WP/13/267

IMF Working Paper

Does Financial Connectedness Predict Crises?

Camelia Minoiu, Chanhyun Kang, V.S. Subrahmanian, Anamaria Berea

INTERNATIONAL MONETARY FUND

 WP/13/267

IMF Working Paper

Research Department

Does Financial Connectedness Predict Crises?[1]

Prepared by Camelia Minoiu, Chanhyun Kang, V.S. Subrahmanian, Anamaria Berea

Authorized for distribution by Andrew Berg

December 2013

Abstract

The global financial crisis has reignited interest in models of crisis prediction. It has also raised the question whether financial connectedness—a possible source of systemic risk—can serve as an early warning indicator of crises. In this paper we examine the ability of connectedness in the global network of financial linkages to predict systemic banking crises. Our results indicate that increases in a country's financial interconnectedness and decreases in its neighbors' connectedness are associated with a higher probability of banking crises after controlling for macroeconomic fundamentals.

JEL Classification Numbers: F01, F36, F47
Keywords: early warning systems, systemic risk, financial networks, banking crises

Author's E-Mail Addresses: cminoiu@imf.org; chanhyun@cs.umd.edu; vs@umiacs.umd.edu; aberea@c4i.gmu.edu

---

[1] Affiliations: IMF (Minoiu); University of Maryland Department of Computer Science (Kang and Subrahmanian); George Mason University C4I Center (Berea). We are grateful for helpful comments from Adrian Alter, Andrew Berg, Charlie Brummit, Eugenio Cerutti, Giovanni Di Iasio, Mathias Drehmann, Grzegorz Halaj, Tom Hurd, Tümer Kapan, Sergio Montagna, Andrea Presbitero, and participants at the "Interlinkages and Systemic Risk" Conference (Ancona, July 2013). We thank Luis Catão for generously sharing his data. Some of the authors have leveraged research funded by the US Army Research Office under grants 11NF0910206, W911NF1160215, W911NF1110344, and W911NF-09-1-0525 (MURI). The views expressed in this paper are those of the authors and do not represent those of funding agencies, the IMF, or IMF policies. All remaining errors are our own.

Contents

Appendices

List of tables

List of figures

1 Introduction

The global financial crisis has underscored the role of financial connectedness as a potential source of systemic risk and macroeconomic instability. It has also highlighted the need to better understand whether an increase in connectedness leads to a higher probability of a financial crisis. In this paper, we contribute to the literature on "early warning systems" (EWS) by investigating whether measures of systemic risk, which are based on modeling the global financial system as a network, can serve as early warning indicators and improve the performance of standard crisis prediction models. In doing so, we use two commonly-employed prediction methods – a data mining algorithm and an econometric crisis incidence model.

We build on the literature that links connectivity in the global banking network ("GBN") to systemic risk and contagion that arise during systemic banking crises.[1] In particular, we compute indicators of connectedness based on the pattern of linkages across banking systems worldwide and test their performance as leading crisis indicators. The EWS literature has traditionally focused on balance of payments and currency crises in emerging market countries,[2] and has identified several macroeconomic factors that are robustly correlated with crises. These include asset price bubbles, low international reserves, real exchange rate appreciation, hyperinflation, and excessive short-term and foreign exchange borrowing. Our paper adds to this literature by focusing on the role of financial connectedness in predicting systemic banking crises. It also speaks to the concern that complex linkages among nations and financial institutions may be a source of systemic risk and accentuate business cycle downturns (Yellen (2013), Stiglitz (2010), Schweitzer et al. (2009)).

Our strategy is two-pronged. First, we assess the ability of a wide range of network-based con-

[1] See Gai et al. (2010), Cont et al. (2012), and Allen et al. (2012).

[2] Seminal contributions include Eichengreen et al. (1995), Kaminsky et al. (1998),Demirguc-Kunt & Detragiache (1998), Kaminsky & Reinhart (1999), and Berg & Pattillo (1999). Berg et al. (2005) caution that the out-of-sample performance of EWS, especially over long horizons, has been mixed. In recent years, the literature has received a new stimulus, with a focus on evaluating factors associated with growth collapses (Dabla-Norris & Gunduz (2012)), fiscal crises (Baldacci et al. (2011)), and financial crises in advanced economies (IMF (2010)). Credit bubbles have been identified as the most robust determinant of banking crises over very long time spans (Schularick & Taylor (2012)).

nectedness measures to predict crises according to a data mining technique, namely a "classification algorithm."[3] Classification algorithms have been used extensively in applications across fields including genetics (Creighton & Hanash (2003)), medical sciences (Karabatak & Ince (2009)), finance (Lu et al. (1998)), meteorology (Feng et al. (2001)), and terrorism research (Subrahmanian et al. (2013)).[4] Second, we employ a standard econometric model of crisis prediction – a probit – and evaluate its performance with and without these indicators. Probit models have been used with varying degrees of success to predict balance of payments, debt, and banking crises.[5]

Our findings indicate that network-based measures of financial connectedness are helpful in predicting systemic banking crises. In particular, increases in a country's interconnectedness and decreases in the connectedness of that country's direct financial partners (or neighbors) are associated with a higher probability of crisis. The results of the classification algorithm suggest that measures of financial connectedness are useful in predicting both crisis-years and the onset of crises. This is especially true when we focus on the most recent wave of crises in advanced economies. The probit analysis reinforces these results. Estimates from a probit model that incorporates connnectedness indicators in addition to standard macroeconomic fundamentals outperforms, in and out of sample, a model based on fundamentals alone.

Our paper adds to a growing literature on the evolution of financial connectedness over the business cycle. In this literature, cross-country linkages are captured in many different ways. Chinazzi et al. (2013) build a GBN using data on cross-country debt flows and equity investments. The authors show that countries that are more connected, or "central", in the GBN prior to the recent crisis experienced a smaller loss of output during 2008-2009. Hale (2012) and Caballero et al. (2009) construct GBNs using granular information on bank-level exposures in the syndicated loan market. Hale (2012) finds that US recessions are associated with lower network connectivity.

[3] The algorithm is described in detail in Section 2.3.

[4] The algorithm we use is patented by one of the authors (Subrahmanian & Ernst (2009)).

[5] See Demirguc-Kunt & Detragiache (2005), Abiad (2003), Edison (2000) and Chui (2002) for reviews of the large econometric literature on EWS.

Caballero et al. (2009) show that the home countries of banks that are centrally located in the GBN experienced lower stock market crashes during 2007-2008 compared to other countries. Hale et al. (2013) argue that direct and indirect financial exposures can act as conduits for the transmission of banking crises internationally. Minoiu & Reyes (2013) construct a GBN based on cross-border bank lending and securities investments and show that connectivity tends to rise before financial debt crises and to fall afterward. This body of work underscores the importance of connectivity for economic performance during crises.

Few studies look directly at the question whether financial connectedness can serve as an early warning indicator. Billio et al. (2012) show that linkages across four US financial sectors – hedge funds, banks, broker/dealers, and insurance companies – have become more significant prior to the 2007-2009 crisis. In particular, the impact of returns of banks and insurers on those of hedge funds and broker/dealers was stronger than the other way around before the crisis. This finding suggests that measures of connectedness from the returns correlation network may be useful out-of-sample indicators of systemic risk. Caballero (2012) measures financial integration based on the network of interbank lending and borrowing in the syndicated loan market, and finds that country-level indicators of financial connectivity, which are proxies for financial integration, have predictive power for the incidence of banking crises during 1980-2007 after controlling for credit growth and other macroeconomic fundamentals.

Our work differs from these studies in several ways. First, while Billio et al. (2012) focus on US financial institutions and hence on systemic risk in the US economy, we look at the ability of *global* financial linkages to predict systemic banking crises in a *cross-country* setting. Second, we investigate the performance of EWS augmented with measures of financial connectivity and include 2008-2010 crises (not covered in Caballero (2012)). Including the latest generation of crises – which account for 74 out of 430 country-years over 1978-2010 – is important because financial connectedness may arguably have played a more important role during the most recent wave of crises compared to earlier ones. Third, while we use less granular information than Hale et al. (2013) and Hale (2012), who employ bank-level information on syndicated loan contracts, our data gives a comprehensive view of cross-border banking *system* exposures – of which syndicated loans are only one component. Finally, we consider a wide range of network-based connectedness indicators – 27 distinct indicators – along with their lags and growth rates.

To construct the GBN we use annual data on cross-border banking system exposures from the BIS locational statistics over 1978-2010. The BIS locational statistics have been used extensively to analyze geographical patterns in financial linkages (Minoiu & Reyes (2013); Hattori & Suda (2007)), financial integration (Kalemli-Ozcan et al. (2013)), and factors associated with capital flows (Blank & Buch (2010, 2007)). The data represent foreign exposures of banks in BIS reporting countries vis-a-vis residents (borrowers) in all other countries. The data are representative of banking activity in each location as reporting banks typically account for more than 90 percent of total banking assets in each BIS reporting country. For each year we construct a directed GBN with countries as "nodes" and cross-border exposures as "edges," and compute network indicators using both the binary and weighted versions of this network (defined in Section 2.1).

Our analysis addresses several related, yet distinct, dimensions of systemic risk. The first one is connectedness narrowly understood as direct exposures among economic agents (such as financial systems or institutions). Connectedness matters for crisis incidence because the failure of an economic agent can lead to the failure of another agent through this direct channel. An indirect channel is contagion, by which an agent's failure leads to panic among agents not directly connected to him/her. It is difficult to empirically separate the effects of connectedness from those of contagion. In this paper we take a view of connectedness that encompasses the two concepts and refers to *both* direct and higher-order (indirect) exposures. We take this concept to the data by mapping financial linkages across countries and extracting information from the pattern of connections by means of network analysis.

The remainder of this paper is organized as follows. In Section II, we introduce the data, describe the construction of the GBN and present discuss long-term trends in financial connectivity. We also describe the classification algorithm. In Section III, we present the results, ranging from simple correlations between connectivity and crises to the performance of the classification algorithm and that of the regression analysis. In Section IV we conclude. Formal definitions for all the network indicators used are provided in Appendix A. Supplemental results are available in an online appendix.[6]

2 Data and Methods

2.1 Data Description

We use data from the BIS bilateral locational statistics – a component of the BIS international banking statistics – that provide a comprehensive view of international banking transactions. The data represent stocks of cross-border assets (such as loans, debt securities, and other assets) held by banking systems in 210 countries during the 1978-2010 period. Similar to balance-of-payments statistics, these data are defined on a *locational* basis, i.e., they capture exposures of reporting banks in a given location regardless of their ownership structure (BIS (2009, 2011)). As such, the BIS locational statistics are useful for measuring financial integration across countries and territories (such as off-shore financial centers) that report financial data.[7]

The sample comprises 29 BIS reporting countries and 181 non-reporting countries. In what

[6]The online appendix is available for download from www.camelia-minoiu.com/crisis-prediction-appendix.pdf.

