Contagion in the Interbank Market with Stochastic Loss Given Default*

Christoph Memmel,[a] Angelika Sachs,[b] and Ingrid Stein[a]
[a]Deutsche Bundesbank
[b]LMU Munich

This paper investigates contagion in the German interbank market under the assumption of a stochastic loss given default (LGD). We combine a unique data set about the LGD of interbank loans with detailed data about interbank exposures. We find that the frequency distribution of the LGD is markedly U-shaped. Our simulations show that contagion in the German interbank market may happen. For the point in time under consideration, the assumption of a stochastic LGD leads on average to a more fragile banking system than under the assumption of a constant LGD.

JEL Codes: D53, E47, G21.

1. Introduction

The collapse of Lehman Brothers turned the 2007/2008 turmoil into a deep global financial crisis. But even before the Lehman default, interbank markets ceased to function properly. In particular, the

*The views expressed in this paper are those of the authors and do not necessarily reflect the opinions of the Deutsche Bundesbank. We thank Gabriel Frahm, Gerhard Illing, Ulrich Krüger, Peter Raupach, Sebastian Watzka, two anonymous referees, and the participants at the Annual Meeting of the European Economic Association 2011, the Annual Meeting of the German Finance Association 2011, the 1st Conference of the MaRs Network of the ESCB, the FSC workshop on stress testing and network analysis, as well as the research seminars of the Deutsche Bundesbank and the University of Munich for valuable comments. Author contact: Memmel and Stein: Deutsche Bundesbank, Wilhelm-Epstein-Strasse 14, D-60431 Frankfurt, Germany; Tel: +49 (0) 69 9566 8531 (Memmel), +49 (0) 69 9566 8348 (Stein). Sachs: LMU Munich, Department of Economics, Ludwigstrasse 28 (Rgb.), D-80539 Munich, Germany; Tel: +49 (0) 89 2180 5673. Author e-mails: christoph.memmel@bundesbank.de, angelika.sachs@lrz.uni-muenchen.de, ingrid.stein@bundesbank.de.

fear of contagion via interbank markets played a crucial role. While banks could gauge their direct losses from exposure to so-called toxic assets, they could not assess their counterparties' losses and creditworthiness and were therefore not willing to lend money to other banks, causing the breakdown of interbank markets. This led to an unprecedented liquidity extension of central banks and government rescue packages (see Stolz and Wedow 2010) which, however, could not avoid deep recessions in many countries of the world. From an economic perspective, it is therefore essential to have a tool allowing to assess potential contagion risks via interbank markets.

Creating such a tool is the aim of this paper. We study contagion in the German interbank market, one of the largest interbank markets in Europe. We carry out a simulation exercise where we assume that a certain bank fails and examine how this failure affects other banks' solvency via direct effects and chain reactions in the banking system. Throughout this paper, our focus is on fourteen large and internationally active German banks and the sectors of savings and cooperative banks.

In this paper, we only investigate the direct, mechanic contagion effects in the interbank market, which means that we analyze the direct (on- and off-balance-sheet) exposure between the banks. What we do not consider are effects like a general loss in confidence among banks, which could lead to a drying up of the interbank market and thereby to a liquidity shortage; contagion due to market perception, i.e., that all banks with a similar business model are subject to distrust when such a bank runs into distress; or herding behavior, where massive sales can drive the price of an asset below its fundamental value and banks using fair-value accounting have to adjust asset values. Hence, our analysis covers only part of the possible contagion effects. However, this analysis is relatively precise because it is based on hard data and not so much on estimated economic relationships.

We investigate in particular the role of the loss given default (LGD) in the contagion process, which is a key factor for the extent of contagion in analyses like this. The LGD, multiplied by the total exposure of a creditor bank to a debtor bank, gives the actual loss of the creditor bank in the event that the debtor bank fails. The LGD can vary between 0 percent (e.g., in the event that the defaulted loan is fully collateralized) and 100 percent (which is equivalent to a zero recovery rate of the defaulted loan). As there is usually only sparse information about recovery rates in the case of bank defaults, the standard approach in the literature on interbank contagion is to assume a fixed value of the LGD and repeat the simulation exercise with different values of this LGD. The literature generally finds that losses in the total banking system crucially depend on the LGD value. Below a certain threshold of LGD, potential losses are minor. However, as soon as the LGD exceeds a certain threshold, there are considerable risks of large parts of the banking system being affected and heavy losses in the banking system occurring (see, e.g., Upper and Worms 2004 and van Lelyveld and Liedorp 2006). Therefore, the standard approach has the considerable drawback that an assessment of contagion risks in the real world is difficult and associated with great uncertainties. In our paper, we overcome this shortcoming by using a unique data set of realized LGDs of defaulted interbank exposures.

Our contributions to the literature are as follows. First, by using this data set of realized LGDs on the interbank market, we are able to investigate the empirical patterns of actual LGDs of bank loans. Second, unlike the vast majority of papers in the literature, we dispose of detailed data about the pairwise interbank exposures and do not need to estimate them. Instead, we are able to precisely quantify interbank exposures (including off-balance-sheet and derivative positions) within the national market. Third, in contrast to most papers in the literature, we conduct the simulation exercise with a *stochastic* LGD derived from the observed distribution of LGDs (instead of a stepwise increase of *constant* values). We thereby obtain a distribution of the number of contagious bank defaults which allows a more realistic assessment of contagion risks.

Our main findings are, first of all, that LGDs follow a markedly U-shaped distribution, which can be reasonably well approximated by a beta distribution. Second of all, using the precise information about interbank exposures and the distribution of LGD, we find that contagion in the German interbank market may happen. Third, for the point in time under consideration, we find that the number of bank defaults increases on average when we assume a stochastic LGD instead of a constant one.

This paper is structured in the following way: In section 2, we give a brief overview of the literature on interbank contagion as well as LGD modeling and state our contribution to the literature. Section 3 deals with the description of the contagion exercise and the structure of the interbank network. Section 4 summarizes how we model the LGD. In section 5, we show the results of the contagion exercise, and in section 6 the conclusion is presented.

2. Literature

Our paper relates to three strands of the literature. The first strand is about empirical simulation studies of interbank contagion (see Upper 2011 for an overview). Especially national European interbank markets have been the focus of empirical studies (see, for instance, van Lelyveld and Liedorp 2006 for the Netherlands, Sheldon and Maurer 1998 for Switzerland, or Mistrulli 2011 for Italy). In addition to studies based on national interbank markets, there are cross-border contagion simulations. These studies are either based on Bank for International Settlements data on consolidated banking statistics (see Degryse, Elahi, and Penas 2010 and Espinosa-Vega and Solé 2010) or analyze international sector interlinkages (see Castrén and Kavonius 2009). Most papers in this strand do not have direct access to information on interbank exposures but either apply statistical methods to derive the bilateral exposures or rely on data which cover only part of the interbank exposures. We have a certain advantage compared with these studies since we are able to precisely quantify the amount of bilateral exposures for a system of fourteen large and internationally active German banks as well as the sectors of the savings and the cooperative banks. Our data set is based on the German credit register and includes off-balance-sheet and derivative positions. It contains all bilateral exposures of the fourteen banks and two sectors above a threshold of EUR 1.5 m. This threshold is not relevant for the purpose of our study since interbank exposures are typically large.

The second strand of literature we contribute to deals with extensions of the usual contagion exercises. Cifuentes, Ferrucci, and Shin (2005) introduce additional stress due to declining asset prices as a result of fire sales; Elsinger, Lehar, and Summer (2006) integrate the interbank contagion model in a stress-testing setting that includes macroeconomic shocks. Chan-Lau (2010) and Espinosa-Vega and Solé (2010) consider not only credit risk but funding risk as well.

