Measuring Option Implied Degree of Distress in the US Financial Sector Using the Entropy Principle[1]

Philipp Matros[2], Johannes Vilsmeier[3]

Abstract

We estimate time series of option implied Probabilities of Default (PoDs) for 19 major US financial institutions from 2002 to 2012. These PoDs are estimated as mass points of entropy based risk neutral densities and subsequently corrected for maturity dependence. The obtained time series are evaluated with regard to their consistency and predictive power and their properties are compared to Credit Default Swap Spreads (CDS). Moreover, we also derive an indicator for the systemic risk in the US financial sector. We find that the PoDs are superior to CDS in identifying the high risk banks prior to the Lehman crisis.

Keywords: Entropy Principle, Risk Neutral Density, Probability of Default, Financial Stability Indicator, Credit Default Swaps
JEL classification: C14, C32, G01, G21

[2]Universität Regensburg, D-93040 Regensburg, Germany, philipp.matros@wiwi.uni-regensburg.de, phone: +49 (0)941 943-3948

[3]Universität Regensburg, D-93040 Regensburg, Germany, johannes.vilsmeier@wiwi.uni-regensburg.de, phone: +49 (0)941 943-5619

1 Introduction

The 2008/2009 global financial crisis and its severe economic repercussions highlighted once again the importance of financial stability surveillance. However, monitoring and assessing national and international financial system soundness in a timely manner is a wide and complex field of work. There is no consensus aggregate indicator which can be used to gauge financial system's resilience to stress situations. Instead, vast research in finance has led to a great variety of non-mutually exclusive financial stability indicators based on different theoretical and empirical grounds (see e.g. IMF (2009)).

This paper derives a market based indicator by estimating option implied risk neutral densities (RNDs) from observed option prices. In contrast to balance sheet based indicators, market based stability measures are inherently forward-looking and available at much higher frequency (see e.g. Knaup (2011)). Assuming that markets are information efficient, prices exhibit all available information concerning the future evolution of a firm's value and, hence, the analysis of market prices offers a very appealing way to evaluate a firm's soundness. While it can be argued to which extent asset markets are information efficient, empirical studies emphasize that derivative markets offer transactional efficiencies over equity and debt markets which expresses in a higher degree of information efficiency (see Swidler and Wilcox (2002); Mayhew (1995)). The reasoning is that the high leverage of derivative products decreases the cost of speculation and portfolio insurance, which in turn can increase the informational content of those markets. Further, a great advantage of option markets over e.g. the market of the frequently used Credit Default Swaps (CDS) is that option markets provide us with several equilibrium prices for just one underlying, such that one obtains a more comprehensive view of investors' expectations. Also, option markets are usually more liquid and prices more widely available than for CDS, as they are exchange traded.

We use the so-called option implied Probability of Default (option iPoD) methodology to derive time series of option based financial stability measures. The approach was suggested in Capuano (2008) and numerically robustified as well as evaluated in Vilsmeier (2011). The framework uses the cross entropy function in order to estimate the option implied RND and allows for a probability mass point in the RND at a value of zero for the underlying. This mass point can be interpreted as Probability of Default (PoD). As opposed to CDS- or bond based PoD estimation approaches, our methodology has the substantial advantage that it requires no assumptions for the recovery rate. Further, we do not need to assume any stochastic process for our risk neutral pricing model, which is the basis for approaches like Distance to Default and for the derivation of option based indicators like Implied Volatility.

Using a non-parametric estimation procedure, our approach provides the entire option implied RND from the observed option prices and is highly flexible regarding the functional form of the implied distribution. We obtain a probability distribution for the future value of the option's underlying at time to maturity, which provides us with a comprehensive image about the investors' expectations and risk preferences.

There exists a large literature on RND estimations, differing by the statistical methods applied to extract the RND from the observed option prices (for an extensive overview see Jackwerth (2004)). So far, though, RND estimates have not been applied to systematically evaluate a specific firm's soundness over a longer period of time; instead they were used for short term applications like evaluations of option pricing models (e.g. Figlewski (2008)) or testing market expectations around certain events (e.g. Capuano (2008); Melick and Thomas (1997)). Possible reasons are that RNDs do not provide a unique and easy to interpret stability measure and that their estimations are often plagued by limited sets of strikes, noisiness and maturity dependence. Moreover, one faces high computational efforts, and numerical instabilities of the statistical approach can make it difficult to obtain consistent RND estimates over long periods of time.

Our empirical implementation of the option iPoD framework overcomes the described problems. We estimate RND- and, hence, PoD time series for 19 of the largest US banks and financial institutions. The considered sample spans over eleven years from 2002 to the beginning of 2012 and comprises the late consequences of the early 2000s recession, the first financial turmoils in mid-2007, the world financial crisis of 2008/2009 as well as the repercussions of the European sovereign debt crisis of 2011/2012. This leaves us with a unique opportunity to evaluate our indicator's properties and forecasting abilities in an environment of huge macroeconomic distress and to compare it with more resilient periods as from 2004 to 2007.

In order to obtain the RND/PoD estimates, we use alternately five-, six-, and seven month call option contracts. Subsequently, we remove the maturity dependence inherent in the original time series by applying a non-parametric quantile regression approach to the pooled PoDs. The time series of the maturity corrected PoDs are evaluated with regard to their consistency and predictive power and their properties are compared to Credit Default Swap Spreads (CDS)[4]. In this context, we derive an indicator for the systemic risk prevailing in the US financial sector by applying a Principal Component Analysis (PCA) to the firm specific PoDs. Thereby, we isolate the systematic components from the idiosyncratic risks of each bank. Further, for a more clear interpretation of the respective levels of CDS and PoDs, we

---

[4]In the following we will use the abbreviation CDS for the term Credit Default Swap Spreads.

check and compare the indicators' properties in relation to the systemic risk, to the most resilient bank and to their own history.

Our results give evidence for the predictive/signalling power of the option iPoD methodology with regard to crises periods and for its ability to identify financially vulnerable institutions in a timely manner. We find that CDS and the option implied PoDs exhibit very similar dynamics, but PoDs being superior to CDS in identifying the high risk banks prior to the Lehman crisis in September 2008.

Our paper contributes to the literature in several ways. We are the very first to empirically apply the option iPoD framework to derive time series of RNDs/PoDs and, hence, are the first that give a comprehensive empirical evaluation of that methodology. It is demonstrated how the option iPoD methodology can be empirically implemented in order to get consistent, smooth and maturity corrected PoDs. This was achieved by the appropriate choice of liquidity weights, the use of a suitable maturity cycle of option contracts and a refined option data filtering technique. With the estimated PoDs we provide highly informative financial risk indicators for the individual firms in our sample, and in addition we use them to derive an indicator for the systemic risk of the overall US financial sector. Moreover, we give empirical evidence for the high informational content of the option iPoD framework, which is shown to be superior to the informational content of CDS. We stress that this is due to the fact that for CDS differing and unknown recovery rates over firms and periods of time complicate the interpretation of the levels of CDS, while iPoDs can be interpreted directly as they are equity based. Finally, our approach provides a sound basis for multivariate financial risk analyses, as the framework provides time series for firm specific PoDs jointly with the corresponding asset distributions.

The remainder of the paper is structured as follows. Section 2 describes the underlying methodology of our empirical framework. This comprises mainly the entropy based estimation of the RND and the PoD. This is followed by the description of the data used in our analysis. Section 4 explains the empirical implementation of our approach on how to obtain a stable and smooth time series of PoD- and RND estimates. The estimation results as well as the comparisons with the CDS are presented in section 5. Finally, section 6 provides concluding remarks and offers prospects for further research.

2 Methodology

The statistical framework used in this paper was originally suggested by Capuano (2008). In Vilsmeier (2011) a more stable objective function for the RND estimation and an alternative PoD determination procedure were proposed. In this paper, we apply this version.

The idea of the framework is to allow for a probability mass point for a value of zero of the stock S in the estimation of a stock option implied RND. A RND is a density function $f(S_T)$ that describes the investors' expectations regarding the value of the underlying at time to maturity T, implied by the observed option prices for different strikes. In order to obtain a RND one uses the continuous risk neutral pricing formula for a call option

$$C_0^{K_i} = e^{-rT} \int\limits_{S_T=K_i}^{\infty} (S_T - K_i) f(S_T) dS_T, \quad i = 1 \ldots B. \tag{1}$$

The formula should be solved with respect to the unknown density $f(S_T)$ for given option prices $C_0^{K_i}$ at different strikes K_i. Equation (1) states that the today's observed option prices must be equal to the discounted expectation of the inner values[5] under risk neutral probability measure (risk neutral pricing), where r denotes the annualized risk free rate and T the time to maturity of the option (measured in years). The number of observable option prices for different strikes K is denoted by B whereat the current stock price S_0 is included as an option with strike $K_1 = 0$. One faces an underdetermined estimation problem, as we do not have an infinite set of strikes (Breeden and Litzenberger (1978)). There are different statistical approaches to determine a unique density $f^*(S_T)$ out of the infinite many that are compatible with the observed prices (see e.g. Jackwerth (2004)). The approach chosen in this paper is to minimize the so-called cross entropy function

$$CE[f(S_T), f^0(S_T)] = \int\limits_{0}^{\infty} f(S_T) \log \frac{f(S_T)}{f^0(S_T)} dS_T, \tag{2}$$

under restrictions imposed on the moments of $f(S_T)$ given by the system of equations (1) and where $f^0(S_T)$ denotes some prior distribution.

The cross entropy function is based on the concept of entropy, whose mathematical representation is: $H[f(S_T)] = -\int_0^{\infty} f(S_T) \log f(S_T) dS_T$. In its information theoretic interpretation of Shannon (1948), the entropy function can be regarded as a measure of average uncertainty in a random variable. In this sense, if the entropy function is maximized subject to

---

[5]An inner value of an option is the payoff of the option given a specific realisation of S_T.

moment constraints on the density, among all densities that are consistent with the moment conditions (1) one identifies this distribution as optimal that implies the most uncertainty regarding a future outcome (Jaynes (1957)). Without constraints the entropy function obtains its maximum on a bounded domain by choosing the uniform distribution, for an unbounded positive real valued domain (for a given mean) by the exponential distribution, and for an unbounded real valued interval (given a mean and a variance) by the the normal distribution.

However, in this framework instead of maximizing the entropy function we minimize the cross entropy function in order to identify an optimal density under moment conditions. Thereby, the so-called entropic distance of $f(S_T)$ to some prior density $f^0(S_T)$ is minimized (see e.g. Cover and Thomas (2006)). Both approaches, though, yield the same optimal solution, if the prior distribution $f^0(S_T)$ in (2) is chosen to be of maximal entropy on the defined domain. We use the cross entropy concept since the prior density, $f^0(S_T)$, is necessary for the determination of the PoD.

If one assumes that a stock price of zero implies default, then a probability mass point for $S_T = 0$ in the RND could be interpreted as the investors' expectation regarding a firm's default between now and time to maturity T of the option. Given our continuous estimation framework, such a mass point cannot be estimated as a 'jump' in the density at $S_T = 0$. Instead, we extend the domain of the RND for S_T such that all realisations within this additional interval of values imply a future stock price of zero. In this way the PoD is not estimated as a mass point but as the integral over the density assigned to a certain sub-domain of the RND. The additional interval of values is obtained by shifting the domain for the future stock value S_T upwards by some constant D, and estimating $f(V_T)$, with $V_T = S_T + D$. For the payoff of the option in T (the inner value) now holds: $C_T = \max(V_T - D - K; 0)$, and any value $V_T \leq D$ will imply an inner value of zero for arbitrary K.[6] The integral of $f(V_T)$ over the interval $[0; D]$ will yield our PoD estimate.

A theoretical interpretation to the described PoD estimation procedure is possible, if one assumes that the so-called structural approach of Merton (1974) applies to a firm's balance sheet. In the structural approach a firm's value of assets is given by the value of its debt plus the value of its equity. The firm defaults if the value of assets does not cover the value of debt. Hence, in the PoD mechanism V_T can be interpreted as value of assets, S_T as the value of equity and D as the value of debt.[7]

In Vilsmeier (2011) some important characteristics of the PoD mechanism were found using

---

[6]Note that before extending the domain of S_T only $S_T = 0$ implied an inner value of zero for arbitrary K.

[7]Note that the assumptions of the structural approach have no implications for the PoD estimation. Any event that will lead in the investors' expectation to $S_T = 0$ will increase the PoD.

numerical experiments. First, the procedure can perfectly estimate a probability mass point for $S_T = 0$ if the length of the interval D is chosen correctly. Second, the estimated PoD is a concavely increasing function in D and flattens out after reaching the optimal D. Third, for arbitrary reasonable RND forms and PoD levels the optimal D lies within the interval $[0; 20]$. Fourth, as an exact rule for the determination of the optimal D could not be detected, averaging the PoDs obtained for RND estimates with Ds ranging from 0 to 20 provides very accurate estimates. The optimal RND is then identified as the one that provides the PoD closest to the average PoD ('averaging approach'). Despite its quite ad hoc nature, this procedure produces highly accurate estimates in numerical experiments covering a great number of reasonable RND/PoD specifications. As will be seen in section 5 the procedure produces also highly plausible results if applied to real option data.[8]

For the RND estimation we use the following system of equations, which we express in terms of V_T and D:

$$CE[f(V_T), f^0(V_T)] = \int_0^\infty f(V_T) \log \frac{f(V_T)}{f^0(V_T)} dV_T \tag{3}$$

$$C_0^{K_i} = e^{-rT} \int_{V_T = D + K_i}^\infty (V_T - D - K_i) f(V_T) dV_T, \quad i = 1 \ldots B \tag{4}$$

$$\int_{V_T = 0}^\infty f(V_T) dV_T = 1. \tag{5}$$

(3) is the cross entropy function for $f(V_T)$ with regard to some prior distribution $f^0(V_T)$. (4) denotes the continuous risk neutral pricing formula and (5) is an additional moment condition that ensures that the density integrates up to one.

Combining (3) to (5), the estimation setup can be written using the Lagrange multiplier

---

[8]The intuition behind the ad hoc procedure is explained in Vilsmeier (2011). A sketch of the intuition is as follows. The absolute level of D does not influence the results regarding the estimated RND/PoD but only the value of V_{min} (minimal possible realisation for V_T) relative to D, i.e. the interval length $[V_{min}, D]$. The 'location' of the interval $[V_{min}, V_{max}]$ does not influence the results as the inner values of the options depend on $(V_T - D)$, and D will increase by the same amount as V_T if we move the interval. The length of the interval $[D, V_{max}]$ does not significantly influence the PoD, as the normalization constant $\mu(\lambda)$ will offset the decreasing $f^0(V_T)$ in the PoD estimation due to an increase in V_{max}, and 'irrelevant' V_T will be assigned a density of virtually zero in the RND estimation. Hence, for the PoD/RND estimation only the interval $[V_{min}, D]$ is of importance, which is addressed in the ad hoc procedure. Moreover, the non-parametric estimation framework is able to estimate arbitrary levels of PoDs for relatively short interval lengths because the framework assigns just enough density to values V_T larger but very close to D that in the 'transition', $V_T = D$, the density is large enough to assign the PoD to the given interval length $[V_{min}, D]$. This is because, if e.g. the PoD is high, the framework has to assign density to values of V_T that do not contribute to the price of the option (very small values and 'negative' equity values).

approach:

$$L = \int_{V_T=0}^{\infty} f(V_T) \left[\log \frac{f(V_T)}{f^0(V_T)} \right] dV_T + \lambda_0 \left[1 - \int_{V_T=0}^{\infty} f(V_T) dV_T \right]$$

$$+ \sum_{i=1}^{B} \lambda_i \left[C_0^{K_i} - e^{-rT} \int_{V_T=D+K_i}^{\infty} (V_T - D - K_i) f(V_T) dV_T \right]. \qquad (6)$$

Optimizing (6) with respect to $f(V_T)$ yields (see e.g. Cover and Thomas (2006)):

$$f^*(V_T) = \frac{1}{\mu(\lambda)} f^0(V_T) \exp \left[\sum_{i=1}^{B} \lambda_i e^{-rT} \mathbf{1}_{V_T > D+K_i} (V_T - D - K_i) \right], \qquad (7)$$

with

$$\mu(\lambda) = \exp(1 - \lambda_0) = \exp(-\lambda_0') = \int_{V_T=0}^{\infty} f^0(V_T) \exp \left[\sum_{i=1}^{B} \lambda_i e^{-rT} \mathbf{1}_{V_T > D+K_i} (V_T - D - K_i) \right] dV_T. \qquad (8)$$

We see from (7) that the optimal solution will be in the family of exponential distributions. Hence, the estimation procedure is highly flexible regarding the underlying shape of the RNDs and is able to approximate almost arbitrary functional forms if we have enough option prices. Further, the estimated PoD will be equal to $\int_0^D \frac{f^0(V_T)}{\mu(\lambda)}$ as the expression in the exponential function will be equal to 0 for all $V_T \leq D$. That means that the estimated PoD and the shape of the RND interact, as the Lagrange multipliers also determine the shape of the RND.