[7]A list of all countries and territories included in our sample is included in the online appendix.

follows, we refer to the former as "core" countries and to the latter as "periphery" countries.[8],[9] The GBN is constructed solely on the basis of data from the core countries. These countries also report liabilities, which we use to impute assets of periphery countries vis-a-vis core countries. To understand how representative these data are of true global linkages, we can imagine the full network as consisting of three layers: the links among core countries, those between core and periphery countries, and those among periphery countries. Since core countries report positions to the BIS, we observe the full extent of linkages within the core. We similarly observe exposures of core countries vis-a-vis periphery countries. The liabilities of core countries vis-a-vis periphery countries are used as the claims of periphery countries vis-a-vis core countries so we also have a measure of linkages from the periphery toward the core.

It is important to note that we do not have data on linkages *among* periphery countries; in other words, connections within the periphery are missing from our GBN. This caveat should be kept in mind when interpreting the results as our network variables may suffer from measurement error and we cannot know if these errors are systematic. To date, data availability remains a common problem in studies of global systemic risk, as discussed, for instance, in Cerutti et al. (2012). Nonetheless, ours is a first attempt to utilize core countries' liability-side data from the BIS locational statistics to obtain as comprehensive a view as possible of cross-country banking linkages.[10]

The GBN is directed: a directional edge is created from country A to country B if banks in

[8] Countries are invited to report financial data to the BIS when their financial sector becomes large. In order to minimize the effect of changes in sample composition on our results – due to countries starting to report to the BIS at different points in time, or because some countries become independent states during the sample period – we also experimented with a GBN based on data from the 13 core countries that report their claims continuously and the 175 countries that exist as independent entities since 1978, but the results were largely unchanged.

[9] We use the "core" and "periphery" labels for the nodes in the GBN based on the structure of the data rather than a formal core-periphery model. For recent work on core-periphery structures of interbank networks, see, e.g., Fricke & Lux (2012).

[10] Note that exposures from asset- and liability-side data are not strictly comparable because the former refer to assets of banking systems in core countries vis-a-vis residents in periphery countries, while the latter refer to assets of residents in periphery countries vis-a-vis banks in core countries. This inconsistency is unlikely to affect our binary GBN, but may lead to some measurement error in the weighted GBN. For purposes of our analysis, we need to obtain as comprehensive a view of cross-country financial linkages as possible; therefore we prefer to use an imperfect measure of periphery-core linkages rather than assume that they do not exist.

country A have non-zero claims vis-a-vis borrowers in country B. In the binary GBN, an edge exists if a country has non-zero banking claims on another country. The data for each year over 1978-2010 is modeled as a separate GBN. We also consider two variants of the GBN with different edge weighting schemes. The first one – the "baseline network" – refers to edge weights given by log-exposures divided by the log-product of GDPs within each pair of nodes. PPP-adjusted GDP data come from the Penn World Tables (Mark 7.1). The advantage of this weighting scheme is that banking system claims are scaled by the economic size of countries, which ensures that increases in weighted indicators of connectedness are not driven solely by economic growth. The second, which we use for robustness purposes and refer to as the "alternative network," uses real log-exposures as edge weights.[11] All the results included in the paper are based on the first variant, but are robust to using the second one.[12]

2.2 Global Connectivity: Stylized Facts

We assess the extent and trend in global connectivity during 1978-2010 using selected indicators of connectedness. (See Appendix A for formal definitions.) The GBNs for 1980 and 2007 are shown in Figure 1. The visualizations depict an increase in global financial connectivity both in the full sample (top panels) and the core (lower panels), with the latter showing a significantly denser network in 2007, before the global financial crisis.

Figure 2 plots network density and total exposures. Network density refers to the share of observed edges in the total number of possible edges, and ranges in the full GBN between 5 percent at the beginning of the sample period and 12 percent in 2007. These figures are markedly larger in the core network, with density rising from 17 percent in 1978 to 46 percent in 2007. While network density has increased continuously over the past three decades, at an annual average rate of 3 percent, cross-border exposures have also risen, but almost three times faster (8 percent). At

[11]Exposures are expressed in real terms using the US CPI.

[12]See online appendix for robustness checks using the alternative network.

the peak of network density in 2007, global exposures amounted to USD 90 trillion – about twice the size of global GDP.

Averages for our network indicators are reported in Table 1 at different points during the sample period. The last column in the table shows the percent increase in each network indicator between 1978 and 2010, indicating that financial integration – proxied by our rich set of measures – has been on the rise since 1978. Average degree and strength, representing the number of financial partners and the average banking system exposures of countries in the GBN, display the largest increases among all indicators, by a factor of 1.5 over 1978-2010.

Clustering coefficients capture the tightness of link formation around a node and are bounded between 0 and 1. The clustering coefficients developed by (Fagiolo (2007)) answer the following question: Given that a node has connections to two neighbors, what is the probability that those neighbors are also connected with each other (i.e., that they form a triangle)? For the 2000s the answer is 80 to 90 percent, which implies a high tendency for nodes to form triangles. The second clustering coefficient we consider, proposed by Lopez-Fernandez et al. (2004), captures the effectiveness with which a node's neighbors are connected with one another and has also risen over the sample period. Overall, the rise in average clustering coefficients – by between 59 and 85 percent during the period of analysis – suggests that not only has the number of bilateral financial links and the size of exposures increased, but countries' financial neighbors have also become more tightly linked with one other.[13]

The next set of indicators refers to connectivity among a country's neighbors. Average nearest-neighbor degree (ANND) is the number of creditors or debtors that each country's immediate neighbors have. For instance, "out-out" ANND is the average number of creditors that a creditor country has. Similarly, "in-in" ANND is the average number of debtors that a debtor country has. ANNS measures take into account the size of cross-border exposures, too. All neighbor connectivity

[13]Recent work argues that clustering coefficients based on different triangle patterns in banking networks have different implications for systemic risk (Tabak et al. (2012, 2011)).

indicators are on the increase until 2007, and fall in the wake of the crisis.

Our last indicator is the Herfindahl-Hirschmann index (HHI), a commonly used measure of market competition. Although this is not a network indicator, it helps assess the degree of portfolio diversification in the GBN. The HHI, computed from the debtor countries' perspective, is equal to the sum of squared shares a country represents in its creditors' portfolios.[14] The HHI indicates a relative lack of concentration in the first half of the period, but increases during the 1990s to levels of moderate concentration. By the end of the sample period the HHI declines again, with debtor countries increasingly diversifying their pool of creditors. This trend coincides with a period of fast economic growth coupled with increased capital account openness and financial system reform in many countries.

In sum, our network measures depict a general trend towards greater financial integration during 1978-2007 – both in terms of the average country's direct connectivity and the tightness of countries' neighbor networks. This trend was interrupted by the global financial crisis. The averages in Table 1 are consistent with the network visualizations in Figure 1 and the evolution of network density and global exposures in Figure 2.

2.3 Classification Algorithm

We begin to evaluate our connectedness measures' effectiveness as early warning indicators with a data mining method, namely a "classification algorithm." The 27 indicators we consider are included in the algorithm in levels lagged up to 5 years, and in growth rates computed over the past 1, 2,.., up to 5 years. We set up the GBN as a matrix whose rows correspond to country-year observations and whose columns correspond to (i) our set of network-based indicators of connectedness, and (ii) the crisis incidence variable `crisis`. The latter takes value 1 if a crisis

[14]For instance, if country A has 50 percent of its liabilities vis-a-vis country B and 50 percent vis-a-vis country C, then A's HHI is equal to 0.5.

occurred in the country during the specified year and 0 otherwise.[15] Crisis incidence data come from Laeven & Valencia (2012).

The output of the classification algorithm is a series of association rules that describe relationships between measures of connectedness and the probability of a crisis (Agrawal et al. (1996)). Association rules are of the form "If condition C holds over the network indicators, then the crisis incidence variable takes value 1." The condition C can take the form $\ell_1 \leq V_1 \leq u_1 \wedge \cdots \wedge \ell_n \leq V_n \leq u_n$ where each V_i is a network indicator and the ℓ_i, u_i's are real numbers.[16]

For an association rule to be acceptable, it needs to have a high *confidence* and a high *support*. Confidence is the conditional probability $\mathbb{P}(\texttt{crisis} = 1|C)$, i.e., the share of correctly-called crises. In most applications, the confidence is required to be above a pre-specified threshold. However, confidence alone does not define "goodness" of an association rule because a rule may have a high probability (e.g., probability of 1) with C being true just once and `crisis` being true on that one occasion. In addition to confidence, we use support, which is usually defined as $|\{r \mid r.\texttt{crisis} = 1 \wedge r.C\}|$ where r is a row in the matrix described above, $r.\texttt{crisis}$ is the value of the crisis entry in row r and $r.C$ is true iff row r satisfies condition C. Thus support is simply the number of times (i.e., rows) for which both C and $\texttt{crisis} = 1$ are simultaneously true for a given country-year pair. As in the case of confidence, support is required to exceed a threshold – i.e., the two conditions must co-occur sufficiently many times – in order for the association rule to be acceptable.

We also require association rules to satisfy the property that "negative probability" be low where negative probability is defined as $\mathbb{P}(\texttt{crisis} = 1|\neg C)$.[17] This computes the probability of a crisis occurring when C is not true (the $\neg$ denotes negation). By ensuring that $\mathbb{P}(\texttt{crisis} = 1|\neg C)$ is below a pre-specified threshold, we ensure that condition C acts like a "canary in a coalmine."

[15] We run the algorithm on the full set of crisis-years (rather than the onset of crises) to preserve the maximum amount of variation in the data and improve the algorithms' chances to learn from it.

[16] Here all network indicators have real valued attributes.

[17] While there are several algorithms around to extract association rules from tabular data, there are fewer which also take negative probabilities into account.

When C is true it predicts a crisis with high probability - when it is not true, it predicts a crisis with very low probability.

Classification algorithms like the one used here have been useful in a range of applications spanning breast cancer detection, prediction of movements in the stock market, meteorological phenomena, and upticks in terrorism. To the best of our knowledge, this is the first application of this algorithm to a macroeconomics question – the prediction of systemic banking crises. Results based on the classification algorithm are presented in Section 3.2.

