Mapping the State of Financial Stability[1]

This version: 5.1.2012

Peter Sarlin[2]
Åbo Akademi University
Turku Centre for Computer Science

Tuomas A. Peltonen[3]
European Central Bank

Abstract

The paper lays out a methodology based upon data and dimensionality reduction for mapping the state of financial stability, and visualizing the sources of systemic risks. The Self-Organizing Financial Stability Map (SOFSM) can be used to monitor macro-financial vulnerabilities by locating a country in the financial stability cycle. Besides of its visualization capabilities, the SOFSM can be used as an Early Warning System that can be calibrated according to policymakers' preferences between Type I and II errors. The SOFSM performs on par with a statistical benchmark model and correctly calls the crises that started in 2007 in the euro area and the United States.

JEL Codes: E44, E58, F01, F37, G01.
Keywords: Self-Organizing Financial Stability Map (SOFSM); systemic financial crisis; systemic risk; visualization; prediction; macroprudential supervision

---

[1] The authors want to thank the anonymous referee, Barbro Back, Tomas Eklund, Kristian Koerselman, Marco Lo Duca, Fredrik Lucander, and seminar participants at the Bank of Finland seminar on 11 November 2011 in Helsinki, the 14th Annual DNB Research Conference 'Complex systems: Towards a better understanding of financial stability and crises' on 3–4 November 2011 in Amsterdam, the ESCB Macro-Prudential Research (MaRs) Network workshop on 14–15 April 2011 in Frankfurt am Main, the International Joint Conference on Artificial Intelligence (IJCAI'11) workshop on Chance Discovery on 16–22 July 2011 in Barcelona, the ECB Financial Stability seminar on 16 September 2011 in Frankfurt am Main, the First Conference of the ESCB MaRs Network on 5–6 October 2011 in Frankfurt am Main, the Bank of Finland Institute for Economies in Transition (BOFIT) and the Data Mining and Knowledge Management Laboratory at Åbo Akademi University for useful comments and discussions. All remaining errors are of our own. The views presented in the paper are those of the authors and do not necessarily represent the views of the European Central Bank or the Eurosystem.
[2] Corresponding author: Peter Sarlin, Department of Information Technologies, Åbo Akademi University, Turku Centre for Computer Science, Joukahaisenkatu 3–5, 20520 Turku, Finland. Email: psarlin@abo.fi tel. +358 2 215 4670.
[3] Tuomas A. Peltonen, Directorate General Financial Stability, Financial Stability Surveillance Division, European Central Bank, Kaiserstrasse 29, 60311 Frankfurt am Main, Germany. Email: tuomas.peltonen@ecb.europa.eu.

1. Introduction

The recent global financial crisis has demonstrated the importance of understanding sources of domestic and global vulnerabilities that may lead to a systemic financial crisis.[4] Early identification of financial stress would allow policymakers to introduce policy actions to decrease or prevent further build up of vulnerabilities or otherwise enhance the shock absorption capacity of the financial system. Finding the individual sources of vulnerability and risk is of central importance since that allows targeted actions for repairing specific cracks in the financial system.

Much of the empirical literature deals with Early Warning Systems (EWSs) that rely on conventional statistical modelling methods, such as the univariate signals approach or multivariate logit/probit models.[5] However, financial crises are complex events driven by non-linearly related and non-normally distributed economic and financial factors.[6] These non-linearities derive, for example, from the fact that crises become more likely as the number of fragilities increase. Potentially due to restrictive assumptions, e.g. on linearity or error distributions, conventional statistical techniques may fail in modelling these events. Novel EWSs attempt to model these complex relationships by applying non-linear techniques (Demyanyk and Hasan, 2010). For example, Peltonen (2006) and Fioramanti (2008) show that a neural network outperforms a probit model in predicting currency and debt crises. However, while the utilization of non-linear techniques may increase *a posteriori* prediction accuracies to a minor extent, Peltonen (2006) and Berg *et al.* (2005) demonstrate that the results of *a priori* predictions of financial crises remain disappointing. Given the changing nature of the occurrences of these extreme events, stand-alone numerical analyzes are unlikely to comprehensively describe them. As a complement, this motivates the development of tools with clear visual capabilities and intuitive interpretability, enabling real human perception.

One reason why the interpretability of the monitoring systems has not been adequately addressed is related to the complexity of the problem. A large number of indicators are often required to accurately assess the sources of financial instability. As with statistical tables, standard two- and three-dimensional visualizations have, of course, their limitations for high dimensions, not to mention the challenge of including a temporal or cross-sectional dimension or assessing multiple countries over time. Although composite indices of leading indicators and predicted probabilities of EWSs enable comparison across countries and over time, these indices fall short in disentangling the sources of vulnerability.[7] The recent work by IMF staff on the Global Financial Stability Map (GFSM) (Dattels *et al.*, 2010) has sought to overcome

[4] Cardarelli *et al.* (2011) show that out of 113 financial stress episodes for 17 key advanced economies, 29 were followed by an economic slowdown and an equal number by recessions.

[5] Logit and probit models have frequently been applied to predicting financial crises. For example, Berg and Pattillo (1999) apply a discrete choice model to predicting currency crises; Schmidt (1984) and Fuertes and Kalotychou (2006) to predicting debt crises; Barrel *et al*. (2010) to predicting banking crises; and Lo Duca and Peltonen (2011) to predicting systemic crises. An alternative method is the univariate non-parametric indicator proposed by Kaminsky *et al.* (1998), and its subsequent versions. See Berg *et al.* (2005) for a comprehensive review.

[6] Fioramanti (2008), Arciniegas and Arciniegas Rueda (2009) and Lo Duca and Peltonen (2011) show that indicators of debt, currency, and systemic crises are non-linearly related.

[7] There exist several composite indices for measuring financial tensions, e.g. Illing and Liu (2006), Cardarelli *et al.* (2011) and Lo Duca and Peltonen (2011). These will be further discussed in Section 2.

this challenge by a mapping of six composite indices.[8] Even here, however, the GFSM spider chart visualization of six indices falls short in disentangling individual sources. Familiar limitations of spider charts are, for example, the facts that area does not scale one-to-one with increases in variables and that the area itself depends on the order of dimensions. In addition, the use of adjustment based on market and domain intelligence, especially during crisis episodes, and the absence of a systematic evaluation gives neither a transparent data-driven measure of financial stress nor an objective anticipation of the GFSM's future precision. Indeed, the GFSM comes with the following caveat: "*given the degree of ambiguity and arbitrariness of this exercise the results should be viewed merely illustrative*".[9]

Methods for exploratory data analysis such as data and dimensionality reduction techniques may help in overcoming these shortcomings by illustrating data structures in easily understandable forms. The Self-Organizing Map (SOM) (Kohonen, 1982; 2001) is a method that combines the aims of data and dimension reduction. It is capable of providing an easily interpretable non-linear description of the multidimensional data distribution on a two-dimensional plane without losing sight of individual indicators. The two-dimensional output of the SOM makes it particularly useful for visualizations, or summarizations, of large amounts of information.

By 2005, over 7700 works had featured the SOM (Pöllä *et al.*, 2009). While extensively applied to topics in engineering and medicine, the literature is short of thorough testing of the SOM for financial stability monitoring. In the emerging market context, Arciniegas and Arciniegas Rueda (2009), Sarlin and Marghescu (2011), Sarlin (2011) and Resta (2009) have applied the SOM to indicators of currency crises, debt crises and general economic and financial performance, respectively. The SOM has not, to our knowledge, been earlier applied to monitoring systemic risk or assessing the global dimensions of financial stability, including global macro-financial proxies as well as individual advanced and emerging market economies.