If we assume some value for D, we obtain our RND and PoD if we are able to determine the λ_i in (7). This could be achieved if one replaces $f(V_T)$ in the Lagrangian by (7), and optimizes regarding the λ_i. The resulting system of equations could be solved using e.g. a multivariate Newton-Raphson algorithm. The search for the roots, though, is numerically very unstable due to near singularities of the involved Jacobian for large domains of the λ_i. Hence, in Vilsmeier (2011), following the suggestions in Alhassid et al. (1978) and Agmon et al. (1979), an objective function for the λ_i was derived that yields efficient and numerically stable optimizations. The derivation is based on the finding that a function can be defined such that for any trial set of parameters λ^{Tr} it provides a theoretical lower bound to the value of the cross entropy of the optimal density. This function has its minimum for the optimal set of λs.

If one assumes some finite maximum and minimum value, V_{max} and V_{min}, for the value of

assets (per share) in the RND domain[9] and a uniform prior, one can solve the integrals involved in the objective function analytically and obtains:

$$
F = \log\left(\frac{1}{V_{max} - V_{min}}\right) + \log\left\{ \exp\left(-\sum_{i=1}^{B} w_i \lambda_i C_0^{K_i}\right)(D - V_{min}) \right.
$$

$$
- \sum_{i=1}^{B-1} \left[\frac{\exp\left(\sum_{j=1}^{i} w_j \lambda_j (e^{-rT}(K_i - K_j) - C_0^{K_j}) - \sum_{k=i+1}^{B} w_k \lambda_k C_0^{K_k}\right)}{e^{-rT}\left(\sum_{j=1}^{i} w_j \lambda_j\right)} \right.
$$

$$
\left. - \frac{\exp\left(\sum_{j=1}^{i} w_j \lambda_j (e^{-rT}(K_{i+1} - K_j) - C_0^{K_j}) - \sum_{k=i+1}^{B} w_k \lambda_k C_0^{K_k}\right)}{e^{-rT}\left(\sum_{j=1}^{i} w_j \lambda_j\right)} \right]
$$

$$
\left. - \left[\frac{\exp\left(\sum_{j=1}^{B} w_j \lambda_j (e^{-rT}(K_B - K_j) - C_0^{K_j})\right) - \exp\left(\sum_{j=1}^{B} w_j \lambda_j (e^{-rT}(V_{max} - D - K_j) - C_0^{K_j})\right)}{e^{-rT}\left(\sum_{j=1}^{B} w_j \lambda_j\right)} \right] \right\},
$$

(9)

where w_i denotes weights that are pre-multiplied to the Lagrange multiplier. The weights will ensure that more liquid option contracts (measured in our approach in terms of open interest) have to be met more closely by the estimated RND. The assignment of the liquidity weights is very important in order to obtain timely consistent and smooth PoD estimates. This issue will be addressed in section 4. The minimization of (9) is numerically highly stable and can be computed in a fast manner, which is prerequisite for our applications as we had to estimate in total about one million RNDs based on options that provide up to 40 strikes a day.

3 Data

Our option and stock data sample comprises 19 US banks and financial institutions and ranges from February 6, 2002 to February 24, 2012. Hence, the late consequences of the early 2000s recession, the first financial turmoils in mid-2007, the world financial crisis of 2008/2009 as well as the repercussions of the European sovereign debt crisis of 2011/2012 are included in our data set. Regarding the level and variance in the degree of financial distress, this data sample provides us with a unique opportunity to evaluate our indicator's properties

---

[9]The value of V_{max} should be large (e.g. ten times the current stock price) but can be arbitrarily chosen as it does not significantly influence the estimates. V_{min} denotes the minimal possible realisation for V_T and is set equal to zero in our applications.

over a highly diversified period of time. As Coffinet et al. (2010) point out, prior to the world financial crisis empirical literature assessing properties of market indicators had to rely on rather soft definitions of default.

Among the 19 covered financial institutions there are 14 banks, namely: Goldman Sachs (GS), Wells Fargo (WFC), Citigroup (C), Bank of America (BAC), JPMorgan Chase (JPM), Morgan Stanley (MS), PNC Bank (PNC), State Street (STT), Bank of New York Mellon (BK), Lehman Brothers (LEH), Bear Stearns (BSC), Wachovia (WB), Merill Lynch (MER) and Washington Mutual (WM).[10] This data set comprises the past and/or present largest US banks. For instance, on November 4, 2011 the Financial Stability Board (FSB) published a list of global systemically important financial institutions, which included the following US banks: GS, WFC, C, BAC, JPM, MS, STT and BK (see FSB (2011)). Thus, our sample can be viewed as an appropriate proxy for the US financial sector. Moreover, banks not covered by the FSB listing had a major impact on the US financial sector before the outbreak of the global financial crisis. LEH, for instance, filed for Chapter 11 bankruptcy protection on September 15, 2008; BSC was sold to JPM on March 16, 2008; MER was purchased by BAC on September 14, 2008; WB was sold to WFC on December 31, 2008 and WM filed for Chapter 11 bankruptcy protection on September 26, 2008 and was shortly after taken over mainly by JPM.[11] LEH, BSC and MER were among the world's leading investment banks and prior to the financial crisis WM was one of the largest US banks.

In addition to the mentioned 14 banks, our sample includes 5 non-banking financial entities with great relevance for the overall US financial sector. The American International Group (AIG) is among the biggest insurance corporations in the world but had to be rescued by the Federal Reserve Bank and the United States Treasury in September 2008 and May 2009. Until its acquisition by BAC on July 1, 2008, Countrywide (CFC) was one of the largest US mortgage banks. Before the financial crisis, the financial services company MBIA (MBI) was the world's largest monoline insurer and highly involved in the market of asset-backed- and mortgage-backed securities. Blackstone (BX) is an asset management and financial services company and currently the largest fund of hedge funds in the world. Finally, Blackrock (BLK) is an investment management corporation and the world's largest asset manager.

As can be seen from the brief discussion above, our data set covers not only a varied time period but also includes a quite heterogeneous sample of financial institutions. Some have performed quite well during the recent financial crisis, some others depended on governmen-

---

[10]We are well aware of the fact that a distinctive classification of the the financial entities is subjected to certain inaccuracies due to the complexity of the underlying business portfolios.

[11]For the sake of completeness, also Citigroup had to be rescued by the US government on November 2008 but did not go bankrupt in the aftermath, though.

tal financial aid during the turmoils and still others went bankrupt. This leaves us with a substantial informational input for our empirical analysis.

The equity option and stock data of the above mentioned financial institutions were obtained from the New York Stock Exchange (NYSE) via the data provider Stricknet. This encompasses on a daily basis the end of day stock- and option price, the trading volume, the strike price, the bid- and ask option price, the open interest (number of outstanding option contracts) as well as the year, month and, for the current option symbol methodology, day of expiration. At first, in order to estimate the option iPoDs we need to extract the relevant information from our vast data set. This is a quite challenging task, due to the option market inherent complexities and due to changes in the option symbolic system in February 2010.

The old option symbol format consisted of a three- to five-character symbol. The first up to three letters denote the root symbol, which is the stock symbol of the respective company if stock- and option trading is located at the same exchange. If not, it can depart. The subsequent letter represents the expiration month of the option[12] and indicates if it is a call (letters A for January to L for December) or put (M to X). The last letter (generally A to T) indicates the strike price. Since there are plenty potential strike prices but only a limited number of letters, each letter stood for more than one strike price. Consequently, large stock price variations may lead to duplications of the option symbols. Therefore, dummy- or extension root symbols were introduced in order to accommodate the limitations of the original symbol. In consequence, one and the same option can have different option symbols whose detection is not straightforward. The identification of option symbols is further complicated by stock splits, odd strike prices and the proliferation of, for instance, LEAPs-, Quarterly-, Weekly-, and FLEX options[13]. For an exemplary overview of such a dataset given the old option symbology see Table 1. Due to the mentioned shortcomings, a more intuitive and flexible option coding system was established. The new symbology consists of a one to six character root symbol, followed by two characters each representing the expiration year, month and day, one character denoting the put (P) or call (C) option, five characters for the strike price and three for possible strike price fractions (e.g. in case of stock splits). This option symbolic method avoids any duplications and can accommodate newer types of options. For an exemplary overview of our extracted dataset given the current option symbolic methodology see Table 2. Most of our data, though, are based on the old option symbology. As a risk

---

[12]Note that specific expirations days have to be obtained from option expiration calenders. Equity options' expiration date is usually every third Saturday in a month.

[13]LEAP stands for Long-Term Equity Anticipation Securities which are options with expiration dates longer than one year. Quarterly options expire at the end of quarters and Weekly options expire after eight days. FLEX options means Flexible Exchange Options which can be modified by investors needs.

Option Symbol	Last	Volume	Strike	Bid	Ask	Open Interest	Expiration Year
JPMRX	5.05	0	44.00	5.15	5.25	4149	2010
JPMRY	6.45	0	46.00	6.45	6.55	1210	2010
JPMRZ	7.00	0	47.00	7.15	7.25	1469	2010
JSAAA	25.55	0	17.50	24.15	24.25	1449	2010
JSAAB	34.45	0	10.00	31.65	31.75	92	2010

Table 1: Extract from the Option Dataset for JPM with old Option Symbology

Option Symbol	Last	Volume	Strike	Bid	Ask	Open Interest	Expiration Year
JPM 110107C00042000	1.67	705	42.00	1.64	1.66	1125	2011
JPM 110107C00043000	0.89	2153	43.00	0.85	0.87	1279	2011
JPM 110107C00044000	0.34	4157	44.00	0.33	0.35	130	2011
JPM 110107C00045000	0.09	351	45.00	0.10	0.11	0	2011
JPM 110107C00046000	0.03	31	46.00	0.03	0.04	0	2011
JPM 110107C00047000	0.02	10	47.00	0.01	0.02	0	2011

Table 2: Extract from the Option Dataset for JPM with new Option Symbology

free interest rate we used the 3-month Treasury Bill secondary market rate obtained from the FRED data base. The CDS spread data were obtained from Markit Group and range from February 6, 2002 to January 18, 2012. The sample covers the same institutions as the stock and option data, except for BK, BX and BLK. Due to enhanced data availability, we extracted 5-year CDS under the credit event of modified restructuring.

The subsequent section will explain how we empirically implemented our estimation approach and how we dealt with problems discussed in this section and the introduction.

4 Empirical Implementation

There are some important issues that have to be taken into account in order to be able to estimate smooth and timely consistent RND time series.

The first issue concerns the so-called maturity dependence of RND estimates. RNDs provide information about the expectation of investors regarding the value of the underlying at the

expiry date of the option contracts. Hence, if one estimates RNDs for subsequent days and uses option contracts with the same expiry date for each of these days, one faces the problem that RNDs closer to the expiry date will exhibit ceteris paribus less uncertainty regarding the future value of the underlying. This characteristic is called maturity dependence of RNDs. The problem arises because traded option contracts exist only for a few expiry dates within a year, such that one cannot extract time series of RNDs with constant time to maturity. To solve this problem we introduce in section 5.1 a regression based procedure to remove the maturity effects from the moments and PoDs of the estimated RNDs. For our procedure to work, though, we need to keep the maturity effect between the estimated RNDs as small as possible, i.e. the used option contracts for different days should exhibit similar time to maturity. Very similar time to maturity would be obtained if one constantly uses option contracts that expire in the subsequent month (as such a contract always exists). Thereby, the maximum difference in time to maturity related to two RNDs would be 31 days. However, this approach has the serious drawback that the estimated RNDs and hence PoDs only display information about the investors' expectations regarding a very short time period. The derived PoDs would indicate the probability of a firm's default within the next few weeks only, and the obtained results would be very erratic as only very imminent risks significantly change the investors expectation. Consequently, one would wish for a risk evaluation for a longer time period. Option contracts with longer time to maturity (e.g. 6 month) are not newly initiated at each month, though. Instead, different firms have different cycles within they initiate contracts with longer time to maturity than one month.

Taking into account the trade-off between maturity dependence and long term risk evaluation, we identified and applied three different 'maturity cycles' for our examined companies and allocated institutions with the same maturity scheme into one group. Hence, in our estimation implementation we considered three sub-samples of financial institutions instead of treating all 19 institutions equally with regard to their maturity intervals. The first group consists of GS, WFC, MS, BLK, BK, LEH, BSC, WB, MER, CFC and WM with, starting in January, a cycle of six-, five- and seven month maturity (i.e. seven month contracts are initiated in March, June, September and December), alternating throughout the years. The second group uses a five-, seven- and six month maturity scheme and covers C, JPM and BX. Finally, the third group comprises BAC, AIG, MBI, PNC and STT and follows a seven-, six- and five month time to maturity cycle. The explained five,- six- and seven month time to maturity schemes leave us with a maximum of roughly 90 days of difference in time to maturity between the applied option contracts. Thereby, over the entire considered time period of 2583 days for which we estimated the RNDs, we obtain repeatedly RNDs with

the same time to maturity. This is the prerequisite for our maturity dependence adjustment procedure, which consists of regressing the pooled (over time and firms) moments and PoDs of the RNDs on their respective time to maturity.

A further important issue that we detected in the estimation of our RND time series is the use of adequate liquidity weights in the optimization procedure as shown in our objective function (9). These weights ensure that more liquid option contracts, which presumably exhibit more information about the future value of the underlying (prices with less noise), are met more closely by our estimated RNDs than illiquid contracts. The weights are calculated by dividing open interest (contracts traded in the past and not exercised or evened up yet) for a specific strike by the sum of open interest over all available strikes for a firm's stock option. We found that the use of liquidity weights based on open interest leads to much smoother and more consistent time series than the use of trading volume, as often suggested in the literature. The reason is that crucial market information are discarded if the sample is weighted by trading volume. As the currently observed price arises as a result of current and past trading, a contemporaneous trading volume of zero does not necessarily imply such a contract is illiquid and has no information about the investors' expectations. Quite the contrary, if there is no trading today but there was high trading in the past (measured in open interest), this means the contract is liquid but the investors' expectations did not change with respect to the previous day(s).

A third issue deals with the old option symbology used previous to February 2010. Under the current option symbol methodology the expiration date of the option is directly inherent in the symbol (as well as the information whether it is a call- or put option) and, hence, the extraction of the correct option contracts from the whole firm specific domain is straightforward. However, the old option symbology causes some difficulties due to problems like dummy- or extension root symbols as outlined in section 3. First, we need to identify the option contracts with the appropriate maturity. Knowing that the second last letter (A to L for call options) of the option symbol denotes the expiration month of the option, together with information about the expiration year (provided by the NYSE in the dataset), we can filter out the options with the adequate time to maturity. Second, we need to deal with the issue of dummy symbols. Therefore, we check how many different option symbols are left in our already filtered dataset and sort the strike prices belonging to each symbol in ascending order. If there is just one option symbol left, we are done. Otherwise, we check if different option symbol groups have different strike ranges. If the corresponding strike ranges do not overlap, we combine these contracts as they contain differing information. If strike ranges overlap, we chose that option symbol corresponding with the most strike prices. Otherwise

we would discard a large set of information in cases where the available strike ranges exceed the number of available letters which denote the respective strike price of an option.

Finally, we have to set some model parameters before we can carry out the RND estimations. The framework described in section 2 is designed such that we can set global parameters which are used for the RND estimations for all institutions and all periods. This is very important in order to be able to estimate such a large number of RNDs. As pointed out in Vilsmeier (2011) the level of D, i.e. our debt value, does not influence our estimation results but only the length of the interval $[V_{min}; D]$. This implies that no matter how large we set D, if V_{min} is always D minus some constant (e.g. 10), the obtained results are exactly the same. Knowing this, we set our $V_{min} = 0$ and D takes integer values 0 to 20 in the averaging approach. For instance, in the first run $D = 1$, in the second $D = 2$ and so on, up to our $D_{max} = 20$. For V_{max}, we choose some arbitrarily large value that will be large enough for the asset value domain of all banks and for all time periods. In our implementation, we used $V_{max} = 850$, which implies that the asset value per share of a firm cannot exceed this value.

