reasonable approach.

In choosing the variable ρ_i we follow the prescription given by BASEL II [19] which sets

$$\rho_i = 0.12\frac{1-\mathrm{e}^{-50\mathrm{PD}_i}}{1-\mathrm{e}^{-50}} + 0.24\left(1 - \frac{1-\mathrm{e}^{-50\mathrm{PD}_i}}{1-\mathrm{e}^{-50}}\right) \approx 0.12\left(1+\mathrm{e}^{-50\mathrm{PD}_i}\right) \quad (4)$$

where PD_i gives the probability of default of firm i over one year, ignoring credit contagion effects. With $p_i = \Phi(-\vartheta_i)$ as the monthly default probability, we have $\mathrm{PD}_i \approx 12\Phi(-\vartheta_i)$.

We still have to specify the form for the economic interactions. We adopt here a probabilistic approach, and investigate a 'synthetic' portfolio, described by its statistical properties. Specifically, we take the interactions to be random quantities of the form

$$J_{ij} = c_{ij}\left[\frac{J_0}{c} + \frac{J}{\sqrt{c}}x_{ij}\right] . \tag{5}$$

Here, the $c_{ij} \in \{0,1\}$ detail the network (presence or absence) of interactions between different firms and we choose these to be randomly fixed according to

$$P(c_{ij}) = \frac{c}{N}\delta_{c_{ij},1} + \left(1 - \frac{c}{N}\right)\delta_{c_{ij},0} , \qquad i<j , \qquad c_{ji} = c_{ij} . \tag{6}$$

We assume that the average connectivity c of each firm is large in the limit of a large economy; this will allow the influence of partner firms to be described by the central limit theorem and the law of large numbers. Concerning the values of the (non-zero) impact parameters, we parametrise them as shown, with x_{ij} assumed to be zero-mean, unit-variance random variables, with finite moments, and which are *pairwise independent*,

$$\overline{x_{ij}} = 0 , \qquad \overline{x_{ij}^2} = 1 , \qquad \overline{x_{ij}x_{ji}} = \alpha , \qquad \overline{x_{ij}x_{kl}} = 0 \ \text{ otherwise} . \tag{7}$$

The parameters J_0 and J determine the mean and variance of the interaction strengths; the scaling of mean and variance with c and $\sqrt{c}$ respectively in (5) is necessary to allow a meaningful large c limit to be taken. Taking $J_0 > 0$ would encode the fact that on average firms have a synergy with their economic partners.

At first sight, specifying the J_{ij} appears to introduce a vast number of parameters into our model, but in fact only the first two moments of the distribution of interaction strengths are sufficient to determine the macroscopic behaviour of the system, and so the model space is not too large.

Let us now turn to the capital required to be held against credit risk. In the BASEL II document [19] the capital requirement for a unit-size loan given to firm i is

$$K_i = \mathrm{LGD}_i\left[\Phi\left(\frac{\sqrt{\rho_i}\,\Phi^{-1}(0.999) + \Phi^{-1}(\mathrm{PD}_i)}{\sqrt{1-\rho_i}}\right) - \mathrm{PD}_i\right] M_i . \tag{8}$$

The first factor, the loss given default LGD_i of firm i, is related to the average fraction of a loan that can be recovered despite default. The last factor, M_i, is related to the maturity (long dated loans are inherently riskier). Adjustments related to liquidity (low liquidity loans are risker) and concentration (fewer, larger loans give a greater variance in returns for given expected return) are occasionally also included in this factor — concentration-adjustments, in fact, are a means to account for reduced granularity in a credit portfolio resulting from the possibility of credit contagion.

The factor inside square brackets in (8) is entirely related to the loss-frequency distribution. The first term is the value of the loss frequency not exceeded with probability $q = 0.999$ under fluctuating macro-economic conditions, with ρ_i describing the dependence of the firm's loss-frequency on the macro-economic factor. The second term is the average loss frequency. The value of the confidence level q is in principle arbitrary, but is related to the target rating of the bank. The risk weighted asset is then found by further multiplying by terms such as the exposure at default (i.e. size of the loan). Thus the capital required for firm i can be viewed as the loss at the 99.9th percentile level of stress, in *excess* of the expected loss, multiplied by a conversion factor. From this structure it is clear that a key ingredient for the capital adequacy requirements is a good model of credit risk that works well into the tail of the loss frequency distribution.