3 Results: Connectivity and Banking Crises

3.1 Exploring the Data: Conditional Correlations

Here we explore the predictive ability of financial connectedness for crisis incidence by examining whether network indicators display a different behavior before vs. after the onset of a crisis. Since the impact of connectedness on the likelihood of crises may be driven by global and country-specific factors, we regress the network indicators (from the baseline GBN) against a full set of country and year dummies and obtain the residuals. We then plot the cross-sectional averages of these residuals – taken across all country-year observations – within a window of five years around the onset of crises. Figure 3 shows the evolution of selected network indicators around banking crises after controlling for unobserved heterogeneity as described above. Note that here we do not account for the impact of any macroeconomic fundamentals (we do so in Section 3.3).

The three panels in Figure 3 present remarkable patterns. In Panel 3a we notice that node degree and strength increase steeply before the onset of a crisis and level off subsequently. The signal here is the rapid growth in interconnectedness captured by these simple network measures. The degree of clustering also rises before banking crises and peaks 1-2 years before their onset, after which it begins to decline (Panel 3b). The pre-crisis decline in clustering reflects a lower probability that a given country's financial partners are connected, and suggests that there is turbulence in the network right before the onset of a crisis. What happens to a country's neighbor connectedness before crises? Panel 3c shows that average neighbor degree and strength (ANND and ANNS) decline 3-4 years before the onset of a crisis. This suggests that there is link destruction among countries' second and higher-degree connections long before the onset of crises; and that the decline in neighbor connectivity can also be informative about the arrival of a crisis.

In the panels of Figure 3 we also show a dotted line that captures most of the information in our indicators. It represents the first principle component (henceforth labelled "1st PC") extracted from the variation in three groups of indicators through Principal Component Analysis (PCA).[18] The first group (Panel 3a) comprises degree and strength indicators. The second group (Panel 3b) includes all the clustering coefficients (both binary and weighted). The third one (Panel 3c) includes all the neighbor degree and strength indicators. The first component explains between 80 and 90 percent of the variation in the underlying indicators, and helps summarize the information from many variables into one single factor. Working with the PCs of several groups of indicators is particularly useful in an EWS if we wish to estimate parsimonious regression models.

Following this preliminary evidence that hints at the ability of financial connectedness to predict crises, we formalize our results with the classification algorithm (Section 3.2) and econometric analysis (Section 3.3).

3.2 Results from the Classification Algorithm

We ran the classification algorithm on the baseline GBN for the full period (1978-2010) and two sub-periods: 1978-2002, which covers first- and second-generation crises mainly affecting developing economies;[19] and 2003-2010, which covers the latest wave of banking crises in advanced economies.

[18] See Appendix B for details.

[19] The only advanced economy crises during this period include the 1991-1992 Scandinavian banking crises and the 1992 Exchange Rate Mechanism crisis.

Table 2 summarizes the in-sample performance of the algorithm in identifying crisis-years for the full sample and separately for core and periphery countries. For each sample and period, we report the number of actual crises (column 1), the number of crises predicted by the algorithm (2), support (3), precision (4), recall (5), and the number of sub-rules (6). Precision and recall are measures of the performance of the algorithm. Precision is defined as the number of correctly predicted crises divided by the number of predicted crises. It takes maximum value when are no "false alarms" (Type 1 error is 0). Recall refers to the number of correctly predicted crises divided by the number of actual crises. It attains a maximum when there are no "missed crises" (Type 2 error is 0).

For the sample of countries included in the baseline GBN there were 409 crisis-years during 1978-2010, of which more than 75 percent occurred before 2003. The algorithm produces 34 association rules (Table 2, column 6) which we can combine into one giant rule to predict crises.[20] To give an example of several sub-rules, the algorithm yields that "A crisis will occur in a given year if (i) out-degree at $t-3$ is between 155 and 171 or (ii) the "in" BCC at $t-5$ is between 0.115 and 0.118 or (iii) in-strength at $t-5$ is between 63.5 and 66.5," etc.[21]

Over the full period, precision is very high, reaching 94 percent for the full sample and 88 percent for core countries (column 4). Recall is relatively low at 11 percent for the full period, but higher by subperiod: 94 percent for 1978-2002 and 40 percent for 2003-2010 (column 5). Precision and recall are lower for core countries compared to periphery countries during the full period, which may be partly due to the lower number of crisis-years experienced by advanced economies overall, and hence the lower ability of the algorithm to learn about them from that sample.[22] By contrast, for the 2003-2010 period which mainly contains crises for core countries, the algorithm has high precision and recall (0.85 and 0.83 respectively), thus performing quite well in identifying crisis-years for the

[20] Each association rule has the form `crisis` if C_1,..., `crisis` if C_n. These rules can be combined into one large rule `crisis` if C_1 or $\cdots$ or C_n saying that a crisis occurs if any of the conditions $C_1, \ldots, C_n$ is true.

[21] See online appendix for all the subrules generated by the algorithm for the core countries, 2003-2010.

[22] Another reason is that four advanced economies that experienced crises during 2007-2008 are in fact part of our periphery because we were unable to obtain banking system claims data for them.

latest wave of banking crises.[23]

The consistently low Type 1 and Type 2 errors produced by the classification algorithm strengthen our preliminary evidence that network-based measures of connectedness can serve as early warning indicators. When we carefully examine the sub-rules, we notice that almost of all the indicators considered are selected by the algorithm to construct the rules, and lagged levels show up more frequently than past growth rates.[24] In the next step, we use this information to construct a relatively parsimonious empirical model that lists as covariates those indicators identified by the algorithm as the most promising. Specifically, we estimate a regression-based EWS that controls for standard macroeconomic fundamentals and adds measures of connectedness. To keep the number of covariates manageable, we group the indicators into three categories – as in Figure 3 – and retain the first two PCs from each group.

3.3 Results from Regression Analysis

We estimate a benchmark binary dependent variable model in the full sample of countries over 1978-2010. The dependent variable is an indicator taking value 1 for the onset of systemic banking crises. [25] We draw on the recent crisis prediction literature to specify a small set of key macroeconomic fundamentals that have previously been identified as leading indicators of crises. Specifically, we control for per capita income (at PPP), net foreign assets (in percentage of GDP), foreign exchange reserves (in percentage of GDP), real exchange rate (RER) misalignment, and periods of sustained capital inflows. The net foreign asset and foreign exchange positions, as well as periods of RER overvaluation are systematically associated with increased likelihood of currency and/or external

[23] Note also that the algorithm produces a larger number of sub-rules when there is less variation in the data to learn from; this is the case, for instance, in the subsample of periphery countries during 2003-2010 when few countries experienced crises; and over the full period when connectedness measures likely have less predictive content than in the more recent years. With less variation in the data, sub-rules have less power to predict crises.

[24] See online appendix for the frequency with which different indicators appear in subrules.

[25] Following Gourinchas & Obstfeld (2012), the dependent variable takes value 1 in the year of crisis onset, is set to missing for the subsequent four years (as countries often lack access to international capital markets in the years of the crisis), and is 0 in non-crisis years.

crises (Catao & Milessi-Ferretti (2013), Kaminsky et al. (1998)) while credit booms fueled by large capital inflows – captured in our model by a dummy variable for capital flows bonanzas (Reinhart & Reinhart (2009)) – play a prominent role in the run-up to banking crises (Schularick & Taylor (2012), Alessi & Detken (2011)). We also include a variable aimed at allowing for the possibility of contagion – a dummy variable for at least one neighbor experiencing a crisis. Our regressors are a mixture of variables that predict financial crises in general, since banking, currency, and debt crises often occur together (Kaminsky & Reinhart (1999)). All regressors are lagged one year.

Estimates from the benchmark model with macroeconomic fundamentals are reported in Table 3 for the probit and logit estimators. We notice that the coefficients on the macroeconomic characteristics have the expected signs across specifications: countries that are net creditors or hold higher foreign exchange reserves are less likely to have banking crises. The probability of a banking crisis goes up when countries experience sustained periods of high capital inflows or an overvalued RER and hence a loss of competitiveness. Having at least one neighbor in crisis also raises the probability of the country itself experiencing a crisis – which hints at the possibility of contagion.

While our preferred specifications do not include country or year fixed effects, for completeness we also show models with them. Country fixed effects control for omitted time-invariant characteristics that may systematically drive a country's predisposition towards crises. Year fixed effects control for global shocks that may cause crises to cluster together (e.g., the 1997-1998 East Asian crises). Including these dummies leaves the coefficient estimates largely unchanged; however, such models are less helpful for crisis prediction. One reason is that we cannot anticipate future global shocks. Another is that estimated country fixed effects are not very informative from a policy perspective. As the probit vs. logit models yield similar results, in the remainder of the analysis we only estimate probits and take as our baseline the model presented in column 4.[26]

[26] As a general check of model performance, we run Kolmogorov-Smirnov tests for the distributions of predicted probabilities across crisis/non-crisis years. For each regression, we reject the null that the two samples of predicted probabilities are drawn from the same distribution, which suggests that the models systematically yield different predicted probabilities across the two regimes.

Next, we add measures of financial connectedness to the baseline model. These are the 1st and 2nd PCs extracted from our rich set of network indicators (as discussed in section 3.1). In each specification we add the PCs from all the lagged levels of the indicators as well as the 1-year lagged growth rates (labelled "growth"). The estimates are presented in Table 4. In column (1) we add the two PCs for a country's own level of connectedness captured in degree, strength, and clustering. In the richest specification we add the PCs for neighbor connectedness as well as the HHI (column 4). The results suggests that an increase in a country's own financial connectedness is associated with a higher probability of crises one year later. By contrast, a *decline* in the country's neighbor connectedness – a possible sign that turbulence is occurring further out in the network and may be transmitted through higher-degree connections – increases the probability of a crisis.[27]

To investigate how much adding the financial connectedness measures improves the predictive power of the benchmark model, we discriminate among probit models based on the area under the Receiver Operating Characteristic (ROC). The ROC depicts the relationship between true and false positives for a range of probability thresholds. The ROC lies above the 45-degree line if the model generates crisis predictions that are superior to random guessing. According to this criterion, the best model is the one that maximizes the area under the curve (AUROC). As seen in Table 3, the baseline model has an AUROC of 70.3 percent, which rises to 75.3 percent for the model augmented with all connectedness measures (Table 4, column 5). Compared to the benchmark of a random guessing model which has an AUROC of 50 percent, adding the connectedness measures improves the predictive performance of the model by 25 percent. In addition, we test that the two areas are equal and reject the null of zero distance between the ROCs with 99 percent confidence (p-value=0.0003).

The ROCs corresponding to these two models are shown in Figure 4. A closer look reveals that the augmented-model ROC lies above the baseline ROC especially at high false positive rates, which

[27]In results not reported, we reject the null hypothesis that all indicators other than the 1st PC jointly a have statistically significant effect on the probability of a crisis.

are generated by setting very low probability thresholds to call crises. Such low probability cutoffs would be chosen by policy-makers who prefer a higher positive rate – i.e., to be alerted often even when there is no crisis – over missing crises. This means the connectedness-augmented model may be particularly useful to such conservative policymakers.