5 Results

This section will show that the empirical implementation of our estimation approach yields consistent and plausible PoD estimates and we will compare their performance to established indicators like CDS. Subsection 5.1 will stress that during institution specific high distress times the levels of our PoD estimates are more reasonable in size than those of CDS. In that context, the issue of recovery rates will be discussed and the superior informational content of our iPoDs emphasized. Subsection 5.2 presents our derived systemic risk indicator which measures the degree of distress of the overall financial sector. Finally, subsection 5.3 interprets our financial stability indicator in relative terms and substantiates the finding that our iPods are often better than CDS in identifying high risk institutions prior to incisive events.

Even though we focus on the PoD time series, it is important to note that the evaluation of the PoDs implicitly evaluates also the entire RND. This is the case because the PoD is a function of all RND shape parameters λ_i. Hence, no consistent and plausible PoDs can be obtained if not a consistent and plausible RND is estimated.

5.1 Time Series of option iPoDs

Applying the option iPoD methodology as described in the previous sections provides us with a RND estimate for each firm at each trading day in our sample period. Figure 1 displays an

example of such a RND, obtained for Goldman Sachs at July 10, 2008. The special feature of this RND estimate is the probability mass point for a future value of assets (per share), V_T, in the amount of the debt value D of the firm. The mass point indicates the probability that the stock value of the firm will be zero at time to maturity of the options and hence corresponds to our PoD estimate. We subsequently want to evaluate the time series of these
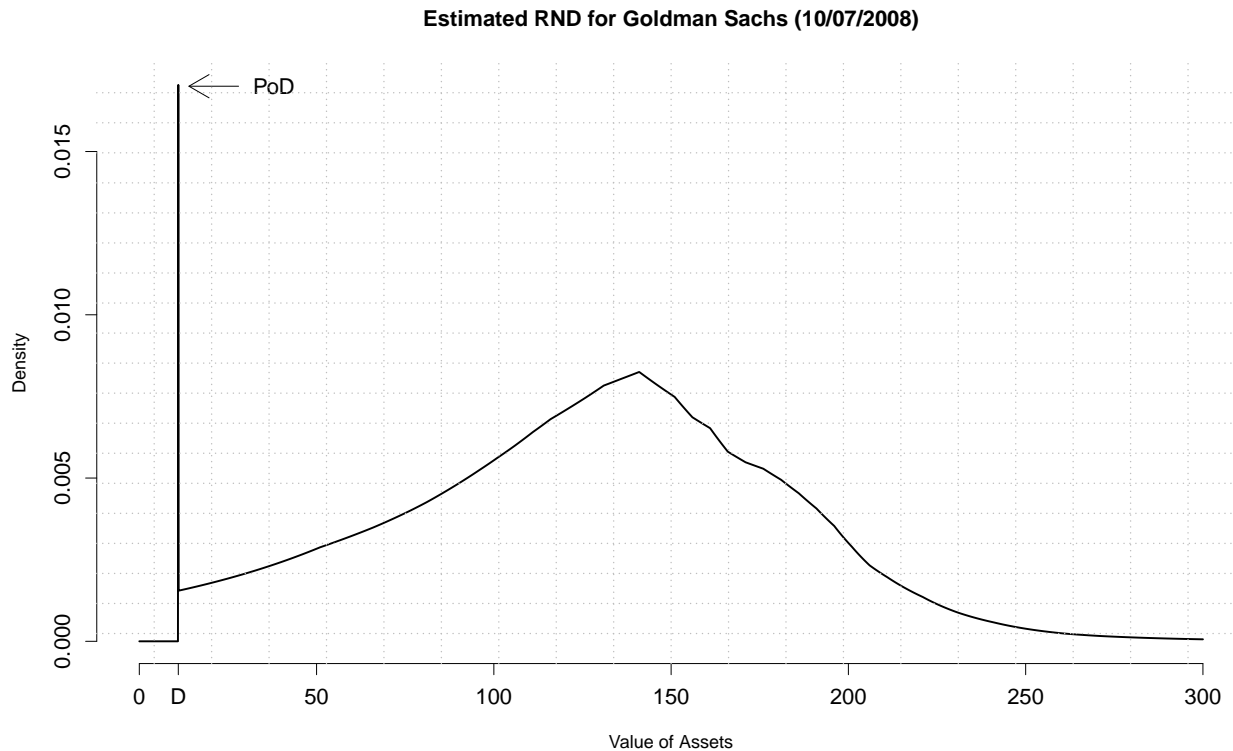
Estimated RND for Goldman Sachs (10/07/2008)

Figure 1: Estimated Risk Neutral Density for the future value of assets (per share) for Goldman Sachs on basis of option prices available at July 10, 2008

PoD estimates obtained for the different financial institutions. Before we can do so, however, we have to remove the maturity dependence from the daily PoD estimates, as outlined in the preceding section.

The basic idea is quite simple. We start by pooling our PoDs over periods of time and firms, and assign to each PoD estimate the time to maturity of the options that was used to estimate the RND. Then we regress the PoDs on the respective time to maturity. As all times to maturity run repeatedly from 130 days to 220 days, we have for each of the time to maturities several PoD estimates such that the regression based approach should yield rough approximations of the true maturity effects. In order to admit for non-linear maturity effects as well as for different effects for different quantiles of the PoD distribution we

apply a non-parametric quantile regression approach. More precisely, we use the methodology of additive quantile regressions (Hastie and Tibshirani (1986); Hastie and Tibshirani (1990)) in which the usual predictor of the quantile regression is augmented with additive non-parametric terms, and smoothing restrictions are imposed onto the fitted function. As smoothing restriction we apply the method of total variation regularization as suggested in Koenker et al. (1994).[14] Further, we restrict our fitted function to have a positive slope, as we expect that a higher time to maturity leads on average to a higher PoD. Figure 11 (Appendix) exemplarily depicts for each time to maturity the 40%- and 90%-PoD quantiles, as well as the respective fitted functions for each quantile. In order to carry out the maturity correction, we obtain the fitted PoDs for each time to maturity and different quantiles (in 5% steps), and calculate the difference between the fitted values of the highest time to maturity and the respective lower time to maturities for each quantile. The obtained differences are our correction factors for the estimated PoDs. The size of the correction factor applied to a specific PoD depends on how large the assigned time to maturity is and to which quantile the respective PoD belongs. After the correction all PoDs have a theoretical time to maturity of 220 days that is the maximum possible time to maturity that we have in our sample.[15]

Figure 12 (Appendix) shows the effects of our maturity correction procedure on the original PoD time series of Lehman Brothers. We can state that the maturity correction changes nothing regarding the dynamics of the time series. PoDs are only systematically higher, which is what we expect if time to maturity increases.

Henceforward, our aim is to check if the estimated PoD time series are plausible and consistent regarding dynamic and level, and if the obtained indicators are able to give some early warning signals in advance of crisis events.

In Figure 2 we see an example for an estimated time series of maturity corrected PoDs, namely for Citigroup. The time series runs from February 2002 to February 2012 and displays some typical dynamics that one can find in the PoD time series of all institutes in our sample (see Figure 13 (Appendix) for a complete overview of all PoDs). PoD levels are elevated in the aftermath of the 2001 recession, followed by a very calm period with low PoDs until mid-2007. Starting with July 2007, PoDs increase continuously until mid-2009 in the course of the US subprime crisis. The bankruptcies of Bear Stearns (BSC) and Lehman Brothers (LEH) are highlighted in Figure 2 and one can discover sharp inclines at the days of these events but also highly elevated risk indication in advance of these events. In mid-

---

[14]The methodology is available in the statistical software R using the 'quanteg' package.

[15]The correction process could also be applied to different moments of the RNDs in order to obtain maturity corrected densities.
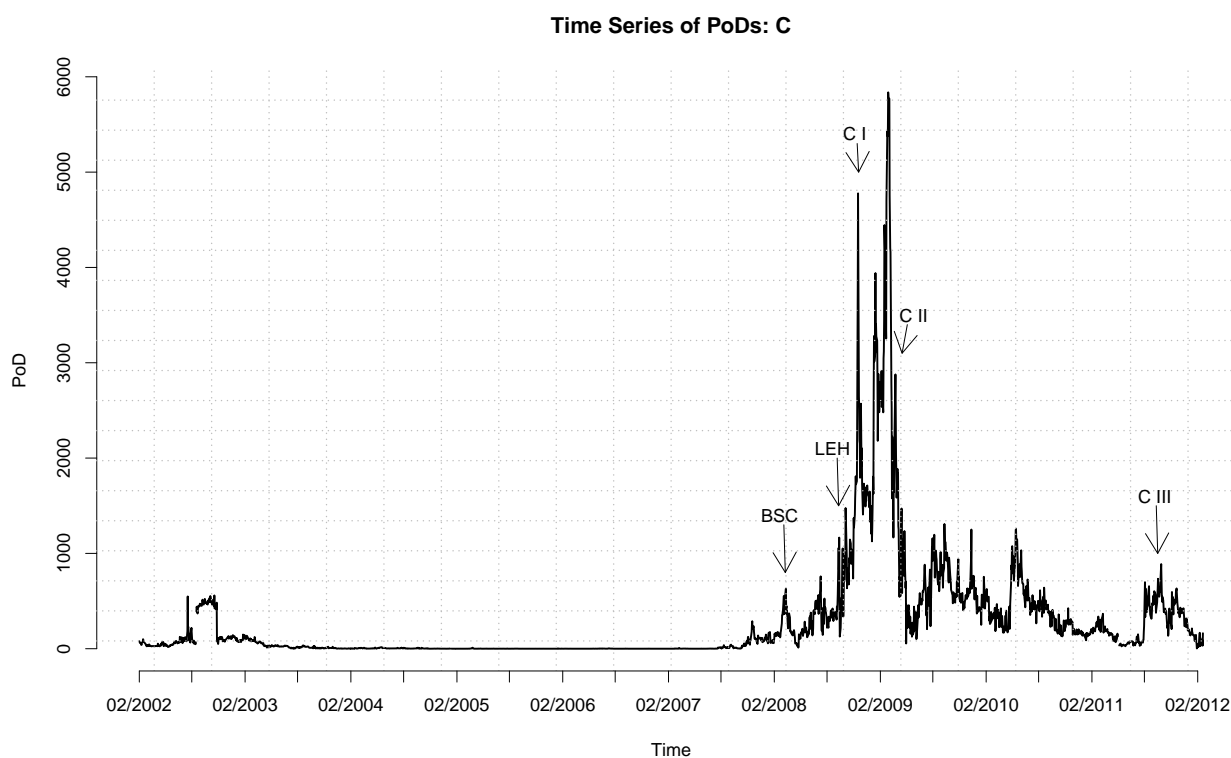
Time Series of PoDs: C

Figure 2: Time Series of maturity corrected PoDs in basis points for Citigroup from February 6, 2002 to February 24, 2012

2009, PoDs return to pre-LEH period levels and in end-2010/beginning-2011 even to levels of the pre-BSC period. In August 2011 PoDs begin to increase again to pre-LEH levels as a consequence of the European sovereign debt crisis.

Firm specific events for Citigroup are highlighted in Figure 2 with C I, C II and C III. C I indicates the date when Citigroup was rescued by the US government, C II indicates the date when large restructuring of the firm was decided at a general meeting, and C III the date when Citigroup was downgraded by several rating agencies. For C I and C III we see that the PoDs already signalled very high risks in advance of these events, with levels of around 1800 basis points (BP) for C I and 800 BP for C III. With regard to C II, the decisions seemed to be expected (as PoDs decreased some days before) and once they were actually made led to a sharp decline in investors' risk perception. Added together, we see that the PoD time series is able to signal risks concerning the entire financial sector (systemic risk) as well as firm specific (idiosyncratic) risks in a timely manner.

To check our indicator's predictive power in comparison to other existing indicators, we next compare our PoD time series to 5-year CDS. CDS are a very commonly used measure to

derive (risk neutral) probabilities of default for firms and countries, such that we cannot only compare dynamics but also to some extent the levels of the two indicators.

In Figure 3, 5-year CDS for Citigroup are plotted against our PoD time series using two different scales. It is obvious that both indicators exhibit a very similar dynamic, i.e. the points of time at which the time series move up or down are virtually the same. The impression that both indicators, CDS and PoDs, exhibit highly similar dynamics is found for all of the time series of the different institutions (see also Figure 4). In Table 3 (Appendix) the Pearson and
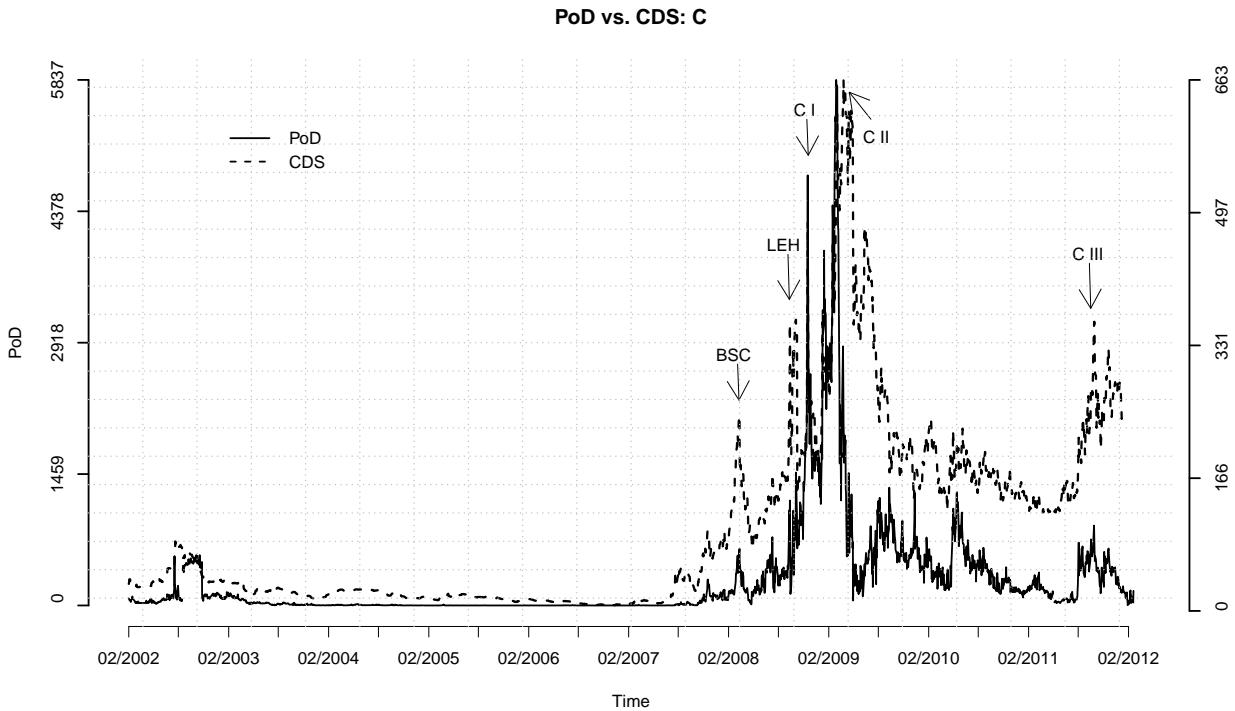

Figure 3: Time series of maturity corrected PoDs versus time series of 5-year CDS for Citigroup, using two different scales

Spearman correlations between CDS and PoDs for the different financial institutions can be found. Correlations with CDS are calculated for maturity corrected and non-corrected PoDs, respectively. After the correction Pearson correlations are on average about 3% higher than before the correction, and are mostly around 70%. Spearman rank correlations are mostly above 80% and remain in general unaffected by the correction. The size of the correlations is extremely high, taking into account that CDS levels imply a default evaluation over one year while our PoDs describe the possibility of default over six month on average. Comparing Spearman and Pearson correlations indicates that there is some non-linear relationship between PoDs and CDS which is not accounted for in the Pearson measure. The maturity

correction seems to linearise the relationship between the two measures to some degree as Pearson correlations increase while Spearman correlations remain unchanged. The effect of the correction is highly reasonable, as we would expect both indicators to exhibit perfect linear correlation (equal size) if there were not the issues of recovery rates (see below), differing risk evaluation horizons and maturity dependence. These issues induce non-linearity to the relationship, which can be accounted for by the Spearman measure but not by the Pearson coefficient. Hence, removing the maturity dependence should increase Pearson correlations and remain Spearman correlations unaffected. We conclude that the correlation analysis does not only tell that PoDs and CDS exhibit quite similar dynamics, but also supports the effectiveness of our maturity correction.