Returning to our description of *default dynamics*, let us first focus on the case of independent firms, with $J_{ij} = 0 \; \forall i, j$, and consider a single epoch for our model with fluctuating forces given by (3) at given macro-economic condition η_0. The probability of a default of firm i with average unconditional monthly default probability p_i occurring during the epoch $t \to t+1$ in our model is given by

$$\langle n_{i,t+1} | n_{i,t} = 0 \rangle = \Phi\left(\frac{\sqrt{\rho_i}\,\eta_0 + \Phi^{-1}(p_i)}{\sqrt{1-\rho_i}} \right) \tag{9}$$

Since the probability of default is increasing with η_0, we can find the probability of default not exceeded at e.g. the 99.9 percent confidence level; it is given by setting $\eta_0 = \Phi^{-1}(0.999)$ in the above equation (recall η_0 is distributed as a zero-mean, unit-variance Gaussian random number). As above, the excess capital required is the loss at the 99.9th percent level minus the expected loss (multiplied by a risk factor). However, when we consider the case of an interacting economy with non-zero J_{ij}, we find that in fact

$$\langle n_{i,t+1} | n_{i,t} = 0 \rangle = \Phi\left(\frac{J_0 m_t + \sqrt{\rho_i}\,\eta_0 - \vartheta_i}{\sqrt{1-\rho_i + J^2 m_t}} \right) , \tag{10}$$

where

$$m_t = \frac{1}{N} \sum_j n_{j,t} \tag{11}$$

is the fraction of firms within the economy that have defaulted up to time t; we also expressed the expected monthly default rate p_i in terms of a 'rescaled initial wealth' ϑ_i, $\Phi^{-1}(p_i) = -\vartheta_i$.

Thus we find that our formulation is very similar to that used in BASEL II. However, we directly take account of the correlations in defaults caused by credit contagion. This introduces two extra parameters into the model but it does markedly change the behaviour in the tails of the loss frequency distribution, and thereby in the tails of the loss distribution itself. Correlation between firms is essentially a dynamic phenomenon — if there is no dynamics, there is no way for one's firms' performance to influence the viability of any other firm. Thus rather than considering firms to be independent over a single epoch which lasts the entire period of any loan, we split the overall time (e.g. one year) into smaller units (e.g. one month) and let the firms evolve over these smaller time units with the default probability adjusted (since the default event in 12 monthly epochs is compounded 12 times as opposed to a single epoch). A firm may default at any point, but will then influence its partner firms for the remainder of the time horizon. The complexity of the theory is merely linear in time, thus it is not a great computational burden to choose this approach.

Following the approach described in [8] it is possible to solve the model in a stochastic manner. Credit contagion within this model is encoded at each time by a single number, the fraction of firms that have defaulted thus far, which evolves according to

$$m_{t+1} = m_t + \left\langle \Big(1 - \langle n_t(\vartheta)\rangle\Big) \Phi\left(\frac{J_0 m_t + \sqrt{\rho(\vartheta)}\,\eta_0 - \vartheta}{\sqrt{1-\rho(\vartheta) + J^2 m_t}}\right)\right\rangle_\vartheta \tag{12}$$

where $\langle n_t(\vartheta)\rangle$ denotes the time-dependent monthly default rate of firms with $\vartheta_i \approx \vartheta$, as influenced by interactions with the economy, and the larger angled brackets with subscript ϑ denote an average over the bare monthly probabilities of default for the ensemble of firms, or equivalently over the distribution $p(\vartheta)$ of their rescaled initial wealth parameters ϑ. A heuristic derivation of this result can be found in the appendix. For a full justification of the assumptions used in the derivation, we refer to [8].

In (12) the Basel II recommendation which links correlations to macro-economic factors with (unconditional) default probabilities, $\rho_i = \rho(p_i) \to \rho(\vartheta_i)$, via (4) is formally taken into account. Note that this correlation was not implemented in [8].

Also, note that credit contagion affects the dynamics of defaults only via two parameters, J_0 and J, which characterise the mean and variance of the impact parameter distribution. Further, the parameter α which quantifies forward-backward correlation of mutual impacts according to (7) does not appear in the final formulation, nor are there any memory-effects in the dynamics, as would normally be expected for systems of this type. The reason for this simplifying feature is in the fact that the defaulted state is taken to be absorbing over the risk horizon of one year.