For policymaking purposes it may be also be useful to monitor a more parsimonious model than that shown in Table 4. If this were the case, then the results from Table 5 would be of interest. In this table we report benchmark and augmented probit models where the latter only include the 1st PC of our three groups of connectedness measures: degree/strength, clustering, and neighbor degree/strength. The estimates in columns 1-2 suggest that over the full period of analysis, an increase in clustering and a decrease in neighbor connectedness are associated with a higher probability of crisis. These two results hold up for the 1978-2002 period (columns 3-4), in which contagion played a prominent role in the unfolding of crises such as the 1980s debt crisis and the 1997-1998 East Asian crises. However, over the most recent period, 2003-2010, it appears that the simplest network indicators, degree and strength, are also the most relevant for crisis prediction.[28]

3.4 In- and Out-of-Sample Performance

We conclude the analysis by looking at the performance of our data mining and regression-based approaches in predicting the most recent wave of systemic banking crises in advanced economies. Fourteen high-income countries experienced a crisis starting in 2007 or 2008.

As shown in Table 6, the classification algorithm correctly predicts in-sample the onset of 5 out of the 14 systemic crises (column 1). The benchmark probit model, which exploits only lagged information about the state of the economy, generates relatively low in-sample crisis probabilities,

[28] By comparing the results across sample periods, we also notice that macroeconomic fundamentals are consistently able to predict crises: surges of capital inflows, an overvalued RER, and having at least a finacial partner in crisis increases the likelihood of a crisis. A large stock of foreign exchange reserves decreses it.

ranging from 2.6 percent for Denmark in 2008 to 10.2 percent for Greece in 2008 (column 3).[29] For all but two countries (Portugal and the US), these probabilities increase when we augment the probit model with variables on financial connectedness. The starkest increase is for the UK, where the predicted probability of a crisis rises from 5.5 to 27.4 percent between the benchmark and augmented model. This suggests that connectedness measures for most countries have substantial predictive content. We find this plausible given the relative centrality of some of these countries' banking systems in the GBN.

When both the algorithm and probit model are re-run on the subsample of core countries over the 2003-2010 period, when connectedness may arguably have paid a more important role as a conduit for financial stress, both techniques yield a better in-sample performance. Now the classification algorithm correctly identifies 11 of 14 onsets of crises (column 2); and the fitted crisis probabilities are larger for both the bechmark and augmented probit models (columns 5-6). The connectedness-augmented probit now consistently outperforms the benchmark model as it generates higher crisis probabilities for all onsets during 2007-2008.

Still focusing on the latest wave of crises, we find that the classification algorithm has a remarkable out-of-sample performance for the subsample of core countries. Table 7 reports the precision and recall associated with different crisis prediction windows. When $k = 1$, we look at the performance of the algorithm in predicting the onset of a crisis in the year ahead; when $k = 2$ the window expands to include the following year, etc. We notice that recall is consistently high across subsamples, which suggests that the algorithm misses few crisis onsets (columns 2, 4, and 6). However, precision – the algorithm's ability to not issue false alarms – is relatively high only for the core countries. This hints at the fact that financial connectedness, not surprisingly, is more useful for crisis prediction for core rather than periphery countries.[30]

[29]Predicted crisis probabilities below 20 percent are a common feature in EWS, as financial crises are relatively rare events.)

[30]Here we focused on the algorithm's ability to predict the onset of crises, but it also does well in predicting all crisis-years (see online appendix).

One last piece of evidence on the out-of-sample performance of our methods is summarized in Table 8, which lists the crises predicted by the algorithm for the 2005-2008 period (columns 1-4); as well as the probit's out-of-sample predicted probabilities for the year 2008 (columns 5-6). We find that the algorithm starts issuing crisis alerts for Ireland, The Netherlands, and Spain as early as 2005; and it issues a maximum number of crisis alerts for the year 2007. In 2007 and 2008 combined, the algorithm correctly anticipates 7 of the 14 crises in advanced economies.[31] It appears that based solely on information about countries' connectedness in the GBN, the classification algorithm would have issued crisis signals since the mid-2000s for many countries that were affected by the global financial crisis.

Looking at the probit's out of sample predicted probabilities, we note that they are relatively small: only those for the UK and US are larger than 10 percent (column 5). For this model to issue a crisis alert for the year 2008, the policymaker would have had to set a relatively low crisis probability threshold. Importantly, the out-of-sample probabilities do increase when the probit incorporates data on financial connectedness, and they turn out larger than 10 percent for 4 countries: Ireland, Spain, the UK and the US (column 6).

4 Conclusions

It is widely believed that growing financial linkages across countries have played an important role in the severity and spread of the global financial crisis. In this paper we examined the performance of financial interconnectedness as an early warning indicator for systemic banking crises. To this end, we used data mining and econometric methods to document the predictive content of a rich set of network-based connectedness measures. We constructed these measures using data on cross-border banking assets from a large sample of countries over the past four decades.

[31] The algorithm also issues false alarms during this period, e.g., predicting banking crises in South Africa in 2006, and in Canada in 2007-2008.

Our results suggest that financial connectedness can predict crises. Higher levels of a country's own connectedness – captured by simple indicators such as the number and intensity of its financial ties in the global banking network – are associated with a higher probability of crises. Lower connectivity among a country's direct financial partners hints at turbulence in the network and potential contagion, and it too is associated with a higher probability of crisis. A classification algorithm that searches for patterns in the data and associates levels and growth rates of connectedness indicators with crisis incidence, has a remarkable ability to predict the onset of crises. A probit model of crisis prediction generally yields higher crisis probabilities, both in and out of sample, when it includes information on country connectedness as opposed to macroeconomic factors alone.

We see these results as a first step towards exploiting the potentially rich informational content of network connectivity indicators for purposes of crisis prediction. While our findings suggest that there is a useful amount of crisis-signal information in the connectedness measures we define, the list of indicators we consider is by no means exhaustive. Future work could expand the scope of the analysis. Our results could be further probed on global banking networks underpinned by alternative data on cross-country linkages. Such work could result in improved early warning systems and more lead time for designing appropriate policy responses.

References

Abiad, A. (2003). Early warning systems: A survey and regime switching approach, IMF Working Paper No. 32.

Agrawal, R., Imieliski, T., & Swami, A. (1996). Mining association rules between sets of items in large databases. *ACM SIGMOD Record*, *22(2)*, ACM.

Alessi, L. & Detken, C. (2011). Quasi real time early warning indicators for costly asset price boom/bust cycles: A role for global liquidity. *European Journal of Political Economy*, *27*, 520–533.

Allen, F., Babus, A., & Carletti, E. (2012). Asset commonality, debt maturity and systemic risk. *Journal of Financial Economics*, *104(3)*, 519–534.

Baldacci, E., Petrova, I., Belhocine, N., Dobrescu, G., & Mazraani, S. (2011). Assessing fiscal stress, IMF Working Paper No. 100.

Berg, A., Borensztein, E., & Pattillo, C. (2005). Assessing early warning systems: How have they worked in practice? *IMF Staff Papers*, *52(3)*, 462–502.

Berg, A. & Pattillo, C. (1999). Are currency crises predictable? a test. *IMF Staff Papers*, *46(2)*, 107–138.

Billio, M., Getmansky, M., Lo, A. W., & Pelizzon, L. (2012). Ecnometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics*, *104*, 535–559.

BIS (2009). Guide to the international financial statistics, Monetary and Economic Department BIS Paper No. 49.

BIS (2011). Guidelines to the international locational financial statistics, Monetary and Economic Department.

Blank, S. & Buch, C. (2007). The euro and cross-border banking: Evidence from bilateral data. *Comparative Economic Studies*, *49(3)*, 389–410.

Blank, S. & Buch, C. (2010). International bank portfolios: short and long-run responses to macroeconomic conditions. *Review of International Economics*, *18(2)*, 289–306.

Caballero, J. (2012). Banking crises and financial integration, IDB Working Paper No. 364.

Caballero, J., Candelaria, C., & Hale, G. (2009). Bank relationships and the depth of the current economic crisis, FRBSF Economic Letter 2009–38, December 14.

Catao, L. & Milessi-Ferretti, G. M. (2013). External liabilities and crisis risk, IMF Working Paper No. 13/113.

Cerutti, E., Claessens, S., & McGuire, P. (2012). *Systemic risk in global banking: What available data can tell us and what more data are needed?*, chapter NBER Systemic Risk and Macro Modeling. Cambridge, MA: University of Chicago Press.

Chinazzi, M., Fagiolo, G., Reyes, J. A., & Schiavo, S. (2013). Post-mortem examination of the international financial network. *Journal of Economic Dynamics and Control*, *37*, 1692–1713.

Chui, M. (2002). Leading indicators of balance-of-payments crises: A partial review, Bank of England Working Paper No. 171.

Cont, R., Moussa, A., & e Santos, E. B. (2012). *Network structure and systemic risk in banking systems*, chapter Handbook of Systemic Risk. New York: Cambridge University Press.

Creighton, C. & Hanash, S. (2003). Mining gene expression databases for association rules. *Bioinformatics*, *19*(1), 79–86.

Dabla-Norris, E. & Gunduz, Y. B. (2012). Exogenous shocks and growth crises in low-income countries: A vulnerabilty index, IMF Working Paper No. 264.

Demirguc-Kunt, A. & Detragiache, E. (1998). The determinants of banking crises in developing and developed countries. *IMF Staff Papers*, *45*.

Demirguc-Kunt, A. & Detragiache, E. (2005). Cross-country empirical studies of systemic banking distress: A survey. *National Institute Economic Review*, *192*.

Edison, H. J. (2000). Do indicators of financial crises work? an evaluation of an early warning system, Board of Governors of the Federal Reserve System International Finance Discussion Paper No. 675.

Eichengreen, B., Rose, A. K., & Wyplosz, C. (1995). Exchange market mayhem: The antecedents and aftermath of speculative attacks. *Economic Policy*, *10(2)*, 251312.

Fagiolo, G. (2007). Clustering in complex directed networks. *Physical Review E*, *76*(2), 026107.

Feng, L., Dillon, T., & Liu, J. (2001). Inter-transactional association rules for multi-dimensional contexts for prediction and their application to studying meteorological data. *Data and Knowledge Engineering*, *37*(1), 85–115.

Fricke, D. & Lux, T. (2012). Identifying a core-periphery structure in the italian interbank market, Department of Economics, University of Kiel, mimeo.

Gai, P., Haldane, A., & Kapadia, S. (2010). Complexity, concentration and contagion. *Journal of Monetary Economics*, *58*, 453–470.