An entirely different picture is obtained if we look at the levels of the PoDs and the CDS, which differ considerably. This is a feature that is true for all of our PoD and CDS time series, and is highlighted in an even more striking way in Figure 4, in which the PoD and CDS time series for Lehman Brothers are plotted at the same scale. At the time of LEH's bankruptcy, PoD values rose to 2500 BP whereas CDS values rose only to 700 BP. This feature is also shown in Figure 3 where at the time of the rescue of Citigroup by the US government on November 23, 2008 (C I) we observe a level of around 4600 BP for the PoDs versus around 250 BP for the CDS. Assuming we interpret the levels of both measures as probabilities of default, the option implied PoDs are almost 18 times as high as CDS implied PoDs in the Citigroup case and more than 3.5 times as high in the case of Lehman Brothers. However, this interpretation is not entirely correct since CDS cannot be directly interpreted as probabilities of default. In order to calculate CDS implied PoDs it is necessary to assume some recovery rate (neglecting liquidity and counterparty risk) for the case of a firm's default. This is because CDS are derivatives on debt based securities. In contrast, our option implied PoDs can be interpreted directly as probabilities of default as we use equity based market instruments. Hence, when comparing the levels of the two indicators one has to be aware that if the investors' assumed recovery rate is unequal to zero, a CDS spread implies a probability of default of a multiple of its amount. This fact may lead to severe problems in the assessment of the resilience of some firm, since an equal CDS level for two firms can imply two totally different probabilities of default. Things are further complicated if recovery rates are not constant over time. In contrast, our PoD estimates have the great advantage that larger values of the indicator always imply higher risk, no matter if comparisons are made over time, firms or both.

Taking into account the interpretation problem of the CDS levels, we subsequently make the simplifying assumption that recovery rates are constant over firms and time in order to
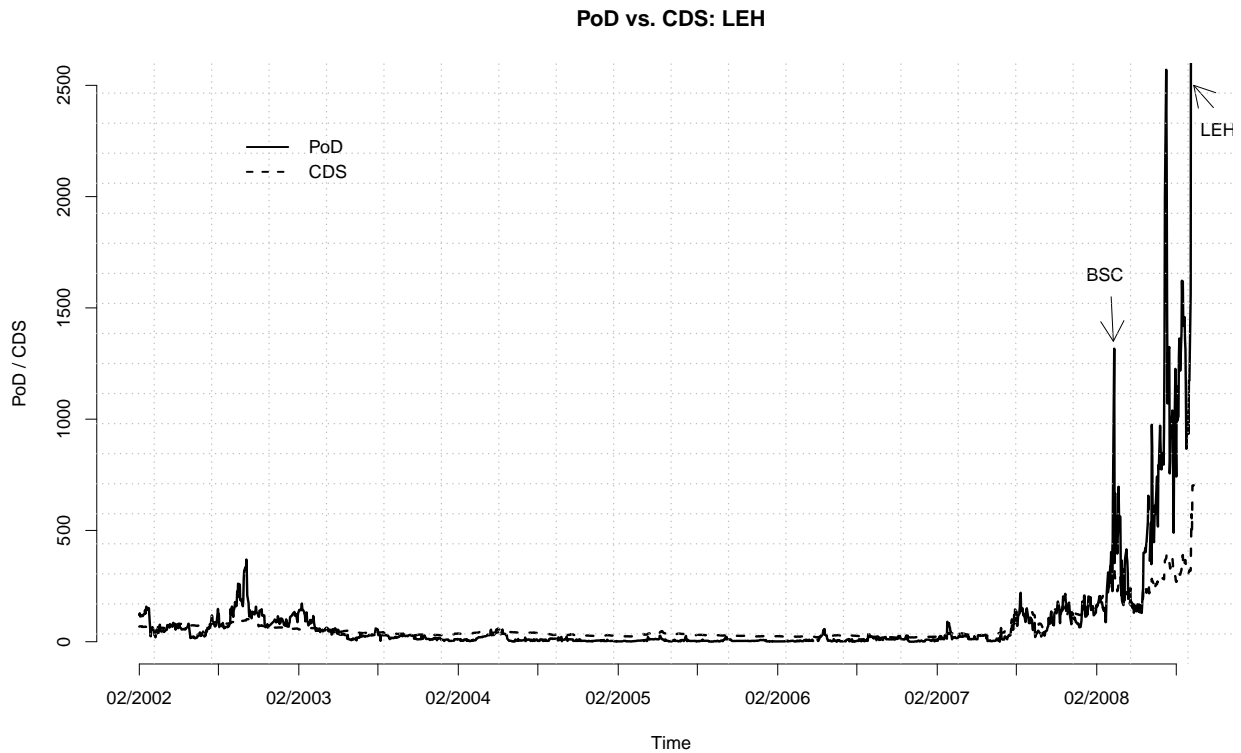
PoD vs. CDS: LEH

Figure 4: Time Series of maturity corrected PoDs versus time series of 5-year CDS for Lehman Brothers, using one scale

be able to compare the indicated risk levels of CDS and PoDs. To compare CDS and PoD levels it is first of all important to know what can be regarded as a high/low value for the respective indicator. From Table 4 (Appendix) one can see that before the crisis (chosen start date for crisis: July 2007) the average PoDs range from 2 BP for WFC to 29 BP for MS and average CDS from 14 BP for BAC to 60 BP for MBI.[16] PoDs are on average lower than CDS spreads, which is due to the shorter implied default evaluation period. The average values after July 2007 range for PoDs from 61 BP for GS to 1135 BP for MBI and for CDS from 89 BP for PNC to 1162 BP for MBI. Knowing this, we can now assess that the PoDs' observed maximum values for LEH of 2500 BP and for C of even 5800 BP are extraordinarily high. In contrary, the maximal values of CDS for LEH of 700 BP and for C of 600 BP are only about half as high as the average CDS value for MBI after the crisis. Given LEH and C faced maximum financial distress at time of the respective maximum values, we obtain a first indication that CDS level are less reasonable in size as option iPoDs. The possible reason for the implausibility of the CDS levels might be that different recovery rates are assumed for the different firms.

---

[16]Only taking into account those banks for which CDS and PoDs are available.

Yet, the indicators' average values alone do not provide us with a very clear picture on how risky a specific firm is at a specific point in time. For this we need relative information. More precisely, information on how large the level of a specific bank's indicator is (i) relative to the indicator level of other banks, (ii) relative to the level of a financially sound bank, and (iii) relative to the bank's indicator level in the past. In the following, we will focus on the calculation and evaluation of these three 'relative' risk indicators.

In the next section, we derive a new, systematic measure of financial risk, which we interpret as a proxy for the systemic risk in the overall US financial sector and which we will use for the first of our three relative risk indicators.

5.2 Systematic/Systemic Risk

Before we proceed with the discussion of the relative risk indicators, it is essential to clarify what we exactly mean by the term systematic risk and how it is derived. As a result, we gain a second type of indicator which, in contrast to our firm specific PoDs, signals the degree of distress of the overall financial sector. This systematic risk indicator can serve as a complementary analysis tool for supervisory authorities. In addition, the new measure will allow us to separate the common dynamic of all firms' indicators from the individual indicator time series. By doing so we can take into account the fact that all banks display historically high indicator levels during crises periods and we can later identify those banks that are especially risky, given the degree of distress in the entire banking sector.

Given the strong similarities in the dynamics of our PoD series (see Figure 13 in the Appendix), there is strong evidence for some latent factor that predominantly drives the pattern of our dataset. We interpret this unobservable joint factor as systemic or financial sector risk. In order to derive a measure for that financial sector risk component, we aim to segregate the systematic risk of our PoD data from the banks' individual idiosyncratic risks. Therefore, we apply a principal component analysis (PCA) and regard the first principal component (PC) as a proxy for the overall financial sector risk. In other words, the first PC represents that joint factor which mainly causes the correlation between the variables and likewise explains the largest part of the variation in the dataset and is therefore regarded as the driving force behind the common pattern underlying our data.

Technically speaking, the first PC can be regarded as a linear combination of the optimally weighted PoD data. Hence, our financial sector risk indicator is a weighted mean of our banks' PoDs. The term optimally refers to the fact that there exists no other set of weights that leads to a PC which accounts for a larger amount of variance in the data. Thereby, we

weigh these banks in the sample the most which exhibit the highest percentage of the total variation and, hence, exert the strongest influence on the overall financial sector risk.

Due to data restrictions we could not use the complete set of banks to calculate our PCs. For the PC of the PoDs we used GS, WFC, C, BAC, JPM, AIG, MS, MBI, PNC and STT. No data are available for LEH, BSC, WB, MER, CFC and WM after their takeover or default and also BLK, BX and BK contain too many missing values. For the CDS we took GS, WFC, C, BAC, AIG, MS and MBI for the same reasons just explained. Consequently, given the data restrictions we decided to use the maximum number of companies available for the respective PC calculation. We found that all eigenvector elements are positive and that the first PC of our PoD data explains 79.21% of the total variation and the first PC of the CDS explains 83.66% of the overall variance in the dataset. Given that high explanatory power, we consider the first PC as an appropriate proxy for the systemic risk inherent in the financial sector.

The PCs of our PoD and CDS series are depicted in Figure 5.[17] It is striking that both graphs exhibit highly similar dynamics but differ significantly in their level values - as already found for the individual bank PoDs and CDS. This shows again that our estimation framework yields very plausible results. Regarding its predictive power, our systemic risk indicator clearly indicates the late consequences of the early 2000s recession, the Bear Stearns takeover by JPM on March 16, 2008 (denoted BSC in the graph), the Lehman Brothers bankruptcy on September 15, 2008 (denoted LEH in the graph) as well as the worsening conditions in the wake of the European sovereign debt crisis (labeled ESDC in the graph). Our approach is very well capable to significantly predict incisive events in US economic history on a sector wide scale and, moreover, indicates magnified distress levels even prior to certain incidents.

Figure 15 (Appendix) shows that the above mentioned properties do not hinge on the exact choice of banks used for the PC determination. In Figure 15, we used the same specification of banks for the PC determination of the PoDs as used for the PC calculation of the CDS.[18] One can clearly see that the main findings remain unchanged.

---

[17]Beforehand, we tested our PoD and CDS series for unit roots (results available upon request). We found that except for the 'crisis banks' which went bankrupt or were taken over during 2008 (LEH, BSC, WB, MER, CFC and WM), the series are significantly stationary. The crisis banks' indicators, though, rose to very high levels shortly before failure and then abruptly ended, hence, pointing towards unit root behaviour. Nonetheless, theoretically it is hard to justify that the CDS and PoD series are non-stationary. If the series were not mean-reverting, very high values lasting over a lengthy periode of time would sooner or later trigger default of the respective company.

[18]The first PC of the PoD for a bank sample equal to the one used for the PC calculation of the CDS explains 72.28% of the total variation in the data, which is still a very high percentage.
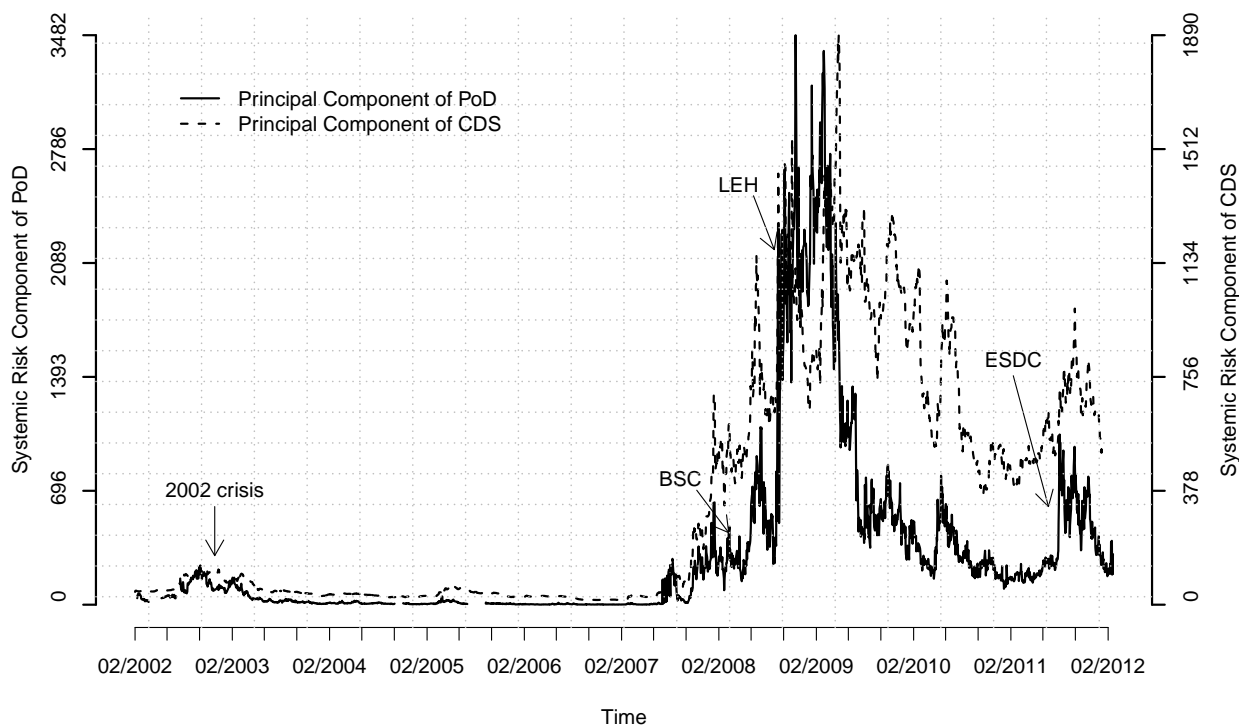
Principal Components of PoD and CDS Series

Figure 5: Systemic risk of PoD and CDS series measured in basis points over the complete sample

In order to have a closer look at the dynamics of our financial sector risk indicator, Figure 6 focuses solely on the crisis period. We highlighted important events during the financial crisis in order to verify the predictive power of our financial sector risk indicator. BSC I marks the Bear Stearns hedge fund troubles of mid-2007, which were the first forerunners of the financial crisis. CFC I represents the announcement of the Countrywide takeover by BAC on January 11, 2008, BSC II denotes the Bear Stearns takeover by JPM on March 16, 2008 and CFC II stands for the Countrywide acquisition by BAC on June 25, 2008. All four events were accompanied by a significant increase in the systemic risk component of our PoDs. More importantly, worsened financial conditions were signalled in advance to the specified events. For instance, long before the CFC II event actually took place, the surge in the systemic risk indicator already pointed to impending financial upheaval. The biggest surge in the indicator took place in the course of the Lehman Brothers bankruptcy (LEH event), the BAC announcement of the Merrill Lynch purchase on September 14, 2008 (MER event) and the bailout of AIG by the Federal Reserve Bank of New York on September 16, 2008 (AIG event). Shortly after the bankruptcy of Washington Mutual on September

23

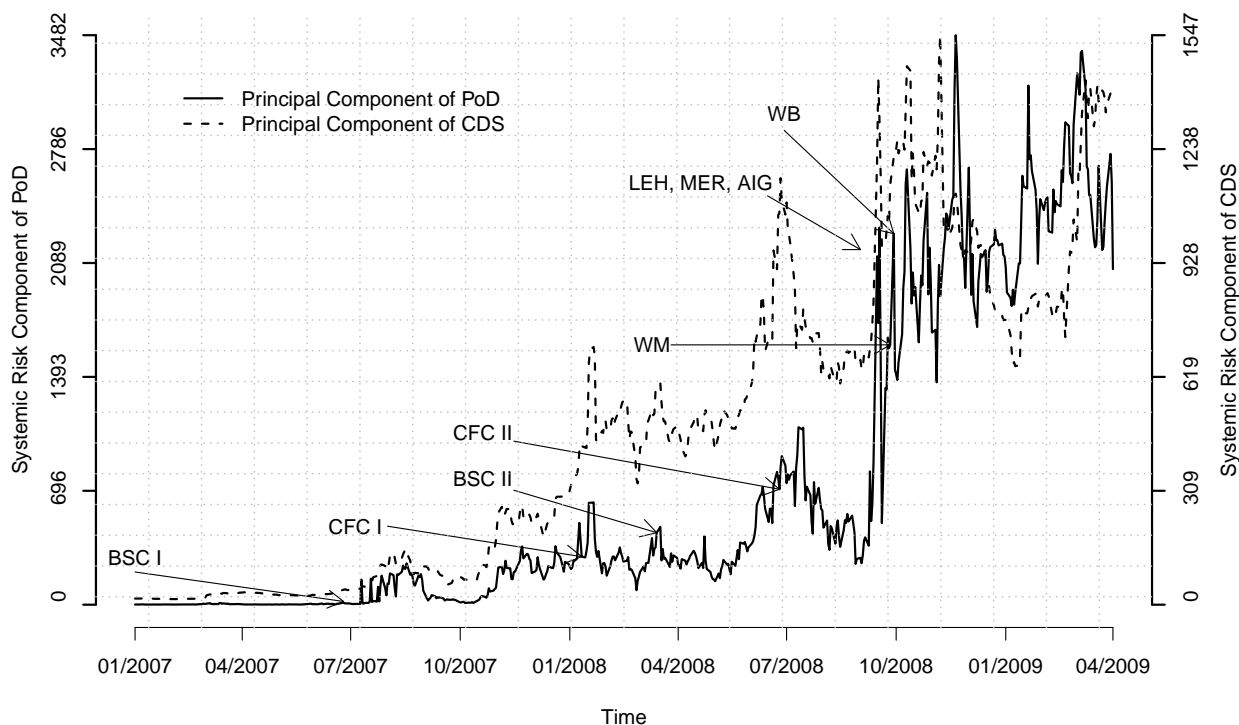**Principal Components of PoD and CDS Series**

Figure 6: Systemic risk of PoD and CDS series measured in basis points from January 2007 to April 2009

26, 2008 (WM event) as well as the announcement of the acquisition of Wachovia by WFC on September 29, 2008 (WB event) market conditions further worsened as indicated by a substantial increase in the systemic risk measure. Again, our financial sector risk indicator rose sharply already prior to these incisive events.