3. Results

We now turn to presenting a few key results of our analysis. Our results concerning default dynamics and loss distributions are obtained for an economy in which the parameters ϑ_i determining unconditional monthly default probabilities p_i according to $\vartheta_i = -\Phi^{-1}(p_i)$, are normally distributed with mean $\vartheta_0 = 2.75$, and variance $\sigma_\vartheta^2 = 0.1$ so that typical monthly bare default probabilities are in the $10^{-5} - 10^{-3}$ range. The couplings to the macro-economic factor are chosen to depend on the expected default probabilities according to the Basel II prescription (4).

In Fig. 1 we we show that renormalisation (with respect to credit contagion) makes little difference to the *typical* default dynamics observed for $\eta_0 = 0$, i.e. for neutral macro-economic conditions. The evolution of the fraction of defaulted firms in interacting economies differs hardly from that of the non-interacting economy with $J_{ij} = 0 \Leftrightarrow (J_0, J) = (0, 0)$. At least, as the differences are of the order of one percent of the portfolio after a twelve month period, the differences are smaller than the uncertainty that would be introduced by using historic default rates to calibrate any model (along with other uncertainties by the model choice). If one were concerned by this difference it would not be unduly taxing (although potentially rather unedifying) to alter the independent default probability so that the mean number of defaulted firms at the end of a year (or any chosen epoch) were independent of $\{J, J_0\}$.

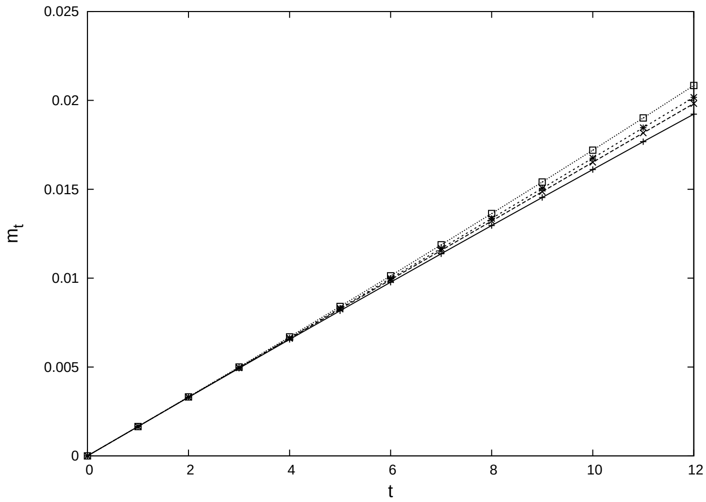

Figure 1. Typical fraction of defaulted companies as a function of time for $(J_0, J) = (0,0), (1,0), (0,1)$, and (1,1) (bottom to top), realized for a neutral macro-economic factor $\eta_0 = 0$.

In marked contrast to this, the tail of loss-frequency distributions is strongly affected by the presence of interactions in the system, as shown in Fig. 2. We also note that the tail of the loss-frequency distribution is more pronounced than in our previous study [8]. This is *solely* due to the fact that in the present paper we followed the Basel II suggestion that relates the coupling of a company to macro-economic factors with its default probability via (4), stipulating that the coupling to the macro-economic factor *decreases* with increasing probability of default. As a consequence, companies with very low unconditional default probabilities will be driven into default *mainly* in rare situations of *extreme* economic distress where interaction generated avalanches of risk-events are likely to occur.

Let us now look at the economy-wide losses per node. Here we will not only evaluate end-of-year results, but exploit the dynamical information contained in (12) to also look at the way losses are accrued as a function of time t.

For a given macro-economic condition η_0, the loss per node at time t is given by

$$L_t(\eta_0) = \frac{1}{N} \sum_i n_{i,t} \ell_i \ . \tag{13}$$

We assume that the ℓ_i are randomly sampled from the loss distribution for node i, are taken to be independent of the stochastic evolution, but are possibly correlated with the bare monthly default probability. In the large N limit this gives

$$L_t(\eta_0) = \lim_{N\to\infty} \frac{1}{N} \sum_i n_t(\vartheta_i) \ell_i = \int \mathrm{d}\vartheta p(\vartheta) \langle n_t(\vartheta) \rangle \overline{\ell}(\vartheta) \tag{14}$$

by the law of large numbers, where $\overline{\ell} = \overline{\ell}(\vartheta)$ is the mean of the loss distribution for a node with default probability $p_d(\vartheta)$.

