City Research Online

City St George's, University of London

Citation: Eichengreen, B., Mody, A., Nedeljkovic, M. & Sarno, L. (2012). How the Subprime Crisis went global: Evidence from bank credit default swap spreads. Journal of International Money and Finance, 31(5), pp. 1299-1318. doi: 10.1016/j.jimonfin.2012.02.002

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: https://openaccess.city.ac.uk/id/eprint/13161/

Link to published version: https://doi.org/10.1016/j.jimonfin.2012.02.002

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to City Research Online policy.

City Research Online: http://openaccess.city.ac.uk/ publications@citystgeorges.ac.uk

How the Subprime Crisis Went Global: Evidence from Bank Credit Default Swap Spreads

Barry Eichengreen, Ashoka Mody, Milan Nedeljkovic and Lucio Sarno*

First version: August 2009 - This version: February 2012

Abstract

How did the Subprime Crisis, a problem in a small corner of U.S. financial markets, affect the entire global banking system? To shed light on this question we use principal components analysis to identify common factors in the movement of banks' credit default swap spreads. We find that fortunes of international banks rise and fall together even in normal times along with short-term global economic prospects. But the importance of common factors rose steadily to exceptional levels from the outbreak of the Subprime Crisis to past the rescue of Bear Stearns, reflecting a diffuse sense that funding and credit risk was increasing. Following the failure of Lehman Brothers, the interdependencies briefly increased to a new high, before they fell back to the pre-Lehman elevated levels – but now they more clearly reflected heightened funding and counterparty risk. After Lehman's failure, the prospect of global recession became imminent, auguring the further deterioration of banks' loan portfolios. At this point the entire global financial system had become infected.

JEL classification: G10; F30.

Keywords: subprime crisis; credit default swap; common factors.

*Eichengreen, Barry. eichengr@econ.berkeley.edu, Department of Economics, University of California, Berkeley, 530 Evans Hall, Berkeley, CA 94720-3880, USA and National Bureau for Economic Research (NBER), Cambridge; Mody, Ashoka. amody@imf.org, International Monetary Fund, 700 19th Street, N.W., Washington, D.C., 2043, USA; Nedeljkovic, Milan. National Bank of Serbia, 12 Kralja Petra, Belgrade, 11000, Serbia, milan.nedeljkovic@nbs.rs; and Sarno, Lucio (corresponding author). Lucio.Sarno@city.ac.uk, Cass Business School, City University, 106 Bunhill Row, London, EC1Y 8TZ, UK, +44 (0) 20 7040 8772, and Centre for Economic Policy Research (CEPR), London.

1 Introduction

One enduring question about the financial turbulence that engulfed the world starting in the summer of 2007 is how problems in a small corner of U.S. financial markets—securities backed by subprime mortgages accounting for only some 3 per cent of U.S. financial assets—could infect the entire U.S. and global banking systems. Moreover, while the banking system became affected in a generalized fashion by the crisis, the fortunes of banks differed substantially in terms of the market assessment (e.g. differentials in the impact on their share prices) and on the scale of government intervention received. In particular, whether the decision to let Lehman Brothers fail was a critical mistake that unleashed a global economic and financial tsunami will be debated for years. Some say that the authorities should have known that investors perceived banks' fortunes as intertwined, so that letting one fail was bound to undermine confidence in the others. Others say that Lehman Brothers was unique and everyone knew it.[1] The crisis that affected the global financial system, in this view, did not reflect the decision to let this one institution fail. Rather it reflected deteriorating global economic and financial conditions that undermined the position of banks as a class.

This paper seeks to shed further light on these issues. We analyze the risk premium on debt owed by individual banks as measured by banks' credit default swap (CDS) spreads, focusing on the CDS spreads of the 45 largest financial institutions in the U.S., the U.K., Germany, Switzerland, France, Italy, Netherlands, Spain and Portugal.[2]

We use principal components analysis (PCA) to extract the common factors underlying weekly variations in the CDS spreads of individual banks. If the spreads for different banks move independently, then we can infer that the risk of bank failure is driven by bank-specific factors. If they move together, then we infer that banks are perceived as subject to common risks. This provides us with the first bit of evidence on how the crisis spread. In addition to estimating the importance of common factors, we attempt to ascertain what they reflect. We examine the association between the

[1] Among other things, whereas other institutions could be saved because they had adequate collateral against which the U.S. Treasury and Federal Reserve could lend, Lehman did not.

[2] These swaps are insurance contracts. The buyer of the CDS makes payments to the seller in order to receive a payment if a credit instrument (e.g. a bond or a loan) goes into default or in the event of a specified credit event, such as bankruptcy. The spreads are, in effect, a measure of the credit risk or the insurance premium charged. This measure has several advantages over the traditional measures which are based on banks' balance sheet information. First, the CDS spreads are forward looking since they encompass available information with respect to expected default risk. Balance sheet data only reflects ex-post information on the institutions' performance. Second, CDS spreads are timely updated without the need to rely on (subjective) interpolation techniques, whereas balance sheet data are only available at quarterly frequency. The CDS spreads also offer advantages over other market measures of risk based on, e.g. bond spreads and stock returns. They are the most actively traded derivatives and lead bond (Blanco, Brennan and Marsh, 2005) and stock (Acharya and Johnson, 2007) markets in price discovery. Also, bond spreads may reflect factors other than the ones related to default risk (due to, for example, different tax treatments) and are sensitive to the choice of the benchmark risk-free rate (Jorion and Zhang, 2007). However, there has been a recent concern that speculative pressure within the CDS market sometimes causes the swaps to become delinked from their function of hedging against default (Soros, 2009). See also Longstaff et al. (2010), who analyze spreads on sovereign CDS, and Zhang, Zhou, and Zhu (2009), who examine the determinants of spreads on corporate CDS spreads.

common factors on the one hand and real-economy influences outside the financial system, transactional relationships among banks, and transactional influences between banks and other parts of the financial system on the other hand.[3]

We reach the following conclusions. The share of common factors was already quite high, at 62 percent, prior to the outbreak of the Subprime Crisis in July 2007. Banks' fortunes rose and fell together to a considerable extent, in other words, even before the crisis. These common factors were associated with U.S. high-yield spreads—the premium paid relative to Treasury bonds by U.S. corporations that had less than investment grade credit ratings—which we take as an indicator of the perceived probability of default by less creditworthy U.S. corporations, and in turn reflects economic growth prospects.[4] For obvious reasons, those defaults and the growth performance that drives them have major implications for the condition of the banking system even in normal times.

The share of the variance accounted for by common factors then rose to 77 percent in the period between the July 2007 eruption of the Subprime Crisis and Lehman's failure in September 2008. This is indicative of a perception that banks as a class faced higher common risks than before. At the same time, the measured association between the common factors and U.S. high-yield spreads declined, while the association with measures of banks' own credit risk and of generalized risk aversion increased (Brunnermeier, 2009; Dwyer and Tkac, 2009). An interpretation is that the Subprime Crisis made investors more wary of the risks in bank portfolios for reasons largely independent of the evolution of the real economy but that lack of detailed information on those risks led them to treat all banks as riskier rather than discriminating among them.

Following Lehman's failure, there was a further brief increase in the share of the variance accounted for by the common components. Then, although the level of CDS spreads remained high, the share of their variance accounted for by the common component fell back relatively quickly to levels below those that prevailed just before the Lehman episode. In other words, the common movements declined from their peaks but remained at the post-Bear Stearns elevated levels. Thus, the perception persisted that the banks' fortunes were linked. The association between the common factors and high-yield corporate spreads also reemerged, evidently reflecting the perception that a global recession was now in train. More importantly, the common component of CDS spreads became more highly related with measures of funding and credit risk as measured by spreads in the asset-backed commercial paper market and LIBOR minus the overnight index swap. An interpretation is that whereas in the July 2007-September 2008 period investors became more aware of systemic risk

[3] To be clear, we do not attempt to identify causality. However, the association measures offer a rich set of stylized characterizations. These characterizations are likely to be the basis for defining and probing more subtle hypotheses.

[4] These high-yield spreads have been found to be good predictors of U.S. GDP growth at horizons of about a year, reflecting a financial-accelerator interaction between credit markets and the real economy (Mody and Taylor, 2003; Mody, Sarno and Taylor, 2007). Because European high-yield spreads are closely correlated with US spreads and, as such, offer no additional information, U.S. high-yield spreads are also a measure of global prospects.

in an unfocused sense, Lehman's failure caused that common risk to be more concretely identified with both developments in the real economy and specific problems in the financial system.

In sum, then, our answer to the question posed in the title is as follows. Banks fortunes rise and fall together even in normal times. But the importance of common factors rose to exceptional levels between the outbreak of the Subprime Crisis and the rescue of Bear Stearns, reflecting increased diffuse sense that credit risk was increasing. The period following the failure of Lehman Brothers then saw a further increase in those interdependencies, reflecting heightened funding and counterparty risk. In addition there were direct spillovers, as opposed to common movements, from the CDS spreads of U.S. banks to those of European banks. After Lehman's failure the prospect of global recession became imminent, auguring the further deterioration of banks' loan portfolios. At this point the entire global financial system had become infected.

It is helpful to be clear about what this paper does not do. It does not pinpoint any one bank or set of banks as systemically important. Rather, the extent of comovement in spreads points to tendencies of the degree to which the system is perceived to be tied to common factors. An individual bank within the set examined may be more or less tied to the common factors to the extent that it has a larger or smaller extent of idiosyncratic risk. Ultimately, then, the methodology outlined here is a guide for policy only to the extent that it highlights overall trends. The task of determining the systemic importance of an individual bank requires examining the data in the books of the banks—or worse, data that should be on the books but is not.

The rest of the paper is organized as follows. Section 2 specifies a dynamic factor model in which common latent factors explain the movement of the CDS spreads of the 45 banks in our sample. The model is estimated using PCA in recursive fashion, allowing the contributions of the components to change over time. In Section 3, we consider the possibility of additional spillovers from inter-bank exposures that go beyond the common movements identified by the latent factors. Then in Section 4, we describe the changing relations between these latent factors and a number of high frequency financial series. We also provide a sensitivity analysis to check the robustness of our results. A final section concludes.

2 Common Factors in CDS Spreads

We start by decomposing the change in CDS spreads of N=45 global banks into common and idiosyncratic components. The term "banks" is used throughout in this paper, although some insurance companies are also included in the sample. The sample runs from July 29, 2002 to November 28, 2008. Thereafter, the intense involvement of the U.S. authorities in managing the short-term vulnerability of the financial sector increasingly reflects the official interventions which, because they

operated differentially across the institutions, limits the validity of a market-driven common set of influences on their CDS. The data are 5-year CDS spreads, as the five-year maturity is the most widely traded. We use end-of-day quotes from the New York market for payment in U.S. dollars based on U.S. dollar-denominated notional amounts. As is customary, the spreads are denominated in basis points (100 basis points equal 1 percentage point). These are averaged over the week to obtain a weekly series, smoothing out sharp daily movements and irregular trading, yielding 331 observations per bank. The data are taken from Bloomberg.