They argue that the banks' funding is hindered when the interbank market does not function properly. Aikman et al. (2009) incorporate various of these aspects into one quantitative model of systemic stability. Degryse and Nguyen (2007) explicitly model the LGDs, deriving them endogenously from the banks' balance sheet composition. Our extension, too, is about LGD modeling. However, we model the LGDs as stochastic.

The third strand of literature deals with the distribution of LGDs. Tarashev and Zhu (2008) and Huang, Zhou, and Zhu (2009) choose a stochastic setting for the LGD. They assume a triangular distribution with the probability mass concentrated in the center of the distribution (more precisely at 55 percent and 50 percent, respectively). Crouhy, Galai, and Mark (2000) model a stochastic LGD with the help of a beta distribution. They estimate the parameters by using bond market data. Their estimations yield the result that the LGD follows a unimodal beta distribution. Our contribution consists in estimating the distribution of the LGDs of *interbank exposures*. We have a unique data set of realized interbank LGDs at our disposal. This data suggests a markedly U-shaped density for the LGD, i.e., a distribution with a vast probability mass at zero and 100 percent. This finding is in line with Dermine and de Carvalho (2006) and Bastos (2010), who use a data set of defaulted loans provided by a large Portuguese bank and find a U-shaped LGD distribution for non-financial firms.

3. Round-by-Round Algorithm

In the event that a bank fails, the banks that have granted loans to this bank suffer losses from their exposures. The contagion process in the interbank market may stop after the first round but may also propagate further through the system. Banks that fell into distress as a consequence of the initial distress may now themselves become a source of contagion. This process will continue round by round until no new banks are affected (possibly leading to a large number of failures in total) or until the supervisory authorities manage to put an end to this process.

In this section we describe a simulation exercise so as to study the extent to which the German banking system may be prone to such a contagious process. We apply the round-by-round algorithm as described in Upper (2011):

(i) Initially, bank i fails exogenously.
(ii) As a result, banks whose exposure to bank i multiplied by the loss given default (LGD) exceeds their buffer of tier 1 capital also fail. We define a bank to be in default in the event that its tier 1 capital ratio is below 6 percent of its risk-weighted assets. This default definition is in line with the new Basel Accord, where the minimum capital requirement is also set at 6 percent.[1] We do not take into account potential reactions of the lender banks. For example, the lender banks may have hidden reserves which they release to raise their tier 1 capital. Instead, we assume that write-offs on interbank loans decrease the lender's tier 1 capital by the same amount.
(iii) Additional banks may fail if their combined exposure to the banks that have failed so far (times the LGD) is greater than their capital buffer.
(iv) The contagious process stops when there is a round with no new failures.

Thus, bank j is in distress if

$$\frac{E_j - \sum_k (LGD_{jk} \cdot x_{jk} \cdot 1_{k \in D})}{RWA_j - 0.2 \cdot \sum_k (x_{jk} \cdot 1_{k \in D})} < 0.06. \tag{1}$$

In this context, E_j is the tier 1 capital of bank j, x_{jk} is the exposure of bank j to bank k, $1_{k \in D}$ is an indicator variable that takes on the value 1 in the event that bank k is in distress (and 0 otherwise), LGD_{jk} is the loss given default associated with the exposure of bank j to bank k, and RWA_j are the risk-weighted assets of bank j. We assume that interbank claims receive a weight of 0.2 in banks' risk-weighted assets.[2] When calculating the tier 1 capital ratio, we also take into account that every claim to a bank that failed completely disappears from the creditor bank's risk-weighted assets.

[1] See Basel Committee on Banking Supervision (2010, paragraph 50).

[2] The risk weight of 0.2 stems from the Basel I and Basel II framework applied to German banks.

One can argue that a bank can serve its debt unless the capital of this bank is totally consumed and becomes negative. However, we use the stricter criterion of a minimum level of 6 percent tier 1 capital. We do this for the following reasons: (i) When a bank is shut down and liquidated (because, for instance, it no longer meets the minimum regulatory capital requirements), it is questionable whether one receives the bank assets' book value, especially because the book value of its illiquid positions may be far higher than the proceeds from a hasty fire sale. (ii) A bank with a sharp drop in its capital ratio will no longer have any access to short-term funding at sustainable rates. Soon afterwards, the unfavorable funding conditions will have consumed what is left of the capital.

We carry out this simulation exercise for each of fourteen large and internationally active banks in Germany (the biggest private commercial banks and the central institutes of the savings and cooperative banks) and the two sectors of the savings and cooperative banks, which we treat as aggregate sectors. Thus, we consider sixteen units in total which potentially have bilateral exposures to each other. We treat the savings and cooperative banks in an aggregate way because single banks that belong to this group are usually very small.[3] Hence, it is almost sure that the default of one of these small banks would not trigger contagious reactions. If, however, the whole sector of savings or cooperative banks were to be hit by an aggregate shock (which is not completely unlikely because of similar business models), a contagious effect on the rest of the banking system is quite possible. To sum up, we cover about 67 percent of the total assets of the German banking system in our paper.

To run the round-by-round algorithm, information is needed on (i) the pairwise exposures between the banks and (ii) the appropriate loss, given a bank fails. Concerning the pairwise exposures, we have detailed information on exposures within the German interbank market. This leaves the question of determining the loss given default. From the literature, we know that this is crucial for the contagion exercises (see, e.g., Upper and Worms 2004). Different solutions are possible:

[3]Craig and von Peter (2010) show that only a small number of banks form the so-called core of the German interbank market and that these core banks act as an intermediary for numerous small banks (like savings and cooperative banks).

(i) *Constant LGD.* The loss given default is exogenously set to a constant value, say 40 percent or 45 percent.[4] To account for the fact that the LGD crucially drives the results, one can vary the constant loss given default over a wide range of values. The contagion exercise is then run for each different value of the LGD.

(ii) *Endogenous LGD.* If information on the actual over-indebtedness of the distressed bank, the bankruptcy cost, and the degree of collateralization were available, it would be possible to endogenously calculate the loss given default.

(iii) *Stochastic LGD.* Our supervisory data concerning the write-offs of interbank loans show that the loss given default varies considerably, with a large portion of the probability mass at 0 percent and at 100 percent. A possible explanation for this quasi-dichotomy may be that the loans are either fully collateralized (as in the repo market) or completely unsecured. This finding is not in line with the assumption of a constant LGD (solution (i)). Solution (i) would, rather, be in line with a distribution of the LGDs concentrated in one point.

In this study, we use the third solution. In contrast to the existing literature that exogenously assumes some constant LGD value, we have a unique data set of actually realized LGDs on the interbank market. This data set provides an empirical frequency distribution of LGDs. The exact properties of the LGD distribution are investigated in section 4.

As outlined above, the first step for running the round-by-round algorithm consists of establishing the matrix of mutual interbank exposures. We use Bundesbank data from the German credit register (MiMiK) to obtain the necessary information.[5] Unlike credit registers in most other countries, the German credit register also includes interbank loans and is not confined to non-financials. This database offers us a certain data advantage compared with other studies since we are able to determine the complete matrix for the

[4]Kaufman (1994) gives an overview of loss given default estimates for bank failures; the estimates vary considerably. James (1991) finds that the average loss of failed U.S. banks during the period of 1985 to 1988 was about 30 percent. In addition, there were direct costs associated with the bank closures of 10 percent of the assets. In our data set, the mean LGD is about 45 percent.

[5]See Schmieder (2006) for more details about this database.

banks. By contrast, balance sheet data only show (for each bank) the aggregate amount lent to or borrowed from all banks. Moreover, payment data or large exposure data are generally less comprehensive than credit register data and include—for example, in the case of payment data—information about short-term lending only.