Given the strong predictive power of our financial sector risk indicator, we are interested in building an indicator which puts more emphasis on the relative intensity of distress in the financial sector. Therefore, to classify the levels of the systemic risk component, we built 'crisis bands' based on four pre-specified events. We took the 50 days average of our PoD and CDS systemic risk components around the early 2000s recession (weakest crisis level), prior to the BSC takeover by JPM (second weakest crisis level) and prior (second strongest crisis level) as well as posterior (strongest crisis level) to the LEH bankruptcy. The four crisis bands are depicted in Figure 16 (Appendix). Based on these crisis thresholds we construct an ordinal scale which indicates on a five step interval the intensity of risk in the US financial sector over time - yielding a scaled magnitude of systemic risk which helps to set current crisis levels in relation to previous ones. That ordinal crisis scale is shown in Figure 7.
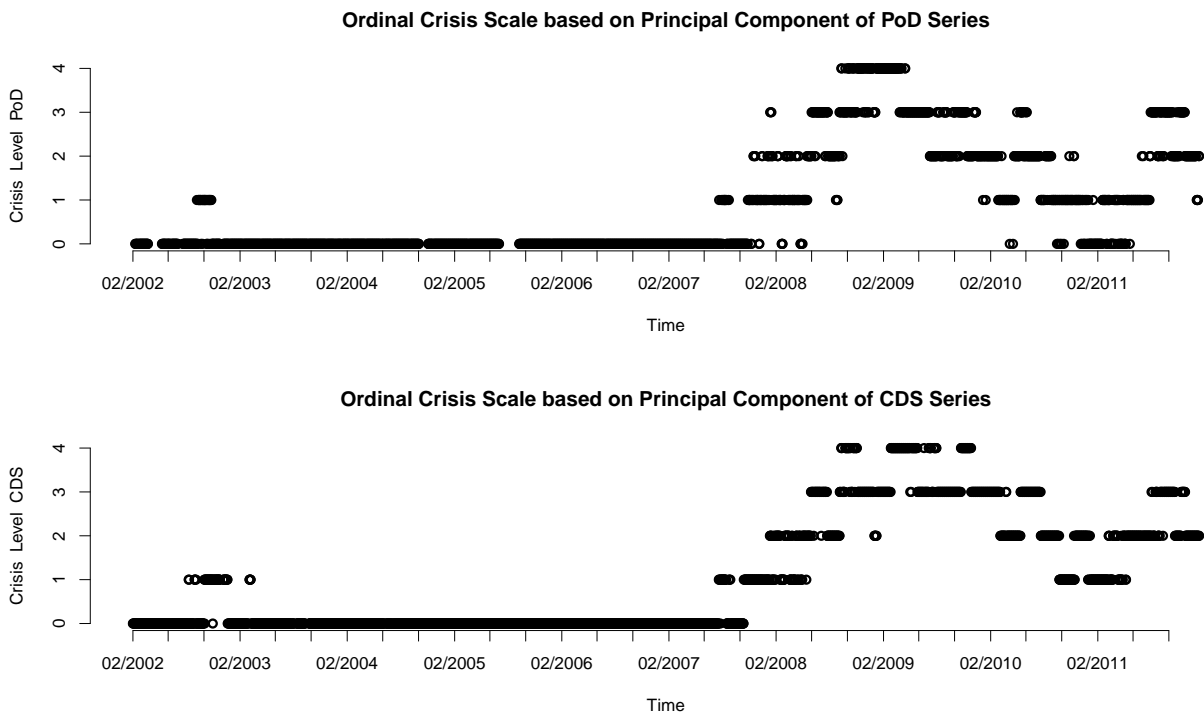
Figure 7: Ordinal crisis indicators for the systemic risk of PoD and CDS series

Again, the dynamics between the PoD and CDS ordinal crisis indicators are highly similar. At first, the indicator signals no crisis at all but soon jumps upon a step one crisis level, associated with the early 2000s recession. This is followed by a relatively long non-crisis period (crisis level zero) before the indicator starts to rise again, triggered by the BSC hedge fund troubles. This initial surge sets in earlier for the PoDs compared to the CDS scale. The ordinal indicator reaches its maximum at the point of the LEH bankruptcy and remains at crisis level four during times of severest financial distress. Due to governmental stability measures, conditions improved and crisis levels decline but never reach pre-crisis levels again. The last surge in the ordinal scale indicates the European sovereign debt crisis.

5.3 Relative Risk Analysis

Relative to Systemic Risk

As we could verify that our proxy for the systemic risk component of the US financial sector provides very conclusive results, we apply this financial sector indicator to our relative risk analysis. Figure 8 depicts exemplarily the PoD and CDS series of LEH and WM relative to the systemic risk. We consider a period of six months prior to the Lehman event in order to

evaluate the predictive power of our financial stability indicator in advance of severe events. Therefore, we subtract the financial sector risk component from the original PoD and CDS series of the respective companies and gain a 'relative risk spread' measured in basis points.[19] This spread indicates whether a company's distress level is considered to be higher or lower compared to the prevailing overall financial sector risk. The upper graph of Figure 8 exhibits
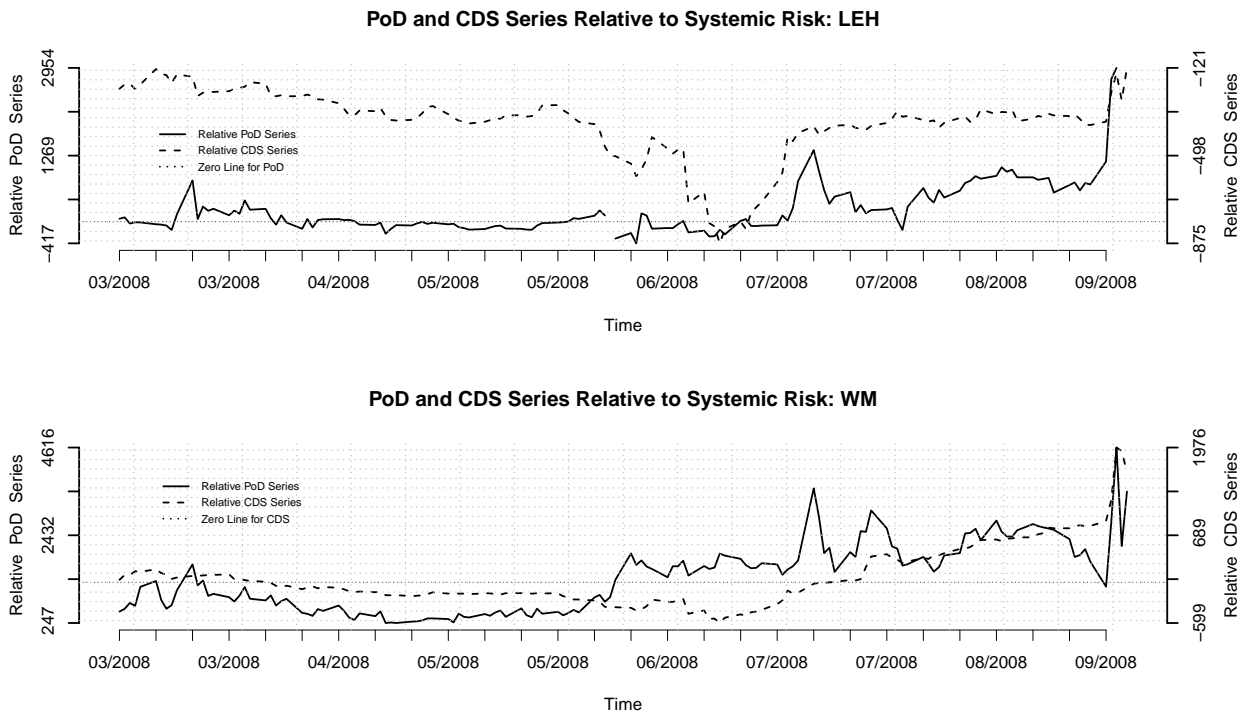

Figure 8: PoD and CDS series relative to systemic risk six month prior to the Lehman event; measured in basis points

the relative risk spread for the PoD and CDS series of LEH. As illustrated by the zero line, the relative PoD time series moves constantly around zero but exceeds that value in July 2008, two and a half months before the actual collapse of LEH. Hence, way in advance of the bankruptcy, our indicator is able to signal in relative terms an increased default risk. One trading day before failure, the risk spread even rose to around 29.5%, implying a default probability which is 29.5 percentage points higher than that of the average overall financial sector. In contrast, the CDS spread remains negative over the complete sample, which implies LEH's distress level to be below the systemic one. This clearly underestimates LEH's inherent default risk especially with regard to the looming insolvency.

The lower graph of Figure 8 depicts the relative PoD and CDS time series for WM. Here,

---

[19]For the sake of consistency, for the systemic risk calculation we used in this case the same bank sample for the PoDs as for the CDS. However, results remain unchanged if we used the PoD sample of section 5.2.

both curves signal a relative risk level above the systemic one in the latter part of the sample, which is plausible since WM collapsed on September 26, 2008. However, the PoD spreads are positive over the complete sample, whereas the CDS curve does not cross the zero line before the end of July 2008. This result speaks again in favor of the predictive power of our option implied financial stability measure. In addition, the absolute values of the PoD spreads are decisively higher than the CDS ones, which is more consistent with the strained financial condition of WM at that time.

Table 5 (Appendix) provides an overview of the relative riskiness of the remaining financial institutions in our sample. We aggregated the information by taking the average of the PoD and CDS spreads for each institution over an interval of one to ten days, ten to twenty days, twenty to forty days and thirty to sixty days prior to the Lehman event. The results substantiate the findings from above that our PoDs are superior to the CDS in identifying the most financially troubled institutions already prior to the Lehman Brothers breakdown (see highlighted LEH, WB, MER, WM and AIG). In a consistent manner, the PoD spreads of the most resilient banks like GS, WFC and JPM exhibit negative spreads since these institutions weathered the financial turmoils quite well.

Relative to the most resilient bank

It is important to consider that the derived systemic risk component used in the analysis above can be mainly driven by one or a few high risk institutions. As a result, a bank's risk level below the systemic risk does not inevitably imply that this bank is riskless. Therefore, we extend our analysis and set our PoDs and the CDS in relation to the most resilient bank. As most stable bank we chose JPM for two reasons: First, PoD and CDS levels are the lowest on average over the complete sample and second JPM has the fewest missing values for both indicators.[20] Figure 9 displays PoD and CDS series for LEH and AIG relative to JPM for a period of six months prior to the Lehman event. In both graphs the PoD and CDS series have a highly similar dynamic structure and spread values are almost exclusively positive. Hence, the relative PoDs as well as the relative CDS point to an increased risk level prior to the Lehman event. Moreover, in the case of LEH even the Bear Stearns crisis of March 2008 is characterized by amplified spread levels. The absolute magnitude of the spread levels, though, differ to a large extent, implying that the PoDs signal a higher degree of distress. This is highly reasonable given the grave financial conditions of both institutions.

In essence, these results are confirmed by Table 6 (Appendix). Again, we aggregated our PoD

---

[20]We tested several reasonable specification for the most resilient bank and found that the main results do not hinge on the exact choice of JPM as most stable institution.
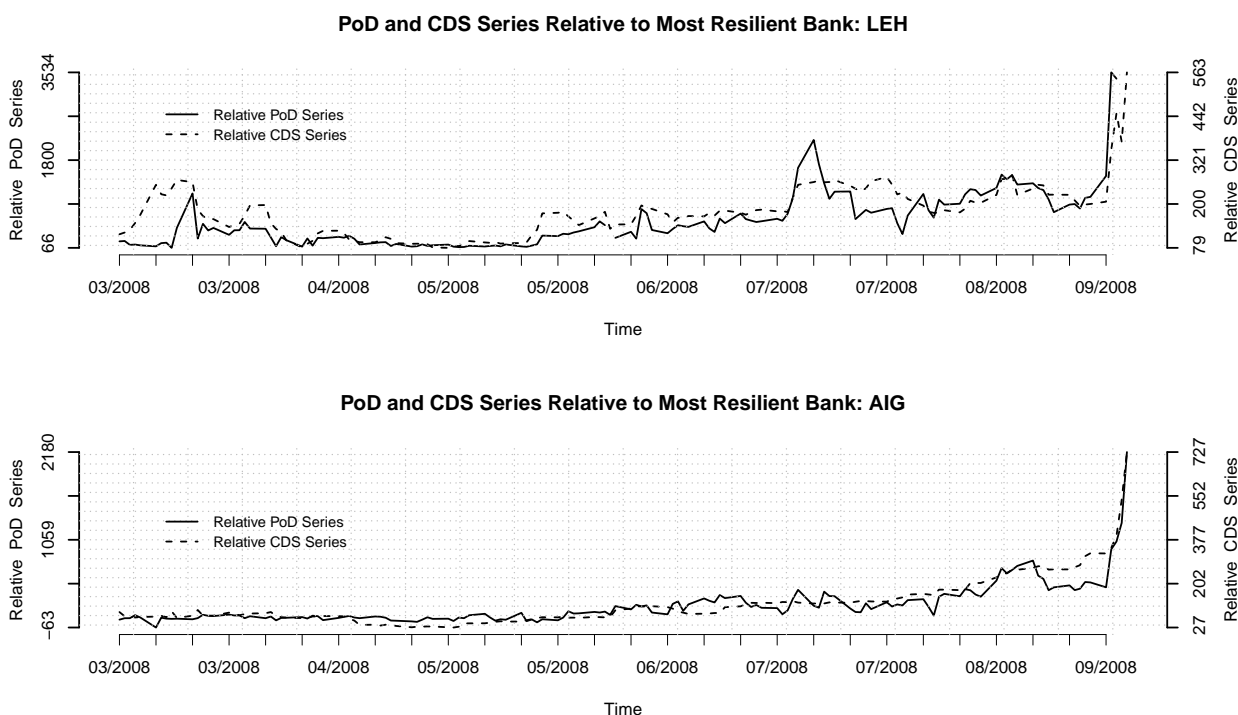
PoD and CDS Series Relative to Most Resilient Bank: LEH

PoD and CDS Series Relative to Most Resilient Bank: AIG

Figure 9: PoD and CDS series relative to the most resilient bank (in both cases JPM) six month prior to the Lehman event; measured in basis points

and CDS spread data for the relevant financial institutions prior to the Lehman event in the same way as done in Table 5. As one would expect, beside very few values, all spreads are positive. More interestingly, for banks like WM, MER, WB and LEH which did not survive the financial crisis, the PoD spread values are consistently higher than the CDS spreads. The capability of our PoDs to act as a financial early warning system prior to disruptive events is thereby further demonstrated. Also institutions like AIG or C, which may have survived the financial turmoils but underwent severe difficulties, are assigned higher spread values for the PoDs than for the CDS.