2.1 Preliminary Data Analysis

Table 1 reports summary statistics on spreads for 45 banks. Average spreads over the period vary significantly across banks (from a low of 17 for Rabobank to a high of 101 basis points for AIG).[5] Our interest is not so much in the cross-sectional variation at this stage, however, as in the variation over time, which has been substantial. Unlike sovereign CDS spreads where the standard deviations are typically smaller than the means, the standard deviations of the CDS spreads of financial institutions studied here are close to the means and sometimes larger. The minimum/maximum values further highlight the considerable time-series variation. For example, the spread for Merrill Lynch ranges from 15 to 473 basis points; in Europe, the range for Commerzbank varies from 8 to 260 basis points.

[Insert Table 1 about here]

Figure 1 tracks the time variation in median spreads for all banks, U.S. banks, and European banks. In 2002, after the tech bubble burst and confidence was challenged by the events of September 11, 2001, CDS spreads were elevated. Some banks were able to purchase protection at relatively low spreads of 20-50 basis points, but others paid more than 100 basis points. Subsequently spreads declined everywhere. The low point was the week of January 17, 2007 when the median spread in the full sample of 45 banks was 7.5 basis points. Thereafter, spreads increased gradually, reaching a median of 12 basis points in the second week of July 2007. But even in that week, the highest spread was 55 basis points for Bear Stearns.[6] In contrast, the subsequent rise in spreads was dramatic with twin peaks corresponding to the Bear Stearns rescue and the Lehman Brothers failure. For U.S. banks, a high of 417 basis points was reached following the severe stress after the Lehman failure during the week of October 1, 2008; the median spread then moderated to 268 basis points in the last week of November 2008. The corresponding numbers for the European banks were 130 and 97 basis points, respectively.

[Insert Figure 1 about here]

[5] This variation is, however, smaller than the 100-fold variation in premia across the Japanese and Brazilian sovereigns in the sample analyzed by Longstaff et al. (2010).

[6] Some, evidently, knew about the extent of its leverage.

2.2 A Dynamic Factor Model of CDS Spreads

The first question we ask is whether the movements in spreads reflected common drivers. To answer that question we estimate the "latent" or unobserved factors generating common movements. The relationship between the unobserved factors (F_t) and the observed spreads ($X_{i,t}$) can be approximately represented by a dynamic factor model (Chamberlain and Rothchild, 1983):

$$X_{i,t} = \Lambda_{i,h} F_{h,t} + \phi_i(L) X_{i,t-1} + \varepsilon_{i,t} \tag{1}$$

where $X_{i,t}$ is a vector of the weekly changes in CDS spreads of each of the 45 banks, "i" refers to a bank and "t" is the time subscript. $\Lambda_{i,h}$ is a vector of factor loadings, and $F_{h,t}$ are latent factors ($h = 1, \ldots, k$). The estimation procedure allows for $\varepsilon_{i,t}$ to be cross-sectionally and time correlated and heteroskedastic.[7] Stacking the terms, specification (1) can be equally represented as:

$$X = F\Lambda' + \Phi(L)X + \varepsilon \tag{2}$$

where X and ε are $T \times N$ matrices, F is a $T \times k$ matrix, and Λ is a $N \times k$ matrix.

The idea here is that the covariance among the series can be captured by a few unobserved common factors. The latent factors and their loadings can be consistently estimated by PCA. As Bai and Ng (2002, 2008) and Stock and Watson (2002) show, the principal component (PC) estimator enables us to identify factors up to a change of sign and consistently estimate the factors space up to an orthonormal transformation.[8] The estimation procedure also provides a measure of the fraction of the total variation explained by each factor, which is computed as the ratio between the k largest eigenvalues of the matrix $X'X$ divided by the sum of all eigenvalues.

The data for each bank are first filtered for autocorrelation since in the presence of serial correlation the covariance matrix of the data will not have the factor structure (exact or approximate) and as such can lead to biased inference. To assess variation over time, the model is estimated recursively after the first 150 data points: as such, the filtering is also performed recursively. At each recursion an AR(p) model is applied to each series, where the order, p is determined using the individual partial autocorrelation function (PACF) and residuals from the AR(p) model are used as the filtered series. It is worth noting that the use of weekly averages of the daily CDS spreads may also introduce a moving average component in the errors of the constructed series. However, since

[7] Note, however, that the contribution of the idiosyncratic covariances to the total variance needs to be bounded (Bai and Ng, 2008), which puts a limit on the amount of time and cross-correlation and heteroskedasticity such that the number of cross-correlated errors can only grow at a rate slower than $\sqrt{N}$ and analogously for the dependence over time.

[8] Note that the consistency is related to the space spanned by the factors and not with respect to the individual factor estimates.

the PC estimation procedure allows for some serial correlation in the errors, this should not affect the results substantially. Finally, all series are standardized at each recursion since PC estimates are not scale invariant.

Some caveats and further motivation are in order with respect to our choice of modeling framework and econometric methods. First, we do not attempt to model the time-variation in conditional pairwise correlations across the CDS analyzed, which one could achieve by estimating a dynamic conditional correlation (DCC) model of the type originally developed by Engle (2002). This is because we are primarily interested in the common drivers across the whole set of banks analyzed, which would be more cumbersome to extract from the richer structure of a DCC model, whereas this task is straightforward in the context of specification (2).

Second, the dynamic factor model considered here does not allow explicitly for time-varying volatility, which could generate biases in the estimation of the principal components and subsequent estimation of the correlation between the principal components and observable economic variables. This is important since changes in correlation between two series could be due to increases in volatility in the common factor, as opposed to changes in the covariance between them. Thus, in addition to the benchmark model for CDS spread changes which ignores the potential biases induced by time-varying volatility, for robustness we also estimate the dynamic factor model with time-varying factor volatility, and the details are given in Section 4.6. In general, a richer dynamic factor model of CDS spreads would allow explicitly for time-varying, stochastic volatility and correlations, and could be estimated by Markov Chain Monte Carlo (MCMC) methods.[9] However, given that this paper provides the first PCA analysis of CDS spreads during the crisis period, the use of a benchmark model that is parsimonious has the advantage of establishing the bare facts in a more accessible fashion. Moreover, the model estimated is comparable to traditional PCA analysis of the kind used, for example, by Longstaff et al. (2010) for sovereign CDS spreads and Collin-Dufrense, Goldstein and Martin (2001) for corporate bonds.[10]

2.3 Estimation Results and Discussion

Figure 2 shows changes over time in the contributions of the common factors to the total variation in the CDS data, obtained from the estimated factors. Recall that up until mid-July 2007, i.e., before the

[9] For examples of MCMC estimation of multivariate stochastic volatility, see Della Corte, Sarno and Tsiakas (2009) and Sarno, Schneider and Wagner (2012).

[10] A related literature focuses on estimating credit spreads, which is a key component in marking-to-market a financial institution's fixed-income investment portfolio. Credit spreads can be estimated using either bond prices (e.g. Campbell and Taksler, 2003; Collin-Dufresne, Goldstein and Martin, 2001; Elton et al., 2001) or, more commonly, CDS spreads, as in Longstaff, Mithal and Neis (2005) and Ericsson, Jacobs and Oviedo (2009). The CDS-based estimates are preferred because of the greater liquidity of the CDS markets, which is due to the fact that CDSs do not require any upfront investment and enable one to more easily short a firm's credit risk. There are two common approaches to modeling: analytical, structural models and statistical, reduced-form models. See Collin-Dufrense and Goldstein (2001), Chen, Collin-Dufrense and Goldstein (2009), and Jarrow and Protter (2005) for a review.

start of the Subprime Crisis, absolute movements in bank CDS spreads were small. The PCA shows that even in this relatively tranquil period the perceived riskiness of different international banks moved together to a considerable extent. The first component contributed just above 40 percent of these movements and the second a further 10 percent. Together the first four common factors explained about 60 percent of movements in CDS spreads in this period.[11] The statistical procedure does not tell us whether these common influences reflected interconnections within the banking system or were the result of common external factors. To explore that distinction, in Section 4, we report evidence on the relations between the common "unobserved" factors and observed variables representing real economic prospects and metrics of banks' funding stress and their systemic credit risk.

[Insert Figure 2 about here]

A new phase commenced in July 2007, when HSBC announced large subprime-mortgage related losses and CDS spreads rose sharply. While spreads retreated somewhat in the months thereafter, they remained at significantly higher levels. Nevertheless, July 2007 saw the start of a tendency towards an increase in the comovement of spreads. The rise in the share of variance explained by the first factor and the first four factors combined trended up even as the average spreads rose and fell. Thus, while the sense of crisis went into remission periodically in late 2007 and early 2008, the perception remained that banks faced common risks. By early March 2008, prior to the rescue of Bear Stearns, the share of the first component had risen by about 15 percentage points to 56 percent.

The importance of the common factors continued to increase following the Bear Stearns rescue, reaching a new high in May 2008, at which point the first common factor explained almost 60 percent of the variance of CDS spreads. Then, the period between May and September 2008 was one of general weakness of financial-market indicators. The share of the variation explained by common components of CDS spreads exhibited some tendency to decline in this period, although it remained not very far below the high level reached in May.

The failure of Lehman Brothers was accompanied by a further increase in the comovement of spreads. The share of the variance accounted for by the first principal component jumped to 65 percent, while the share accounted for by the first four reached 80 percent. However, there was moderation of the comovement by the first week of October, as the share of the variation explained by recursively computed common factors fell back to pre-Lehman levels, implying that the importance

[11] Factors other than the first four individually explained less than 4 percent of the variation. Thus, our analysis below focuses on the PC estimates obtained using the four factor "space." This is supported by the results obtained by running the Onatski's (2010) criterion for determination of the number of factors in the data through a grid of parameter values. The Bai and Ng (2002) information criterion does lead to the possibility of more than four factors, but this criterion tends to overestimate the number of factors in samples with relatively small cross-sections and high cross-correlations (Anderson and Vahid, 2007, Onatski, 2010). As a practical matter, note that our results remain the same with either 3 or 5 factors.

of factors decreased to below pre-Lehman levels over this period.