Gourinchas, P.-O. & Obstfeld, M. (2012). Stories of the twentieth century for the twenty-first. *American Economic Journal: Macroeconomics*, *4(1)*, 226–265.

Hale, G. (2012). Bank relationships, business cycles, and financial crises. *Journal of International Economics*, *88(2)*, 312–325.

Hale, G., Kapan, T., & Minoiu, C. (2013). Crisis transmission in the global banking network, International Monetary Fund, mimeo.

Hattori, M. & Suda, Y. (2007). *Developments in a cross-border bank exposure network*, volume 29, chapter Research on Global Financial Stability: The Use of BIS International Financial Statistics, (pp. 16–31). Committee on the Global Financial System Publications.

IMF (2010). The imf-fsb early warning exercise—design and methodological toolkit.

Kalemli-Ozcan, S., Papaioannou, E., & Peydro, J.-L. (2013). Financial regulation, financial globalization and the synchronization of economic activity. *Journal of Finance*, *xx*, xx–xx.

Kaminsky, G., Lizondo, S., & Reinhart, C. M. (1998). Leading indicators of currency crises. *IMF Staff Papers*, *45(1)*, 1–48.

Kaminsky, G. & Reinhart, C. (1999). The twin crises: The causes of banking and balance-of-payments crises. *American Economic Review*, *xxx*, 473–500.

Karabatak, M. & Ince, M. C. (2009). An expert system for detection of breast cancer based on association rules and neural network. *Expert Systems with Applications*, *36*(2), 3465–3469.

Laeven, L. & Valencia, F. (2012). Systemic banking crises: An update, IMF Working Paper No. 12/163.

Lopez-Fernandez, L., Robles, G., & Barahona, J. (2004). Applying social network analysis to the information in cvs repositories. *26th International Conference on Software Engineering - W17S Workshop.*

Lu, H., Han, J., & Feng, L. (1998). Stock movement prediction and n-dimensional inter-transaction association rules. In *Proceedings of the ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery*, (pp.$\tilde{1}$2).

Minoiu, C. & Reyes, J. A. (2013). A network analysis of global banking: 1978-2010. *Journal of Financial Stability*, *9*, 168–184.

Reinhart, C. M. & Reinhart, V. R. (2009). Capital flow bonanzas: An encompassing view of the past and present, NBER Working Paper No. 14321.

Schularick, M. & Taylor, A. M. (2012). Credit booms gone bust: Monetary policy, leverage cycles, and financial crises, 1870-2008. *American Economic Review*, *102(2)*, 1029–1061.

Schweitzer, F., Fagiolo, G., Sornette, D., Vega-Redondo, F., Vespignani, A., & White, D. R. (2009). Economicnetworks:thenewchallenges. *Science*, *325*, 422.

Stiglitz, J. (2010). Risk and global economic architecture: Why full financial integration may be undesirable, NBER Working Paper No. 15718.

Subrahmanian, V. & Ernst, J. (2009). Method and system for optimal data diagnosis, US Patent Nr. 7474987 dated January 6, 2009.

Subrahmanian, V., Mannes, A., Sliva, A., Shakarian, J., & Dickerson, J. (2013). Computational analysis of terrorist groups: Lashkar-e-taiba. *Information Systems and Applications*, *XIV*, Springer.

Tabak, B. M., Takami, M., Rocha, J., & Cajueiro, D. (2011). Directed clustering coefficient as a measure of systemic risk in complex banking networks, Banco Central de Brasil Research Department Working Paper No. 249.

Tabak, B. M., Takami, M., Rocha, J., Cajueiro, D., & Souza, S. (2012). Interpreting directed clustering coeffcients in complex banking networks, Banco Central de Brasil Research Department, mimeo.

Yellen, J. L. (2013). Interconnectedness and systemic risk: Lessons from the financial crisis. Remarks at the AEA/AFA Joint Luncheon, San Diego, CA, January 4.

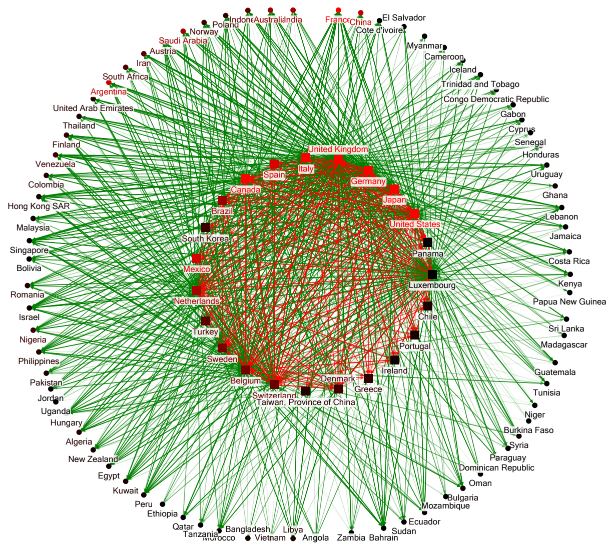

(a) Full network, 1980

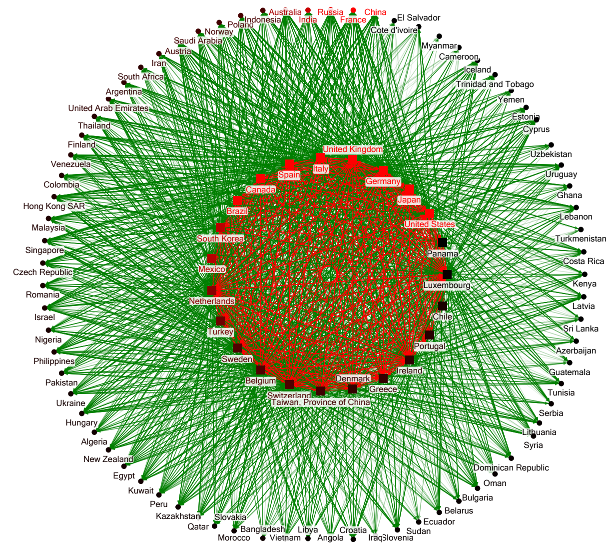

(b) Full network, 2007

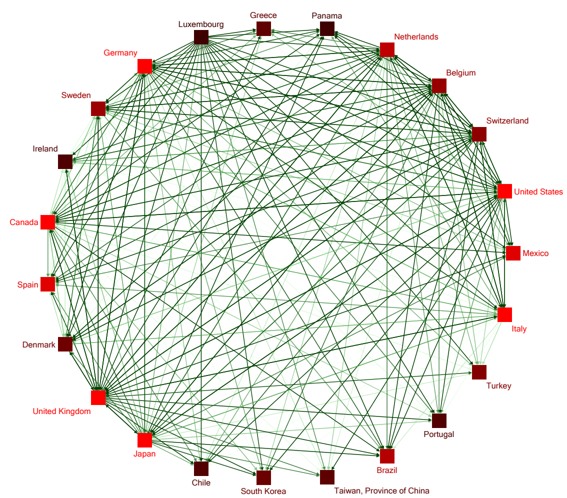

(c) Core network, 1980

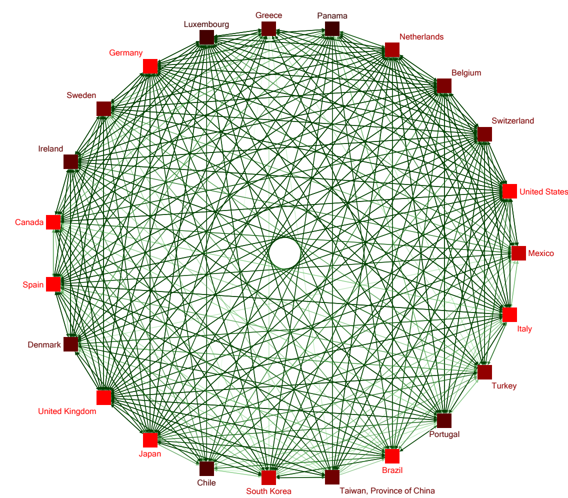

(d) Core network, 2007

Figure 1: *Global banking network in 1980 vs. 2007*: Visualization of the baseline GBN for the largest 100 countries by GDP. Node color is proportional to GDP (black is lower, red is higher). Edge weight is proportional with the intensity of the edge color (from light to dark green). In panels (a)-(b), the nodes on the inner ring are core countries and the nodes on the outer ring are periphery countries; red edges connect core countries and green edges connect core with periphery countries.

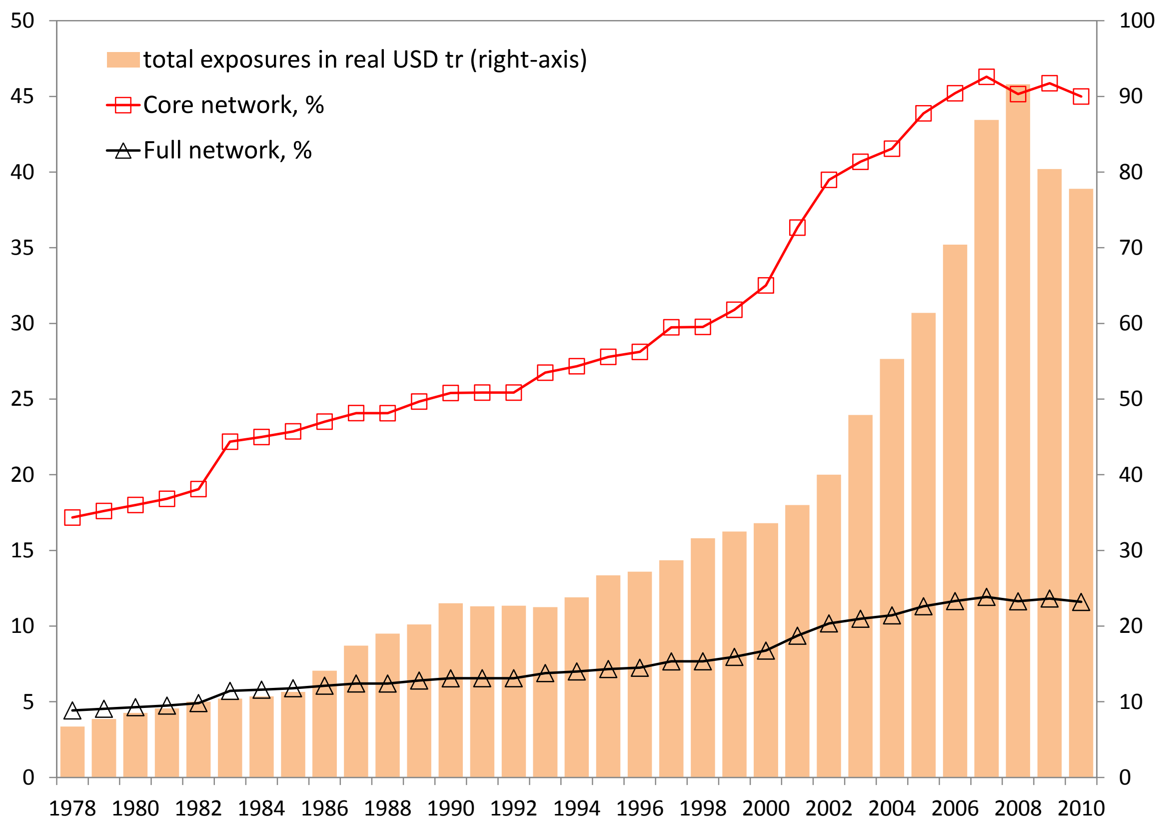

Figure 2: *Network density and total exposures*: Network density is defined as the number of edges in the GBN divided by the number of possible edges. Total exposures, representing cross-border banking claims, are expressed in constant (2005) USD trillion.