Relative to the own history

For a complementary risk analysis one also needs to evaluate a bank's performance in relation to its own history. Especially, in calmer periods prior to financial turmoil spreads relative to the most resilient bank and to the financial sector risk might be rather small. This is due to the fact that investors' risk perceptions do not diverge much in financially stable periods. If for instance the two previously discussed relative risk indicators did not signal elevated degrees of distress that would not inevitably imply that a bank's financial condition has

not worsened yet. Hence, a firm's PoD relative to its own history can be regarded as a first warning sign, whereas the indicators relative to the most stable bank and the systemic risk follow with a lag.

Therefore, in Figure 10 we consider a period of six months prior to the Bear Stearns event of March 14, 2008 (first takeover arrangements by JPM induced by the Federal Reserve Bank of New York), reflecting a rather tranquil financial environment. It is about to see if our relative PoD series is able to identify the increasing risk emerging at the horizon. The PoD and CDS spreads are determined by taking the difference between the original series and its mean over a moving window of the last three years. As initial interval for our mean calculation we chose January 2004 to January 2007. Figure 10 depicts the PoD and CDS spreads for Bear Sterns and Lehman Brothers relative to their own history. In both graphs,
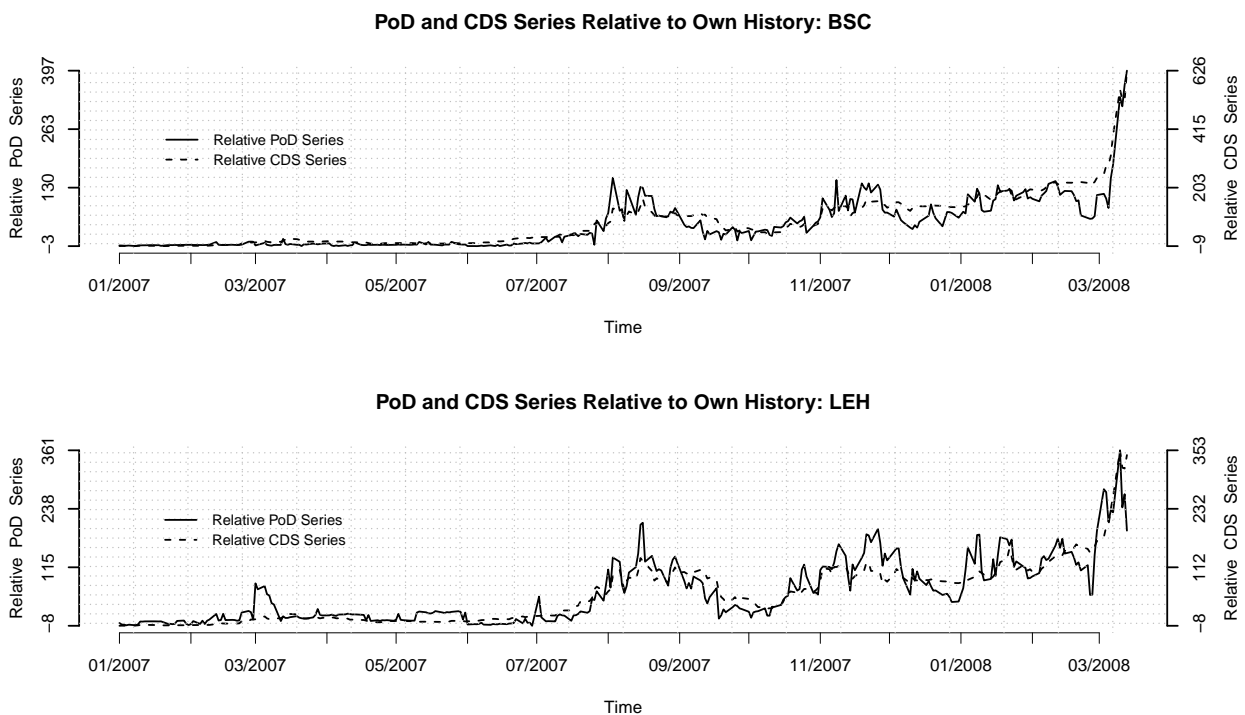
PoD and CDS Series Relative to Own History: BSC

PoD and CDS Series Relative to Own History: LEH

Figure 10: PoD and CDS series relative to the own history six months prior to the Bear Stearns event; measured in basis points

the PoD and CDS spreads exhibit highly similar dynamics and are also quite similar with regard to the magnitude of their spread values. This is not very surprising since - as already mentioned in section 5.1 - CDS and PoD levels are relatively close in financially solid periods. More strikingly, the PoD as well as the CDS spreads for Bear Stearns and Lehman Brothers clearly indicate the Bear Stearns hedge fund crisis of mid 2007, which was the first signifier

of the looming financial crisis. In the following, the spread series remain on their elevated levels and surge shortly before the Bear Stearns event.

Table 7 (Appendix) aggregates the results for the remaining financial institutions of the sample. This time, however, we chose different intervals for our mean calculation in order to evaluate if and how the relative risk measures signal the different distress levels over time. We took the mean over an interval of: 1 to 10 days (to capture the period close to the event), 30 to 136 days (September 2007 to February 2008 as the period between the hedge fund crisis and the final event), 137 to 180 days (July 2007 to September 2007 as the period of the hedge fund crisis itself) and 181 to 225 days (May 2007 to July 2007 to represent the financially resilient period) prior to the Bear Stearns event.

Almost exclusively all institutions' relative risk indicators rise over time. Starting from very low and partly even negative values, all series increase to noticeable higher amounts in the wake of the hedge fund crisis. The averaged spread values further rise which points to worsening financial conditions prior to the Bear Stearns event. This is especially the case for WM and CFC whereas spreads of more stable banks like GS and JPM increase only moderately. BSC's and LEH's spread values also rose rather slightly which suggests that their collapse came quite surprisingly relative to their own history. This involved not only investors from equity- but also from debt markets. Shortly before the final event, almost all companies' spreads increased sharply given that the moving window average already accounts for past risen values. Hence, results presented in Table 7 are fully in line with economic events which occurred at that time.

This section showed that also on a relative scale, our option implied financial stability indicator provides plausible and consistent results, and gives proof of its high predictive power. As consequence, these properties also speak in favor of our RND estimation procedure in general.

6 Conclusion

In this paper we applied the so-called option iPoD methodology to a dataset of option prices for 19 of the largest US financial institutions, ranging from February 2002 to February 2012. We showed how to empirically implement the framework in order to obtain consistent and smooth time series of option implied Probability of Defaults (PoDs). This was achieved by the appropriate choice of liquidity weights, the use of a suitable maturity cycle of option contracts and a refined option data filtering technique. The option implied PoDs are esti-

mated as probability mass point for a future stock price of zero in a non-parametric risk neutral density (RND) estimation framework. To obtain the time series of RND/PoD estimates, alternately five-, six-, and seven month call option contracts were used. Subsequently, maturity dependence in the time series was removed by applying a non-parametric quantile regression approach to the pooled PoDs.

The time series of PoDs for the different financial institutions were comprehensively evaluated regarding their plausibility and consistency in size and regarding their signalling/predictive power in advance and during crises periods. To do so, we contrasted our indicators to historical events and to time series of 5-year Credit Default Swap Spreads (CDS). We found that the PoD estimates are very well capable to signal the occurrence of adverse shocks to the financial sector as a whole as well as to specific financial institutions - in most cases even way in advance of the actual events. Correlation analyses between 5-year CDS and the option implied PoDs showed that both indicators exhibit highly similar dynamics and gave strong indication for the effectiveness of our maturity dependence correction methodology. Comparing the levels of CDS and PoDs, we found clearly differing values, PoDs being very much larger in advance and during times of distress than CDS. However, in order to derive CDS implied probability of defaults, recovery rates for the different firms and points of time have to be assumed which in turn severely complicates the direct interpretation of the levels of CDS as a risk measure. Hence, we stress the valuable advantage of the iPoD framework, which provides crises indicators that can be interpreted in a direct way as (risk neutral) probabilities of default and, hence, allows for a simple comparison of risk levels over firms and points of time.

For a more clear interpretation of the respective levels of CDS and PoDs, we further calculated three different 'relative' risk indicators. We analysed the spreads of the CDS/PoDs for the different banks relative to the prevailing systematic risk in the financial sector, relative to the risk of the most resilient bank in the sample, and relative to the respective bank's risk in the past (its own history). In this context, we calculated a proxy for the prevailing systematic risk in the financial sector. Therefor, we applied a Principal Component Analysis to the time series of PoDs/CDS. We identified that principal component with the most explanatory power regarding the variation in the sample of the original time series as the latent factor that drives the risk of all of the firms to some extent. The proxy provides a self contained financial stability measure, that can be used to gauge the stability of the US financial sector as a whole. In the analyses of the dynamics of the systematic risk component we could confirm its high signalling and predictive power.

The evaluation of the relative risk indicators showed that the iPoD approach was able to

identify the high risk banks in advance of the Lehman Brother's bankruptcy. In this context, we could emphasize the clearly superior signalling/predictive power of the iPoDs compared to CDS. During this period, the most informative signals were given by the spreads relative to systematic risk, followed by those relative to the most resilient bank. The spreads relative to its own history, in turn, clearly signalled the first turmoils in the course of the US sub-prime crisis in mid-2007. This is true for the CDS and PoD derived relative indicators.

Finally, given the results from our comprehensive empirical analyses we stress that the option iPoD framework provides highly informative risk indicators, which in many cases are superior to CDS with regard to its signalling/predictive power. Importantly, the amount of information provided by the methodology is also superior to almost any commonly used financial risk framework, since one does not only obtain PoDs, but at the same time the corresponding (compatible) asset value distributions for the different firms over time. By evaluating the plausibility/consistency of PoDs in the paper, we implicitly also evaluated the plausibility/consistency of the respective asset distributions (RNDs), as the PoDs are a function in the RND shape parameters. This makes the framework highly attractive for multivariate risk analyses, in which joint PoDs and asset distributions of several banks are derived. In the most straightforward way, such a approach could be implemented using the entropy based Consistent Implied Multivariate Density Optimization (CIMDO) methodology suggested by Segoviano (2006).

References

Agmon, N., Alhassid, Y., and Levine, R. (1979). An algorithm for finding the distribution of maximal entropy. *Journal of Computational Physics*, 30:250–258.

Alhassid, Y., Agmon, N., and Levine, R. D. (1978). An upper bound for the entropy and its applications to the maximal entropy problem. *Chemical Physics Letters*, 53:22–26.

Breeden, D. and Litzenberger, R. (1978). Prices of state-contingent claims implicit in option prices. *Journal of Business*, 51:621–651.

Capuano, C. (2008). The option-ipod. the probability of default implied by option prices based on entropy. *IMF Working Paper*, 08(194).

Coffinet, J., Pop, A., and Tiesset, M. (2010). Predicting financial distress in a high-stress financial world: The role of option prices as bank risk metrics. *Banque De France Working Paper*.

Cover, T. M. and Thomas, J. A. (2006). *Elements of Information Theory*. John Wiley & Sons, 2nd edition.

Figlewski, S. (2008). *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*, chapter Estimating the Implied Risk Neutral Density for the U.S. Market Portfolio. Oxford University Press.

FSB (2011). Policy Measures to Address Systemically Important Financial Institutions. *Press Release*.

Hastie, T. and Tibshirani, R. (1986). Generalized additive models. *Statistical Science*, 1:297–310.

Hastie, T. and Tibshirani, R. (1990). *Generalized Additive Models*. Chapman-Hall.

IMF (2009). Global financial stability report: Responding to the financial crisis and measuring systemic risks. *World Economic and Financial Surveys, IMF*.

Jackwerth, J.-C. (2004). *Option implied Risk Neutral Distributions and Risk Aversion*. Research Foundation of AIMR (CFA Institute).

Jaynes, E. T. (1957). Information theory and statistical mechanics. *Physical Review*, 106:620–630.

Knaup, M. (2011). *Market-Based Measures of Bank Risk and Bank Aggressiveness*. PhD thesis, Tilburg University.

Koenker, R., Ng, P., and Portnoy, S. (1994). Quantile smoothing splines. *Biometrika*, 81:673–680.

Mayhew, S. (1995). Implied volatility. *Financial Analysts Journal,*, 51(4):8–20.

Melick, W. and Thomas, C. (1997). Recovering an asset's implied pdf from option prices: An application to crude oil during the gulf crisis. *Journal of Financial and Quantitative Analysis*, 32:91–115.

Merton, R. C. (1974). On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance*, 29:449–470.

Segoviano, M. A. (2006). Consistent information multivariate density optimizing methodology. *IMF Working Paper*.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423.

Swidler, S. and Wilcox, J. A. (2002). Information about bank risk in option prices. *Journal of Banking & Finance*, 26:1033–1057.

Vilsmeier, J. (2011). Updating the option implied probability of default methodology. *BGPE Working Paper*.
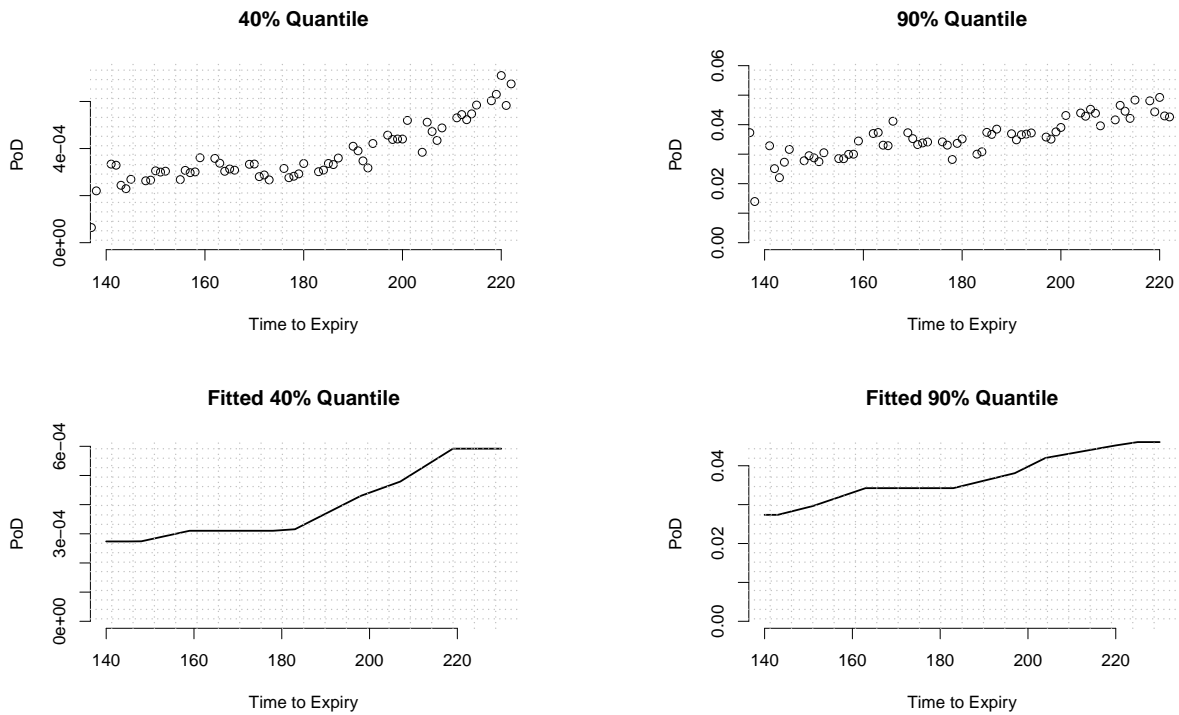
Appendix

Figure 11: 40%- and 90%-Quantile of the PoDs for different time to maturities and the respective fitted functions