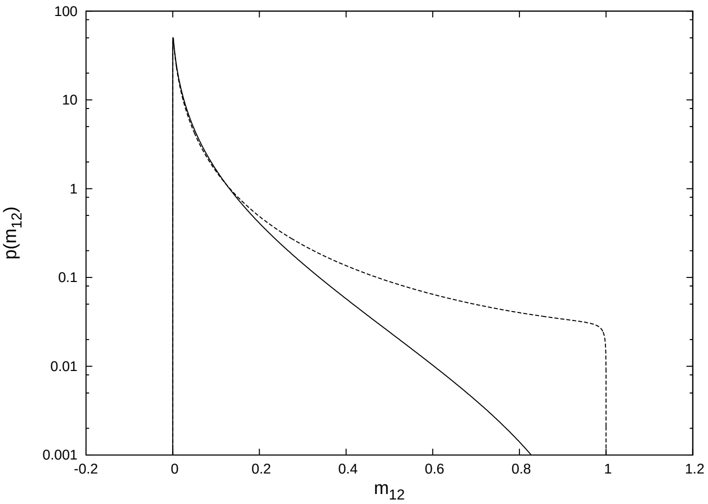

Figure 2. Probability density function of the fraction of defaulted companies at $t = 12$ months, for $(J_0, J) = (0, 0)$ (bottom) and $(J_0, J) = (1, 1)$ (top).

The distribution of the economy-wide fraction of defaulted nodes and thereby the distribution of the losses per node is driven by the distribution of the macro-economic factor η_0. Noting that the fraction of defaulted nodes $m_t = m_t(\eta_0)$ (and thereby the loss accrued up to time t) are monotone increasing functions of η_0 which is itself assumed to be $\mathcal{N}(0, 1)$, one has

$$\text{Prob}[m \le m_t(\eta_0)] = \Phi(\eta_0(m_t)) \; , \tag{15}$$

which entails $\text{Prob}[L \le L_t(\eta_0)] = \Phi(\eta_0(L_t))$ for the loss distribution. The corresponding probability density functions are obtained via a single numerical differentiation.

A typical result for the loss distribution is shown in Fig. 3, for which we consider an economy where average losses are inversely proportional to the unconditional default probabilities $p_i = p_d(\vartheta_i) = \Phi(-\vartheta_i)$,

$$\overline{\ell}(\vartheta) = \frac{\ell_0}{\varepsilon + p_d(\vartheta)} \tag{16}$$

with a parameter $\varepsilon > 0$ as a regularizer preventing divergence as $p_i \to 0$. In this way, the contribution to the total losses will be approximately uniform over the bands of different default probabilities. We make no claim that this choice is in any way singled out on any form of a-priori grounds and there is relatively large freedom of modelling choice in (16) — one just requires a reasonably behaved function of ϑ. Indeed, if we make the simpler assumption that the loss distributions are independent of default probabilities, and take $\overline{\ell}(\vartheta) \equiv 1$, then the distribution of losses per node simply replicates the distribution of the fraction of defaulted companies as shown in Fig 2 at $t = 12$ months. For the results in Fig. 3 we took $\ell_0 = 1$ and the regularizer $\varepsilon = 0.005$.

In Fig. 3 we also show the loss distribution at half term, to illustrate how losses build up over the full risk horizon of 12 months. Economic interactions are seen to strongly affect the tail of the loss distribution at large losses, which is due to the possibility of avalanches of loss events in times of extreme economic stress. Indeed a comparison of the loss distribution of the interacting system at half term and at full term reveals that the avalanches that create the mass of the loss distribution at the extreme end have at half term not yet been able to sweep through the entire economy, while at full term they have.