To get a sense of whether the degree of commonality we observe for international banks is high or low, we can compare these results with those of Longstaff et al. (2010) for sovereign CDS spreads. Longstaff et al. find that the fraction of the variance of spreads on the CDS of 26 sovereigns explained by the first component varies between 32 and 48 percent. Note that they use monthly (rather than weekly) changes in spreads and the higher share is for months in which observations were available for all sovereigns. These features of their analysis smooth out some part of the idiosyncratic movements and, to that extent, might overstate the commonality. As such, the initial—pre-subprime—contribution of the first component of banks' spreads at about a 40 percent share is at least in a comparable range, and possibly implies greater commonality. That share of the first component for the banks, as we have reported, then rose to approximately 60 percent just before the Lehman failure and further to 65 percent just after. When considering the first four components, Longstaff et al. (2010) find that they explain 60 to 70 percent of the share of sovereign spreads variation, which is comparable to our range of 60 to 80 percent. Another implication is that for sovereign spreads, the second component and beyond have a more substantial contribution than is the case for banks, implying greater variety of global common influences on sovereign spreads.

3 Additional Spillovers

In this section, we investigate whether, once the common factors are considered, the CDS spread of a particular bank is further influenced by current and lagged changes in CDS spreads of other banks. In other words, we ask whether there is significant information in CDS spreads of other banks over and above that contained in the common factors.

These spillover tests are performed by regressing the change in the (non-filtered) CDS spread of bank "i" on its own lags, the common factors, and the (lagged and current) changes in CDS spreads of another bank (bank "j"), as in equation (3). Since the CDS data exhibit a potential break in the last week of July 2007 we also include a dummy variable for the sub-prime period, where $D_t = X_{j,t-1} I\left(t \geq \tau\right)$:[12]

$$X_{i,t} = \Lambda_i F_{h,i} + \phi_i(L) X_{i,t-1} + \lambda_j(L) X_{j,t-1} + \delta D_t + u_{i,t}. \tag{3}$$

Equation (3) is estimated for each pair of banks, which yields a total of 2025 regressions. In each case, we test for the statistical significance of the coefficients λ_j by computing the heteroskedascity-

[12] Bai and Ng (2006a) showed that these so-called "factor-augmented" regressions yield consistent estimates of the parameters as $T, N \to \infty$ provided that $\sqrt{T/N} \to 0$. Of course, in a finite sample the estimation error will not disappear completely.

robust Lagrange multiplier (LM) test of specification (3). The restrictions are tested using a standard χ^2 test.[13]

In only about 2.5 percent of the 2025 regressions is the CDS spread of another bank significant at the 5 percent significance level. This supports the validity of the common factors estimated in the preceding section. The relatively few additional spillovers we identify are illustrated in Figure 3(a). Following the start of the sub-prime crisis, the incidence of such spillovers declined, implying that the commonality is well captured by the latent factors. However, the additional spillovers increased notably starting in mid-July 2008 and reached their maximum during the Lehman Brothers crisis. An interpretation is that not just global economic drivers (which are presumably being picked up by the common factors) but also counterparty risk and other similarities in a few banks' portfolios (which are being captured by the additional spillovers) figured importantly in individual CDS spreads around the time of the Lehman Brothers failure.[14] The banks for which additional spillovers matter tend to be well-known names: they include ING, Royal Bank of Scotland and UBS in Europe, and Bank of America, J.P. Morgan and Morgan Stanley in the U.S.

[Insert Figure 3 about here]

Our procedure also allows us to glean some evidence of international spillovers. Here we consider the percentage of banks in the U.S. significantly influenced by a bank in the European Union (Figure 3b) and vice versa (Figure 3c). Note that we are not looking here at pairwise influences; rather the question is what percentage of banks has *at least* one cross-border relationship with evidence of additional spillovers.[15] The period before the sub-prime crisis is relatively stable with moderate evidence of spillovers from the European Union to the U.S. (Figure 3b) and little evidence of spillovers from U.S. banks to European banks (Figure 3c). Following the start of the crisis in July 2007, however, the incidence of additional spillovers from the U.S. system to Europe increased, particularly in periods of high distress (in August 2007 and from the beginning of July 2008 onwards). Interestingly, the magnitude of additional spillovers from European to U.S. banks declined in this period. This is consistent with the view that developments in U.S. banks were the stronger source of perceived financial risk starting in the early stages of the Subprime Crisis.

[13] The heteroskedasticity-robust estimate of the covariance matrix is obtained using the Davidson and MacKinnon (1985)'s transformation of the squared residuals. Simulation results in Clark and McCracken (2005) and Rossi (2005) establish that this test does not lose power in establishing the statistical significance of variables in OLS regressions subject to a structural break provided that the relationship existed for any subsample (i.e. before or after the break). However, it is well known that in the presence of breaks and heteroskedasticity, all classical tests may be oversized (Hansen, 2000). Therefore, we also run the robust bootstrap LM test based on the fixed design wild bootstrap (Hansen, 2000) and the recursive wild bootstrap (Goncalves and Kilian, 2004), but the results are not significantly different. In addition, we also perform a battery of Monte Carlo experiments using the LM and Wald statistics (with and without bootstrapping), and while both tests had similar power properties, the LM test is found to display better size properties. Full details are available upon request.

[14] The importance of counterparty problems due to the failure of Lehman Brothers is emphasized by, inter alia, Brunnermeier (2009), Dwyer and Tkac (2009) and Jones (2009).

[15] Clearly, this leads to a larger fraction of banks than where the assessment is on a pairwise basis.

4 Correlating Latent Factors with Observed Financial Variables

The next step is to examine the relation between the latent factors identified in Section 2 and the observed financial variables. While the exact association of a financial variable with any one of the estimated factors is hard to define due to non-uniqueness of the factor estimates, we can measure the association of financial variables with the entire set of estimated factors and investigate under which conditions correlations with individual factors are still informative.

4.1 Some Statistical Considerations

Let G_t be an M-dimensional vector of observed financial variables. Bai and Ng (2006b) develop statistical criteria which can be used to investigate whether any of the candidate series yields the same information that is contained in the factors. The general idea behind their tests is to examine whether any of the candidate series can be represented as a linear combination of the latent factors allowing for a (limited) degree of noise in the relationship, so that $G_{m,t} = \beta'_m F_t + \xi_{m,t}$, where $G_{m,t}$ is an observable financial time series. Defining the OLS estimates from this regression as $\widehat{\beta}_m$, the residual as $\widehat{\xi}_m$ and the predicted value $\widehat{G}_{m,t} = \widehat{\beta}_m \widehat{F}_t$, a convenient measure that allows us to compare $G_{m,t}$ with $\widehat{G}_{m,t}$ is defined as:

$$R^2(m) = \frac{\widehat{V}\left(\widehat{G}_{m,t}\right)}{\widehat{V}\left(G_{m,t}\right)} \tag{4}$$

where $\widehat{V}(\cdot)$ denotes the sample variance and $\widehat{V}\left(\widehat{G}\right)$ is computed using the sample analogue of the factors' asymptotic covariance matrix, which is calculated to be robust to non-zero cross-correlations and time-series heteroskedasticity (see Bai and Ng, 2006a,b for further details). By definition $R^2(m)$ is bounded between 0 and 1: it equals 1 if there is exact association between the observed variable G_m and the factor space, and is close to 0 in the absence of any relation.[16]

Instead of examining the relationship with the entire factor space, the observed series can also be associated with a particular factor subspace, including one particular factor. Ahn and Perez (2010) proposed moment selection criteria that consistently determine the number of factors that can be related to the set of observable series of interest.[17] The criteria resemble the well-known likelihood-based selection criteria BIC and HCC, using the GMM J-statistic for testing the over-identifying restrictions. Once this number of factors is determined, the individual correlations between factors and the observed series can be examined. The Ahn-Perez analysis makes the relatively strong assumption of independence between the idiosyncratic part of the movement in CDS spreads and the

[16] An alternative measure defined in Bai and Ng (2006b) with an analogous intuition is $NS(m) = \frac{\widehat{V}\left(\widehat{\xi}_{m,t}\right)}{\widehat{V}\left(\widehat{G}_{m,t}\right)}$. The obtained results with this measure are equivalent to those obtained with $R^2(m)$.

[17] This procedure is essentially an extension of Andrews and Lu's (2001) general approach to model and moment selection in the generalized method of moment (GMM) estimation.

observed series. If this assumption is invalid, then the results will be biased since the rejections of the null hypothesis of no correlation may occur for any number of factors related to the observables. To obtain robust results we therefore adopt a pre-step estimation procedure to validate the series of observables we use.[18]

Among other statistical considerations that are worth noting, with regard to the Bai and Ng's (2006b) R^2 criterion, it is not straightforward to determine the threshold that would signal a "matching" between the factor space and individual series of interest when the relationship is contaminated by some degree of noise. Also, with regard to correlations with individual factors, the concern is whether these correlations uniquely identify the relationships with a particular factor. For instance, finding a correlation of 0.5 between the first factor and a series of interest may genuinely reflect a correlation but may also be spuriously picking up a correlation between the second factor and the series. This is a direct consequence of the non-consistency of the individual PC estimates - the first principal component can be a good approximation for the first factor but it can also be a linear combination of (all) factors. Hence, looking at the individual correlations between the series and principal components can spuriously pick up the correlation between the series and other factors. In general, the literature has not fully dealt with these issues. It is commonplace to report the criteria and the correlations without acknowledging their non-uniqueness. To evaluate the seriousness of these limitations, we used a Monte Carlo experiment to investigate how these procedures behave and to reassure ourselves that the results are meaningful and robust.[19]

[18] Given the structure of the procedure and if the true number of factors is known and equal to k, then for m observed series it follows that if we reject the null hypothesis that $m(N-k)$ moment conditions are zero, this may happen only if the instruments (observed series) are correlated with a subset of the idiosyncratic errors. If the observed series were correlated with all idiosyncratic errors, then this would imply the existence of another factor in the data. The useful pre-step therefore is to compute a J-test using a standard optimal HAC weighting matrix for $m(N-k)$ moment conditions and test whether the null is rejected. Following the non-rejection of the null, Ahn and Perez (2010)'s model selection procedure can be applied; otherwise the set of observed series needs to be reconsidered. Note that the pre-step procedure can be seen as leading to a conservative selection of correlates since we exclude all observable series that are correlated with both the factors and idiosyncratic variations: the selected observables are then very robust.

[19] We perform the following experiment, designed to resemble the characteristics of our CDS data (not reported but available upon request). We allow for the data to be cross-correlated and heteroskedastic and subject to a break in volatility. In the first experiment the factors explain a roughly equal percentage of total variation, whereas the second experiment captures the situations when the first factor explains the largest part of the overall variance and when its importance increases after the break point. The "observable" series are generated through a linear relationship with factors with a varying degree of noise. The R^2 criterion and simple correlations between observable series and the estimates of the first three factors are obtained using 5000 simulations. The Ahn and Perez's (2010) GMM-based BIC criterion was computed using 1000 simulations and 100 randomizations to save on computing time. The main findings from the experiment are the following. First, the GMM-based BIC criterion performs fairly well selecting in all cases the correct number of factors. The proposed pre-step estimation captures whenever the series are correlated with the idiosyncratic errors. Second, the R^2 estimates are significantly lower than those proposed by Bai and Ng (2006b) when there is some noise in the relation and breaks in volatility. In particular, the R^2 estimates we highlight below are meaningful measures of the relationships of interest under moderate levels of noise. Third, the signal-to-noise ratio from using correlations as a proxy improves with the difference in the importance of factors.