The German credit register contains quarterly data on large exposures of banks to individual borrowers or single borrower units (e.g., groups). Banking institutions located in Germany are required to report if their exposures to an individual borrower or the sum of exposures to borrowers belonging to one borrower unit exceeds the threshold of EUR 1.5 m at least once in the respective quarter. We think that the threshold of EUR 1.5 m does not cause a serious bias since the typical interbank loan is relatively large and exceeds the threshold of EUR 1.5 m.

The credit register applies a broad definition of a loan. Loans in this sense include traditional loans, bonds, off-balance-sheet positions, and exposures from derivative positions. However, trading-book positions are excluded. We start by analyzing gross on- and off-balance-sheet exposures as a benchmark case. In section 5.2, we run simulations considering on-balance-sheet exposures only. Furthermore, we investigate the case of netting. It is, however, by far not clear whether netting can be enforced in case of a bank failure.[6]

For the simulation exercise, we use data from the fourth quarter of 2010. The resulting matrix of interbank exposures gives some interesting insight into the German interbank market. As we consider fourteen large and internationally active German banks as well as the savings and cooperative sector, we obtain a 16×16 matrix of interbank exposures. Not surprisingly, we almost have a complete network; i.e., most of the off-diagonal elements of this matrix are non-zero. To be precise, only two off-diagonal elements are zero; i.e., only 2 of the 240 possible interbank relations do not exist.

Table 1 shows some summary statistics of the interbank network we consider in this paper. To capture the inequality of how banks spread their interbank assets/liabilities among their counterparties and thus to evaluate how banks differ in terms of their connections to other banks, we calculated for each bank the normalized Herfindahl-Hirschman Index (HHI) of the share of single interbank

[6]See Mistrulli (2011) for this and other arguments concerning the simulation method.

Table 1. Summary Statistics of the Interbank Network

	p25	**Median**	**p75**	***N***
HHI(AIB)	0.05	0.10	0.28	16
HHI(LIB)	0.11	0.26	0.38	16
AIB/E	0.04	0.12	0.24	240
LIB/E	0.05	0.10	0.26	240
off-bs(AIB)	0.06	0.09	0.24	16
off-bs(LIB)	0.03	0.05	0.10	16

Note: HHI = normalized Herfindahl-Hirschman Index; HHI(AIB)/HHI(LIB) = bank-specific normalized HHI of the share of single interbank exposures to total exposures on the asset/liability side of the bank's balance sheet; AIB/E (LIB/E) = ratio of single interbank exposures on the asset (liability) side of the bank's balance sheet to its tier 1 capital; off-bs(AIB)/off-bs(LIB) = bank-specific ratio of off-balance-sheet exposures to total exposures on the asset/liability side of the bank's balance sheet; p25 = 0.25 percentile; p75 = 0.75 percentile; N = number of observations.

assets/liabilities to the bank's total interbank assets/liabilities (that are included in our analysis). The maximum HHI of 1 would indicate that a bank has all its interbank assets/liabilities towards one single counterparty. The minimum HHI of 0 would indicate that a bank spreads its interbank assets/liabilities as equally as possible. The large difference between the 25 percent quantile (with an HHI of 0.05 and 0.11, respectively) and the 75 percent quantile (with an HHI of 0.28 and 0.38, respectively) implies that results differ substantially among banks. The reason is that the central institutions of the savings and cooperative banks concentrate their exposures on the savings and cooperative sector, respectively; large private banks, however, do not. It is also remarkable that banks in our network tend to spread their interbank assets more equally than their interbank liabilities. The relative size of interbank exposures (i.e., interbank assets/liabilities over tier 1 capital) is already quite large at the 25 percent quantile, with 4 percent and 5 percent, respectively. The median of the relative size is at 12 percent and 10 percent, respectively. This confirms that interbank exposures are a considerable source of contagion.[7]

[7]Generally, a single loan must not exceed 25 percent of a bank's liable capital. Exceptions are, however, exposures between banks within the associations of savings and cooperative banks, respectively.

4. Loss Given Default (LGD)

As stated in the previous section, another key component for the contagion exercise is the loss given default (LGD). We have some information about the loss rate banks face in the event that a debtor bank defaults. More precisely, our LGD data are assigned to the respective lender bank and not—as is usually done—to the debtor bank. Although we do not know the LGD of the lender bank for a default of a *specific* debtor bank, we know for each lender bank the average LGD of interbank exposures (at an annual frequency). We have data on the volume of non-performing interbank loans and on the corresponding write-downs. For each bank and each year, two figures are provided: the amount (in euro) of interbank loans for which provisions have been made and the amount of these provisions. We interpret the ratio of these two figures as a realization of the stochastic LGD of a single interbank relationship, not of an average of two or more LGDs. It is important to have realizations of LGDs of single interbank relationships because realizations of average LGDs tend to be biased towards unimodal distributions; the average of, let's say, fifty LGDs is by virtue of the central limit theorem approximately normally distributed, even if the distribution of single LGDs is markedly U-shaped.[8] The data are taken from the quantitative supervisory reports collected by the Bundesbank on banks in Germany.[9] Based on this data, we can estimate the distribution of LGDs.[10]

Looking at figure 1, we see that the empirical distribution of the LGDs is markedly U-shaped. This characteristic and the nature of the LGDs, especially its range between 0 and 1, suggest modeling the LGD distribution with the beta distribution. Figure 1 also displays the probability density function of a beta distribution with the estimated parameters. Compared with the empirical frequency distribution, only small deviations can be observed. Statistical tests confirm this observation. The null hypothesis of a χ^2 goodness-of-fit

[8] We will discuss this topic in more detail later in this section.

[9] For more details on these data, see Memmel and Stein (2008).

[10] Note that the LGD data and the exposure data are not fully compatible: Whereas the LGD data refers to unconsolidated accounts and includes both the trading and the banking book, the exposure data refers to consolidated accounts and does not include the trading book. We believe, however, that this lack of compatibility does not call in question the use of the data.

Figure 1. Relative Frequency of the Loss Given Default for Interbank Loans

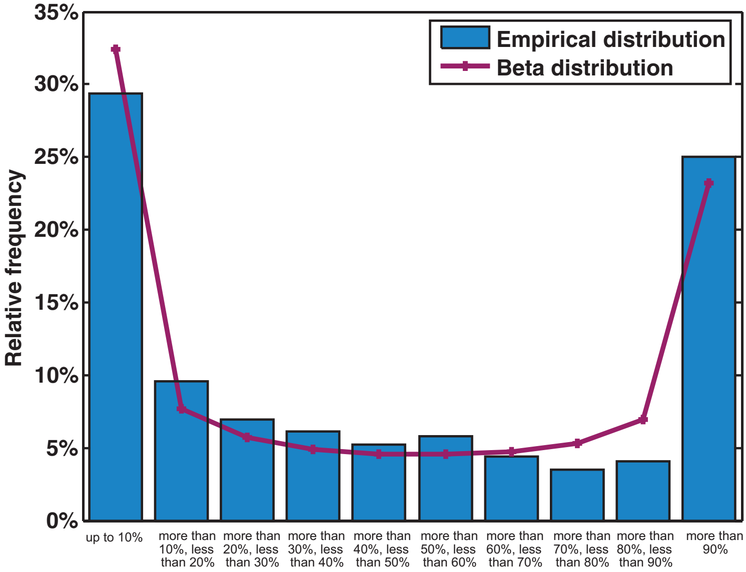

Note: Relative frequency was derived from data on German private commercial banks and the central institutions of the savings and cooperative banks. There were 344 observations for the period 1990–2008.

test on whether our data follow a beta distribution with the estimated parameters $\hat{\alpha}$ and $\hat{\beta}$ cannot be rejected on a 5 percent significance level. Choosing ten equidistant intervals and comparing the observed frequency with the expected frequency within the intervals yields a p-value of ≈ 0.075.[11]

[11]The result of this test gives strong evidence that the assumed distribution is very close to the observed distribution, as the test is very sensitive due to the large number of observations. To illustrate the correlation between the number of observations and the sensitiveness of the test, we run simulations with a sample randomly drawn from a beta distribution. Drawing 344 observations from a beta(0.28,0.35) distribution 10,000 times and testing each sample against a beta(0.18,0.25) distribution yields a probability of making a type II error (i.e., the error of falsely accepting the null hypothesis) of around 18 percent. Repeating this exercise for only half of the sample (i.e., drawing 172 observations each time) leads to a probability of making a type II error of 62 percent. Thus, the larger the sample, the more sensitive the test becomes to only small deviations from the distribution tested.