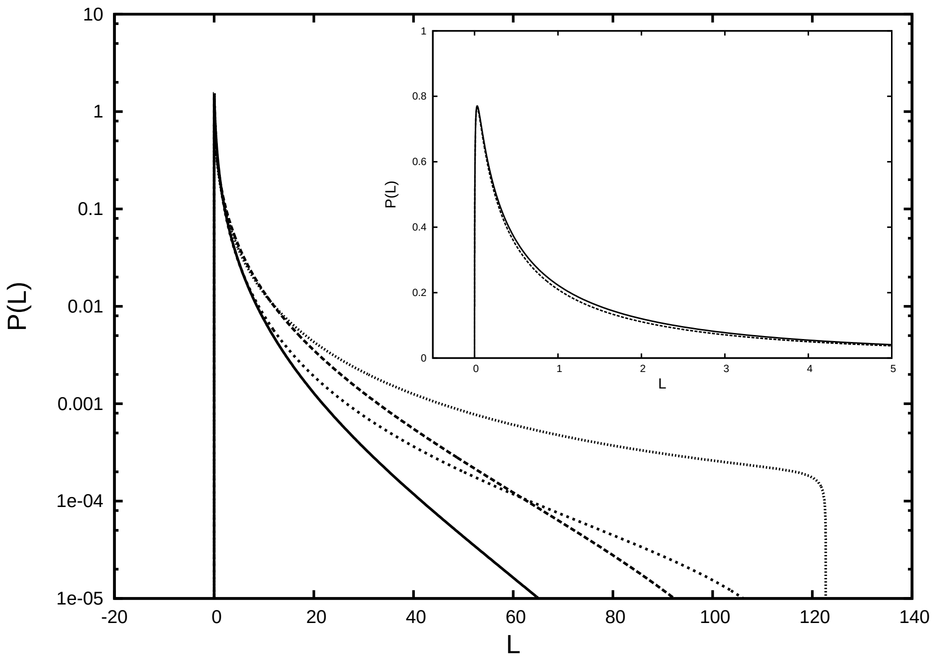

Figure 3. Probability density function of losses for the non-interacting system with $(J_0, J) = (0, 0)$ at 6 months (full line) and 12 months (long-dashed), as well as for the interacting economy with $(J_0, J) = (1, 1)$ at 6 months (short-dashed) and 12 months (dotted). Note the logarithmic scale. The inset reveals (for the 12 months risk horizon) that the loss distribution of the non-interacting system (full line) is slightly larger than that of the interacting system (dashes) in the region of small losses.

If one were to consider a *finite* portfolio of M bonds selected from the economy (which is still considered large), one would have normal Gaussian fluctuations around the mean (14) with a variance for the losses per node inversely proportional to the portfolio size [8].

Note that we have been dealing here with "synthetic" parameter distributions for averages of loss distributions, as well as for the bare monthly failure probabilities. These could be replaced by realistic ones without affecting the general set-up.

4. Conclusion

In this paper we have looked to incorporate the risk due to credit contagion into the internal ratings based approach discussed in BASEL II. While the mathematical subtleties are discussed in full detail elsewhere [8], essentially the large number of neighbours assumed for firms means that the law of large numbers and central limit theorems apply to the interactions, meaning that our theory requires only two more parameters than the BASEL II approach. In terms of risk, one of the striking results is that while the effect of interactions is relatively weak in typical economic scenarios, it is pronounced in times of large economic stress, which leads to a significant fattening of the tails of the portfolio loss distribution. This has implications on the fitting of loss distributions to historical data, where care must be taken not only to fit the average behaviour but also to take care with the more extreme events.

We have not considered the question of pricing in the current paper. One interesting facet that immediately drops out of these results is that in terms of expected returns, the effect of contagion does not affect the outlook much (and if variables are renomalised, then arguably not at all). By examining figure 3 we can see that in this scenario if the pricing regime only requires risk tolerance up to the 95th percentile or even the 99th percentile then the loss distribution will not change things much from the case where there is no contagion. However, for a bank or other institution that really does require capital adequacy at the 99.9th percentile, the pricing of risk will change markedly due to the increased capital requirements to insure against these extreme conditions.