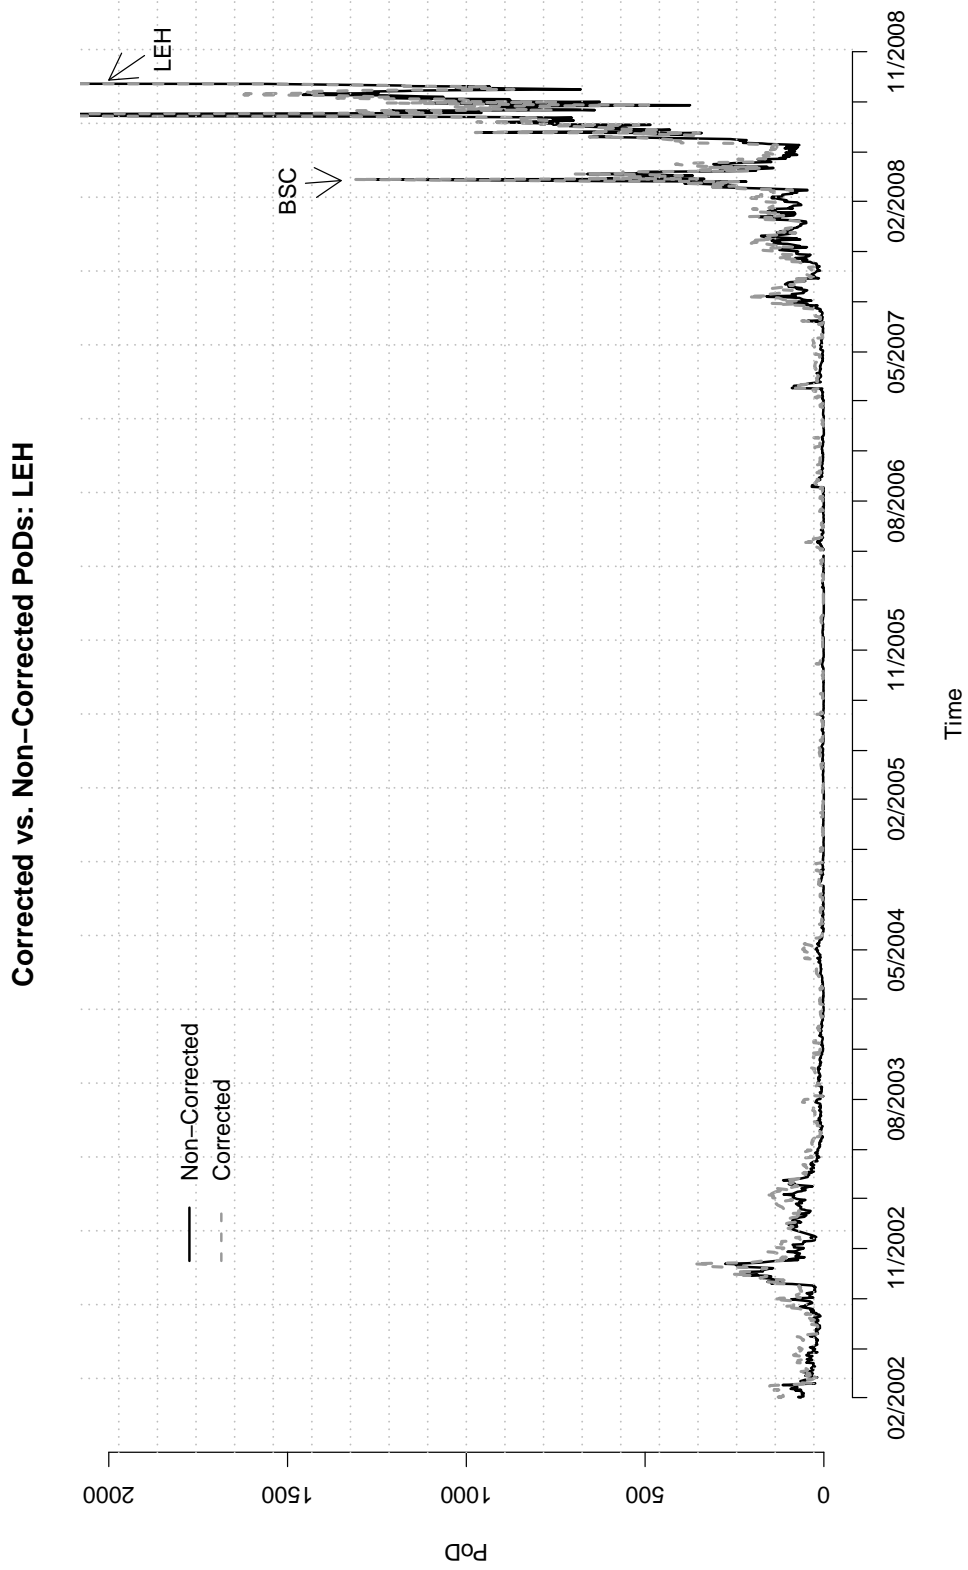**Figure 12:** Maturity corrected and non-maturity corrected PoD time series for LEH measured in basis points over the complete sample
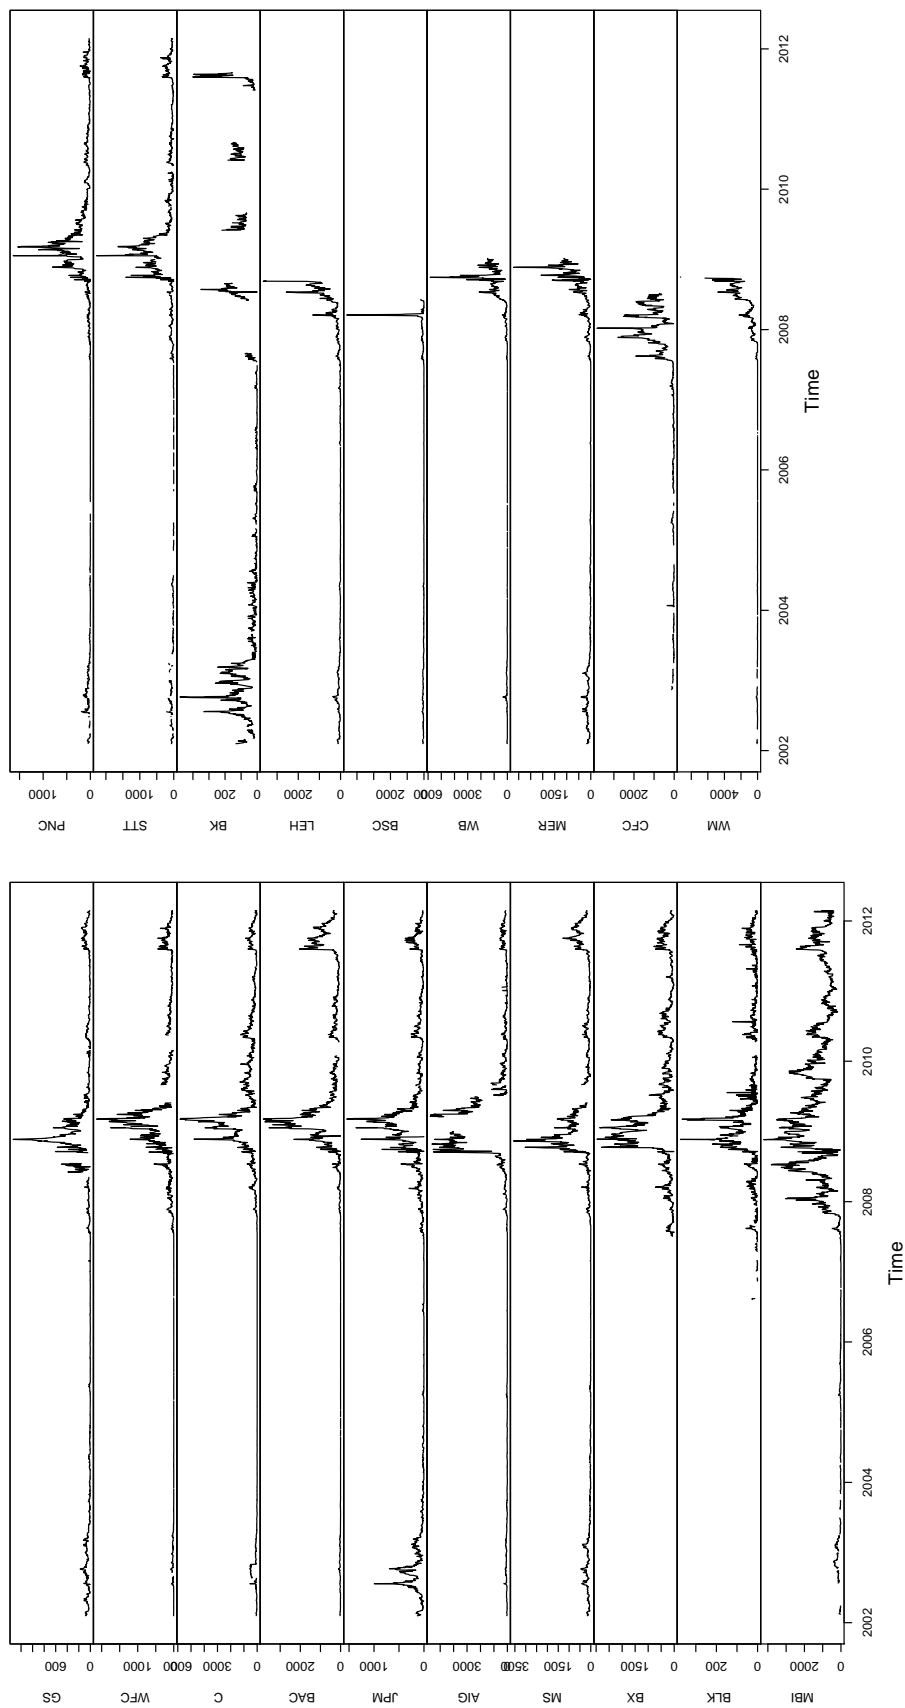
Figure 13: PoD time series for all financial institutions over the complete sample, measured in basis points

37
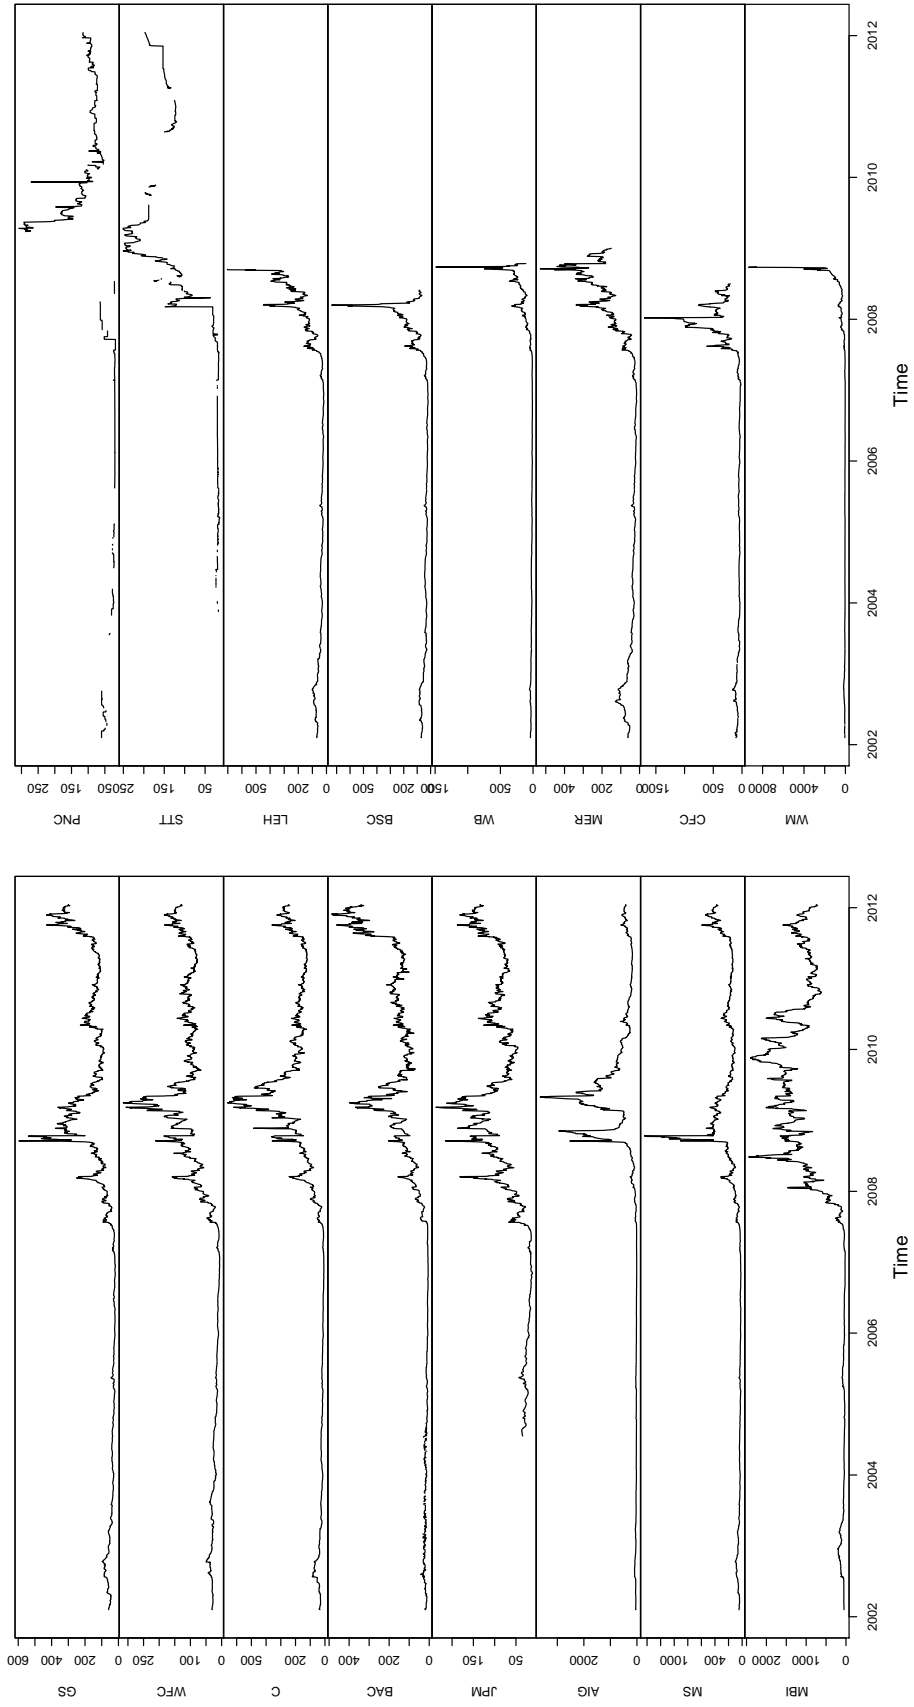
Figure 14: CDS time series for all financial institutions over the complete sample, measured in basis points