4.2 Correlates

We limit our attention to U.S. variables, since the corresponding European variables are highly correlated with U.S. series. A first set of variables representing the real economy includes the corporate default risk measured by the high-yield spread (HYS), risk aversion (VIX), and returns on the S&P500 stock index.[20] A second set of variables representing the banks' financial risks includes the credit spread (LIBOR minus overnight index swap), the liquidity spread (overnight index swap minus the Treasury bill yield), and spreads on asset backed commercial paper (ABCP).[21]

The GMM-based model and moment selection criterion of the validity of the observed series is performed first and the results are presented in Table 2. The test is performed for the full sample and two subsamples (up to July 2007 and up to May 2008) in order to examine whether the most recent period (with possible outliers) influences the results. In the first column of Table 2 we can see that none of the proposed series is correlated with the idiosyncratic part of CDS spreads since the frequency of rejections of the null among all randomizations of the data is very small for all samples. This implies that we can use the moment selection criteria to investigate the relationship between the observed series and the factors. In turn, the results from full and subsample estimation of the criteria suggest that the information in the set of observed series can be associated with the three or four factor subspace.[22]

[Insert Table 2 about here]

4.3 The Real Economy Prior to the Subprime Crisis

In the "real economy" group, we consider three correlates. High-yield spreads (HYS, spreads on bonds issued by less-than-investment-grade issuers) reflect increased corporate default probabilities and are known to do well in predicting short-term GDP growth (Mody and Taylor, 2003). The S&P 500 average reflects the market's perception of the economic outlook, while the VIX is a measure of economic volatility embedded in stock price movements.

Figure 4 shows the evolution of (median) CDS spreads in relation to these observed variables. Prior to the start of the Subprime Crisis, the HYS and the VIX trended down along with the median CDS spreads. As the figure suggests and we show below, the real economy as represented by the stock market bore less short-term relationship with the movement in CDS spreads. The HYS, VIX,

[20] VIX is the implied volatility on the S&P 100 option and is a widely used measure of global risk aversion.

[21] As with the CDS data, all series are recursively filtered and are standardized prior to the estimation of the correlations.

[22] For the full sample, the various criteria (namely, BIC1, BIC2, BIC3 and HQQ) suggest the existence of a relationship between the series and four factors. For the subsamples, BIC1 suggests a relationship with only one factor, whereas three or four factors are suggested by other criteria. Given that apart from the first factor, all other factors may be perceived as weak, BIC1 may underselect the true number of relationships; however, BIC3 and HQQ tend to overselect the true number of relationships (Ahn and Perez, 2010). As such we base our inference primarily on BIC2. See Ahn and Perez (2010) for formal definitions of the BIC1, BIC2, BIC3 and HQQ criteria.

and CDS spreads were all at low levels prior to July 2007. While they rose subsequently, their short-term movements became less correlated between the start of the Subprime Crisis and the failure of Lehman Brothers.

[Insert Figure 4 about here]

In Figure 5, panel A, we show the association of the first four factor space with HYS; this was relatively high prior to the Subprime Crisis. The R^2 criterion gives a value of 0.5 on the eve of the Subprime Crisis. In contrast, the association with the S&P 500 returns is low—less than 0.1. Notice that the R^2 for the VIX lies in between, but at 0.2 is at the lower end of the range.

[Insert Figure 5 about here]

The implication is that perceptions of banks' risk in that tranquil period were shaped by a global factor that is best summarized by corporate default risk. This is reasonable not just because of the banks' direct exposure to default risk but also because HYS has a proven track record as a predictor of economic prospects (e.g. Mody and Taylor, 2003). In particular, HYS movements capture the operation of the financial accelerator: high spreads imply high expected default rates (and hence lower collateral), lower credit supply, reduced growth prospects, and hence higher spreads. HYS has also been found to be a significant explanatory variable of emerging market spread differences across countries and their movements over time (see Eichengreen and Mody, 2000, and Longstaff et al., 2010, who find it to be the most potent of their candidate variables). In contrast, stock returns include both upside and downside movements: while high stock returns presumably lower risk to a degree, banks' risks are apparently more clearly defined by downside risks as reflected in HYS. The fact that the correlation with the VIX is significantly smaller than for HYS suggests that a higher generalized risk aversion does not necessarily translate into banks' risk premia.

These interpretations are supported by the correlations with specific factors reported in Panel B of Figure 5. Up through the start of the Subprime Crisis, the HYS was most highly correlated with the first principal component of CDS spreads.[23] In contrast, there was almost no correlation with the second factor. The same is true for the VIX (panel C). In contrast, S&P 500 returns had a higher correlation with the second factor (panel D). An interpretation is that the first factor reflects global perceptions of downside risks, while the second gives more weight to general movements in expected future profitability. Note, though, that the second factor explains a much smaller fraction of the overall variance of CDS spreads. Hence returns had a much weaker association with spreads' movements.

[23] Hereafter, all the correlations are expressed in absolute terms for ease of comparison.

4.4 The Emergence of Financial Factors

Thus, prior to the Subprime Crisis, global economic factors as summarized in HYS were the main drivers of the commonality in CDS spreads of international banks. Following the onset of the crisis and through the Bear Stearns bailout, however, the association with the HYS declined (Figure 5, panel A). The decline in the overall association between the HYS and the factor space reflected a decline in correlation with the increasingly important first factor of CDS spreads and occurred despite some increase in correlation with the second factor (panel B). Evidently, there was some dissociation with the real economy despite the fact that banks' prospects appear to have been perceived as increasingly correlated with one another in this period. Intuitively, the common risk did not obviously emanate from a sense of worsening of economic prospects. Moreover, an initial sharp increase in association with the VIX, reflecting generalized risk aversion, died down to pre-crisis levels by the time of Bear Stearns rescue (Figure 5, panels A and C). The small rise in the R^2 between S&P 500 returns and the CDS factor space probably reflects the fact that fears about the stability of the banking system were driven more by problems in the banks' positions in securities than by the ability of their corporate customers to stay current on their loans.

Once the Subprime Crisis started, the relevance of financial variables, particularly those related to banks' credit and funding risks, acquired greater prominence. A common metric of banks' credit risks is the TED spread, or the difference between the interest rates on interbank loans (we use the US dollar, 3-month LIBOR) and short-term U.S. government debt (3-month US Treasury bills). This captures the risk premium on bank borrowing, since LIBOR is the rate at which banks borrow and Treasury bills (T-bills) are commonly considered risk-free.

However, the TED spread reflects not just banking sector credit risk but also includes liquidity or flight-to-quality risk. These two categories of risks can be approximately decomposed. TED=(LIBOR-OIS)+(OIS-"T-bill"), where the OIS is an "overnight index swap" which measures the expected daily average Federal Funds rate over the next 3 months. Thus, the TED spread can be decomposed into the banking sector credit risk premium (LIBOR-OIS) and liquidity or flight-to-quality premium (OIS-T-bill).[24]

The TED spread rose sharply in the post-Lehman-crisis period as Figure 6 shows. While the liquidity premium (the OIS-T-bill differential) also increased, the more substantial increase was in credit risk (the LIBOR-OIS differential). Note also the spike after the start of the Subprime Crisis in the spread on ABCP. Since banks use these instruments for their short-term funding, the rise in this spread proxies the risks associated with rollover in short-term funding. That the trading in

[24] Of course, the OIS-T-Bill spread picks up some credit risk, but most analysts view the "general collateral" repo rate as the "risk free" rate (e.g. Longstaff, 2000; Della Corte, Sarno and Thornton, 2008), and that plots very closely to the OIS rate. The LIBOR-OIS spread is analyzed by Taylor (2009) and critiqued by Jones (2009).

market for ABCP issued by banks and conduits decreased substantially within days of the Lehman bankruptcy is well known; e.g. see Dwyer and Tkac (2009) for an overview of events in fixed-income markets before March 2009.[25]

[Insert Figure 6 about here]

The measured association of the common factors with these financial variables, which had been historically close to zero, rose following the start of the crisis (Figure 7). Thus, perceived bank risk, which had previously stemmed mainly from the development of the real economy, now stemmed more from banks' own internal credit and funding risks. However, while liquidity risk (as captured by the OIS-T-bill spread) and the ABCP spread showed some correlation with the CDS factor space initially, those correlations were not sustained. In contrast, the association with credit risk was sustained. Credit risk has the largest R^2 of these three financial variables with the factor space after the start of the crisis; in particular, the correlation with the first factor rose steadily up until Bear Stearns.

[Insert Figure 7 about here]

While the change in the pattern of relationships clearly points to greater emphasis on the internal workings of the banking system (rather than to the broader global economy), the estimates we find even at their elevated levels during this phase are small. To the extent that banks' credit risk and funding risk did become more important, our measures for these features may not be sufficiently encompassing. In addition other concerns, such as lack of transparency of the complex asset holdings, may have also acquired greater prominence in assessing bank risks.

Not much changed between the Bear Stearns rescue and the Lehman failure. The relation with HYS stabilized at below pre-crisis levels and the association with credit risk remained significantly above pre-crisis levels. The values of R^2 for other variables remained low, near their pre-crisis levels. Thus while corporate default risk remained a salient factor determining banks' risk, there was a shift as this traditionally-dominant factor lost ground and the risk that banks themselves may not be able to honor their obligations gained prominence.

4.5 After Lehman

The immediate post-Lehman phase is remarkable for the unprecedented alignment of risks. The association between the CDS factor space and all of the observed variables rose, according the R^2 criterion, with the association with the financial variables increasing most sharply. The increase in credit and funding risk premia reflected the stress faced by banks. In addition, these developments presumably contributed to a revision of prospects of the real economy that further undermined

[25] Note also from Figure 6 that the LIBOR-OIS spread moved rather closely with the spreads on ABCP, often used to proxy the banks' costs of funding since banks issue such paper to fund their investments.

confidence in the condition of and prospects for the banking system.

By the R^2 criterion, the association between the space of common factors in bank CDS spreads and HYS increased following the failure of Lehman, reversing its decline in the preceding four quarters. The association with the S&P 500 returns showed a particularly large increase as global economic prospects were seen as increasingly tied to the fortunes of banks. However, despite the high level of VIX during this period, the association between the CDS factor space and VIX increased relatively little. In contrast, there was an especially sharp increase in the association between the financial variables and the common factor in banks' CDS spreads.