The conclusions concerning loss distributions due to contagion are broadly in line with those of previous studies, [4, 5, 6, 7, 8, 9]. Detailed comparison is difficult as setups and underlying assumptions vary significantly between studies. Whereas [5] and to some extent [9] allow credit quality migration, the other studies do not. Both [5] and [7, 8] allow fully heterogeneous economies, whereas in other studies [4, 6] there is a large degree of homogeneity; the small portfolios studied in [9] would hardly allow one to make a distinction between homogeneity and heterogeneity at all. Clearly the very specific properties of the voter model studied in [4] including its extreme version of homogeneity and regularity are crucially responsible for many of the findings in that study, including, e.g., the anomalous scaling of the variance of loss distributions with portfolio size, which is related to the fact that the average fraction of low-liquidity vertices in the model is conserved under its dynamics.

Let us briefly mention the issue of model calibration of the model, which is discussed in much greater detail in [7]. We note that our model requires bare default probabilities and conditional default probabilities as inputs. Historical data, however, only contain interaction-renormalised default probabilities, and thus the problem arises of how to disentangle the two effects. Concerning typical behaviour, Fig. 1 shows that the effect of interactions is fairly small, and interaction-renormalised default probabilities can, to a first approximation within this model, be taken as substitutes for the bare ones. Concerning conditional default probabilities, these would have to be obtained from refined rating procedures as described in [7]. An important lesson to learn from the many particle perspective, however, is to realize that there is no need to get conditional default probabilities for *individual* pairs of companies correct, as *only the low order statistics* of these is needed to describe the collective macroscopic dynamics of the system. Thus the calibration task appears to be much less daunting than expected at first sight. The effect of interactions manifests itself only in situations of economic stress, generating fat tails in portfolio loss distributions.

The model we have proposed is relatively simple in two important respects. Firstly, we do not take into account credit quality migration but have just two states for our firms, solvent or defaulted. The model could be extended to allow for more states for each firm, although the full complexity of non-Markovian dynamics would resurface in an attempt to take credit quality migration along these lines into account. Secondly, the firms and their environment are still fairly homogeneous — local connectivities being on average $c >> 1$, with $\mathcal{O}(\sqrt{c})$ fluctuations — which in practical situations is of course an approximation. This approximation has been made for convenience rather than out of necessity; the techniques described in [8] can be adapted so as to treat situations with more heterogeneity in local environments. We intend to work on some of these possible model generalisations in the future.

One advantage of our simple model is that it is exactly solvable and the solution itself is not overly involved theoretically or computationally, and we only need to introduce two extra parameters to quantify the effect of economic interactions — compared to the BASEL II approach, which ignores credit contagion altogether.

Acknowledgements We thank Peter Neu for useful discussions and helpful remarks. This paper has been prepared by the authors as a commentary on the topic as at September 2006 and is not a definitive analysis of the subject matter covered. It does not constitute advice or guidance in this area and should not be relied upon as such. The paper has not been prepared by J.H. in his capacity as an employee of Hymans Robertson LLP and the views expressed do not represent those of Hymans Robertson LLP. Neither the authors nor Hymans Robertson LLP accept liability for any errors or omissions

Appendix A. Probabilistic solution of our model

In this appendix we show how the law of large numbers and the central limit theorem can be applied to yield a solution to our model. First we note that the complications arise in the dynamic equations (1) due to interactions between the different firms. The effect of other firms in the wider economy on firm i is described by the *local field*: $h_{i,t}(\mathbf{n}) = \sum_{j\neq i} J_{ij} n_{j,t}$. The variables $n_{j,t}$ are correlated, indeed, any variables $n_{j,t}$ in the set with $J_{ij} \neq 0$ will explicitly depend on $n_{i,t'}$ for $t' < t$. Thus the local field on firm i at time t will depend on the state of firm i at times $t' < t$. However, the model is set up so that if firm i is solvent, $n_{i,t} = 0$ and firm i does not effect its neighbours (and hence the correlation described above is not present). Further, if firm i defaults and $n_{i,t} = 1$, then the firms interacting with i will have correlated states due to them all experiencing i's default, but it will no longer have any effect on firm i - the default state is absorbing over the time horizon we consider. Using the definition (5) the local field defined above is given by:

$$h_{it} = \sum_j J_{ij} n_{jt} = \frac{J_0}{c} \sum_j c_{ij} n_{jt} + \frac{J}{\sqrt{c}} \sum_j c_{ij} x_{ij} n_{jt} \ , \tag{A.1}$$