	1978	1985	1995	2007	2010	% increase 1978-2007
Degree and strength						
Degree-in/out	9.3	12.2	15.4	23.7	23.5	154%
Strength-in/out	3.3	4.3	5.4	8.3	8.2	152%
Binary clustering coefficients						
BCC Lopez-Fernandez 2004	0.56	0.66	0.85	0.89	0.89	59%
BCC Fagiolo 2007 (“cycle”)	0.49	0.59	0.78	0.89	0.89	84%
BCC Fagiolo 2007 (“middleman”)	0.49	0.59	0.78	0.90	0.89	82%
BCC Fagiolo 2007 (“in”)	0.53	0.62	0.78	0.88	0.88	66%
BCC Fagiolo 2007 (“out”)	0.51	0.61	0.80	0.89	0.89	75%
Weighted clustering coefficients						
WCC Lopez-Fernandez 2004	0.23	0.27	0.36	0.37	0.36	64%
WCC Fagiolo 2007 (“cycle”)	0.18	0.21	0.28	0.32	0.32	85%
WCC Fagiolo 2007 (“middleman”)	0.18	0.22	0.28	0.33	0.32	83%
WCC Fagiolo 2007 (“in”)	0.19	0.23	0.28	0.32	0.31	64%
WCC Fagiolo 2007 (“out”)	0.18	0.22	0.29	0.33	0.32	78%
Average nearest neighbor degree and strength						
ANND (out-in)	54.5	66.5	92.9	117.0	110.6	115%
ANND (out-out)	52.2	66.7	87.3	114.6	110.8	119%
ANND (in-in)	50.9	63.6	106.7	119.1	113.5	134%
ANND (in-out)	49.9	65.5	102.8	124.1	122.4	149%
ANNS (out-in)	19.9	24.6	33.5	42.0	39.5	111%
ANNS (out-out)	19.1	24.7	31.5	41.3	39.8	117%
ANNS (in-in)	18.2	22.8	37.9	43.1	40.4	136%
ANNS (in-out)	17.8	23.6	36.4	45.0	43.8	152%
Herfindahl-Hirschmann index						
HHI	0.11	0.13	0.20	0.12	0.13	12%

Table 1: *Average network indicators over time*: The table reports average values (across nodes) for selected network indicators.

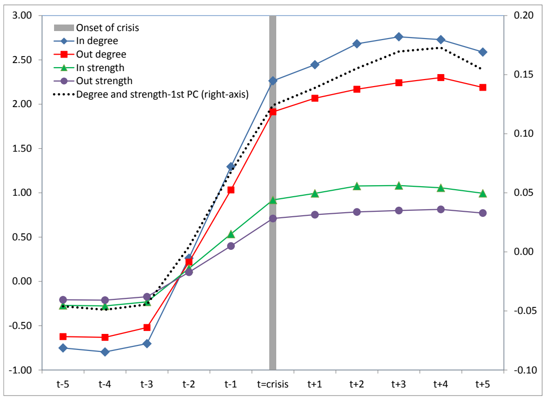

(a) Degree and strength

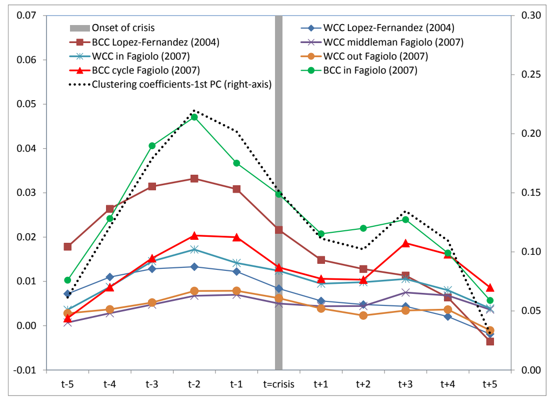

(b) Clustering

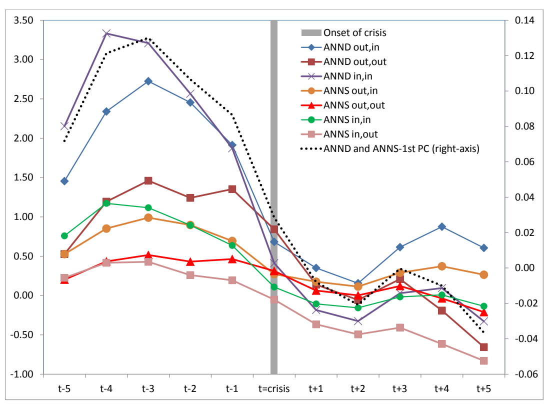

(c) Neighbor connectedness

Figure 3: *Financial connectedness around systemic banking crises*: The chart depicts average levels of selected network indicators during 5 years before and after systemic banking crises (t=`crisis`). The indicators are conditional on country and year fixed effects. The dotted line represents the first principal component (labeled "1st PC") extracted from each of the three groups of indicators (see Section 3.1 and Appendix B for details).

	(1)	(2)	(3)	(4)	(5)	(6)
Sample, period	# crisis-years	# predicted crisis-years	Support	Precision	Recall	# sub-rules
Full, 1978-2010	409	49	46	0.94	0.11	34
Full, 1978-2002	329	342	309	0.90	0.94	481
Full, 2003-2010	80	35	32	0.91	0.40	18
Core, 1978-2010	87	34	30	0.88	0.34	33
Core, 1978-2002	45	59	43	0.73	0.96	76
Core, 2003-2010	42	41	35	0.85	0.83	24
Periphery, 1978-2010	322	310	283	0.91	0.88	335
Periphery, 1978-2002	284	9	9	1.00	0.03	1
Periphery, 2003-2010	38	53	38	0.72	1.00	420

Table 2: *Classification algorithm in-sample performance*: The table summarizes the in-sample performance of the classification algorithm. Columns 3-5 report the support (correctly predicted crises), precision (correctly predicted crises/total predicted crises), and recall (correctly predicted crises/total actual crises). Column 6 reports the number of association sub-rules identified for each case.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Probit	Probit	Probit	Probit	Logit	Logit	Logit	Logit
Log-per capita GDP	0.009	0.290	-0.006	-0.019	0.027	0.670	0.006	-0.037
	(0.026)	(0.285)	(0.030)	(0.028)	(0.059)	(0.663)	(0.064)	(0.061)
Net foreign assets/GDP	-0.092*	-0.240**	-0.085	-0.092	-0.152	-0.449*	-0.133	-0.148
	(0.055)	(0.109)	(0.054)	(0.056)	(0.118)	(0.244)	(0.115)	(0.117)
Capital inflows bonanza	0.379***	0.356***	0.353***	0.390***	0.845***	0.744***	0.730***	0.876***
	(0.095)	(0.116)	(0.103)	(0.096)	(0.209)	(0.236)	(0.224)	(0.211)
Forex reserves/GDP	-0.018***	-0.030***	-0.019***	-0.019***	-0.044***	-0.067***	-0.046***	-0.045***
	(0.005)	(0.011)	(0.005)	(0.005)	(0.012)	(0.023)	(0.014)	(0.012)
RER misalignment	1.100**	1.185*	0.994**	1.122**	2.010*	2.391	1.597	2.082**
	(0.505)	(0.653)	(0.505)	(0.504)	(1.028)	(2.061)	(0.986)	(1.015)
At least 1 neighbor in crisis				0.209**				0.489***
				(0.082)				(0.186)
Country fixed effects		yes				yes		
Year fixed effects			yes				yes	
Observations	3,949	2,319	3,139	3,949	3,949	2,319	3,139	3,949
AUROC	0.696	0.729	0.769	0.703	0.692	0.720	0.763	0.699

Table 3: *Benchmark probit/logit model*: The table reports the estimates of a benchmark crisis prediction model estimated on the full sample of countries over 1978-2010. The dependent variable takes value 1 for the onset of systemic banking crises. A constant is included, but the coefficient is not shown. Standard errors are clustered on country. *** indicates statistical significance at 1 percent, ** at 5 percent, and * at 10 percent.

	(1)	(2)	(3)	(4)
Degree and strength-1st PC	0.203***	0.142**	0.139**	0.121*
	(0.054)	(0.055)	(0.055)	(0.067)
Degree and strength-2nd PC	0.055	-0.288**	-0.319**	-0.371**
	(0.108)	(0.130)	(0.136)	(0.163)
Degree and strength, growth-1st PC	-0.001	-0.004	-0.031	0.014
	(0.038)	(0.036)	(0.044)	(0.054)
Degree and strength, growth-2nd PC	-0.001	-0.005	0.009	0.064
	(0.041)	(0.040)	(0.042)	(0.070)
Clustering-1st PC	0.226***	0.412***	0.437***	0.426***
	(0.072)	(0.081)	(0.083)	(0.086)
Clustering-2nd PC	0.796	0.897*	0.863*	0.965*
	(0.488)	(0.483)	(0.487)	(0.501)
Clustering, growth-1st PC	-0.004	0.001	0.017	0.016
	(0.009)	(0.008)	(0.017)	(0.019)
Clustering, growth-2nd PC	-0.029	-0.026	-0.025	-0.014
	(0.052)	(0.050)	(0.055)	(0.062)
Neighbor connectedness-1st PC		-0.287***	-0.317***	-0.326***
		(0.063)	(0.068)	(0.070)
Neighbor connectedness-2nd PC		0.246	0.245	0.255
		(0.218)	(0.243)	(0.251)
Neighbor connectedness, growth -1st PC			-0.059	-0.062
			(0.041)	(0.040)
Neighbor connectedness, growth -2nd PC			-0.002	-0.016
			(0.044)	(0.051)
Herfindahlindex				-0.639
				(0.702)
Herfindahlindex, growth				0.004
				(0.004)
Observations	3,368	3,368	3,368	3,368
AUROC	0.725	0.748	0.750	0.753

Table 4: *Augmented probit results (full version)*: The table reports the estimates of a standard crisis prediction model augmented with network-based measures of connectedness and estimated on the full sample of countries over 1978-2010. The dependent variable takes value 1 for the onset of a systemic banking crisis. All macroeconomic variables from Table 3 (column 4) and a constant are included, but the coefficients are as expected and not shown. Standard errors are clustered on country. *** indicates statistical significance at 1 percent, ** at 5 percent, and * at 10 percent.