Differences in the correlations between these real and financial variables and the first and second common factors provide further intuition. The financial stress indicators became more correlated with the first factor of CDS spreads; that correlation rose to levels that had not been seen before. In contrast, the variables that had moved the first factor in the past, particularly the HYS, declined sharply in importance in the immediate post-Lehman phase. Instead, the correlation between the real economy and the common CDS movements shifted to the smaller, second factor.

4.6 Sensitivity Analysis

The consistency of the PC factor space estimates which is established in a series of papers by Bai and Ng constitutes the basis for our empirical analysis. Consistency is obtained under the assumptions of a limited degree of cross-section and time correlation and heteroskedasticity in the data. Moreover, the PC estimation does not take into account time-varying volatility of individual factors, which may lead to biases in the estimation of the factors and subsequent estimation of the correlation between the principal components and observable economic variables. To assess the seriousness of these limitations we use three additional methods of estimation.

Makarov and Papanikolaou (2009) recently proposed an extension of specification (1) that explicitly allows for time-varying factor volatility such that:

$$F_t = \Sigma_t^{1/2} u_t \tag{5}$$

where $E\Sigma_t = I$ and $u_t \sim N(0, I)$. The model still implies an unconditional factor structure as in (1), but the individual factors are now allowed to display time-varying volatility. The method is especially useful in cases when the relative volatility of factors varies over time. Estimation of the model consists of two steps. In the first step standard PCA estimation is performed. In the second step, the first-step estimates are corrected using the estimated rotation matrix such that the computed factors are also conditionally uncorrelated. The rotation matrix is estimated as the solution that minimizes the squared off-diagonal elements of the realized factor correlation matrix.

The realized factor correlation matrix is computed using a rolling window of daily factor realizations over the previous eight weeks.[26] To control for possible changes in the volatility of the observed correlates each series is filtered for volatility in the same way using the rolling window of daily realizations over the previous eight weeks.

To assess the impact that the non-spherical idiosyncratic errors may have on the results, we also perform two alternative estimations: the weighted principal component (WPC) estimation proposed by Boivin and Ng (2006), and the robust PCA estimation proposed by Pison, Rousseeuw, Filzmoser and Croux (2003). The former controls for non-negligible cross-correlations and heteroskedasticity by exploiting the information from the sample error covariance matrix to improve the efficiency of the PC estimator through a two-step procedure. The robust PC controls for the presence of outliers through computing the minimum covariance determinant (MCD) estimator, which is an outlier resistant estimate of the data covariance matrix. The PC estimator is then obtained in a classical way using the MCD estimate.

To save space, we only give the description of the results from this further analysis, whereas full details are available on request. The results from all three methods generally support the main findings from the PCA discussed earlier. In particular, if we allow for time-varying volatility in the factors, the R^2 estimates are virtually identical to the ones we recorded with the standard PCA. This is not surprising since it can be expected that the PC estimates of the factor space remain consistent even in the presence of time-varying volatility in the factors, given that the factors unconditional covariance matrix is the same regardless of whether there is time variation in factors' volatility. However, individual factors' correlations exhibit more fluctuations when the procedure is employed recursively, although the general pattern of correlations is consistent with that observed in Figures 5 and 7 for the standard PCA.

Results from weighted PC estimation confirm the previously obtained interpretation of the factors, both quantitatively and qualitatively. The dynamics of correlations computed from robust PC estimation also remain unchanged, the only difference being in the level of correlations, which is slightly lower when computed using the robust factor estimates.

5 Conclusions

We have analyzed common factors in bank credit default swaps both before and during the credit crisis that broke out in July 2007 in order to better understand how this crisis spread from the subprime segment of the U.S. financial market to the entire U.S. and global financial system. We showed first that common factors in CDS spreads are present even in normal times; they reflect the

[26] We have also experimented with different choices of the window length, but the results were very similar.

impact of the macroeconomy—the ultimate common factor from this point of view—on banks as a group. But the importance of the common factor increased significantly between the eruption of the Subprime Crisis in July 2007 and the failure of Lehman Brothers in September 2008. This increase in the common factor seems to have been associated with a proxy for the banking-sector credit risk premium, especially in the period prior to the rescue of Bear Stearns. In contrast, the association with the state of the real economy, which had been evident prior to the crisis, appears to have been somewhat attenuated. In other words, in this abnormal period investors were not yet concerned so much with the prospect of a global recession that would impact the banks' loan books as with other credit risks affecting the banks – connected, presumably, with their investments in subprime related securities.

After the failure of Lehman Brothers the importance of the common factor remained elevated. But where movements in that factor had previously been related to diffuse measures of generalized banking-sector credit risk, they now became increasingly linked to measures of funding risk. In addition, the association of the common factor with the real economy reasserted itself, as evidence of the deepening recession mounted.

What does this evidence imply for policy decisions taken in this period? With benefit of hindsight (which is what a retrospective statistical analysis permits), we can see a substantial common factor in banks' CDS spreads that could have alerted the authorities to the risks of allowing a major financial institution to fail. The further increase in that common factor in the period between the outbreak of the Subprime Crisis and the critical decision concerning Lehman Brothers should have implied further caution in this regard. It was not the implications of any impending economic slowdown about which investors were primarily worried in this period; rather the concern was about the state of the banks' asset portfolios and, presumably, their investments in securities in particular. The heightened comovement at least in part reflected incomplete knowledge about the magnitude of toxic asset positions in this relatively early stage of the crisis and, hence, raised the possibility that instability could spread more quickly and widely than assumed in the consensus view. In the event, Lehman Brothers was allowed to fail, after which the sensitivity of the CDS spreads of global banks as a group experienced heightened sensitivity to the whole range of economic and financial variables. As those variables deteriorated, the result was a perfect storm.

Acknowledgements

This paper was partly written while Milan Nedeljkovic was in the Winter Internship program at the International Monetary Fund. Sarno acknowledges financial support from the Economic and

Social Research Council (No. RES-062-23-2340). The authors are indebted to Jim Lothian (Editor), an anonymous Referee, Pierre Collin-Dufrense, Charlie Kramer, Ian Marsh, Marcos Perez, Giovanni Urga and to participants at the Yale/RFS Financial Crisis Conference for comments and suggestions, to Igor Makarov for sharing his code for estimating dynamic factor models with time-varying factor volatility, and to Susan Becker and Anastasia Guscina for valuable research assistance. The views expressed here are those of the authors and should not be attributed to the National Bank of Serbia or the International Monetary Fund, its management or Executive Directors. The authors alone are responsible for any errors and for the views expressed in the paper.

References

Acharya, V., Johnson, T., 2007. Insider Trading in Credit Derivatives. Journal of Financial Economics 84, 110-141.

Ahn, S., Perez, M.F., 2010. GMM Estimation of the Number of Latent Factors. Journal of Empirical Finance 17, 783-802.

Anderson, H., Vahid F., 2007. Forecasting the Volatility of Australian Stock Returns: Do Common Factors Help? Journal of Business and Economic Statistics 25, 76-90.

Andrews, D.W.K., Lu, B., 2001. Consistent Model and Moment Selection Procedures for GMM Estimation with Application to Dynamic Panel Data Models. Journal of Econometrics 101, 123-164.

Bai, J., Ng, S., 2002. Determining the Number of Factors in Approximate Factor Models. Econometrica 70, 191–221.

Bai, J., Ng, S., 2006a. Confidence Intervals for Diffusion Index Forecasts and Inference with Factor-Augmented Regressions. Econometrica 74, 1133–1150.

Bai, J., Ng, S., 2006b. Evaluating Latent and Observed Factors in Macroeconomics and Finance. Journal of Econometrics 113, 507–537.

Bai, J., Ng, S., 2008. Extremum Estimation when the Predictors are Estimated from Large Panels. Annals of Economics and Finance 9, 201-222.

Blanco, R., Brennan, S., Marsh, I.W., 2005. An Empirical Analysis of the Dynamic Relation between Investment-Grade Bonds and Credit Default Swaps. Journal of Finance 60, 2255-2281.

Boivin, J., Ng, S., 2006. Are More Data Always Better for Factor Analysis? Journal of Econometrics 132, 169–194.

Brunnermeier, M.K., 2009. Deciphering the Liquidity and Credit Crunch 2007-08. Journal of Economic Perspectives 23, 77-100.

Campbell, J., Taksler G., 2003. Equity Volatility and Corporate Bond Yields. Journal of Finance 58, 2321-2349.

Chamberlain, G., Rothschild, M., 1983. Arbitrage, Factor Structure and Mean-Variance Analysis in Large Asset Markets. Econometrica 51, 1281–1304.

Chen, L., Collin-Dufrense, P., Goldstein, R.S., 2009. On the Relation Between the Credit Spread Puzzle and the Equity Premium Puzzle. Review of Financial Studies 22, 3367-3409.

Clark, T.E., McCracken M.W., 2005. The Power of Tests of Predictive Ability in the Presence of Structural Breaks. Journal of Econometrics 124, 1-31.

Collin-Dufrense, P. Goldstein, R.S., 2001. Do Credit Spreads Reflect Stationary Leverage Ratios? Journal of Finance 56, 1929-1958.

Collin-Dufrense, P., Goldstein, R.S., Martin, S.J., 2001. The Determinants of Credit Spread Changes. Journal of Finance 56, 2177-2208.

Davidson, R., MacKinnon J., 1985, Heteroskedasticity-robust Tests in Regression Directions. Annales de l'INSEE 59, 183–218.

Della Corte, P., Sarno, L., Thornton D.L., 2008. The Expectation Hypothesis of the Term Structure of Very Short-term Rates: Statistical Tests and Economic Value. Journal of Financial Economics 89, 158-174.

Della Corte, P., Sarno, L., Tsiakas I., 2009. An Economic Evaluation of Empirical Exchange Rate Models. Review of Financial Studies, 22, 3491-3530.

Dwyer, G.P., Tkac, P., 2009. The Financial Crisis of 2008 in Fixed-Income Markets. Journal of International Money and Finance 28, 1293-1316.

Eichengreen, B., Mody, A., 2000. What Explains Changing Spreads on Emerging Market Debt?, in Sebastian Edwards (ed.), Capital Flows and the Emerging Economies, Chicago. University of Chicago Press, pp. 107-136.

Elton E., Gruber, M., Aggrawal, D., Mann, C., 2001. Explaining the Rate Spread on Corporate Bonds. Journal of Finance 56, 247-277.

Engle, R.F., 2002. Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models. Journal of Business and Economic Statistics 20, 339-350.

Ericsson J., Jacobs, K., Oviedo, R., 2009. The Determinants of Credit Default Swap Premia. Journal of Financial and Quantitative Analysis 44, 109-132.