which is a sum of random quantities (with randomness both due to the Gaussian fluctuating forces, and due to the heterogeneity of the environment). The first contribution is a sum of terms of non-vanishing average. By the law of large numbers this sum converges to the sum of averages in the large c limit,

$$h^0_{it} \equiv \frac{J_0}{c} \sum_j c_{ij} n_{jt} \to \frac{J_0}{c} \sum_j \overline{c_{ij} \langle n_{jt} \rangle} \simeq \frac{J_0}{c} \sum_j \overline{c_{ij}}\ \overline{\langle n_{jt} \rangle} = J_0 \frac{1}{N} \sum_j \overline{\langle n_{jt} \rangle}$$

in which angled brackets $\langle \ldots \rangle$ denote an average over the fluctuating forces, and the overbar $\overline{(\ldots)}$ an average over the J_{ij}, i.e., the c_{ij} and the x_{ij}. An approximation is made by assuming negligible correlations between the c_{ij} and the $\langle n_{jt} \rangle$ induced by the heterogeneity of the interactions. The second contribution to (A.1) is a sum of random variables with zero mean, which we have argued are sufficiently weakly correlated for the central limit theorem to apply for describing the statistics of their sum. Thus the sum

$$\delta h_{it} \equiv \frac{J}{\sqrt{c}} \sum_j c_{ij} x_{ij} n_{jt}$$

is a zero-mean Gaussian whose variance follows from

$$\begin{aligned} \overline{\langle (\delta h_{it})^2 \rangle} &= \frac{J^2}{c} \sum_{jk} \overline{c_{ij} c_{ik} x_{ij} x_{ik} \langle n_{jt} n_{kt} \rangle} \simeq \frac{J^2}{c} \sum_{jk} \overline{c_{ij} c_{ik} x_{ij} x_{ik}}\, \overline{\langle n_{jt} n_{kt} \rangle} \\ &= J^2 \frac{1}{N} \sum_j \overline{\langle n_{jt} \rangle} \end{aligned}$$

An approximation based on assuming negligible correlations has been made as for the first contributions. Thus the local field h_{it} is a Gaussian with mean h^0_{it} and variance $\overline{\langle (\delta h_{it})^2 \rangle}$ both scaling with the average fraction of defaulted nodes in the economy. By the law of large numbers this average fraction will be typically realized in a large economy, i.e. we have

$$m_t = \frac{1}{N} \sum_j n_{jt} \to \frac{1}{N} \sum_j \overline{\langle n_{jt} \rangle} \tag{A.2}$$

in the large N limit. The dynamics of the fraction of defaulted nodes then follows from (1),

$$m_{t+1} = \frac{1}{N} \sum_i n_{it+1} = m_t + \frac{1}{N} \sum_i (1 - n_{it}) \Theta\Big(h_{it} - \vartheta_i + \sqrt{\rho(\vartheta_i)} \eta_0 + \sqrt{1 - \rho(\vartheta_i)}\ \xi_{it} \Big)\ . \tag{A.3}$$

The sum in (A.3) is evaluated as a sum of averages over joint n_{it}, h_{it}, and ξ_{it} distribution by the law of large numbers. We exploit the fact that n_{it}, ξ_{it} and h_{it} are uncorrelated.

Noting that the sum $h_{it} + \sqrt{1-\rho(\vartheta_i)}\ \xi_{it}$ is Gaussian with mean $J_0 m_t$ and variance $1-\rho(\vartheta_i) + J^2 m_t$, and taking into account that n_{it}-averages, depend on i through ϑ_i, $\langle n_{it}\rangle = \langle n_t\rangle_{(\vartheta_i)}$, we find

$$m_{t+1} = m_t + \frac{1}{N}\sum_i (1-\langle n_t\rangle_{(\vartheta_i)})\Phi\left(\frac{J_0 m_t + \sqrt{\rho(\vartheta_i)}\ \eta_0 - \vartheta_i}{\sqrt{1-\rho(\vartheta_i)+J^2 m_t}}\right)$$