The main contribution of this paper is to lay out a methodology for mapping the state of financial stability on a two-dimensional plane. The methodology uses five elements for constructing a Self-Organizing Financial Stability Map (SOFSM): (1) data and dimensionality reduction based upon the SOM, (2) identification of systemic financial crises (3) macro-financial indicators of vulnerabilities and risks, (4) an evaluation framework for assessing model performance, and (5) a model training framework. As an enhancement to the GFSM proposed by the IMF, the SOFSM not only allows disentangling the individual sources of vulnerability, but also performs well as an EWS in predicting out-of-sample systemic financial crises. Robustness of the SOFSM is tested by varying the SOM parameters, thresholds of the models, policymaker preferences, and the forecast horizons. In addition, when assessing a topologically

[8] The GFSM has appeared quarterly in the Global Financial Stability Report (GFSR) since April 2007.
[9] The authors state that the definitions of starting and ending dates of the assessed crisis episodes are arbitrary. Similarly, the assessed crisis episodes are arbitrary, as some episodes in between the assessed ones are disregarded, such as Russia's default in 1999 and the collapse of Long-Term Capital Management. Introduction of judgment based on market intelligence and technical adjustments are motivated when the GFSM is "*unable to fully account for extreme events surpassing historical experience*", which is indeed an obstacle for empirical models, but also a factor of uncertainty in terms of future performance since nothing assures manual detection of vulnerabilities, risks and triggers.

ordered SOFSM, we use the concept of a financial stability neighborhood for assessing contagion through similarities in macro-financial conditions. That is, a crisis in one position on the map indicates propagation of financial distress to adjacent locations. This type of representation may help in identifying the changing nature of crises. Further, inspired by Minsky's (1982) and Kindleberger's (1996) vindicated financial fragility view of a credit or asset cycle, we introduce the notion of the financial stability cycle. We show how the SOFSM can be used to monitor macro-financial vulnerabilities by locating a country in the financial stability cycle: being it either in the pre-crisis, crisis, post-crisis or tranquil state. We visualize samples of the panel dataset, cross-sectional and temporal country-level data, as well as different levels of aggregation, such as world, emerging market and advanced economies. The SOFSM enables disentangling the specific threats, risks and triggers, and should be treated as a starting point rather than an ending point for financial stability surveillance. The paper has also some more technical contributions compared to the earlier SOM literature. Of the above applications, only Sarlin and Marghescu (2011) perform as thorough and systematic evaluation of the model's predictive capabilities as the one in this paper. In addition, we account for policymakers' preferences regarding type 1 and 2 errors when evaluating. This paper also implements a semi-supervised SOM rather than the standard unsupervised SOM in Sarlin and Marghescu (2011).

The paper is structured as follows. Section 2 introduces the five elements necessary for creating the SOFSM, while Section 3 presents its training and evaluation as well as robustness checks. Section 4 illustrates how the SOFSM can be used for detecting signs of vulnerabilities and potential for contagion and for mapping the state of financial stability over time and across countries as well as for different levels of aggregation. Section 5 concludes.

2. Methodology

This section describes the five elements that are necessary for constructing the Self-Organizing Financial Stability Map (SOFSM): (1) data and dimensionality reduction based upon the Self-Organizing Map (SOM), (2) identification of systemic financial crises (3) macro-financial indicators of vulnerabilities and risks, (4) a model evaluation framework for assessing performance, and (5) a model training framework.

Self-Organizing Maps (SOMs)

Exploratory data analysis concerns illustrating data structures in easily understandable forms. Two main groups of methods for multivariate exploratory data analysis are those attempting data and dimensionality reduction. While data reduction through clustering is common for enabling analysis of fewer mean profiles, dimensionality reduction through non-linear projection, e.g. multidimensional scaling and its variants (Cox and Cox, 2001), is common for representing high-dimensional data in a lower dimension. Dimensionality reduction is, however, particularly problematic as all information in a high dimensional space cannot be preserved in a lower dimension. Thus, methods often differ in the properties of data they attempt to preserve, such as global interpoint distances, local interpoint distances and local neighborhood relations.

The SOM is a unique method in that it performs a simultaneous data and dimensionality reduction (Kohonen, 1982; 2001). It differs from non-linear projection techniques like multidimensional scaling by attempting to preserve the neighbourhood relations in the data space Ω on a k-dimensional array of units (represented by reference vectors) instead of attempting to preserve absolute distances in a continuous space. Due to the fact that all information cannot be preserved in dimensionality reductions, the neighborhood preservation of the SOM has been shown to be more trustworthy than that of alternative methods (Venna and Kaski, 2001). This is particularly important in this study as we mainly attempt to visualize individual data points in the reduced space. When interpreting a two-dimensional SOM grid, it is important to note that the numbers on the x- and y-axes do not carry a numeric meaning in a parametric sense; they represent positions in the data space of the map, where each of these positions (x,y) is a mean profile. The vector quantization capability of the SOM performs this data reduction into mean profiles (i.e. reference vectors or units). It models from the continuous space Ω, with a probability density function $f(x)$, to the grid of units, whose location depend on the neighbourhood structure of the data Ω. A second-level clustering can be applied on the reference vectors of the SOM, i.e. separation of data into units and units into clusters. Vesanto and Alhoniemi (2000) show that, compared to other clustering methods, the two-level SOM enhances the clustering through greater robustness on non-normally distributed data and the dual advantage of efficiency and speed. In Marghescu (2007), the data visualization features of the two-level SOM have been reviewed as better than those of other techniques. Information products of two-level SOMs have also been evaluated as superior than currently used methods by end-users within the domain of financial analysis (Eklund *et al.*, 2008).

The intuition of the basic SOM algorithm is presented here. See the Annex for further details on the SOM implementation used in this paper and Kohonen (2001) for a broad overview of the SOM. We use the standard batch SOM algorithm with a Euclidean metric. The SOM grid consists of a user specified number of units m_i (where i=1,2,…,M), which are reference vectors representing the same dimensions (number of variables) as the actual dataset Ω. Generally, the SOM algorithm operates according to the following steps:[10]

1. Initialize the reference vector values using the two principal components
2. Compare all data vectors x_j with all reference vectors m_i to find for each data the nearest reference vector m_b (i.e., best-matching unit, BMU)
3. Update each reference vector m_i to averages of the attracted data, including with diminishing weight data located in a specified neighborhood
4. Repeat steps 2 and 3 a specified number of times
5. Group reference vectors into a reduced number of clusters using Ward's (1963) hierarchical clustering.

The SOM parameters are radius of the neighbourhood σ, number of units M, map format (ratio of x and y dimensions), and number of training iterations t. Large radii result in stiff maps that stress topology preservation at the cost of quantization

[10] See the Annex for details of the steps and Ben Omrane and de Bodt (2007) for a simple example of the functioning of the SOM.

accuracy, while updates based upon solely attracted data (σ=0) leads to a standard *k*-means clustering with no topology preserving mapping.

For the purpose of this analysis, the output of the SOM algorithm is visualized on a two-dimensional plane. The rationale for not using a one-dimensional map is differences within clusters. A three-dimensional map, while adding a further dimension, impairs the interpretability of data visualizations. Here, the multidimensional space of the grid is visualized through layers, or "feature planes".

For each individual element of x_j, a feature plane represents the distribution of its values on the two-dimensional map. As the feature planes are different views of the same map, one unique point represents the same unit on all planes. We produce the feature planes in colour. Cold to warm colours represent low to high values of the indicator according to a colour scale below each feature plane. Shading on the two-dimensional map indicates the distance between each reference vector and its corresponding second-level cluster centre, i.e. those close to the centre have a lighter shade and those farther away have a darker shade.

The quality of the map is usually measured in terms of quantization error, distortion measure and topographic error (see e.g. Kohonen, 2001). As we have class information, we mainly use classification performance measures for evaluating the quality of the map.

Identifying systemic financial crises

The dataset used in this paper is an updated version of that in Lo Duca and Peltonen (2011). It consists of a database of systemic events and a set of vulnerability and risk indicators. The quarterly dataset consists of 28 countries (10 advanced and 18 emerging economies) for the period 1990:1–2011:2.[11] Hence, data vector $x_j \in \Re^{18}$ is formed of the class variables $x_{cla} \in \Re^4$ and the indicator vector $x_{ind} \in \Re^{14}$ for each quarter and country in the sample. The data are retrieved from Haver Analytics, Bloomberg and Datastream. This section explains how the systemic financial crises are identified and how the class variables are defined for enabling assessment of the financial stability cycle.