Goncalves, S., Kilian, L., 2004. Bootstrapping Autoregressions with Conditional Heteroskedasticity of Unknown Form. Journal of Econometrics 123, 89-120.

Hansen, B., 2000. Testing for Structural Change in Conditional Models. Journal of Econometrics 97, 93-115.

Jarrow, R., Protter, P., 2004. Structural Versus Reduced Form Models: A New Information Based Perspective. Journal of Investment Management 2, 1-10.

Jorion, P., Zhang, G., 2007. Good and Bad Credit Contagion: Evidence from Credit Default Swaps. Journal of Financial Economics 84, 860-883.

Longstaff, F.A., 2000. The Term Structure of Very Short Term Rates: New Evidence for the Expectations Hypothesis. Journal of Financial Economics 58, 397-415.

Longstaff, F.A., Mithal, S., Neis, E., 2005. Corporate Yield Spreads: Default Risk or Liquidity? New Evidence from the Credit Default Swap Market. Journal of Finance 60, 2213-2253.

Longstaff, F.A., Pan, J., Pedersen, L.H., Singleton, K.J., 2010. How Sovereign Is Sovereign Credit Risk? American Economic Journal: Macroeconomics 3, 75–103.

Makarov, I., Papanikolaou, D., 2009. Sources of Systematic Risk. London Business School, mimeo.

Mody, A., Taylor, M.P., 2003. The High Yield Spread as a Predictor of Real Economic Activity: Evidence of a Financial Accelerator for the United States. IMF Staff Papers 50, 373-402.

Mody, A., Sarno, L., Taylor, M.P., 2007. A Cross-Country Financial Accelerator: Evidence from North America and Europe. Journal of International Money and Finance 26, 149-165.

Onatski, A., 2010. Determining the Number of Factors from Empirical Distribution of Eigenvalues. Review of Economics and Statistics 92, 1004-1016.

Pison, G., Rousseeuw, P.J, Filzmoser, P., Croux, C., 2003. Robust Factor Analysis. Journal of Multivariate Analysis 84, 145–172.

Rossi, B., 2005. Optimal Tests for Nested Model Selection with Underlying Parameter Instability. Econometric Theory 21, 962-990.

Sarno, L., Schneider, P., Wagner, C., 2012. Properties of Foreign Exchange Risk Premiums. Journal of Financial Economics, forthcoming. DOI: 10.1016/j.jfineco.2012.01.005

Soros, G., 2009. The Game Changer. Financial Times, January 28.

Stock, J., Watson, M., 2002. Forecasting Using Principal Components from a Large Number of Predictors. Journal of the American Statistical Association 97, 1167–1179.

Taylor, J.B., 2009. Getting Off Track: How Government Actions and Interventions Caused, Prolonged, and Worsened the Financial Crisis, Stanford, CA: Hoover Institution Press.

Zhang, B.Y.-B., Zhou, H., Zhu, H., 2009. Explaining Credit Default Swap Spreads With the Equity Volatility and Jump Risks of Individual Firms. Review of Financial Studies 22, 5099-5131.

Table 1. The Sample of Financial Institutions: Descriptive Statistics

	Mean	Standard deviation	Median	Minimum	Maximum
Abbey	26.29	27.36	15.16	4.51	137.93
Barclays	28.45	37.81	12.21	5.65	251.89
HBOS	35.64	54.12	13.78	4.84	385.86
HSBC	23.73	23.88	14.01	4.99	134.96
Lloyds TSB	22.39	27.24	12.30	3.87	184.40
RBS	28.84	39.71	13.02	4.00	289.11
Standard Chartered	31.65	30.37	18.59	5.87	176.73
Allianz	36.19	27.76	24.92	6.03	136.59
Commerzbank	45.81	41.82	27.05	8.14	221.83
HVB	43.24	36.41	29.64	6.32	167.81
Deutsche Bank	31.38	28.54	18.29	10.11	161.68
Dresdner Bank	34.73	30.22	22.00	5.52	154.83
Hannover Rueckversicherung	44.88	30.08	34.67	8.17	138.25
Münchner Hypoth.	32.31	20.56	25.56	5.78	120.19
Monte dei Paschi	30.90	24.31	21.03	6.17	137.93
UniCredit	28.27	25.28	16.06	7.65	133.66
AXA	49.19	46.95	27.93	9.11	235.14
BNP Paribas	20.51	19.01	12.31	5.51	103.83
Credit Agricole	24.19	27.47	13.27	6.08	145.56
LCL	24.58	27.61	12.58	6.16	150.27
Société Générale	24.85	27.11	13.25	5.97	138.93
ABN AMRO	26.79	29.92	15.03	5.26	172.69
ING	26.54	31.12	15.44	4.45	163.93
Rabobank	16.83	22.49	8.35	3.12	131.92
Banco Santader	31.46	30.47	16.89	7.68	143.98
Credit Suisse	38.98	35.44	20.87	9.04	169.57
UBS	29.04	43.34	11.63	4.55	276.23
Banco Comerc. Port.	33.20	28.23	21.86	8.20	145.36
American Express	59.97	90.51	25.84	9.07	603.16
AIG	101.24	310.18	24.87	8.57	2624.15
Bank of America	35.03	33.93	21.34	8.47	191.45
Bear Sterns	60.14	62.51	34.47	19.01	574.31
Chubb	39.28	29.26	29.94	10.01	153.70
Citibank	43.19	54.79	20.83	7.52	351.93
Fed. Mortgage	22.81	14.09	20.67	6.35	82.55
Freddie Mac	21.77	14.77	19.03	5.28	83.10
Goldman Sachs	57.72	64.67	34.33	18.95	437.37
Hartford	67.82	111.88	36.36	10.82	826.17
JP Morgan	44.05	31.96	31.37	11.81	174.98
Lehman Brothers	86.00	127.03	36.74	19.49	641.91
Merrill Lynch	68.94	77.04	34.60	15.61	417.10
Met Life	62.58	108.22	30.46	11.20	790.42
Morgan Stanley	73.21	117.44	34.42	18.59	1153.09
Safeco	46.01	34.38	32.88	18.04	181.00
Wachovia	52.70	84.18	21.42	10.40	741.79

Table 2. The Validity of the Correlates

Variable set	Sample size	J_{N-4}	$BIC1$	$BIC2$	$BIC3$	HQQ
I_1	$29/07/2002 - 28/11/2008$	0.01	4(0.52)	4(0.77)	4(0.91)	4(0.74)
I_1	$29/07/2002 - 30/04/2008$	0.01	1(0.52)	4(0.64)	4(0.85)	4(0.525)
I_1	$29/07/2002 - 18/07/2007$	0.01	1(0.74)	3(0.47)	4(0.87)	4(0.475)

Notes: The test formulas are defined in Ahn and Perez (2008). The number of randomizations of the cross-section ordering was set to 500. J_{N-4} shows the percentage of rejections of the moment condition across randomizations when the true number of factors is 4 (test for independence of idiosyncratic errors and observable series). BIC1, BIC2, BIC3 and HQQ show the number of factors suggested by each and, in parenthesis, the empirical frequency of the selected number of factors over all randomizations. The set of correlates, I_1 includes the high-yield spread, the VIX, the S&P 500 returns, the spread on asset-backed commercial paper, the spread of the LIBOR over the overnight index swap, and the spread of the overnight index swap over the T-bill rate.

Figure 1: Evolution of Spreads on Credit Default Swaps (median, in basis points)