Therefore, we use the beta distribution for further analysis. The density of the beta distribution is given by

$$f(x) = \frac{1}{B(\alpha, \beta)} x^{\alpha-1}(1-x)^{\beta-1} \quad x \in (0,1) \tag{2}$$

with

$$B(\alpha, \beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}, \tag{3}$$

where $\Gamma(\cdot)$ is the gamma function. The parameters $\alpha > 0$ and $\beta > 0$ determine the shape of this distribution.[12] The beta distribution is especially suited to model the LGD because (i) the domain is confined to the economic sensible interval from 0 to 1, (ii) it is highly flexible, and (iii) it nests other distributions.[13] For instance, when both parameters equal 1, then the beta distribution becomes a uniform distribution. When both of the parameters are smaller than 1, the probability density function is U-shaped, with a large portion of the probability mass close to 0 and 1. For parameter values close to 0, this distribution converges to the binomial distribution. By contrast, the density is unimodal in the case of both parameters α and β being greater than 1. For very large parameter values, it converges to the degenerate distribution, where the entire probability mass is concentrated on one point. The expectation and the variance of a random variable X following a beta distribution are functions of the parameters α and β:

$$E(X) =: \mu = \frac{\alpha}{\alpha+\beta} \tag{4}$$

and

$$var(X) =: \sigma^2 = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}. \tag{5}$$

[12]Figure 2 summarizes the possible shapes of the probability density function dependent on the parameter values.

[13]See, e.g., Hahn and Shapiro (1967, p. 91).

Given estimates for the expectation and the variance, estimators for the parameters α and β are obtained by solving the equations (4) and (5) for α and β, respectively:[14]

$$\hat{\alpha} = \hat{\mu}\left(\frac{\hat{\mu}(1-\hat{\mu})}{\hat{\sigma}^2} - 1\right) \tag{6}$$

$$\hat{\beta} = (1-\hat{\mu})\left(\frac{\hat{\mu}(1-\hat{\mu})}{\hat{\sigma}^2} - 1\right). \tag{7}$$

We calculate the sample mean and variance of the distribution of the LGD for the whole sample and different sub-samples and then estimate the parameters α and β (see tables 2 and 3). It is noteworthy that the average LGD of savings banks being the creditors (58 percent) is well above the average LGD of the total sample (38 percent). Cooperative banks, whose business model is comparable, however, suffer only from a rather low LGD of 24 percent on average in the event they incur losses on the interbank market. The average LGD incurred by large and internationally active banks is in between (45 percent). Furthermore, we see that the LGD tends to be higher the larger the lender bank (measured as the lender bank's total assets). Irrespective of the sub-sample under consideration, we observe a U-shaped distribution. We explicitly test the null hypothesis that the beta distribution is not U-shaped, i.e., that $\alpha \geq 1$ or $\beta \geq 1$ (see figure 2). We do this by applying the delta method.[15] The result is that we can reject the null hypothesis on a 1 percent and 5 percent significance level, respectively, in all cases. Thus, we can conclude that, irrespective of the banking group and size of the lender banks, we can assume a U-shaped distribution of the LGD.[16]

As our analysis focuses mostly on large and internationally active banks in Germany, it would be obvious to use this sub-sample to estimate the parameters of the LGD distribution. This would give us 101

[14] This procedure is called the method of matching moments; see, e.g., Hahn and Shapiro (1967, p. 95). We do not use maximum-likelihood estimation because there is a considerable amount of observations which equal exactly 0 and 1 and for which, therefore, the likelihood function is not defined.

[15] For details on the delta method, see the appendix.

[16] In the literature, however, the LGD is often modeled by using a unimodal distribution (which implies that $\alpha > 1$ and $\beta > 1$) or as a constant. Hence, these results may also have further implications for this literature.

Table 2. Mean and Standard Deviation of the Empirical Frequency Distribution of the LGD and Estimated Parameters of the Respective Beta Distribution Dependent on Different Samples of Lender Banks

		Loss Given Default		Beta Distribution	
Sample	***N***	**Mean**	**Standard Dev.**	α	β
All Banks	667	0.38	0.39	0.20***	0.33***
Sample Used in This Paper	344	0.45	0.39	0.28***	0.35***
Large and Internationally Active Banks	101	0.45	0.32	0.62***	0.76**
Savings Banks	50	0.58	0.38	0.42***	0.30***
Cooperative Banks	222	0.24	0.37	0.08***	0.24***

Note: N = number of observations; "Large and Internationally Active" includes large private commercial banks and central institutions of the savings and cooperative banks, "Sample Used in This Paper" includes large and internationally active banks and all private commercial banks; **/*** means significantly < 1 on the 5%/1% level.

Table 3. Mean and Standard Deviation of the Empirical Frequency Distribution of the LGD and Estimated Parameters of the Respective Beta Distribution Dependent on the Lender Banks' Size (lender banks' balance sheet total)

	Loss Given Default		Beta Distribution	
Size Group	**Mean**	**Standard Dev.**	α	β
Smallest 20%	0.26	0.40	0.05***	0.14***
2nd Quintile	0.35	0.41	0.12***	0.22***
3rd Quintile	0.38	0.40	0.17***	0.28***
4th Quintile	0.48	0.38	0.33***	0.36***
Largest 20%	0.42	0.33	0.53***	0.73**

Note: Whole sample of 667 observations; **/*** means significantly < 1 on the 5%/1% level.

Figure 2. Shapes of the Probability Density Function of the Beta Distribution Dependent on the Values of α and β

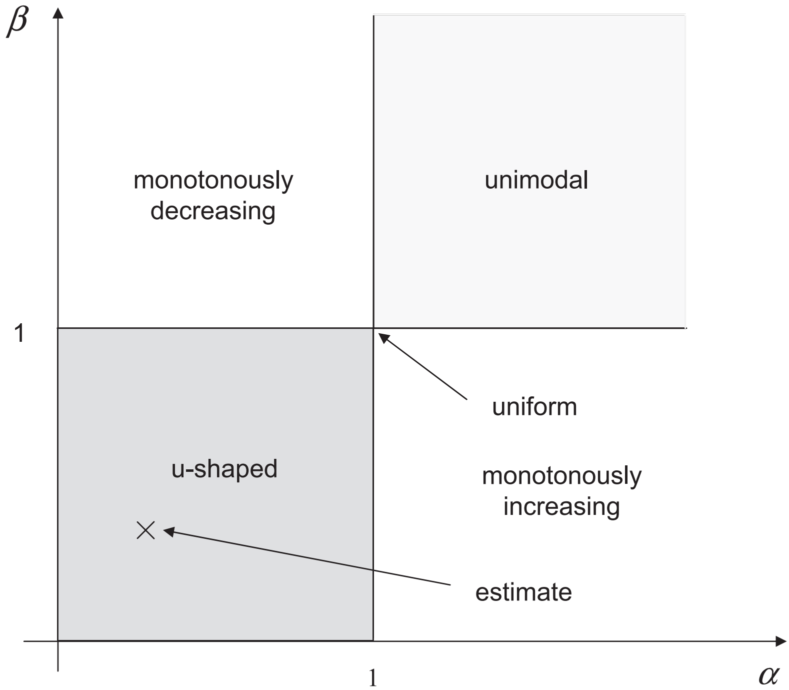

observations of realized LGDs. However, the problem is that these LGDs are probably due not only to a single credit relationship but to several credit relationships in distress. If we assume that the number of defaulted credit relationships per bank and year follows a Poisson distribution, we obtain for the case that exactly one debtor bank defaulted, given the bank reports non-zero write-downs, a probability of only 71 percent. Thus, we have the problem that the reported LGD value of that bank is often (in about 29 percent of the cases) just an average of several LGD values, and the LGD distribution is therefore biased towards unimodal distributions.[17]