Following Lo Duca and Peltonen (2011), the identification of systemic financial crises is done using a Financial Stress Index (FSI). This approach provides an objective criterion for the definition of the starting date of a systemic financial crisis.[12] The rationale behind the FSI is that the larger and broader the shock is (i.e. the more systemic the shock), the higher the co-movement among variables reflecting tensions in different market segments. By aggregating variables to an index that measures

[11] The advanced economies are Australia, Denmark, euro area, Japan, New Zealand, Norway, Sweden, Switzerland, the United Kingdom, and the United States. The emerging market economies are Argentina, Brazil, China, Czech Republic, Hong Kong, Hungary, India, Indonesia, Malaysia, Mexico, the Philippines, Poland, Russia, Singapore, South Africa, Taiwan, Thailand and Turkey.

[12] There are several composite indices for measuring financial tensions. For example, Illing and Liu (2006) and Hakkio and Keeton (2009). Cardarelli *et al.* (2011) and Balakrishnan *et al.* (2009) constructed financial stability indices for a broad set of advanced and emerging economies. The focus of the FSI is, however, on systemic events.

stresses across market segments, the FSI captures the starting and ending points of a systemic financial crisis. The FSI is a country-specific composite index that covers the main segments (money market, equity market and foreign exchange market) of the domestic financial market: (1) the spread of the 3-month interbank rate over the 3-month government bill rate (Ind_1); (2) negative quarterly equity returns (Ind_2); (3) the realized volatility of the main equity index (as average daily absolute changes over a quarter) (Ind_3); (4) the realized volatility of the nominal effective exchange rate (Ind_4); and (5) the realized volatility of the yield on the 3-month government bill (Ind_5).[13] Each indicator *j* (Ind_j) of the FSI for country *i* at quarter *t* is transformed into an integer from 0 to 3 according to the quartile of the country-specific distribution, while the transformed variable is denoted as $q_{j,i,t}(Ind_{j,i,t})$. For example, a value for indicator *j* falling into the third quartile of the distribution would be transformed to a "2". The FSI is computed for country *i* at time *t* as a simple average of the transformed variables as follows:

$$FSI_{i,t} = \frac{\sum_{j=1}^{5} q_{j,i,t}(Ind_{j,i,t})}{5} \qquad (1)$$

To define systemic financial crises, the FSI is first transformed into a binary variable. In order to capture the systemic nature of the financial stress episodes, we focus on episodes of extreme financial stress that have led in the past (on average) to negative consequences for the real economy. In practice, we create a binary "crisis" variable, denoted as C0 that takes a value 1 in the quarter when the FSI moves above the predefined threshold of the 90th percentile of its country-specific distribution and 0 otherwise. This approach identifies a set of 94 systemic events over 1990–2011.

To describe the financial stability cycle, we create a set of other class variables besides to the crisis variable. First, a "pre-crisis" class variable C18 is created by setting the binary variable to 1 in the 18 months preceding the systemic financial crisis, and to 0 in all other periods. The pre-crisis variable mimics an ideal leading indicator that perfectly signals a systemic financial crisis in the 18 months before the event. In order to evaluate robustness for different horizons, we also create other pre-crisis class variables, by setting the binary variables C24, C12 and C6 to 1 in the 24, 12 and 6 months before the systemic event and zero otherwise. Similarly, we create "post-crisis" class variables P6, P12, P18 and P24 that are set to 1 in the 6, 12, 18 and 24 months after the systemic event. Finally, all other time periods are "tranquil" periods denoted as T0. The class vector $x_{cla} \in \Re^4$ consists of the benchmark horizons C18, C0, P18 and T0.

Macro-financial indicators of vulnerabilities and risks

To analyze the sources of systemic risk and vulnerability, we use the same indicators as in Lo Duca and Peltonen (2011). The set of indicators consists of commonly used metrics in the macroprudential literature for capturing the build up of vulnerabilities

[13] When the 3-month government bill rate is not available, the spread between interbank and T-bill rates of the closest maturity is used. The equity returns are multiplied by minus one, so that negative returns increase stress, while positive returns are set to 0. When computing realized volatilities for components $Ind_{3\text{-}5}$, average daily absolute changes over a quarter are used.

and imbalances in the domestic and global economy (e.g. Borio and Lowe, 2002; 2004; Alessi and Detken, 2011). Our key variables are asset price developments and valuations, and variables proxying for credit developments and leverage. In addition, traditional variables (e.g. government budget deficit and current account deficit) are used to control for vulnerabilities stemming from macroeconomic imbalances.[14]

Following the literature, we construct several transformations (e.g. annual changes and deviations from moving averages or trends) of the indicators (in total more than 200 transformations) to proxy for misalignments and a build up of vulnerabilities. To proxy for global macro-financial imbalances and vulnerabilities, we calculate a set of global indicators by averaging the transformed variables for the United States, the euro area, Japan and the United Kingdom.[15] The indicator vector $x_{ind} \in \Re^{14}$ consists of the best-performing transformation per indicator in terms of their univariate performance in predicting systemic events. The indicators and their summary statistics are shown in Table 1.

Statistical properties of the chosen indicators (Table 1) reveal that the data are significantly skewed and non-mesokurtic, and thus do not exhibit normal distributions. To take into account cross-country differences and country-specific fixed effects, we follow Kaminsky *et al.* (1998) by measuring indicators in terms of country-specific percentiles. While such outlier trimming is unnecessary for the clustering of the SOM, an even distribution is highly desirable for visualization.

Finally, the analysis is conducted in a real-time fashion to the extent possible. Thus, we take into account publication lags by using lagged variables. For GDP, money and credit related indicators, the lag ranges from 1 to 2 quarters depending on the country. We also de-trend variables and measure indicators in terms of country-specific percentiles using the latest available information. To test the predictability of the 2008–2009 financial crisis, we split the sample into two sub-samples: the training set spans 1990:4–2005:1, while the test set spans 2005:2–2009:2.

(INSERT TABLE 1 HERE)

Model evaluation framework

For comparing the performance of models, we need an evaluation framework that computes the usefulness of models in terms of predicting systemic financial crises. As we have class information, we mainly use classification performance measures for finding the optimal model rather than the traditional SOM quality measures. We classify the outcomes into combinations of predicted and actual classes using a contingency matrix.

[14] Even though Peltonen and Lo Duca (2011) include interaction terms of both domestic and global vulnerability indicators, we do not replicate them since they are included in the SOM processing *per se*.
[15] Qualitatively similar results are obtained when global variables are constructed as simple averages of variables of all countries in the sample.

		Actual class	
		1	-1
Predicted class	1	*True positive (TP)*	*False positive (FP)*
	-1	*False negative (FN)*	*True negative (TN)*

Based on the elements of the matrix, we compute ratios for measuring performance: recall, precision, False Positive (FP), True Positive (TP), False Negative (FN) and True Negative (TN) rates, and overall accuracy.[16] Due to unbalanced class sizes and differences in class importance, the above measures are sometime unsuited to summarize evaluations of crisis predictions. By assigning every object to the tranquil class, we would achieve a useless classifier for policy action, but still a high proportion of correct predictions (80%). This motivates using a common measure in information retrieval for evaluating performance on unbalanced class sizes. Matthews Correlation Coefficient (MCC) (Matthews, 1975) measures the correlation between the actual and predicted classes. It is defined in the range [-1,1], where -1 represents an inverse prediction and 1 a perfect prediction.[17] The costs of FPs and FNs (type 1 and 2 errors) might, however, be asymmetric, where the weight depends on the policymakers' preferences between giving false signals of crisis and tranquil periods. To calibrate an optimal model and threshold for policy action, we adapt the approach pioneered in Demirgüç-Kunt and Detragiache (2000) with the technical implementation suggested by Alessi and Detken (2011) that also accounts for differences in class size. The loss function of the policymaker is thus defined as:

$$L(\mu) = \mu(FN/(FN+TP)) + (1-\mu)(FP/(FP+TN)), \qquad (2)$$

where the parameter μ represents the relative preference of the policymaker between FPs and FNs (type 1 and 2 errors), and the errors are related to their class size. When $\mu = 0.5$, the policymaker weights equally the FN and FP ratios. She is less concerned about issuing false alarms when $\mu > 0.5$ and more concerned when $\mu < 0.5$. To find out the usefulness of our predictions, we subtract the loss from the best-guess of the policy maker. This is given by $Min(\mu, 1-\mu)$, i.e., the expected value of a guess with the given preferences. From this, we obtain the usefulness of the model:

$$U = Min(\mu, 1-\mu) - L(\mu). \qquad (3)$$

When using the above framework with a predefined preference parameter value, we classify crisis and tranquil events by setting the threshold on the probability of a crisis as to maximize the usefulness of the model for policy action. We do not explicitly assess the extent to which policymakers might be more or less concerned about failing to identify an impending crisis than issuing a false alarm. Missing a crisis may often, however, be more expensive than an internal alarm for further in-depth investigation

[16] Recall positives = TP/(TP+FN), Recall negatives = TN/(TN+FP), Precision positives = TP/(TP+FP), Precision negatives = TN/(TN+FN), Accuracy = (TP+TN)/(TP+TN+FP+FN), TP rate = TP / (TP + FN), FP rate = FP/(FP+TN), FN rate = FN/(FN+TP) and TN rate = TN/(FP+TN).