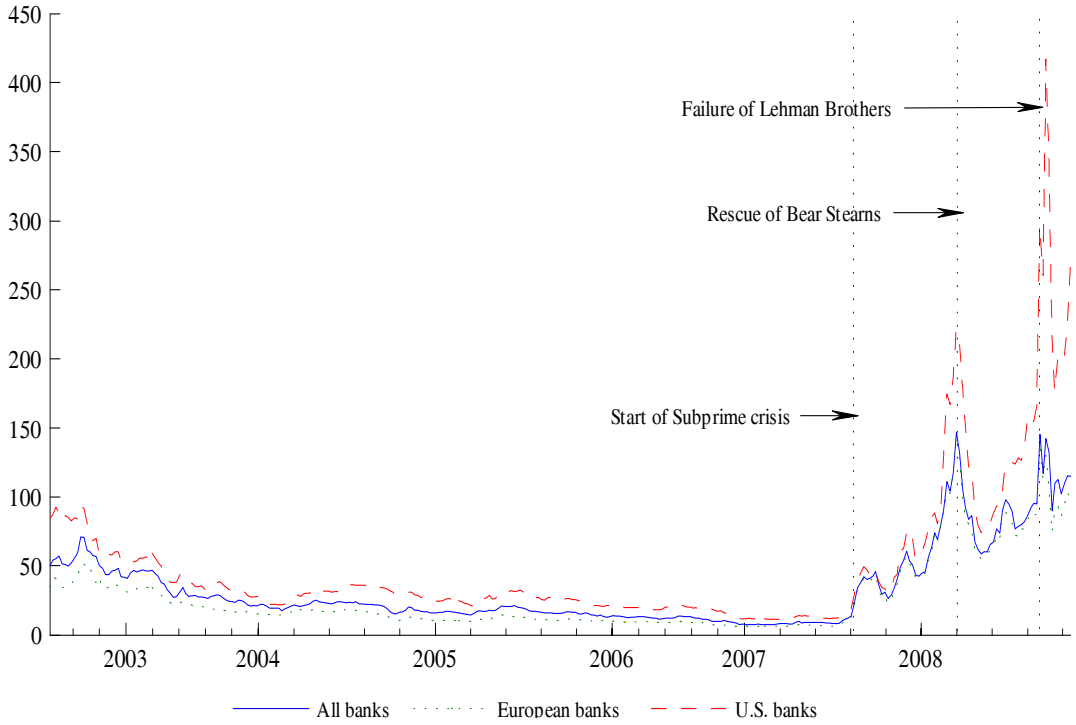

Figure 2: Share of CDS Spreads' Variation Explained by the Four Factors

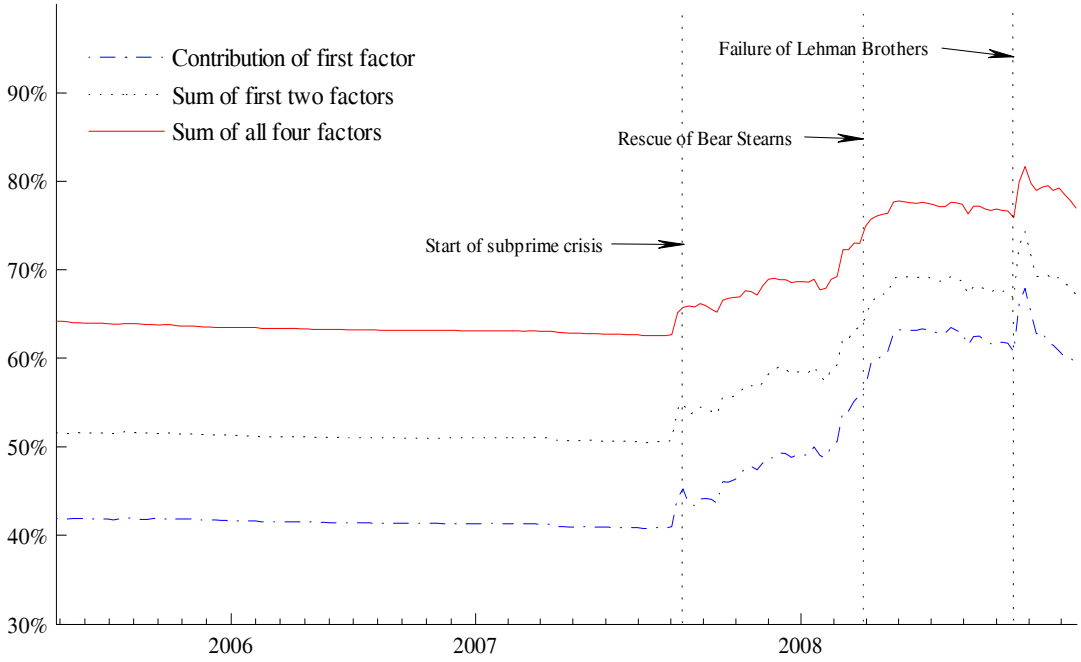

Figure 3. Additional Spillovers - (a) Fraction of Pairs of Banks Incurring "Additional" Spillovers

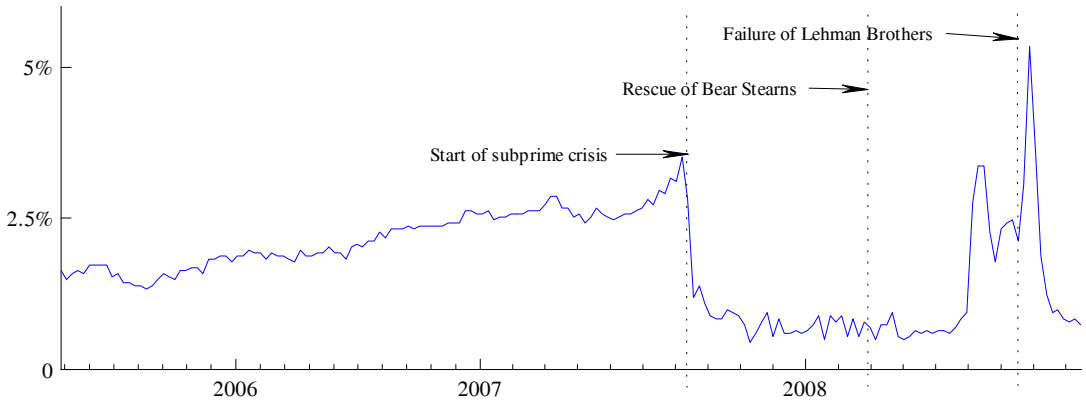

(b) Fraction of European Banks Facing Additional Spillovers from at least one American Bank

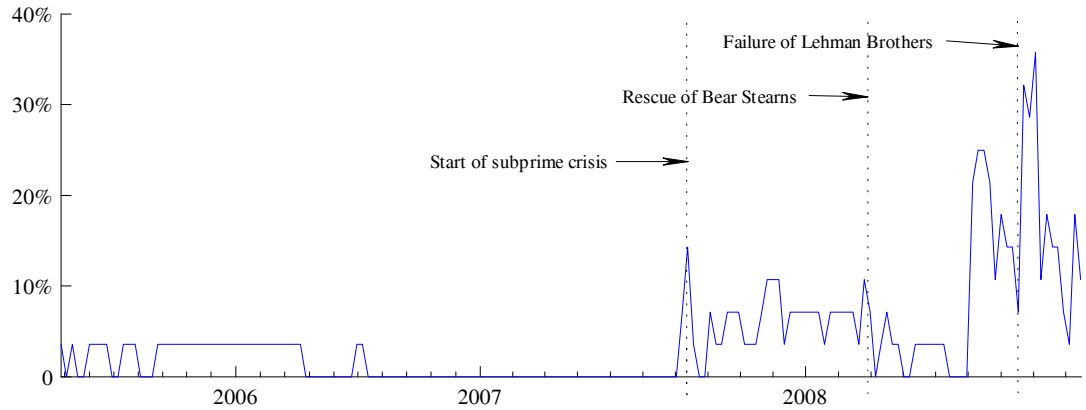

(c) Fraction of American Banks Facing Additional Spillovers from at least one European Bank

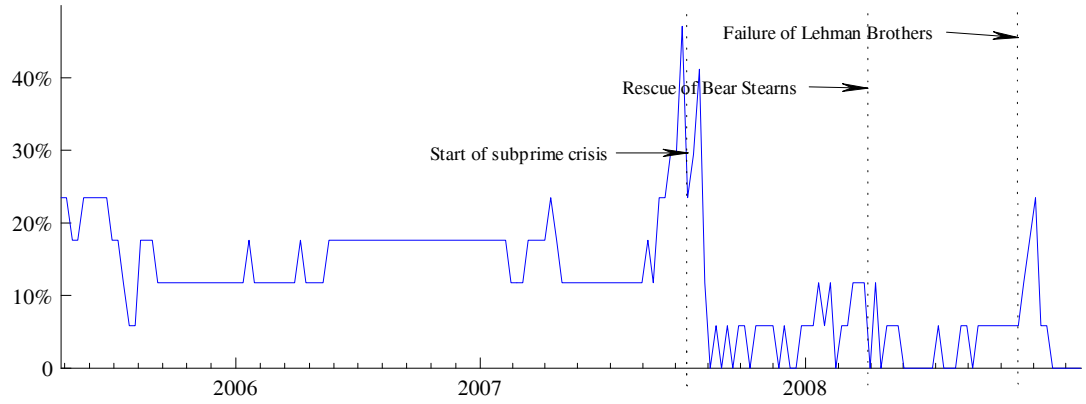

Figure 4. CDS Spread and the "Real Economy"

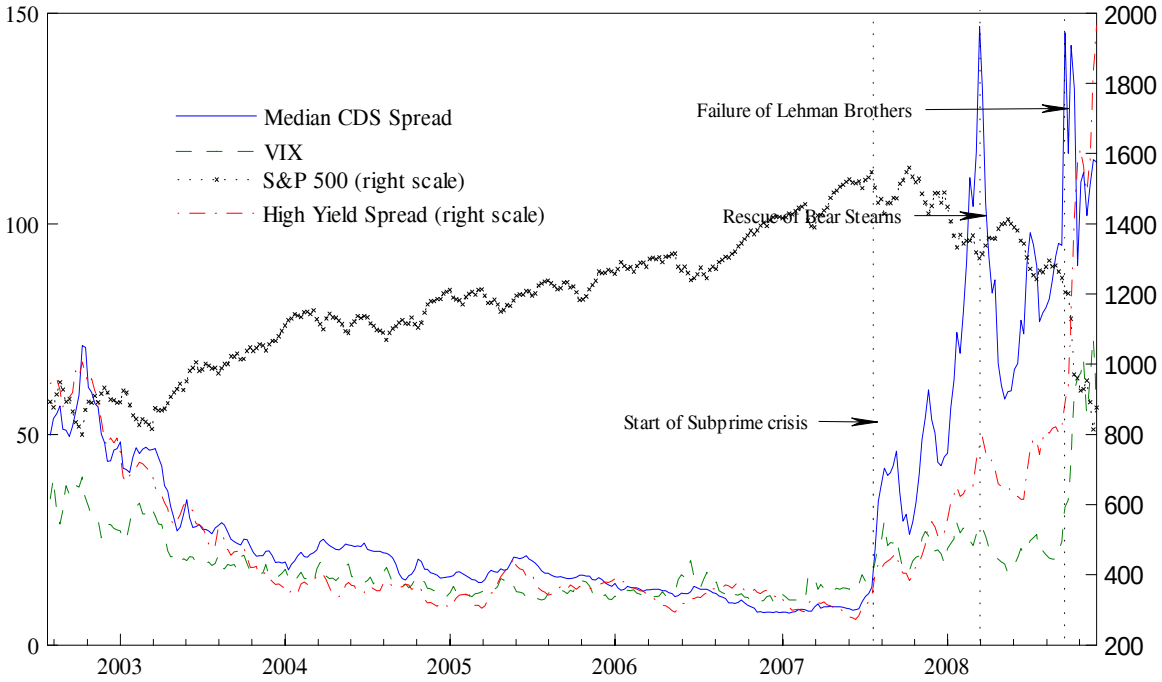

Figure 5. Association of Common Components with the "Real Economy"