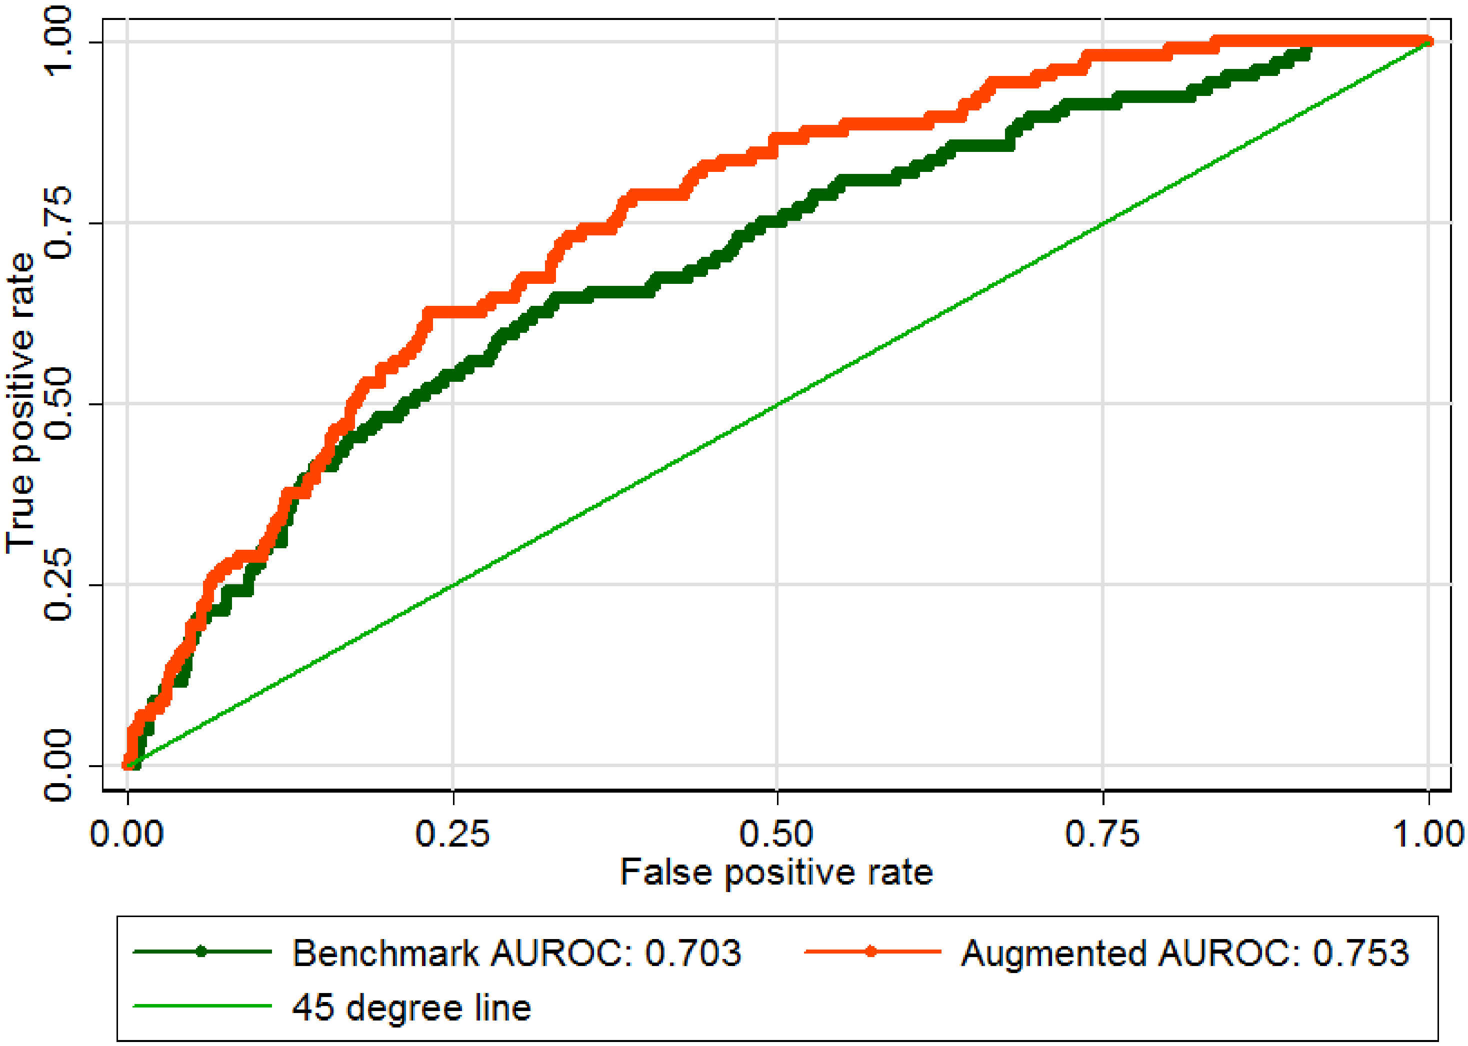

Figure 4: *ROCs for benchmark vs. augmented model (full version)*: The chart shows the ROCs corresponding to the benchmark model (Table 3, column 4) vs. the model fully augmented with network indicators (Table 4, column 4).

	(1)	(2)	(3)	(4)	(5)	(6)
	1978-2010		1978-2002		2003-2010	
Log-per capita GDP	-0.019	-0.069**	-0.103***	-0.107***	0.364***	0.210*
	(0.028)	(0.034)	(0.036)	(0.041)	(0.105)	(0.116)
Net foreign assets/GDP	-0.092	-0.095*	-0.091	-0.080	-0.079	-0.151
	(0.056)	(0.058)	(0.060)	(0.058)	(0.133)	(0.157)
Capital inflows bonanza	0.390***	0.382***	0.300***	0.290***	0.798***	0.844***
	(0.096)	(0.097)	(0.110)	(0.112)	(0.220)	(0.227)
Forex reserves/GDP	-0.019***	-0.014***	-0.008*	-0.008*	-0.027**	-0.017
	(0.005)	(0.005)	(0.004)	(0.004)	(0.012)	(0.012)
RER misalignment	1.122**	1.087**	0.971**	0.927**	4.500***	4.745***
	(0.504)	(0.490)	(0.470)	(0.462)	(1.509)	(1.722)
At least 1 neighbor in crisis	0.209**	0.202**	0.124	0.155	0.910***	0.879**
	(0.082)	(0.091)	(0.101)	(0.106)	(0.313)	(0.352)
Degree and strength-1st PC		0.028		-0.022		0.233***
		(0.018)		(0.028)		(0.073)
Clustering-1st PC		0.128**		0.121**		0.357
		(0.055)		(0.055)		(0.252)
Neighbor connectedness-1st PC		-0.134**		-0.110*		-0.103
		(0.053)		(0.056)		(0.216)
Observations	3,949	3,949	2,799	2,799	1,150	1,150
AUROC	0.703	0.722	0.684	0.696	0.921	0.931

Table 5: *Augmented probit results (parsimonious version)*: The table reports estimates of a parsimonious version of the crisis prediction models presented in Table 4 in which we retain only the 1st PC of each group of connectedness measures. The dependent variable takes value 1 for the onset of a systemic banking crisis. The model is estimated on the full sample of countries over 1978-2010 (columns 1-2), 1978-2002 (columns 3-4) and 2003-2010 (columns 5-6). For completeness we also report the specifications with macroeconomic variables only (columns 1, 3, 5). A constant is included, but the coefficient is not shown. Standard errors are clustered on country. *** indicates statistical significance at 1 percent, ** at 5 percent, and * at 10 percent.

	(1)	(2)	(3)	(4)	(5)	(6)
	Algorithm in-sample predictions		Probit in-sample predicted crisis probabilities			
	Full, 1978-2010	Core, 2003-2010	Full, 1978-2010		Core, 2003-2010	
			Benchmark	Augmented	Benchmark	Augmented
Belgium	2008	2008	3.5%	12.0%	3.0%	6.1%
Denmark	2008	2008	2.6%	4.7%	1.4%	2.3%
Germany		2008	3.7%	15.9%	3.6%	7.2%
Greece		2008	10.2%	13.9%	23.1%	25.4%
Ireland			9.3%	11.3%	19.9%	36.7%
Italy		2008	4.1%	6.1%	6.0%	12.7%
Luxembourg	2008	2008	.	.	.	.
Netherlands		2008	.	.	.	.
Portugal		2008	5.0%	3.1%	10.8%	19.2%
Spain		2008	10.1%	17.9%	23.2%	47.9%
Sweden			3.5%	4.0%	4.4%	6.7%
Switzerland	2008	2008	.	.	.	.
United Kingdom	2007	2007	5.5%	27.4%	19.6%	40.3%
United States			5.7%	2.5%	24.4%	28.6%

Table 6: *Classification algorithm and probit model in-sample performance*: The table summarizes the in-sample performance of the classification algorithm and probit model, focusing on the 2007-2008 crises. Column headings indicate the sample and period over which the algorithm and probit are run. Columns 1-2 report the year for which the algorithm predicts a crisis. Columns 3-6 report predicted crisis probabilities from the benchmark and augmented probit. In columns 3-6, missing values refer to countries for which data on at least one macroeconomic variable is missing.

	(1)	(2)	(3)	(4)	(5)	(6)
	Core		Full sample		Periphery	
prediction k years ahead	Precision	Recall	Precision	Recall	Precision	Recall
k=0	0.12	0.50	0.02	0.15	0.01	0.17
k=1	0.26	0.64	0.03	0.23	0.02	0.25
k=2	0.37	0.71	0.06	0.42	0.03	0.33
k=3	0.44	0.71	0.09	0.46	0.05	0.42
k=4	0.49	0.71	0.10	0.50	0.07	0.58
k=5	0.53	0.71	0.13	0.58	0.08	0.58

Table 7: *Classification algorithm out-of-sample performance*: The table summarizes the out-of-sample performance of the classification algorithm in predicting the onset of crises for the full sample, core, and periphery countries. The algorithm is run on the 2003-2010 period. The measures of performance are precision (correctly predicted crises/total predicted crises) and recall (correctly predicted crises/total actual crises).