[17] The MCC is computed as follows: $MCC = \frac{TP * TN - FP * FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}}$.

of the vulnerabilities and risks. In contrast, given the risks associated with self-fulfilling prophecies, a publicly reported false alarm can have costs on par with failure to not identify a crisis. We use as a benchmark model with $\mu = 0.5$, but test model robustness by varying the preference parameter. The preference parameter of 0.5 belongs to a policymaker who weights equally the FN and FP ratios.

Using receiver operating characteristics (ROC) curves and the area under the ROC curve (AUC), we measure the global performance of the models. The ROC curve shows the trade-off between the benefits and costs of choosing a certain threshold. When two models are compared, the better model has a higher benefit (expressed in terms of TP rate on the vertical axis) at the same cost (expressed in terms of FP rate on the horizontal axis).[18] In this sense, as each FP rate can be associated with a threshold for classifying crisis and tranquil events, the measure shows performance over all thresholds. The size of the AUC is estimated using trapezoidal approximations. It measures the probability that a randomly chosen crisis observation is ranked higher than a tranquil one. A random ranking has an expected AUC of 0.5, while a perfect ranking has an AUC equal to 1.

Model training framework

In the analysis, we employ a semi-supervised SOM by using data vector $x_j \in \Re^{18}$, including class variables (C18, C0, P18 and T0), in training. In contrast to Sarlin and Marghescu (2011), where only the indicator vector $x_{ind} \in \Re^{14}$ is used in determining the best-matching units (BMUs), we also let the class vector $x_{cla} \in \Re^4$ have an impact when determining the BMUs. By including the class variables in the topology preservation the projection better separates the classes, which yields the benefit of easier interpretation of the stages of the financial stability cycle.

We obtain the predictive feature of the model by assigning to each data point $x_{ind} \in \Re^{14}$ the C18 (as well as C6, C12 and C24 when testing robustness) value of its BMU.[19] The performance of a model is evaluated using the framework introduced earlier based on the usefulness criterion for a policymaker. The performance is computed using static and pooled models, i.e. the coefficients or maps are not re-estimated recursively over time and across countries. Following Fuertes and Kalotychou (2006), it can be assumed that by not varying the specification over time or across countries, the parsimonious models better generalize in-sample data and predict out-of-sample data. Although static models have the drawback of ignoring the latest available information, they allow for more thorough in-sample evaluation for setting the SOM parameters as well as better generalization for out-of-sample

[18] In general, the ROC curve plots, for the whole range of measures, the conditional probability of positives to the conditional probability of negatives: $ROC = \frac{P(x|positive)}{P(x|negative)}$.

[19] As discussed in the Annex in detail, the BMU is the node that has the shortest Euclidean distance to a data point. When evaluating an already trained SOM model, we project all data onto the map using only the indicator vector $x_{ind} \in \Re^{14}$. For each data point, probabilities of a crisis in 6, 12, 18 and 24 months are obtained by retrieving the values of C6, C12, C18 and C24 of its BMU.

prediction. We account for a possible adjustment process that economic variables go through in between crisis and tranquil periods, i.e. a post-crisis bias (Bussière and Fratzscher, 2006), as the post-crisis class variable (P18) is included in SOM training.

The training framework and choice of the SOM is implemented with respect to three aspects: (1) the model does not overfit the in-sample data (parsimonious); (2) the framework does not include out-of-sample performance (objective); and (3) visualization is taken into account (interpretability). For a parsimonious model that avoids overfitting, we estimate a benchmark logit model similar to the one in Lo Duca and Peltonen (2011). Aiming at a parsimonious, objective and interpretable model, we employ the following training framework.

1. Train and evaluate in terms of in-sample usefulness models for $\sigma \in \{0.0001,0.3,0.5,0.75,1.0,1.5,2.0\}$ and $M \in \{50,100,150,200,250,300,400,500,600,1000\}$.[20] For each model, set the threshold on the probability of a crisis such that the usefulness is maximized. For each *M*-value, order the models in a descending order.
2. Find for each *M*-value the first model with in-sample usefulness equal to or better than that of the benchmark logit model. Choose none of the models if for an *M*-value all or none of the models' usefulness exceed that of the logit model.
3. Evaluate the interpretability of the models chosen in Step 2. Choose the one that is easiest to interpret and has the best topological ordering.[21]

The above evaluation framework results in a performance matrix with positions for each *M*-σ combination, highlights first models per *M* to outperform the logit model and uses information on topological ordering and interpretability for choosing the final model.

To partition the map into a reduced number of clusters, the units are grouped using Ward's clustering. By performing the clustering on the class variables (C18, C0, P18 and T0), the map is partitioned according to the four stages in the financial stability cycle. This creates four crisp so-called class clusters. However, the clustering given by lines on a map should only be interpreted as an aid in finding the four stages of the financial stability cycle rather than four distinct clusters.

[20] We keep constant the map format (75:100) and the training length. As is recommended by Kohonen (2001) for a stable orientation, this particular map format is close to the ratio of the two largest eigenvalues. To have a comparable training length for different parameters, we use an implementation in SOMine with an increasing function of map size and decreasing of data points, among other things. The varied parameters, *M* and tension σ, have the following effect on performance: an increase in the *M* value increases the in-sample usefulness, where $U \to 0.5$ when $M \to \infty$, but decreases out-of-sample usefulness. In fact, if *M* equals the cardinality of *x*, then perfect in-sample performance may be obtained by each m_i attracting one data. This would, however, be an overfitted model for out-of-sample prediction. Increases in tension decrease quantization accuracy, and thus in-sample usefulness, but do not have a direct effect on out-of-sample performance.

[21] Due to no consensus on a single topology-preservation metric of the SOM projection, it is evaluated following Kaski *et al.* (2000). The nodes m_i are projected into two- and three-dimensional spaces using Sammon's (1969) mapping, a non-linear mapping from a high-dimensional input space to a lower dimension. Topology preservation is defined to be adequate if the map is not twisted at any point and has only adjacent nodes as neighbours in Euclidean space. Interpretability is a subjective measure of the SOM visualization defined by the user.

3. Self-Organizing Financial Stability Map (SOFSM)

This section presents the creation of the Self-Organizing Financial Stability Map **(**SOFSM). We first use the training and evaluation frameworks for constructing the SOFSM and then perform thorough robustness checks.

Training and evaluating the Self-Organizing Financial Stability Map

The model training phase starts by estimating a pooled logit model as a benchmark. The logit model is estimated using the quarterly in-sample panel data for 28 countries from 1990:4–2005:1. The estimates are reported in Table 2 and are later used for predicting out-of-sample data from 2005:2–2009:2. On the in-sample data, the usefulness of the pooled logit model is 0.25. The training of the SOFSM is performed on the same panel data set and the evaluation results are shown in Table 3. For $M \in \{50,400,500,600,1000\}$ no model is chosen for analysis, as they never or always exceed the usefulness of the logit model (U=0.25). Finally, of the five highlighted models, we select the one with M=150 and $\sigma = 0.5$ (shown in bold) for its interpretability and topological ordering.

(INSERT TABLE 2–3 HERE)

The chosen model has 137 units on an 11x13 grid and is trained with a tension of 0.5. Henceforth, this model is referred to as the Self-Organizing Financial Stability Map (SOFSM). Figures 1–3 present the two-dimensional grid of the SOFSM, the feature planes for the 14 indicators, and the feature planes for the class variables. The feature planes in Figure 3 show the real distribution of the classes on the SOFSM. [22] When maximizing the usefulness for policymakers with different preferences, Figure 4 shows how the map is classified into two parts, where the shaded area represents early warning units and the rest tranquil units.