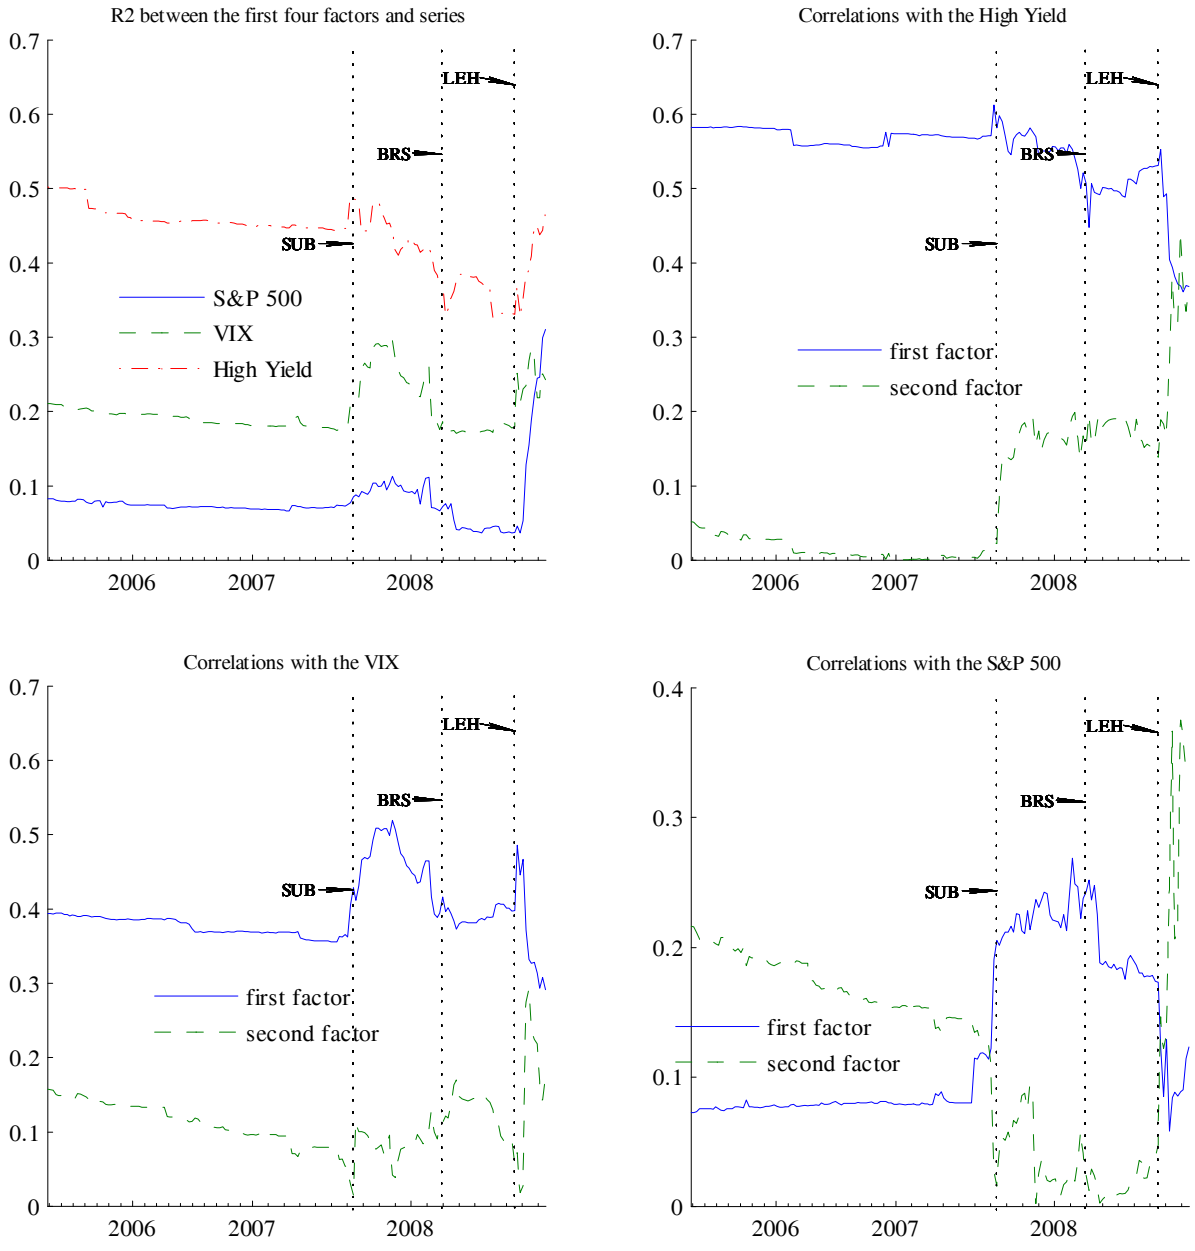

Notes: Vertical line indicators: SUB: Start of Subprime crisis; BRS: Rescue of Bear Stearns; LEH: Failure of Lehman Brothers

Figure 6. CDS Spread and Costs of Funding (basis points)

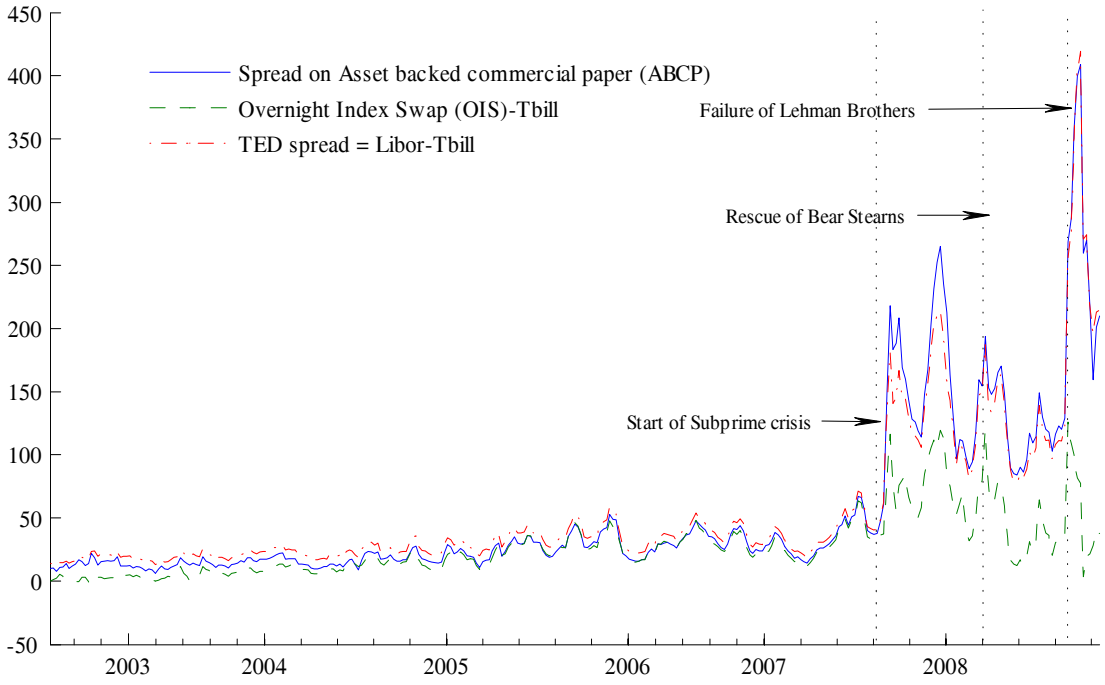

Figure 7. Association of Common Components with the Costs of Funding

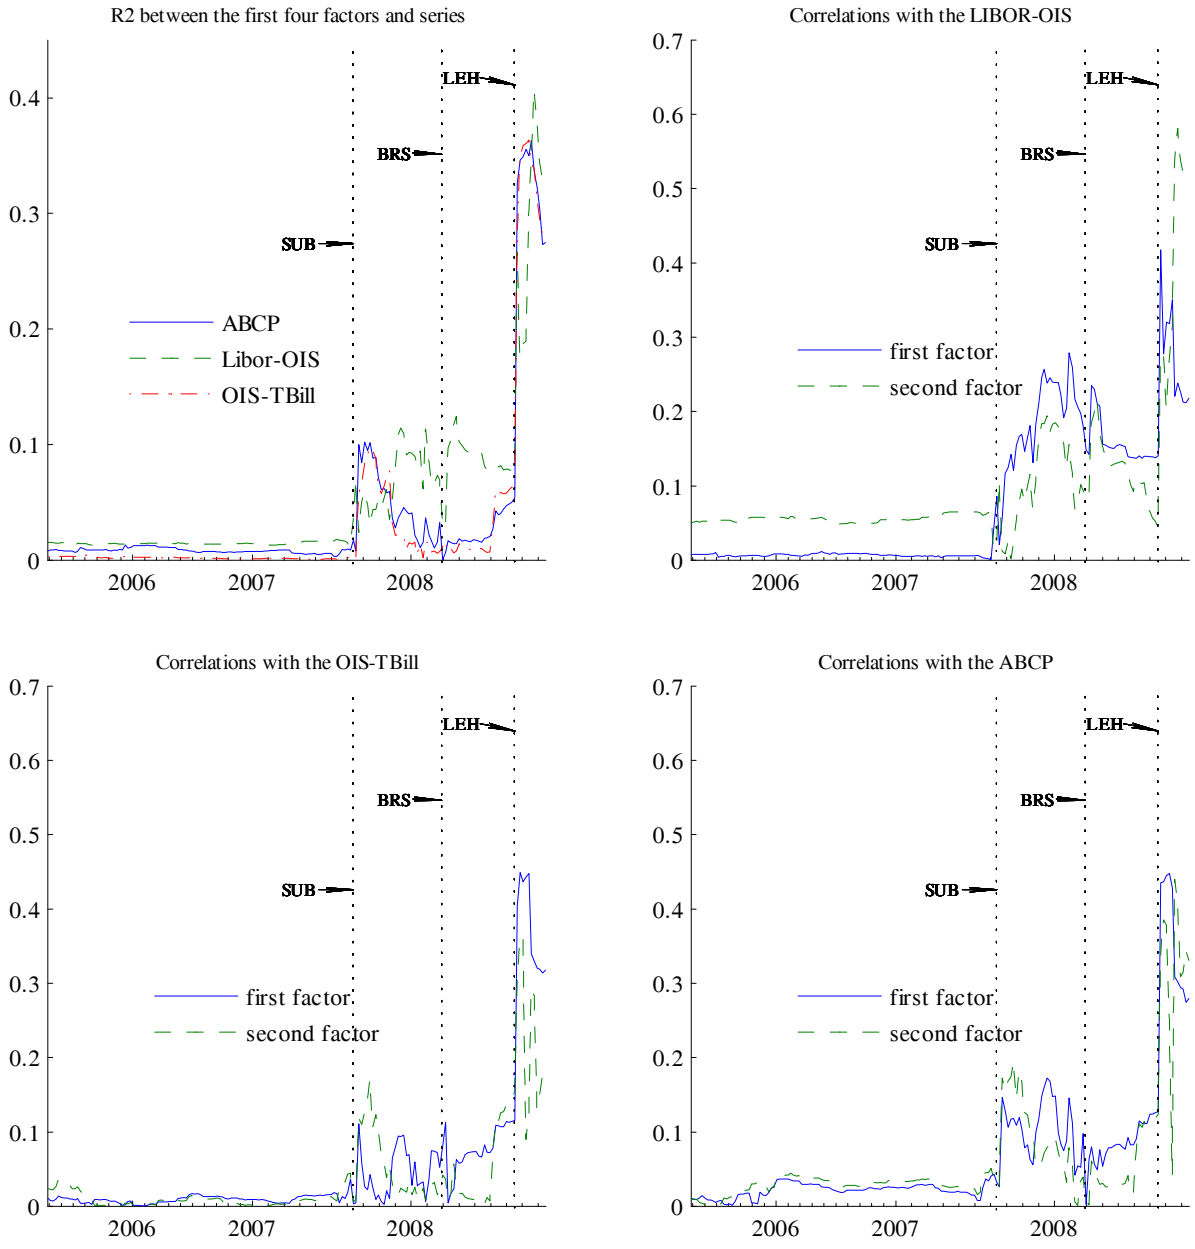

Notes: Series: OIS: Overnight Index Swap; ABCP: Spread on Asset Backed Commercial Paper; Tbill: 3-month Treasury Bills. Vertical line indicators: SUB: Start of Subprime crisis; BRS: Rescue of Bear Stearns; LEH: Failure of Lehman Brothers;