[17]The higher standard deviation of the sample used for the contagion analysis compared with the sample of large and internationally active banks confirms our assumption; see table 2.

We can mitigate this problem by including additional banks in our sample for which it is reasonable to assume that exactly one credit relationship is in distress, given the bank reports non-zero write-downs. Thus, we include all private commercial banks in our sample, which yields a probability that exactly one single credit relationship defaulted (given the bank reported non-zero write-downs) of 93 percent. Regional savings banks and cooperative banks, which are generally small and medium-sized, are not included in our sample. The reason is that we consider their position in the German interbank market as less representative for our stability analysis because these banks' interbank market activities are very much characterized by relationships to their central institutes. This is not the case for the smaller private banks. In addition, the mean LGD, which is not affected by the aforementioned problem, is quite similar (around 45 percent) in the sample we chose and in the sample of the large and internationally active banks (see table 2). We therefore believe that our sample is a balanced compromise between statistical properties (a high share of single default events) and economic fit (similarity of the banks in the contagion exercise and estimation of the LGD distribution).

As our LGD data are applied to situations of severe stress in the interbank market, an important point to investigate is how LGDs change in crisis time compared with normal times. Our data enable this, as they contain the period from 1998 to 2008 and thus include the crisis year 2008. By comparing the mean LGD of the pre-crisis years (1998–2007) to the crisis year 2008, we obtain the surprising result that the LGDs in 2008 are, on average, lower compared with the period before. For the sample including all banks in 2008, we obtain 251 observations with a mean LGD of 0.27 (compared with 0.44 in pre-crisis years). The sub-sample of all private commercial banks and the central institutions of the savings and cooperative banks (twenty-eight observations) yields a mean LGD of only 0.22 (compared with 0.47 in pre-crisis years). A possible explanation for this fact is that banks become more cautious in times of stress and, e.g., demand more collateral for interbank lending. Thus, we conclude that the potential rise of LGDs in times of severe stress (e.g., due to reduced asset values in banks' balance sheets) is counteracted by more precautious lending by banks.

Our final sample of LGD observations consists of 344 observations in the period from 1998 to 2008. Figure 1 shows the frequency distribution of the LGDs. Using the sample mean $\hat{\mu}$ and variance $\hat{\sigma}^2$ as an estimator for the population mean and variance, we obtain $\hat{\mu} = 0.45$ and $\hat{\sigma}^2 = 0.15$. Inserting $\hat{\mu}$ and $\hat{\sigma}^2$ into equation (6) and (7) yields $\hat{\alpha} = 0.28$ and $\hat{\beta} = 0.35$. These parameter values indicate a U-shaped distribution (see figure 2).

As stated above, LGDs of banks can in theory be derived endogenously from their balance sheet composition. However, we do not apply this solution because we would have to make a lot of additional assumptions in our contagion exercise.[18] For example, we would have to make assumptions about who has to bear the losses that arise from the bank failures. A standard assumption in a case like this is that losses are distributed pro rata among creditors (see the clearing algorithm of Eisenberg and Noe 2001), which is definitely a strong assumption. For instance, we find that the mean LGD for totally unsecured interbank exposures is 64 percent, whereas it is only 24 percent for those that are at least partly collateralized. Additionally, it would be necessary to model losses due to fire sales of assets of distressed banks. A detailed contagion analysis with an endogenous LGD is thus not feasible since we lack the necessary data. Besides, our data on realized LGDs suggest that the borrower banks' balance sheet composition and other bank specific variables only explain a small fraction of the LGD variation. We carried out a variance decomposition of the LGDs and we find that most of the variation is due to the lender bank and due to the nature of the relationship; i.e., the variation owing to the balance sheet composition of the borrower bank is less important. Furthermore, endogenizing the LGD disregards the time dimension. Upper (2011) cites the default of Bankhaus Herstatt as an example for the observation that the LGD varies across the time horizon; i.e., the LGD decreases when the recovery horizon becomes longer. This observation is backed up by Bastos (2010), who shows that the recovery rate ($= 1 - LGD$) (though for defaulted loans to non-financials) increases steadily with the recovery horizon. Thus, in our opinion, the best approach is to use the U-shaped frequency distribution of the LGD data that are derived from actual write-downs following the default of a bank.

[18]See Upper (2011) for an overview of these assumptions.

Figure 3. Frequency Distribution of Bank Failures

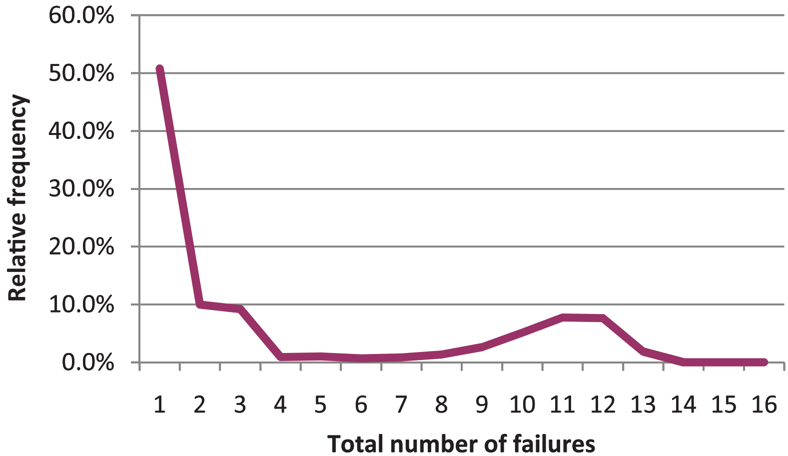

5. Results

5.1 Benchmark Case

The initial assumption for our simulations is that one of the sixteen banks/banking sectors[19] fails. This could trigger a cascade of failures if the ratio of tier 1 capital to risk-weighted assets of one of the creditor banks falls below 6 percent. The simulations (based on a stochastic LGD) are run by drawing from a beta distribution with parameters $\alpha = 0.28$ and $\beta = 0.35$. This means that for each exposure of a creditor bank to a bank in distress, we randomly draw an LGD value from the beta distribution estimated in section 4. We repeat this exercise by varying the bank that fails first from bank number 1 to 16. In contrast to simulations based on a constant LGD, the approach with a stochastic LGD yields for each of the sixteen banks a *distribution* of the number of banks in distress (and not only one single number of subsequent failures). We repeat the contagion exercise 100,000 times for each bank, with a different one of the sixteen banks starting the contagious process each time.

Figure 3 indicates the relative frequency of the number of bank failures, assuming that the probability of the initial failure is the

[19]For notational convenience, we will call the sectors in the following just banks.