By performing Ward's clustering on the class variables, four class clusters are created according to the stages of the financial stability cycle. The upper left cluster represents the pre-crisis cluster (Pre crisis), the lower left represents the crisis cluster (Crisis), the centre and lower-right cluster represents the post-crisis cluster (Post crisis) and the upper right represents the tranquil cluster (Tranquil). It is important to note that the crisp clustering given by the lines that separate the map into four class clusters should only be interpreted as an aid in finding the four stages of the financial stability cycle, not as completely distinct clusters.

To compare the performance of the semi-supervised SOFSM with a standard unsupervised SOM (as the one in Sarlin and Marghescu (2011)), we train an unsupervised SOM model with the same specifications as the SOFSM. For the semi-supervised SOFSM, unsupervised SOM and the logit model, the in-sample and out-of-sample performance with the benchmark specifications ($\mu = 0.5$ and C18) are shown in Table 4. As anticipated, the unsupervised SOM performs to some extent better than the SOFSM along all measures, but it still lacks the separation of classes,

[22] The feature plane PPC0, with a high frequency on the border between the post- and pre-crisis cluster, represents the co-occurrence of pre- and post-crisis periods. In this case, the cycle need not include the tranquil stage if a new pre-crisis period is entered directly after the previous event.

which is necessary for interpreting the stages of the financial stability cycle. Hence, below we only focus on comparing the SOFSM and the logit model. For the benchmark models, the overall performance is similar between the SOFSM and the logit model. On the train set, the SOFSM performs slightly better than the logit model in terms of usefulness, recall positives, precision negatives and the AUC measure, while the logit model outperforms on the other measures. The classification of the models are of opposite nature, as the SOM issues more false alarms (FP rate=31%) than it misses crises (FN rate=19%), whereas the logit model misses more crises (31%) than it issues false alarms (19%). That explains also the difference in the overall accuracy, since the class sizes are unbalanced (around 20% crisis and 80% tranquil periods). The performance of the models on the test set differs, in general, similarly as the performance on the train set, except for the SOM having slightly higher overall accuracy. This is, in general, due to the higher share of crisis episodes in the out-of-sample dataset.

(INSERT FIGURES 1–4 HERE)
(INSERT TABLE 4 HERE)

Robustness checks

We test the robustness of the SOFSM with respect to policymakers' preferences ($\mu = 0.4$ and $\mu = 0.6$), forecast horizon (6, 12 and 24 months before a crisis) and thresholds (with the AUC measure). The results of the robustness tests are shown in Tables 5–6 and Figure 5. Table 5 shows the performance over different policymakers' preferences, Table 6 over different forecast horizons and Figure 5 and the second last column of Tables 5–6 over all possible thresholds.

For a policymaker, who is less concerned about issuing false alarms ($\mu = 0.6$), the performance of the models are similar, except for higher usefulness of the SOFSM compared to the logit model. This confirms that the SOM better detects the rare crisis occurrences. For a policymaker, who is less concerned about missing crises ($\mu = 0.4$), the usefulness of the models is similar, but the nature of the prediction is reversed; the SOM issues less false alarms than it misses crises, whereas the logit model issues more false alarms than misses crises.

Over different forecast horizons, the in-sample performance is generally similar. However, the out-of-sample usefulness, with the exception of forecast horizon of 12 months (C12), is better for the SOFSM than for the logit model. Interestingly, the logit model fails to yield any usefulness (U=0.02) at a forecast horizon of 6 months. Finally, the AUC measure, which summarizes the performance of a model over all thresholds, can be computed for all models by calculating the areas under the ROC curves, such as those shown in Figure 5 for the benchmark models ($\mu = 0.5$ and C18). It is the only measure to consistently show superior performance for the SOFSM. A caution regarding the AUC measure is, however, that parts of the ROC curve that are not policy relevant are included in the computed area. When comparing usefulness for each pair of models, the SOFSM shows consistently equal or superior performance except for a single out-of-sample evaluation with a forecast horizon of 12 months. To sum up, the SOM performs, in general, as well as or better than a logit model in both

classifying the in-sample data and in predicting out-of-sample the global financial crisis that started in 2007.

(INSERT TABLES 5–6 HERE)
(INSERT FIGURE 5 HERE)

4. Mapping the State of Financial Stability

In this section, we use the SOFSM for detecting signs of vulnerabilities and potential for contagion, and for mapping state of financial stability using macro-financial conditions. We map samples of the panel dataset by showing cross-sectional and temporal data on the two-dimensional SOFSM. We also compute aggregates for groups of countries for exploring states of financial stability globally, in advanced economies and in emerging economies. Data points are mapped onto the grid by projecting them to their best-matching units (BMUs) using only the indicator vector $x_{ind} \in \Re^{14}$. Consecutive time-series data are linked with lines. While the mappings are here performed on the evaluated SOFSM that was estimated with the in-sample data, recursive re-estimations in real-time fashion would computation-wise be feasible.

Detecting signs of vulnerabilities and potential for contagion on the SOFSM

In contrast to EWSs using binary classification methods, such as discrete choice techniques, the SOFSM enables simultaneous assessment of the correlations with all four stages of the financial stability cycle, i.e. class clusters. Thus, new models need not be derived for different forecast horizons or definitions of the dependent variable. By assessing the feature planes in Figures 2 and 3 of the SOFSM in Figure 1, the following strong correlations are found, for example. First, we can differentiate between "early" and "late" signs of a crisis by assessing differences within the pre-crisis cluster. The strongest early signs of a crisis (upper right part of the cluster) are high domestic and global real equity growth and equity valuation, while most important late signs of a crisis (lower left part of the cluster) are domestic and global real GDP growth, and domestic real credit growth, leverage, budget surplus, and CA deficit. Second, the highest values of global leverage and real credit growth in the crisis cluster exemplify the fact that increases in some indicators may reflect a rise in financial stress only up to a specific threshold. Increases beyond that level are, in these cases, more concurrent than preceding signals of a crisis. Similarly, budget deficits characterize the late post-crisis and early tranquil periods. The characteristics of the financial stability states are summarized in Table 7.

(INSERT TABLE 7 HERE)

The topological ordering of the SOFSM enables assessing, in terms of macro-financial conditions, neighbouring financial states of a particular position on the map. Transmission of financial contagion is often defined by other types of neighbourhood measures such as financial or trade linkages, proxies of financial shock propagation, equity market co-movement or geographical relations (see for example Dornbusch *et al.* (2000) and Pericoli and Sbracia (2003)). When assessing the SOFSM, the concept of neighborhood of a country represents the similarity of the current macro-financial conditions. Thus, a crisis in one position on the map indicates propagation of financial

instabilities to adjacent locations. This type of representation may help in identifying events surpassing historical experience and the changing nature of crises.

Cross-sectional and temporal analysis on the SOFSM

For a simultaneous temporal and comparative analysis, we map the state of financial stability based on the evolution of macro-financial conditions for the United States and the euro area in Figure 6. The data for both economies represent the first quarters of 2002 to 2011 and the final point of the sample, 2011Q2. Without a precise empirical treatment for accuracy, the map well recognizes for both countries the pre-crisis, crisis and post-crisis stages of the financial stability cycle by circulating around the map during the analyzed period. The early warning units in Figure 4 confirm that even a policymaker with $\mu \leq 0.6$ would have correctly predicted crises in both economies. Interestingly, the euro area is located in the tranquil cluster in 2010Q1. This indicates that the aggregated macroprudential metrics for the euro area as a whole did not reflect the elevated risks in the euro area periphery at that point in time. However, it also coincides with a relatively low FSI for the aggregate euro area. This can be explained by the weaknesses and financial stress in smaller economies being averaged out by improved macro-financial conditions in larger euro area economies, highlighting the importance of country-level analysis. As the SOFSM is flexible with respect to input data, it is of central importance that the included set of vulnerability indicators capture the particular events of interest. The macro-financial vulnerabilities currently used are best suited for capturing the build-up of vulnerabilities in the form of boom-bust cycles. However, they are less useful in identifying situations, where, for example, bank funding constraints or counterparty risks in a post-crisis recovery phase cause elevated financial stress that feeds back to the real economy, increasing the probability of a financial crisis. Furthermore, by using the traditional macro-financial vulnerabilities, it is rather difficult to capture situations where, as in the ongoing debt crisis, self-fulfilling expectations drive the equilibrium outcomes. Nevertheless, the euro area has moved to the border of the pre-crisis cluster in 2010Q4, and to an adjacent unit in 2011Q1 and Q2. This reflects the ongoing sovereign and banking crises as with $\mu = 0.4$ this particular location is an early warning unit (see Figure 4). The United States is located in the post-crisis cluster in 2010Q1 and in the tranquil cluster in 2011Q2.