This version can be understood as an average over the ϑ distribution (i.e. the default probability distribution)

$$p(\vartheta) = \frac{1}{N}\sum_i \delta(\vartheta - \vartheta_i)\ ,$$

which maps onto a distribution of unconditional default probabilities as discussed above. Denoting that average by $\langle \ldots \rangle_\vartheta$ we finally get the following evolution equation for the macroscopic fraction of defaulted companies in the economy

$$m_{t+1} = m_t + \left\langle (1-\langle n_t\rangle_{(\vartheta)})\Phi\left(\frac{J_0 m_t + \sqrt{\rho(\vartheta)}\ \eta_0 - \vartheta}{\sqrt{1-\rho(\vartheta)+J^2 m_t}}\right)\right\rangle_\vartheta \qquad \text{(A.4)}$$

which is just the stated result (12). We have thus an explicit dynamic equation for the macroscopic fraction of defaulted nodes in the economy. It involves first propagating ϑ-dependent default probabilities via

$$\langle n_{t+1}\rangle_{(\vartheta)} = \langle n_t\rangle_{(\vartheta)} + (1-\langle n_t\rangle_{(\vartheta)})\Phi\left(\frac{J_0 m_t + \sqrt{\rho(\vartheta)}\ \eta_0 - \vartheta}{\sqrt{1-\rho(\vartheta)+J^2 m_t}}\right)\ , \qquad \text{(A.5)}$$

which depends only on m_t, thereafter performing an integral over the ϑ distribution to obtain the updated fraction m_{t+1} of defaulted nodes given in (A.4). Note that all the assumptions and approximations used in this appendix can be fully justified by using path integral techniques as shown in [8].

References

[1] M Davis and V Lo 2001, Quant. Finance **1**, 382-387
[2] R Jarrow and F Yu 2001, J. Finance **56**, 1765-1799
[3] E Rogge and P Schönbucher, *Modelling Dynamic Portfolio Credit Risk*, Working Paper, Imperial College, London, ABN Amro Bank, London, and ETH, Zürich (February 2003)
[4] K Giesecke and S Weber 2006, J. Econ. Dyn. Control **30**, 741-767
[5] D Egloff, M Leippold, and P Vanini 2004, EFA Maastricht Meetings Paper No. 1142
[6] R Frey and J Backhaus, *Interacting defaults and counterparty risk; a Markovian approach*, Working Paper, University of Leipzig (July 2003).
[7] P Neu and R Kühn 2004, Physica A **342**, 639-655
[8] J P L Hatchett and R Kühn 2006, J. Phys. A **39**, 2231-2251
[9] H Kraft and M Steffensen 2006, *Bancruptcy, Counterparty Risk, and Contagion*, preprint
[10] T M Liggett 1999, *Stochastic Interacting Systems: Contact, Voter and Exclusion Processes* Springer-Verlag: Berlin
[11] T K Siu, W K Ching, E S Fung and M K Ng 2005, Quantitative Finance **5**, 543-556
[12] J P Morgan Global Research 1997 CreditMetrics^{TM}: The Benchmark for Understanding Credit Risk *Technical Document* (New York) (www.creditmetrics.com)

[13] Credit Suisse First Boston 1997 Credit Risk+: A Credit Risk Management Framework *Technical Document* (New York) (www.csfb.com/creditrisk)
[14] M B Gordy 2000, J. Bank. Finance **24**, 119-149
[15] R Merton 1974, J. Finance **29**, 449-470
[16] F Yu 2007, Math. Finance **17**, 155-173
[17] S R Das, D Duffie, N Kapadia, and L Saita 2007 Journal of Finance **62**, 93-117.
[18] J P L Hatchett, B Wemmenhove, I Pérez Castillo, T Nikoletopoulos, N S Skantzos, and A C C Coolen 2004, J. Phys. A **37**, 6201-6220; I Pérez Castillo, B Wemmenhove, J P L Hatchett, A C C Coolen, N S Skantzos, and T. Nikoletopoulos 2004 J. Phys. A **37**, 8789-8799; T Nikoletopoulos, A C C Coolen, I Pérez Castillo, N S Skantzos, J P L Hatchett and B Wemmenhove 2004, J. Phys. A **37**, 6455-6475
[19] Basel Committee on Banking Supervision 2005 *International Convergence of Capital Measurement and Capital Standards: A Revised Framework* (Basel) (http://www.bis.org) Section III: Credit Risk - The Internal Ratings Based-Approach
[20] R Kühn and P Neu 2003, Physica A **322**, 650-666