Figure 4. Distribution Function of Bank Failures

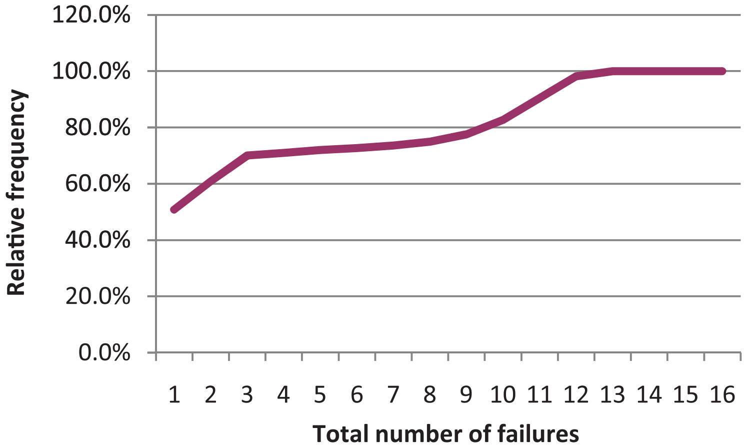

same for all of the sixteen banks. The figure shows that in 51 percent of the 1,600,000 simulation runs, no further failure occurs. In 8 percent of the cases, however, eleven subsequent bank failures occur. On average, we observe 3.06 subsequent bank failures (i.e., 4.06 bank failures in total) in our simulations. Figure 4 shows, for example, that in almost 18 percent of the cases more than ten banks fail. These results indicate that there is a considerable risk of interbank contagion.

5.2 Robustness Checks

We carry out robustness checks concerning six issues. The first three checks consider the robustness of the results to different LGD specifications. First, we investigate if drawing from the beta distribution estimated from our data set is a good approximation for the empirical distribution. Second, we examine if results change significantly when using different LGD distributions for the savings and cooperative sector. Third, we compare the simulation results of a stochastic LGD with the results under the assumption of a constant LGD that is equal to the mean of our data set. The next two checks consider the specification of the matrix of interbank exposures. With one special feature of our data set (compared with most of the existing literature) being the inclusion of off-balance-sheet exposures, we thus additionally run simulations excluding off-balance-sheet exposures and compare the results. In our fifth robustness check, we examine if netting of interbank assets and liabilities between counterparties can solve the problem of contagion. Finally, we check whether the number of bank failures is a good indicator for the stability of the system, as it is, of course, not only important how many banks fail but also how many assets are affected by failure. Thus, as a last robustness check, we take the balance sheet total of failing banks into account when judging the severity of contagion.

5.2.1 Drawing from the Empirical LGD Distribution

To investigate the sensitivity of our results with respect to the assumed distribution, we draw from the discrete distribution observed by the data instead of the beta distribution. For this purpose, one observed LGD value is randomly allocated to each exposure of a creditor bank to a bank in distress. Compared with drawing the LGD from a beta distribution, the results of this exercise do not differ much. The average amount of bank failures is 4.11 (compared with 4.06 in section 5.1). Furthermore, if we look at the relative frequency distribution as well as the cumulative distribution function of the total number of bank failures, there are virtually no differences to the results of the simulations with the beta-distributed LGD. We can therefore conclude that drawing from the beta distribution is a good approximation for our observed LGD values.

5.2.2 Different LGD Distribution for Savings and Cooperative Banks

As our interbank network consists not only of private commercial banks but also of savings and cooperative banks, an obvious question is how results are driven by the parameters of the beta distribution the LGDs are drawn from. As table 2 shows, LGDs corresponding to write-downs of savings and cooperative banks rather resemble a beta(0.42,0.30) and beta(0.08,0.24) distribution, respectively. Thus, we run our simulations by drawing from the respective distributions for exposures of the savings and cooperative sector and from the "standard" beta(0.28,0.35) distribution for the exposures of the remaining banks.

The results of the contagion analysis differ only slightly from our benchmark results. The overall expectation of bank failures is now at 4.24 (compared with 4.06 in the benchmark case). On the bank level, it is not clear whether the system is more stable than in the benchmark case. The initial default of five of the sixteen banks triggers more failures in the benchmark case; for the initial default of nine of the sixteen banks, fewer failures occur in the benchmark case and two of the sixteen banks do not trigger any further bank failure in any case. The deviations of the expected number of bank defaults, given that one specific bank fails, are, however, rather small regarding the benchmark case and do not exceed 0.95.

5.2.3 Stochastic versus Constant LGD

As the standard assumption in the existing literature is a constant LGD, we compare our simulation results generated under the assumption of a stochastic LGD with results under the assumption of a constant LGD. We set the constant LGD equal to the mean of our LGD data set (0.45; see table 2). Contrary to the case of the stochastic LGD, where we receive for each trigger bank a whole distribution of results, we obtain for each trigger bank one single number of failures under the assumption of a constant LGD. For only four of sixteen initial bank failures, a constant LGD yields a more unstable system (compared with the average number of bank failures under the assumption of a stochastic LGD). In total, we obtain on average 2.69 bank failures under the assumption of a constant LGD (compared with the average of 4.06 bank failures under the assumption of a stochastic LGD). Thus, we conclude that there is a certain risk of underestimating the effects of a bank failure on financial stability if the distribution of the LGD is not considered.

5.2.4 On-Balance-Sheet Exposures Only

Additionally, we examine the impact of including off-balance-sheet positions in our simulations. Most literature on interbank contagion ignores off-balance-sheet exposures due to data restrictions, while we have considered them in our above simulations. We therefore repeat the simulation exercise by excluding off-balance-sheet positions. According to our data set, the ratio of off-balance-sheet

Figure 5. Difference between the Relative Frequency Distributions of Bank Failures Considering Total Exposures and Balance Sheet Exposures Only

exposures to total exposures varies considerably between banks. Table 1 shows that 25 percent of the banks hold less than 6 percent of total interbank assets (3 percent of total interbank liabilities) off balance sheet. There are, however, also 25 percent of the banks that have a share of more than 24 percent (10 percent) of off-balance-sheet interbank assets (liabilities).

Not surprisingly, banks with a high amount of off-balance-sheet positions on their liability side trigger much fewer bank failures when ignoring these exposures. In total, the average amount of bank failures is only 3.47 (compared with 4.06 when considering all exposures).

To elaborate the differences between the simulation results with and without off-balance-sheet exposures, we calculate the difference between the two relative frequency distributions of bank failures (see figure 5). Figure 5 shows, for example, that the overall relative frequency of observing only one bank failure (i.e., contagion effects not occurring) is 5 percentage points higher when only considering balance sheet exposures. For high numbers of bank defaults, the result is reversed. For instance, the overall relative frequency of observing twelve bank failures is more than 5 percentage points higher when off-balance-sheet exposures are considered. Thus, figure 5 shows that the inclusion of off-balance-sheet exposures leads to a higher frequency of observing extreme events and therefore captures tail risk in a more adequate way. Therefore, we can conclude that off-balance-sheet exposures considerably contribute to the interdependence of banks and possibly change the results of the stability analysis in a remarkable way.

5.2.5 Netting

As a next robustness check, we examine how netting affects our results. Thus, we assume that banks net their exposures to each other. Technically, this means that we calculate the difference of element (i, j) and (j, i) of the matrix of interbank exposures and change all negative entries to zeros. The outcome is a matrix of net interbank exposures.

The result is that, of course, far fewer bank failures occur. This could be seen easily by looking at our simulation method. A significant reduction of interbank exposures necessarily induces less-contagious bank failures. What is surprising, however, is that contagion could still occur. Our simulations show that in almost 13 percent of the 1,600,000 simulation runs, a second-round effect occurs. On average, 1.16 banks fail, which is naturally much less than in our benchmark case (4.06 bank failures on average).