Figure 7 represents a cross-section mapping of the state of financial stability for all countries in 2010Q3 and in 2011Q2, which is the latest data point in the analysis. In 2010Q3, the countries are divided into three groups of financial stability states. The map indicates elevated risks in several emerging market economies (Mexico, Turkey, Argentina, Brazil, Taiwan, Malaysia and the Philippines), while most of the advanced economies are in the lower right corner of the map (post-crisis and tranquil cluster). Three countries (Singapore, South Africa and India) are located on the border of the tranquil and pre-crisis clusters, which is an indication of a possible future transition to the pre-crisis cluster. Interestingly, in 2011Q2, most economies are located in the tranquil cluster, while the euro area has the highest financial stress by being located close to the pre-crisis cluster.

For this type of cross-sectional data, the topological ordering of the SOFSM enables assessing potential propagation of financial instabilities and contagion to adjacent macro-financial locations. When the SOFSM does not account for events surpassing

historical data, as empirical models of non-stationary processes may do, this type of representation may help in identifying the changing nature of crises. For the cross section in 2010Q3 (Figure 7), a crisis in, say, Argentina and Brazil would as well indicate possible financial distress in neighbouring countries (e.g. Taiwan, Mexico and Turkey).

(INSERT FIGURE 6–7 HERE)

Exploring aggregate financial stability states on the SOFSM

In this section, we map the financial stability states for three aggregates: the world, emerging market economies and advanced economies. We compute the state of financial stability for the aggregates by weighting the indicators for the countries in our sample using stock market capitalization to proxy their financial importance. Hence, an aggregated data vector is computed as follows: $x_{agg}(i,t) = \sum_{i=1}^{I}(w(i,t)/W(t))x_j(i,t)$, where $x_j(i,t)$ is a data vector for country i at time t, $w(i,t)$ is stock market capitalization, W is aggregated stock market capitalization and I represents all countries. These aggregates can, like any data point, be projected onto the map to their BMUs.

Figure 8 shows the evolution of global macro-financial conditions in the first quarters of 2002 to 2011. The global state of financial stability enters the pre-crisis cluster in 2006Q1 and the crisis cluster in 2007Q1. It moves via the post-crisis and tranquil cluster back to the post-crisis cluster in 2011Q1. This coincides with the global evolution of the financial stress index (FSI). More interestingly, the model signals out of sample a global financial crisis as early as in 2006Q1. The separation of the global aggregate into emerging market and advanced economies is shown in Figure 9. The mapping of the advanced economy aggregate is very similar to the one of the world aggregate, which is mainly a result of the high share of stock market capitalization of the advanced economies. Notably, the movements of the financial stability states of the emerging markets are also similar to those in the advanced economies, illustrating the global dimension of the current crisis. While the emerging market cycle moves around that of the advanced economies, it does not indicate significant differences in the timeline or strength of financial stress.

(INSERT FIGURE 8–9 HERE)

5. Conclusions

This paper creates a Self-Organizing Financial Stability Map (SOFSM) based upon data and dimensionality reduction methods for mapping the state of financial stability and visualising the sources of systemic risks. The SOFSM is a two-dimensional representation of a multidimensional financial stability space that allows disentangling the individual sources of vulnerabilities impacting on systemic risks. The model can be used to monitor macro-financial vulnerabilities by locating a country in the financial stability cycle: being it either in the pre-crisis, crisis, post-crisis or tranquil state. Moreover, the SOFSM can be used as an early warning system, and to analyse contagion on the basis of similarities in macro-financial vulnerabilities across countries. Our results indicate the SOFSM performs as well or better than a logit

model in classifying in-sample data and predicting out of sample the global financial crisis that started in 2007. The SOFSM makes an out-of-sample prediction identifying the onset of the global financial crisis as early as the first quarter of 2006. The robustness of the SOFSM is tested by varying the SOM parameters, thresholds of the models, the policymakers' preferences, and the forecast horizon.

Annex: The SOM Algorithm

This description of the SOM algorithm follows that in Sarlin (2011). This study uses the Viscovery SOMine 5.1 package.[23] In addition to an easily interpretable visual representation and interaction features, it attempts to reduce computational cost by some extensions to the basic SOM. It employs the batch training algorithm, and thus processes data in batches instead of sequences. Important advantages of the batch algorithm are the reduction of computational cost and reproducible results (given the same initialization). The training process starts with initialization of the reference vectors set to the direction of the two principal components of the input data. The principal component initialization not only further reduces computational cost and enables reproducible results, but is also shown to be important for convergence when using the batch SOM (Forte *et al.*, 2002). Following Kohonen (2001), this is done in three steps:

1. Determine two eigenvectors, v_1 and v_2, with the largest eigenvalues from the covariance matrix of all data Ω.
2. Let v_1 and v_2 span a two-dimensional linear subspace and fit a rectangular array along it, where the two dimensions are the eigenvectors and the center coincides with the mean of Ω. Hence, the direction of the long side is parallel to the longest eigenvector v_1 with a length of 80% of the length of v_1. The short side is parallel to v_2 with a length of 80% of the length of v_2.
3. Identify the initial value of the reference vectors $m_i(0)$ with the array points, where the corners of the rectangle are $\pm 0.4v_1 \pm 0.4v_2$.

Following the initialization, the batch training algorithm operates a specified number of iterations 1,2,…,t in two steps. In the first step, each input data vector x is assigned to the best-matching unit (BMU) m_c:

$$\|x - m_c(t)\| = \min_i \|x - m_i(t)\|. \tag{A.1}$$

We employ a semi-supervised version of the SOM by also including class information when determining the BMU. In the second step, each reference vector m_i (where i=1,2,…,M) is adjusted using the batch update formula:

$$m_i(t+1) = \frac{\sum_{j=1}^{N} h_{ic(j)}(t) x_j}{\sum_{j=1}^{N} h_{ic(j)}(t)} \tag{A.2}$$

[23] There are several other implementations of the SOM. The seminal packages – SOM_PAK, SOM Toolbox for Matlab, Nenet, etc – are not regularly updated or adapted to their environment. Out of the newer implementations, Viscovery SOMine provides the needed techniques for interactive exploratory analysis (Moehrmann *et al.*, 2011). For a thorough discussion of SOM software and the implementation in Viscovery SOMine, see Deboeck (1998a; 1998b).

where index j indicates the input data vectors that belong to unit c, and N is the number of the data vectors. The neighbourhood function $h_{ic(j)} \in (0,1]$ is defined as the following Gaussian function:

$$h_{ic(j)} = \exp\left(-\frac{\|r_c - r_i\|^2}{2\sigma^2(t)}\right), \tag{A.3}$$

where $\|r_c - r_i\|^2$ is the squared Euclidean distance between the coordinates of the reference vectors m_c and m_i on the two-dimensional grid, and the radius of the neighbourhood $\sigma(t)$ is a monotonically decreasing function of time t. The radius of the neighbourhood begins as half the diagonal of the grid size ($\sigma = (X^2 + Y^2)^2 / 2$), and goes monotonically towards the specified tension value $\sigma(t) \in (0,2]$.

Second-level clustering is done using an agglomerative hierarchical clustering. The following modified Ward's (1963) criterion is used as a basis for measuring the distance between two candidate clusters:

$$d_{kl} = \begin{cases} \dfrac{n_k n_l}{n_k + n_l} \cdot \|c_k - c_l\|^2 & \text{if } k \text{ and } l \text{ are adjacent} \\ \infty & \text{otherwise} \end{cases}, \tag{A.4}$$

where k and l represent two clusters, n_k and n_l the number of data points in the clusters k and l, and $\|c_k - c_l\|^2$ the squared Euclidean distance between the cluster centres of clusters k and l. The Ward clustering is modified only to merge clusters with other topologically neighbouring clusters by defining the distance between non-adjacent clusters as infinitely large. The algorithm starts with each unit as its own cluster and merges units for all possible numbers of clusters using the minimum Ward distance (1,2,…,M).