	(1)	(2)	(3)	(4)	(5)	(6)
	Algorithm				Probit	
	2008	2007	2006	2005	Benchmark	Augmented
Belgium	xx	x	x		0.7%	0.9%
Denmark	xx	x			0.3%	0.5%
Germany		x			0.8%	1.0%
Greece					4.9%	3.7%
Ireland	xx	x	x	x	5.8%	10.8%
Italy					0.9%	1.8%
Luxembourg					.	.
Netherlands		x	x	x	.	.
Portugal			x		1.3%	2.8%
Spain	xx	x	x	x	5.6%	12.5%
Sweden	xx	x	x		0.9%	1.4%
Switzerland	xx	x			.	.
United Kingdom					10.9%	15.1%
United States		xx			12.9%	15.0%

Table 8: *Classification algorithm and probit model out-of-sample performance*: The table summarizes the out-of-sample performance of the classification algorithm and probit model, focusing on the 2007-2008 crises. Columns 1-4 report “x” if the classification algorithm predicts a crisis in the year indicated as column heading. “xx” labels crisis predictions that are accurate. The algorithm is run on the subsample of core countries on a rolling basis (over 1978-2004 for 2005 prediction, 1978-2005 for 2006 prediction, etc.). Columns 5-6 report crisis probabilities predicted for the year 2008 by the augmented probit model (parsimonious version) estimated on the subsample of core countries over 1978-2007.

Appendix A Network Indicator Definitions

This appendix provides formal definitions for the network indicators employed in the analysis. Consider a weighted directed network $G = (V, E, w)$ where:

- V is a finite set of nodes (countries)
- $E \subseteq V \times V$ is a finite set of directed edges
- $w : E \rightarrow [0, 1]$ assigns a weight to each edge.

Edge weights represent log-transformed real cross-border banking exposures (scaled by the log-product GDP of the two nodes in the baseline network; and unscaled in the alternative network). In what follows, adjacent nodes of a node v are given by $N_v = N_v^{in} \cup N_v^{out}$ where $N_v^{in} = \{v' | w(v', v) > 0\}$ and $N_v^{out} = \{v' | w(v, v') > 0\}$ and $b_{v,v'} = 1$ if $w(v, v') > 0$ and 0 otherwise. We compute the following network indicators:

1. In-degree and Out-degree. The in-degree of a node v is denoted d_v^{in} and represents the total number of a node's creditors. The out-degree of a node v is denoted d_v^{out} and represents the total number of a node's debtors.

$$d_v^{in} = \sum_{v' \in V} b_{v',v} \tag{1}$$

$$d_v^{out} = \sum_{v' \in V} b_{v,v'} \tag{2}$$

2. In-strength and Out-strength. The in-strength of a node v is denoted $str^{in}(v)$ and refers to the total weight of in-coming edges (a node's liabilities). The out-strength of a node v is denoted $str^{out}(v)$ and represents the total weight of out-going edges (a node's assets or exposures).

$$str^{in}(v) = \sum_{(v',v) \in E} w(v', v) \tag{3}$$

$$str^{out}(v) = \sum_{(v,v') \in E} w(v, v') \tag{4}$$

3. A^{in} and A^{out}. These are measures of in-strength and out-strength that are normalized by the total strength of each nodes' neighbors.

$$A^{in}(v) = \frac{str^{in}(v)}{\sum_{(v',v) \in E} str^{in}(v')} \tag{5}$$

$$A^{out}(v) = \frac{str^{out}(v)}{\sum_{(v,v')\in E} str^{in}(v')} \tag{6}$$

4. AN^{in} and AN^{out}. These indicators denote the average A^{in} and A^{out} values normalized across all nodes in the GBN.

$$AN^{in}(v) = \frac{\sum_{(v',v)\in E} A^{in}(v')}{\sum_{v''\in V} A^{in}(v'')} \tag{7}$$

$$AN^{out}(v) = \frac{\sum_{(v,v')\in E} A^{out}(v')}{\sum_{v''\in V} A^{out}(v'')} \tag{8}$$

5. Average nearest node degree (ANND). The ANND denotes the average in-degree (or out-degree) of neighbor nodes connected toward (or from) a node v.

$$ANND^{in,in}(v) = \frac{\sum_{(v',v)\in E} d_{v'}^{in}}{d_v^{in}} \tag{9}$$

$$ANND^{out,in}(v) = \frac{\sum_{(v',v)\in E} d_{v'}^{out}}{d_v^{in}} \tag{10}$$

$$ANND^{in,out}(v) = \frac{\sum_{(v,v')\in E} d_{v'}^{in}}{d_v^{out}} \tag{11}$$

$$ANND^{out,out}(v) = \frac{\sum_{(v,v')\in E} d_{v'}^{out}}{d_v^{out}} \tag{12}$$

6. Average nearest node strength (ANNS). The ANNS denotes the average in-strength (or out-strength) of neighbor nodes connected to (or from) a node v.

$$ANNS^{in,in}(v) = \frac{\sum_{(v',v)\in E} str^{in}(v')}{d_v^{in}} \tag{13}$$

$$ANNS^{out,in}(v) = \frac{\sum_{(v',v)\in E} str^{out}(v')}{d_v^{in}} \tag{14}$$

$$ANNS^{in,out}(v) = \frac{\sum_{(v,v')\in E} str^{in}(v')}{d_v^{out}} \tag{15}$$

$$ANNS^{out,out}(v) = \frac{\sum_{(v,v')\in E} str^{out}(v')}{d_v^{out}} \tag{16}$$

7. Binary and weighted local clustering coefficients. We use two types of clustering coefficients for various types of triangle-type relationships respectively introduced in Lopez-Fernandez et al. (2004) and Fagiolo (2007).

$$BCC^1(v) = \frac{\sum_{v',v'' \in N_v} b_{v',v''}}{|N_v| \times (|N_v| - 1)} \tag{17}$$

$$WCC^1(v) = \frac{\sum_{v',v'' \in N_v} w(v', v'')}{|N_v| \times (|N_v| - 1)} \tag{18}$$

as defined in Lopez-Fernandez et al. (2004), and:

$$BCC^2_{Cycle}(v) = \frac{\sum_{v',v'' \in N_v} \sqrt[3]{b_{v,v'} b_{v',v''} b_{v'',v}}}{d_v^{in} \times d_v^{out} - d_v^{\leftrightarrow}} \tag{19}$$

$$BCC^2_{Middleman}(v) = \frac{\sum_{v',v'' \in N_v} \sqrt[3]{b_{v,v'} b_{v'',v'} b_{v'',v}}}{d_v^{in} \times d_v^{out} - d_v^{\leftrightarrow}} \tag{20}$$

$$BCC^2_{In}(v) = \frac{\sum_{v',v'' \in N_v} \sqrt[3]{b_{v',v} b_{v',v''} b_{v'',v}}}{d_v^{in} \times (d_v^{in} - 1)} \tag{21}$$

$$BCC^2_{Out}(v) = \frac{\sum_{v',v'' \in N_v} \sqrt[3]{b_{v,v'} b_{v',v''} b_{v,v''}}}{d_v^{out} \times (d_v^{out} - 1)} \tag{22}$$

$$WCC^2_{Cycle}(v) = \frac{\sum_{v',v'' \in N_v} \sqrt[3]{w(v,v') w(v',v'') w(v'',v)}}{d_v^{in} \times d_v^{out} - d_v^{\leftrightarrow}} \tag{23}$$

$$WCC^2_{Middleman}(v) = \frac{\sum_{v',v'' \in N_v} \sqrt[3]{w(v,v') w(v'',v') w(v'',v)}}{d_v^{in} \times d_v^{out} - d_v^{\leftrightarrow}} \tag{24}$$

$$WCC^2_{In}(v) = \frac{\sum_{v',v'' \in N_v} \sqrt[3]{w(v',v) w(v',v'') w(v'',v)}}{d_v^{in} \times (d_v^{in} - 1)} \tag{25}$$

$$WCC^2_{Out}(v) = \frac{\sum_{v',v'' \in N_v} \sqrt[3]{w(v,v') w(v',v'') w(v,v'')}}{d_v^{out} \times (d_v^{out} - 1)} \tag{26}$$

where $d_v^{\leftrightarrow} = |N_v^{in} \cap N_v^{out}|$ as defined in Fagiolo (2007). See online appendix for a graphical representation of the triangle taxonomies.

8. Herfindahl-Hirschmann index (HHI). This is the standard measure of competition or market power concentration from the international organization literature and is not strictly a network indicator. The HHI helps gauge the degree of diversification (or concentration) of a

node's cross-border banking portfolios.

$$HHI(v) = \sum_{v' \in N_v^{in}} \left(\frac{\sum_{(v',v)\in E} w(v',v)}{\sum_{(v'',v)\in E} w(v'',v)}\right)^2 \tag{27}$$

Appendix B Principal Component Analysis (PCA)

We group 26 network indicators into three groups – degree and strength, clustering, and neighbor connectedness – and extract the first and second principal components (1st PC and 2nd PC) through PCA. All indicators are computed on the baseline GBN. Each group include the following indicators:

Group of indicators	List of indicators
1) Degree and strength	
level 1st PC: 90.8 percent	d^{in}, d^{out}
level 2nd PC: 3.7 percent	str^{in}, str^{out}
growth 1st PC: 43.5 percent	A^{in}, A^{out}
growth 2nd PC: 32.0 percent	AN^{in}, AN^{out}
2) Clustering	
level 1st PC: 88.3 percent	BCC^1, Lopez-Fernandez et al. (2004)
level 2nd PC: 7.0 percent	WCC^1, Lopez-Fernandez et al. (2004))
growth 1st PC: 69.0 percent	BCC^2_{Cycle}, Fagiolo (2007)
growth 2nd PC: 14.1 percent	$BCC^2_{Middleman}$, Fagiolo (2007)
	BCC^2_{In}, Fagiolo (2007)
	BCC^2_{Out}, Fagiolo (2007)
	WCC^2_{Cycle}, Fagiolo (2007)
	$WCC^2_{Middleman}$, Fagiolo (2007)
	WCC^2_{In}, Fagiolo (2007)
	WCC^2_{Out}, Fagiolo (2007)
3) Neighbor connectedness	
level 1st PC: 89.7 percent	$ANND^{in,in}$
level 2nd PC: 9.7 percent	$ANND^{out,in}$
growth 1st PC: 58.1 percent	$ANND^{in,out}$
growth 2nd PC: 36.8 percent	$ANND^{out,out}$
	$ANNS^{in,in}$
	$ANNS^{out,in}$
	$ANNS^{in,out}$
	$ANNS^{out,out}$