38
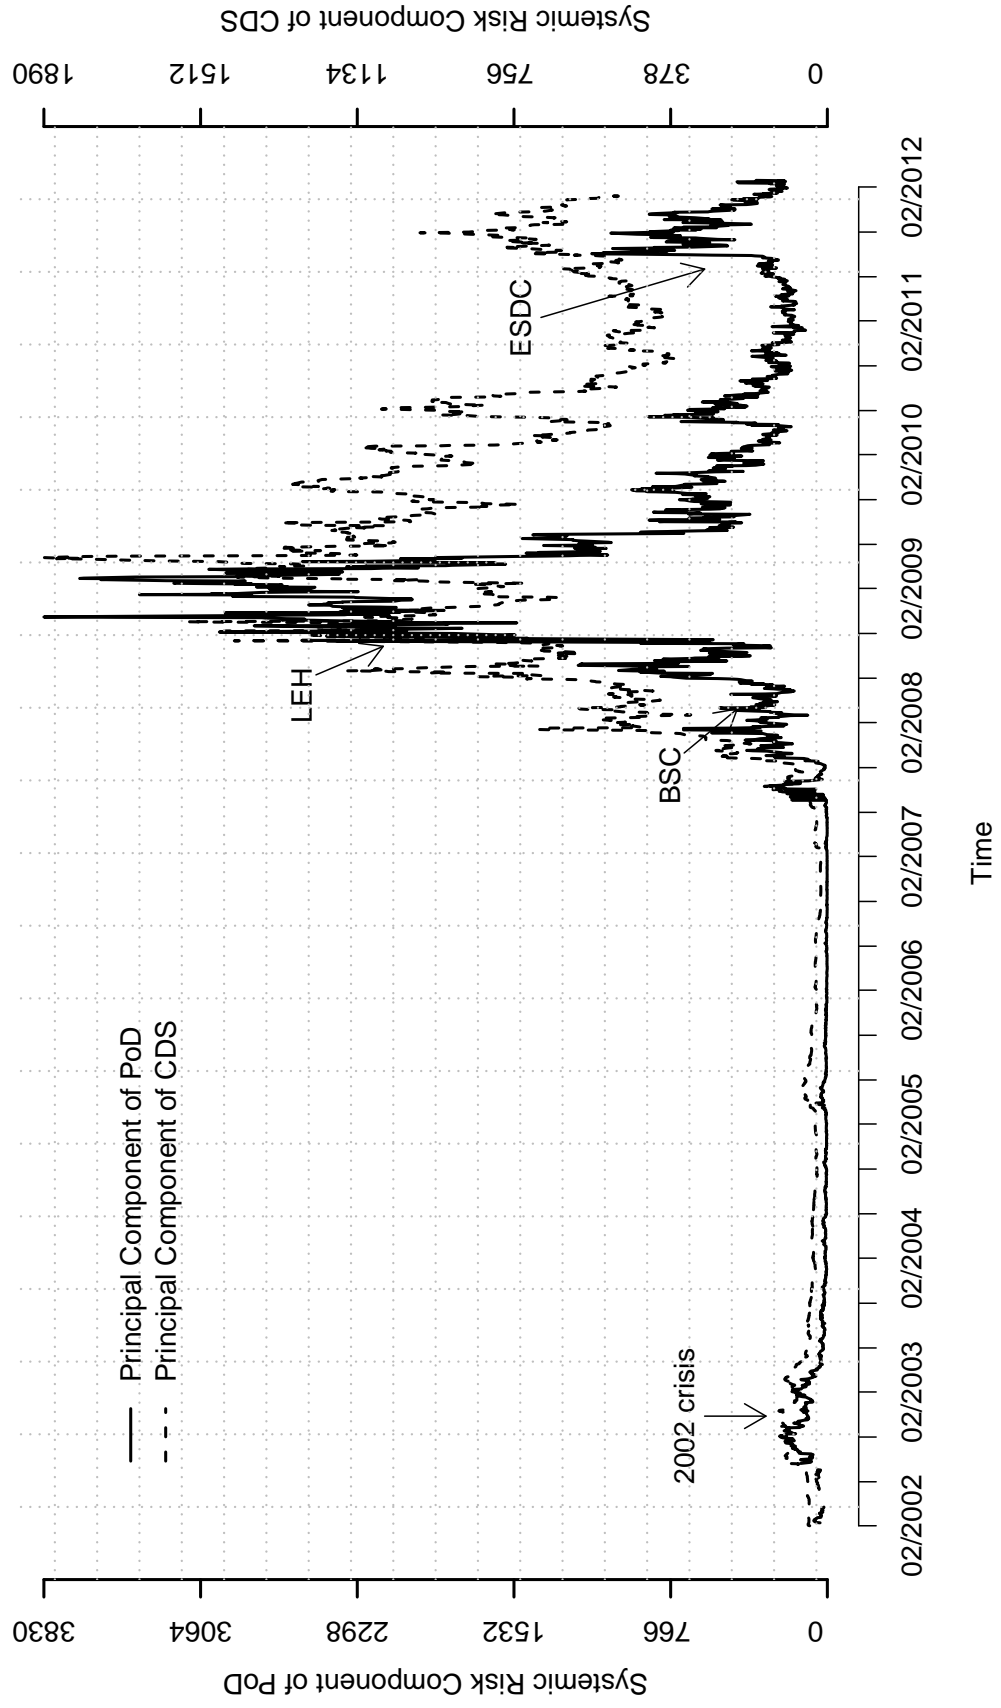
Figure 15: Systemic risk of PoD and CDS series measured in basis points over the complete sample; same banks used for principal components of PoD and CDS series
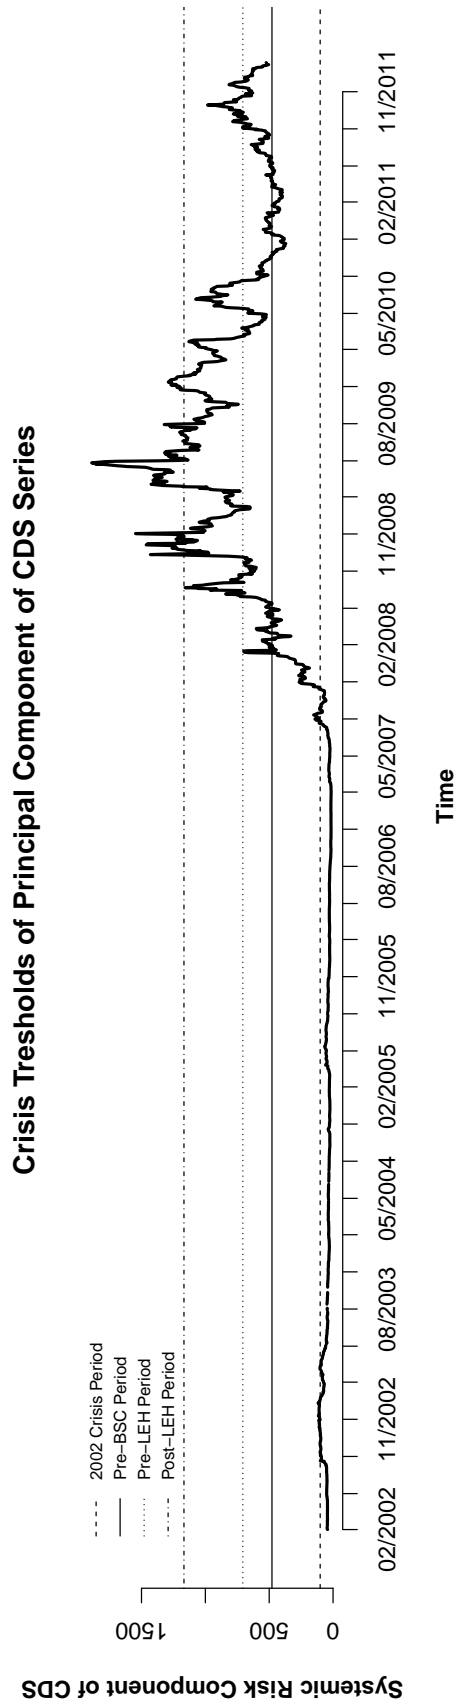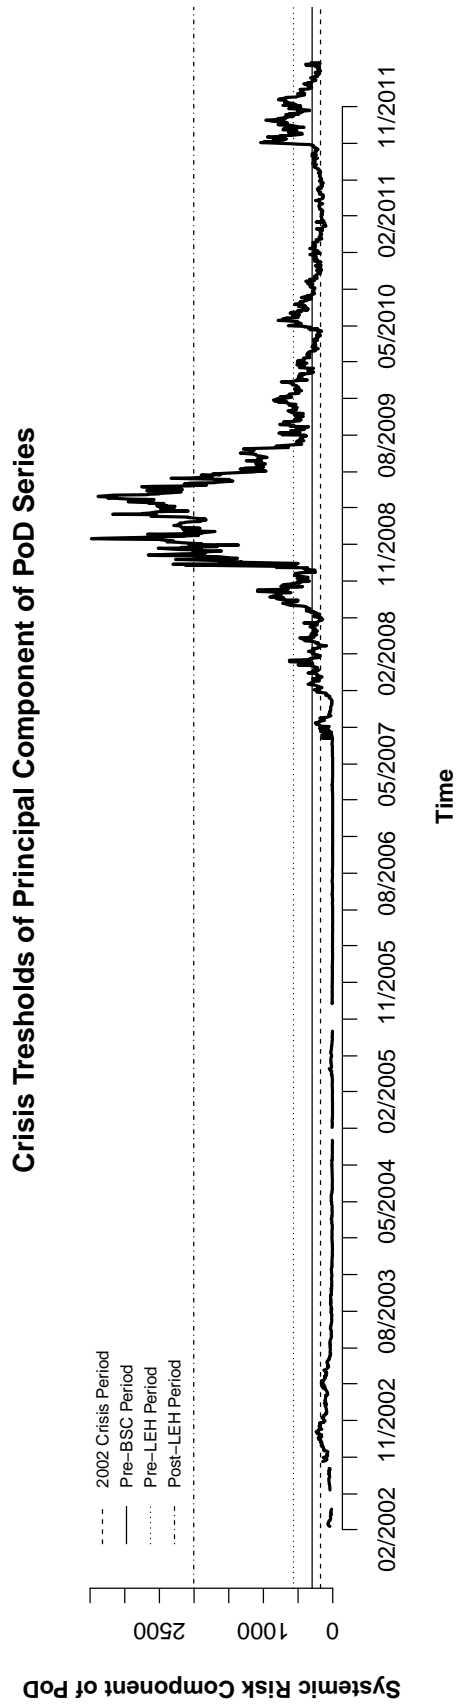
Figure 16: User specified crisis thresholds for the systemic risk of PoD and CDS series

Company	Pearson non-corrected	Spearman non-corrected	Pearson corrected	Spearman corrected
GS	0.60	0.78	0.62	0.78
WFC	0.70	0.91	0.73	0.91
C	0.66	0.94	0.69	0.94
BAC	0.67	0.92	0.69	0.92
JPM	0.65	0.88	0.70	0.89
AIG	0.78	0.90	0.78	0.90
MS	0.72	0.89	0.73	0.89
MBI	0.85	0.94	0.86	0.94
PNC	0.81	0.83	0.85	0.82
STT	0.62	0.83	0.64	0.82
LEH	0.82	0.84	0.84	0.83
BSC	0.45	0.78	0.49	0.77
WB	0.73	0.77	0.74	0.79
MER	0.65	0.90	0.68	0.90
CFC	0.92	0.69	0.92	0.71
WM	0.82	0.79	0.82	0.77

Table 3: Pearson and Spearman correlations between PoD and CDS before and after maturity correction

	Avg.PoDs	Avg.CDS	Avg.PoDs (before)	Avg.CDS (before)	Avg.PoDs (after)	Avg.CDS (after)
GS	33.10	96.76	9.41	36.51	61.33	166.53
WFC	78.81	61.30	1.89	20.37	172.91	108.64
C	242.16	103.72	24.26	22.82	488.33	197.47
BAC	187.13	80.26	2.42	14.13	402.87	155.62
JPM	69.33	66.80	26.41	22.40	114.45	95.36
AIG	277.52	263.05	13.27	28.78	590.85	534.54
MS	154.60	123.87	28.76	36.54	306.87	225.02
BX	345.03	NA	NA	NA	345.03	NA
BLK	16.26	NA	0.72	NA	17.70	NA
MBI	585.63	570.69	26.24	60.38	1134.84	1161.69
PNC	45.63	60.48	5.92	27.40	87.00	89.49
STT	81.30	80.81	11.53	18.96	137.68	133.96
BK	29.16	NA	19.03	NA	74.63	NA
LEH	74.90	67.94	17.32	39.59	328.62	192.04
BSC	23.30	56.34	7.32	37.78	113.75	163.75
WB	112.96	48.54	4.15	21.29	484.12	160.13
MER	82.49	71.11	22.73	38.87	290.66	184.04
CFC	205.47	117.23	21.93	57.14	896.36	432.82
WM	167.47	149.36	3.84	49.75	378.26	570.68

Table 4: Average PoDs and CDS over entire sample period, before July 2007 (before) and after July 2007 (after)

42

Company	PoD (1-10)	CDS (1-10)	PoD (10-20)	CDS (10-20)	PoD (20-40)	CDS (20-40)	PoD (30-60)	CDS (30-60)
GS	-938.50	-535.36	NA	-519.02	NA	-528.68	-698.31	-699.39
WFC	-401.63	-571.10	-327.75	-552.55	-439.10	-555.66	-641.59	-720.22
C	-138.22	-529.06	-122.76	-508.60	-259.61	-518.29	-414.85	-694.79
BAC	-393.72	-568.17	-416.62	-550.93	-517.67	-551.79	-689.16	-723.82
JPM	-437.54	-579.70	-382.36	-568.95	-449.17	-564.82	-649.29	-730.84
AIG	386.14	-203.66	213.20	-313.80	-174.74	-402.92	-399.65	-606.62
MS	-403.90	-473.37	-315.92	-452.86	-391.65	-441.07	-608.45	-617.60
MBI	133.66	475.78	319.97	534.20	1142.67	592.61	1888.70	836.88
PNC	-492.93	NA	-421.74	NA	-521.31	NA	-782.65	-927.36
STT	-473.26	-581.62	-403.34	-568.91	-528.81	-533.36	-751.63	-743.98
LEH	1294.84	-282.63	846.63	-324.13	445.47	-348.05	203.89	-514.05
WB	424.74	-376.51	635.05	-388.34	373.34	-400.79	230.02	-583.92
MER	12.27	-366.53	25.49	-355.08	-117.97	-381.52	-437.81	-558.62
WM	2423.84	1172.43	2561.60	694.21	2172.71	362.79	1991.85	-102.63

Table 5: Average PoD and CDS spreads with regard to systemic risk prior to Lehman event

Company	PoD (1-10)	CDS (1-10)	PoD (10-20)	CDS (10-20)	PoD (20-40)	CDS (20-40)	PoD (30-60)	CDS (30-60)
GS	-65.05	44.35	NA	49.94	NA	36.15	19.34	31.45
WFC	35.91	8.60	54.61	16.41	10.06	9.16	7.69	10.62
C	299.32	50.64	259.60	60.35	189.55	46.53	234.43	36.05
BAC	43.83	11.53	-34.26	18.03	-68.50	13.03	-39.87	7.02
AIG	823.68	376.04	595.57	255.16	274.43	161.90	249.63	124.22
MS	33.64	106.34	66.45	116.09	57.52	123.75	40.83	113.24
MBI	571.21	1055.48	702.33	1103.16	1591.83	1157.43	2537.99	1567.72
STT	-35.72	-1.92	-20.98	0.04	-79.64	26.83	-104.81	36.64
LEH	1656.72	297.07	1228.99	244.82	894.64	216.78	853.17	216.79
WB	862.29	203.19	1017.41	180.61	822.50	164.03	879.31	146.93
MER	449.81	213.17	407.85	213.87	331.20	183.31	211.47	172.22
WM	2861.39	1752.13	2943.96	1263.16	2621.88	927.61	2641.13	628.21

Table 6: Average PoD and CDS spreads with regard to the most resilient bank prior to Lehman event

	PoD (1-10)	CDS (1-10)	PoD (30-136)	CDS (30-136)	PoD (137-180)	CDS (137-180)	PoD (181-225)	CDS (181-225)
GS	21.15	166.27	20.73	38.93	18.12	39.42	-1.71	1.09
WFC	127.06	112.41	58.97	34.71	23.02	16.27	0.55	-4.55
C	349.59	176.01	70.53	44.83	8.03	20.92	-0.53	-4.15
BAC	95.90	99.60	25.52	36.23	4.79	16.33	0.57	-1.92
JPM	131.14	102.21	39.86	26.81	29.07	21.44	-1.29	-6.88
AIG	169.07	175.69	57.20	42.32	15.75	21.11	-6.01	-8.33
MS	56.58	205.72	11.99	56.31	-3.99	39.98	-4.26	3.75
MBI	581.06	726.47	681.53	380.20	121.97	109.88	1.08	4.46
PNC	47.79	36.91	14.21	28.47	4.90	-0.77	0.05	-0.89
STT	45.01	69.06	55.17	9.15	30.25	0.32	2.19	-1.76
LEH	248.41	255.30	90.56	83.27	74.38	67.16	7.12	3.44
BSC	216.33	399.24	68.56	115.72	58.39	78.41	0.12	6.71
WB	198.19	235.25	46.96	61.86	17.19	19.82	0.37	-3.21
MER	136.01	215.68	63.91	77.97	31.58	42.95	-2.79	2.61
CFC	1349.10	477.21	1003.49	491.70	450.70	170.76	12.14	7.60
WM	969.10	508.09	347.06	221.88	71.31	53.94	-0.58	-1.57

Table 7: Average PoD and CDS spreads with regard to the own history prior to Bear Stearns event