References

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(3), 520–533.

Arciniegas Rueda, I.E., Arciniegas, F., 2009. SOM-based data analysis of speculative attacks' real effects. *Intelligent Data Analysis* 13(2), 261–300.

Balakrishnan, R., Danninger, S., Elekdag, S., Tytell, I., 2009. The Transmission of Financial Stress from Advanced to Emerging Economies. IMF Working Paper, WP/09/133.

Barrell, R., Davis, P.E., Karim, D., Liadze, I., 2010. Bank regulation, property prices and early warning systems for banking crises in OECD countries. *Journal of Banking & Finance* 34(9), 2255–2264.

Ben Omrane, W., de Bodt, E., 2009. Using self-organizing maps to adjust for intra-day seasonality. *Journal of Banking & Finance* 31(6), 1817–1838

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

Berg, A., Pattillo, C., 1999. Predicting currency crises – the indicators approach and an alternative. *Journal of International Money and Finance* 18, 561–586.

Borio, C., Lowe, P., 2002. Asset Prices, Financial and Monetary Stability: Exploring the Nexus. BIS Working Papers, No. 114.

Borio, C., Lowe, P., 2004. Securing Sustainable Price Stability: Should Credit Come Back from the Wilderness?. BIS Working Papers, No. 157.

Bussière, M., Fratzscher, M., 2006. Towards a new early warning system of financial crises. *Journal of International Money and Finance* 25(6), 953–973.

Cardarelli, R., Elekdag, S., Lall, S., 2011. Financial stress and economic contractions. *Journal of Financial Stability* 7(2), 78–97.

Cox, T.F., Cox, M.A.A., 2001. *Multidimensional Scaling*. Chapman & Hall/CRC, Florida.

Dattels, P., McCaughrin, R., Miyajim, K., Puig, J., 2010. Can you Map Global Financial Stability?. IMF Working Paper, WP/10/145.

Deboeck, G., 1998a. Software Tools for Self-Organizing Map, in: Deboeck, G., Kohonen, T., (Eds.), *Visual Explorations in Finance with Self-Organizing Maps*, Springer-Verlag, Berlin, pp. 179–194.

Deboeck, G., 1998b. “Best practices in data mining using self-organizing maps, in: Deboeck, G., Kohonen, T., (Eds.), *Visual Explorations in Finance with Self-Organizing Maps*, Springer-Verlag, Berlin, pp. 201–229.

Demirgüç-Kunt, A., Detragiache, E., 2000. Monitoring Banking Sector Fragility. A Multivariate Logit. *World Bank Economic Review* 14(2), 287–307.

Demyanyk, Y.S., Hasan, I., 2010. Financial crises and bank failures: a review of prediction methods. *Omega* 38(5), 315–324.

Dornbusch, R., Park, Y.C., Claessens, S., 2000. Contagion: How it Spreads and How it can be Stopped. *World Bank Research Observer* 15, 177–197.

Eklund, T., Back, B., Vanharanta, H., Visa, A., 2000. Evaluating a SOM-based financial benchmarking tool. *Journal of Emerging Technologies in Accounting* 5(1), 109–127.

Illing, M., Liu, Y., 2006. Measuring financial stress in a developed country: An application to Canada. *Journal of Financial Stability* 2(3), 243–65.

Fioramanti, M., 2008. Predicting sovereign debt crises using artificial neural networks: a comparative approach. *Journal of Financial Stability* 4(2), 149–164.

Forte, J.C., Letrémy, P., Cottrell, M., 2002. Advantages and drawbacks of the Batch Kohonen algorithm, in: Verleysen, M., (Ed.), *Proceedings of the 10th European Symposium on Neural Networks*, Springer-Verlag, Berlin, pp. 223–230.

Fuertes, A.M., Kalotychou, E., 2006. Early Warning System for Sovereign Debt Crisis: the role of heterogeneity. *Computational Statistics and Data Analysis* 5, 1420–1441.

Hakkio, C.S., Keeton, W.R., 2009. Financial Stress: What is it, How can it be measured and Why does it matter?. Federal Reserve Bank of Kansas City Economic Review, Second Quarter 2009, 5–50.

Kaminsky, G., Lizondo, S., Reinhart, C., 1998. Leading Indicators of Currency Crises. *IMF Staff Papers* 45(1), 1–48.

Venna, J., Kaski, S., 2001. Neighborhood preservation in nonlinear projection methods: An experimental study, in Dorffner G., Bischof, H., Hornik, K., (Eds.), *Proceedings of the International Conference on Artificial Neural Networks*, Springer Verlag, Vienna, Austria, pp. 485–491.

Kaski, S., Venna, J., Kohonen, T., 2000. Coloring that reveals cluster structures in multivariate data. *Australian Journal of Intelligent Information Processing Systems* 6, 82–88.

Kindleberger, C., 1996. Maniacs, Panics, and Crashes. Cambridge University Press, Cambridge.

Kohonen, T., 1982. Self-organized formation of topologically correct feature maps. *Biological Cybernetics* 66, 59–69.

Kohonen, T., 2001. *Self-Organizing Maps*, 3rd edition. Springer-Verlag, Berlin.

Lo Duca, M., Peltonen, T.A., 2011. Macro-Financial Vulnerabilities and Future Financial Stress – Assessing Systemic Risks and Predicting Systemic Events. ECB Working Paper, No. 1311.

Marghescu, D., 2007. Multidimensional Data Visualization Techniques for Exploring Financial Performance Data, in: *Proceedings of 13th Americas Conference on Information Systems*, Keystone, Colorado, USA.

Matthews, B.W., 1975. Comparison of the predicted and observed secondary structure of T4 phage lysozyme. *Biochimica et Biophysica Acta (BBA) – Protein Structure* 405(2), 442–45.

Minsky, H., 1982. Can "it" Happen Again?: Essays on Instability and Finance. M.E. Sharpe, Armonk, N.Y.

Moehrmann, J., Burkovski, A., Baranovskiy, E., Heinze, G.A., Rapoport, A., Heideman, G., 2011. A Discussion on Visual Interactive Data Exploration Using Self-Organizing Maps, in: Laaksonen, J., Honkela, T., (Eds.), *Proceedings of the 8th International Workshop on Self-Organizing Maps*, Springer-Verlag, Berlin, pp. 178–187.

Peltonen, T.A., 2006. Are emerging market currency crises predictable? A test. ECB Working Paper, No. 571.

Pericoli, M., Sbracia, M., 2003. A Primer on Financial Contagion. *Journal of Economic Surveys* 17, 571–608.

Pöllä, M., Honkela, T., Kohonen, T., 2009. Bibliography of Self-Organizing Map (SOM) Papers: 2002-2005 Addendum. TKK Reports in Information and Computer Science, Helsinki University of Technology, Report TKK-ICS-R24.

Resta, M., 2009. Early Warning Systems: an approach via Self Organizing Maps with applications to emergent markets, in: Apolloni, B., Bassis, S., Marinaro, M. (Eds.), *Proceedings of the 18th Italian Workshop on Neural Networks*, IOS Press, Amsterdam, pp. 176–184.

Sammon Jr., J.W., 1969. A Non-Linear Mapping for Data Structure Analysis. *IEEE Transactions on Computers* 18(5), 401–409.

Sarlin, P., 2011. Sovereign Debt Monitor: A Visual Self-Organizing Maps Approach, in: *Proceedings of the IEEE Symposium on Computational Intelligence for Financial Engineering & Economics*, IEEE Press, Paris, pp. 357–364.

Sarlin, P., Marghescu, D., 2011. Visual Predictions of Currency Crises using Self-Organizing Maps. *Intelligent Systems in Accounting, Finance and Management* 18(1), 15–38.

Schmidt, R., 1984. Early warning of debt rescheduling. *Journal of Banking & Finance* 8(2), 357–370.

Ward Jr., J.H., 1963. Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association* 58, 236–244.