5.2.6 Balance Sheet versus Number of Banks

Of course, one could argue that the number of bank failures is not a good indicator for financial stability, as the size of the defaulted bank also matters. Hence, as an additional robustness check, we use the ratio of assets that belong to banks that fail in reaction to the trigger event to total assets that could theoretically fail as an indicator for the contagious effects. A value of 0 would thus mean that only the trigger bank fails and no subsequent bank failures occur. A value of 1 means that all banks in the system default.

In figure 6 we compare the distribution function of the total number of bank failures (which is the same as in figure 4) with the distribution function of the share of assets that belong to failing

Figure 6. Distribution Function of Bank Failures (N) and Share of Assets that Belong to Failing Banks (BT)

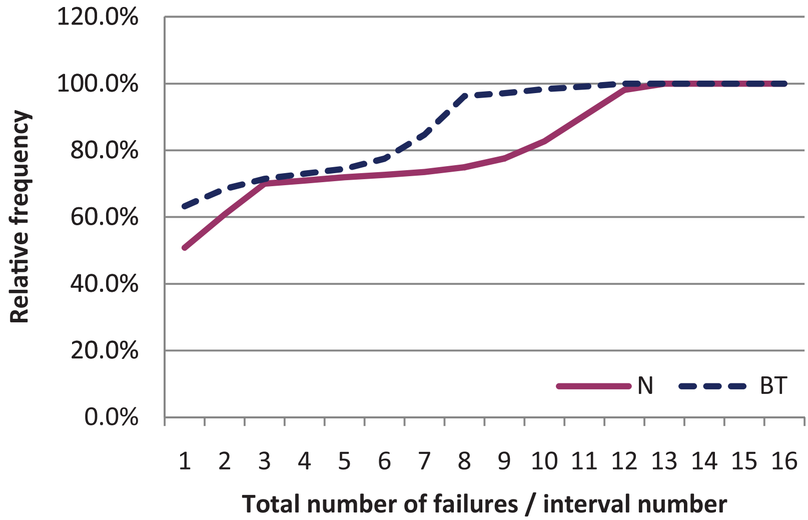

banks (without the trigger bank). To make these two functions comparable, we divided the share of assets that belong to failed banks into sixteen intervals of the same size and counted the frequency of results being in a particular interval. It is now easy to see that it is, e.g., more likely to observe less than 50 percent of the total assets of the banking system (without trigger bank) failing than it is to observe fewer than eight bank failures.

Furthermore, our simulations show that on average 14 percent of assets in the remaining banking system (without the trigger bank) are affected by bank failure. By comparing this result with the average share of banks that fail subsequently ($3.05/15 \approx 20$ percent), we can conclude that the banks that fail in our simulations belong on average to the smaller banks of our sample.

6. Conclusion

In this paper, we investigate contagion risk in the German interbank market. We have access to a unique data set on loss given defaults (LGDs) of interbank exposures. Our data reveal that the frequency distribution of the LGD is markedly U-shaped; i.e., defaults of interbank loans often imply either a low or a high loss. This markedly U-shaped distribution stands in contrast to the assumption of a unimodal LGD distribution in the literature.

Next, we run simulations investigating the extent of potential contagion in the German interbank market. For this purpose, we focus on fourteen systematically relevant German banks and the sectors of the savings and cooperative banks. We run simulations under the assumption of a stochastic LGD by drawing from a beta distribution. The shape of the beta distribution is derived from our LGD data set.

The result of our simulations is that contagion in the German interbank market may happen. For the period of time under review (end 2010), we find that the contagion exercise under the assumption of a stochastic LGD yields on average a more vulnerable system than under the assumption of a constant LGD. Furthermore, banks' off-balance-sheet exposures considerably contribute to the interdependence of banks and change the results of the stability analysis in a remarkable way.

An additional open question for future research is to compare the loss distribution at different points in time and to develop an indicator showing by how far the interbank market is prone to contagious processes.

Appendix. Delta Method to Test for the U-Shape of the Beta Distribution

Our goal is to explicitly test whether the observed LGD distribution is significantly U-shaped; i.e., we test the null hypothesis that $\alpha \geq 1$ or $\beta \geq 1$. We carry out a sequence of two t-tests with the two null hypotheses $\alpha \geq 1$ and $\beta \geq 1$, respectively. In the event that we can reject both null hypotheses, we accept the hypothesis $\alpha < 1$ and $\beta < 1$. Given the same significance level in both t-tests, the significance level of the joint hypothesis $\alpha < 1$ and $\beta < 1$ is at least as strong (see Frahm, Wickern, and Wiechers 2010). Using the delta method and the relations given in equations (6) and (7), we derive the asymptotic distribution of the estimates for $\hat{\alpha}$ and $\hat{\beta}$, respectively. Using a first-order Taylor expansion, the delta method gives us a relation between the variance-covariance matrix of the estimators $\hat{\mu}$ and $\hat{\sigma}^2$, and the variance-covariance matrix of $\hat{\alpha}$ and $\hat{\beta}$:

$$Var\begin{pmatrix}\hat{\alpha}\\ \hat{\beta}\end{pmatrix} \approx \nabla\begin{pmatrix}f_1(\hat{\mu},\hat{\sigma}^2)\\ f_2(\hat{\mu},\hat{\sigma}^2)\end{pmatrix}^T \cdot Var\begin{pmatrix}\hat{\mu}\\ \hat{\sigma}^2\end{pmatrix} \cdot \nabla\begin{pmatrix}f_1(\hat{\mu},\hat{\sigma}^2)\\ f_2(\hat{\mu},\hat{\sigma}^2)\end{pmatrix} \tag{8}$$

with $f_1(\hat{\mu},\hat{\sigma}^2) = \hat{\alpha} = \hat{\mu}(\frac{\hat{\mu}(1-\hat{\mu})}{\hat{\sigma}^2} - 1)$ and $f_2(\hat{\mu},\hat{\sigma}^2) = \hat{\beta} = (1 - \hat{\mu})(\frac{\hat{\mu}(1-\hat{\mu})}{\hat{\sigma}^2} - 1)$.

The variance-covariance matrix of $\hat{\mu}$ and $\hat{\sigma}^2$ is given by

$$Var\begin{pmatrix}\hat{\mu}\\ \hat{\sigma}^2\end{pmatrix} = \begin{pmatrix}\sigma^2_{\hat{\mu}} & \sigma_{\hat{\mu},\hat{\sigma}^2}\\ \sigma_{\hat{\mu},\hat{\sigma}^2} & \sigma^2_{\hat{\sigma}^2}\end{pmatrix} = \begin{pmatrix}\frac{1}{N}\sigma^2 & \frac{1}{N}\mu_3\\ \frac{1}{N}\mu_3 & \frac{1}{N}\left(\mu_4 - \frac{N-3}{N-1}\sigma^4\right)\end{pmatrix}, \tag{9}$$

where μ_3 and μ_4 denote the third and fourth central moments, respectively.[20] For implementation purposes, we replace the true moments with their estimators; i.e., $\hat{\sigma}^2$, $\hat{\mu}_3$, and $\hat{\mu}_4$ are given by $\frac{1}{N-1}\sum_{i=1}^N(x_i-\hat{\mu})^2$, $\frac{1}{N}\sum_{i=1}^N(x_i-\hat{\mu})^3$, and $\frac{1}{N}\sum_{i=1}^N(x_i-\hat{\mu})^4$, respectively.[21] From equations (8) and (9), we see that the variances of $\hat{\alpha}$ and $\hat{\beta}$ are linear combinations of $\sigma^2_{\hat{\mu}}$, $\sigma_{\hat{\mu},\hat{\sigma}^2}$, and $\sigma^2_{\hat{\sigma}^2}$:

$$Var(\hat{\alpha}) = \left(\frac{\partial f_1}{\partial\hat{\mu}}\right)^2 \cdot \sigma^2_{\hat{\mu}} + 2\cdot\left(\frac{\partial f_1}{\partial\hat{\mu}}\right)\cdot\left(\frac{\partial f_1}{\partial\hat{\sigma}^2}\right)\cdot\sigma_{\hat{\mu},\hat{\sigma}^2} + \left(\frac{\partial f_1}{\partial\hat{\sigma}^2}\right)^2\cdot\sigma^2_{\hat{\sigma}^2} \tag{10}$$

$$Var(\hat{\beta}) = \left(\frac{\partial f_2}{\partial\hat{\mu}}\right)^2 \cdot \sigma^2_{\hat{\mu}} + 2\cdot\left(\frac{\partial f_2}{\partial\hat{\mu}}\right)\cdot\left(\frac{\partial f_2}{\partial\hat{\sigma}^2}\right)\cdot\sigma_{\hat{\mu},\hat{\sigma}^2} + \left(\frac{\partial f_2}{\partial\hat{\sigma}^2}\right)^2\cdot\sigma^2_{\hat{\sigma}^2}. \tag{11}$$

Calculations based on our sample (i.e., all private commercial banks and the central institutions of the savings and cooperative banks) yield $Var(\hat{\alpha}) = 0.0007$ and $Var(\hat{\beta}) = 0.0013$. As a next step, we use these values to calculate the test statistics T for the t-test with the null hypothesis that $\alpha \geq 1$ and $\beta \geq 1$. The results $T_\alpha \approx -27$ and $T_\beta \approx -18$ clearly show that the null hypothesis can be rejected.

[20]See, for example, Mood, Graybill, and Boes (1974, p. 228) and Zhang (2007) for the variances and covariances of the estimators $\hat{\mu}$ and $\hat{\sigma}^2$.

[21]See Hahn and Shapiro (1967, p. 48).

Thus, we can conclude that, contrary to the common assumption of a unimodal LGD distribution in the literature, our data set of the LGD follows a U-shaped distribution.

References

Aikman, D., P. Alessandri, B. Eklund, P. Gai, S. Kapadia, E. Martin, N. Mora, G. Sterne, and M. Willison. 2009. "Funding Liquidity Risk in a Quantitative Model of Systemic Stability." Working Paper No. 372, Bank of England.

Basel Committee on Banking Supervision. 2010. "Basel III: A Global Regulatory Framework for More Resilient Banks and Banking Systems." Bank for International Settlements.

Bastos, J. A. 2010. "Forecasting Bank Loans Loss-Given-Default." *Journal of Banking and Finance* 34 (10): 2510–17.

Castrén, O., and I. K. Kavonius. 2009. "Balance Sheet Interlinkages and Macro-Financial Risk Analysis in the Euro Area." ECB Working Paper No. 1124.

Chan-Lau, J. A. 2010. "Balance Sheet Network Analysis of Too-Connected-to-Fail Risk in Global and Domestic Banking Systems." IMF Working Paper No. 10/107.

Cifuentes, R., G. Ferrucci, and H. S. Shin. 2005. "Liquidity Risk and Contagion." *Journal of the European Economic Association* 3 (2–3): 556–66.

Craig, B., and G. von Peter. 2010. "Interbank Tiering and Money Center Banks." BIS Working Paper No. 322.

Crouhy, M., D. Galai, and R. Mark. 2000. "A Comparative Analysis of Current Credit Risk Models." *Journal of Banking and Finance* 24 (1–2): 59–117.

Degryse, H., M. A. Elahi, and M. F. Penas. 2010. "Cross-Border Exposures and Financial Contagion." *International Review of Finance* 10 (2): 209–40.

Degryse, H., and G. Nguyen. 2007. "Interbank Exposures: An Empirical Examination of Contagion Risk in the Belgian Banking System." *International Journal of Central Banking* 3 (2): 132–71.

Dermine, J., and C. N. de Carvalho. 2006. "Bank Loan Losses-Given-Default: A Case Study." *Journal of Banking and Finance* 30 (4): 1219–43.

Eisenberg, L., and T. H. Noe. 2001. "Systemic Risk in Financial Systems." *Management Science* 47 (2): 236–49.

Elsinger, H., A. Lehar, and M. Summer. 2006. "Risk Assessment for Banking Systems." *Management Science* 52 (9): 1301–14.

Espinosa-Vega, M. A., and J. Solé. 2010. "Cross-Border Financial Surveillance: A Network Perspective." IMF Working Paper No. 10/105.

Frahm, G., T. Wickern, and C. Wiechers. 2010. "Multiple Tests for the Performance of Different Investment Strategies." Discussion Paper No. 5/10, Department of Economic and Social Statistics, University of Cologne.

Hahn, G. J., and S. S. Shapiro. 1967. *Statistical Models in Engineering*. New York: John Wiley & Sons.

Huang, X., H. Zhou, and H. Zhu. 2009. "A Framework for Assessing the Systemic Risk of Major Financial Institutions." *Journal of Banking and Finance* 33 (11): 2036–49.

James, C. 1991. "The Losses Realized in Bank Failures." *Journal of Finance* 46 (4): 1223–42.

Kaufman, G. G. 1994. "Bank Contagion: A Review of the Theory and Evidence." *Journal of Financial Services Research* 8: 123–50.

Memmel, C., and I. Stein. 2008. "European Data Watch: The Deutsche Bundesbank's Prudential Database (BAKIS)." *Schmollers Jahrbuch* 128 (2): 321–28.

Mistrulli, P. E. 2011. "Assessing Financial Contagion in the Interbank Market: Maximum Entropy versus Observed Interbank Lending Patterns." *Journal of Banking and Finance* 35 (5): 1114–27.

Mood, A. M., F. A. Graybill, and D. C. Boes. 1974. *Introduction to the Theory of Statistics.* Third ed. McGraw-Hill.

Schmieder, C. 2006. "European Data Watch: The Deutsche Bundesbank's Large Credit Database (BAKIS-M and MiMiK)." *Schmollers Jahrbuch* 126 (4): 653–63.

Sheldon, G., and M. Maurer. 1998. "Interbank Lending and Systemic Risk: An Empirical Analysis for Switzerland." *Swiss Journal of Economics and Statistics* 134 (4.2): 685–704.

Stolz, S., and M. Wedow. 2010. "Extraordinary Measures in Extraordinary Times: Public Measures in Support of the Financial Sector in the EU and the United States." Discussion Paper Series 1, No. 2010/13, Deutsche Bundesbank.

Tarashev, N., and H. Zhu. 2008. "Specification and Calibration Errors in Measures of Portfolio Credit Risk: The Case of the ASRF Model." *International Journal of Central Banking* 4 (2): 129–73.

Upper, C. 2011. "Simulation Methods to Assess the Danger of Contagion in Interbank Markets." *Journal of Financial Stability* 7 (3): 111–25.

Upper, C., and A. Worms. 2004. "Estimating Bilateral Exposures in the German Interbank Market: Is There a Danger of Contagion?" *European Economic Review* 48 (4): 827–49.

van Lelyveld, I., and F. Liedorp. 2006. "Interbank Contagion in the Dutch Banking Sector: A Sensitivity Analysis." *International Journal of Central Banking* 2 (2): 99–133.

Zhang, L. 2007. "Sample Mean and Sample Variance: Their Covariance and Their (In)dependence." *The American Statistician* 61 (2): 159–60.