Vesanto, J., Alhoniemi, E., 2000. Clustering of the self-organizing map. *IEEE Transactions on Neural Networks* 11(3), 586–600.

Table 1: Statistical properties of the dataset

Type	Variable	Abbreviation	Mean	SD	Min.	Max.	Skew.	Kurt.	KSL	AD
Domestic	Inflation[a]	Inflation	0.89	5.17	-10.15	42.53	4.80	26.72	0.29*	263.90*
Domestic	Real GDP[b]	Real GDP growth	3.73	3.76	-17.54	14.13	-0.86	3.16	0.06*	11.34*
Domestic	Real credit to private sector to GDP[b]	Real credit growth	234.07	4724.00	-69.42	101870.34	20.76	429.59	0.51*	Inf*
Domestic	Real equity prices[b]	Real equity growth	5.93	33.01	-84.40	257.04	0.99	4.31	0.05*	7.28*
Domestic	Credit to private sector to GDP[a]	Leverage	3.48	51.64	-62.78	1673.04	22.76	673.35	0.29*	Inf*
Domestic	Stock market capitalisation to GDP[a]	Equity valuation	3.90	28.32	-62.79	201.55	0.77	2.41	0.03*	3.86*
Domestic	Current account deficit to GDP[c]	CA deficit	-0.02	0.07	-0.27	0.10	-0.98	0.73	0.09*	33.12*
Domestic	Government deficit to GDP[c]	Government deficit	0.01	0.05	-0.19	0.22	-1.09	3.46	0.09*	35.90*
Global	Inflation[a]	Global inflation	0.03	0.64	-1.33	2.29	0.71	1.28	0.08*	12.12*
Global	Real GDP[b]	Global real GDP growth	1.84	1.59	-6.34	4.09	-3.02	11.74	0.20*	122.16*
Global	Real credit to private sector to GDP[b]	Global real credit growth	3.87	1.68	-0.23	7.20	-0.21	-0.31	0.07*	8.82*
Global	Real equity prices[b]	Global real equity growth	2.31	19.08	-40.62	37.77	-0.57	-0.68	0.15*	41.90*
Global	Credit to private sector to GDP[a]	Global leverage	1.15	2.79	-2.79	11.21	1.84	3.40	0.22*	105.26*
Global	Stock market capitalisation to GDP[a]	Global equity valuation	0.89	17.41	-40.54	27.46	-0.50	-0.43	0.09*	19.11*

Notes: Transformations: [a], deviation from trend; [b], annual change; [c], level. KSL: Lilliefors' adaption of the Kolmogorov-Smirnov normality test. AD: the standard Anderson-Darling normality test. Significance levels: 1%, *.

Table 2: The evaluation of the SOFSM over *M* and σ values (μ=0.5 and forecast horizon 18 months)

σ (tension) / *M* (# units)	0.001	0.3	0.5	0.75	1	1.5	2
50 (52)	0.24	0.23	0.22	0.21	0.21	0.20	0.20
100 (85)	0.27	0.25	0.23	0.22	0.21	0.21	0.21
150 (137)	0.29	0.24	**0.25**	0.23	0.21	0.23	0.21
200 (188)	0.29	0.29	0.29	0.24	0.23	0.22	0.21
250 (247)	0.30	0.29	0.29	0.24	0.25	0.21	0.22
300 (331)	0.32	0.33	0.30	0.28	0.25	0.23	0.22
400 (408)	0.40	0.40	0.38	0.33	0.30	0.27	0.27
500 (493)	0.42	0.40	0.40	0.36	0.33	0.28	0.27
600 (609)	0.43	0.43	0.41	0.36	0.33	0.28	0.27
1000 (942)	0.46	0.46	0.44	0.41	0.36	0.31	0.30

Notes: Over the neighborhood radii σ, first models to outperform the logit model (*U*=0.25) per *M* value are highlighted in gray and the chosen map is shown in bold. The real number of units is shown in parenthesis since fulfilling the map ratio (75:100) affects the number of units.

Table 3: The estimates of the logit model (μ=0.5 and forecast horizon 18 months)

Variable	Estimate	Error	Z	Sig.
Intercept	-6.744	0.612	-11.024	0.000 ***
Inflation	-0.100	0.300	-0.334	0.738
Real GDP growth	0.076	0.334	0.229	0.819
Real credit growth	-0.001	0.001	-0.613	0.540
Real equity growth	1.791	0.382	4.685	0.000 ***
Leverage	0.003	0.001	3.204	0.001 ***
Equity valuation	0.002	0.001	2.689	0.007 ***
CA deficit	1.151	0.308	3.741	0.000 ***
Government deficit	0.076	0.342	0.223	0.823
Global inflation	0.207	0.341	0.608	0.543
Global real GDP growth	1.156	0.419	2.761	0.006 ***
Global real credit growth	0.685	0.381	1.799	0.072 *
Global real equity growth	0.832	0.419	1.985	0.047 **
Global leverage	0.712	0.427	1.668	0.095 *
Global equity valuation	0.959	0.472	2.029	0.042 **

Notes: Significance levels: 1%, ***; 5 %, **; 10 %, *.

Table 4: Performance of the benchmark models on in-sample and out-of-sample data (μ=0.5 and forecast horizon 18 months).

							Positives		Negatives					
Model	**Data set**	**Threshold**	**TP**	**FP**	**TN**	**FN**	**Precision**	**Recall**	**Precision**	**Recall**	**Accuracy**	**U**	**AUC**	**MCC**
Logit	Train	0.72	162	**190**	**830**	73	**0.46**	0.69	0.92	**0.81**	**0.79**	0.25	0.81	0.44
SOFSM	Train	0.60	190	314	706	45	0.38	0.81	0.94	0.69	0.71	0.25	0.83	0.40
SOM	Train	0.58	**215**	319	701	**20**	0.40	**0.91**	**0.97**	0.69	0.73	**0.30**	**0.88**	**0.48**
Logit	Test	0.72	77	**57**	**249**	93	0.57	0.45	0.73	**0.81**	0.68	0.13	0.72	0.28
SOFSM	Test	0.60	112	89	217	58	0.56	0.66	0.79	0.71	0.69	0.18	0.75	0.36
SOM	Test	0.58	**139**	95	211	**31**	**0.59**	**0.82**	**0.87**	0.69	**0.74**	**0.25**	**0.76**	**0.49**

Notes: The table reports results for the logit, semi-supervised SOFSM and unsupervised SOM on the train and test datasets and the optimal threshold. The thresholds are chosen to maximize usefulness with μ=0.5 and forecast horizon 6 quarters. To assess the performance of the models, the table also reports in columns the following measures: TP = True positives, FP = False positives, TN= True negatives, FN = False negatives, Precision positives = TP/(TP+FP), Recall positives = TP/(TP+FN), Precision negatives = TN/(TN+FN), Recall negatives = TN/(TN+FP), Accuracy = (TP+TN)/(TP+TN+FP+FN), U = Usefulness (see formulae 2 and 3), AUC = area under the ROC curve (TP rate to FP rate, see Section 2 and Figure 5) and MCC = (TP*TN-FP*FN)/√((TP+FP)(TP+FN)(TN+FP)(TN+FN)). The best accuracy measure, as per data set and evaluation measure, is shown in bold.

Table 5: Robustness tests on in-sample and out-of-sample data for different μ values (forecast horizon 18 months)

								Positives		Negatives					
Model	**Data set**	**μ**	**Threshold**	**TP**	**FP**	**TN**	**FN**	**Precision**	**Recall**	**Precision**	**Recall**	**Accuracy**	**U**	**AUC**	**MCC**
Logit	Train	0.4	0.72	**162**	190	830	**73**	0.46	**0.69**	**0.92**	0.81	0.79	**0.16**	0.81	**0.44**
SOFSM	Train	0.4	0.75	153	**166**	**854**	82	**0.48**	0.65	0.91	**0.84**	**0.80**	**0.16**	**0.83**	**0.44**
Logit	Train	0.5	0.72	162	**190**	**830**	73	**0.46**	0.69	0.92	**0.81**	**0.79**	**0.25**	0.81	**0.44**
SOFSM	Train	0.5	0.60	**190**	314	706	**45**	0.38	**0.81**	**0.94**	0.69	0.71	**0.25**	**0.83**	0.40
Logit	Train	0.6	0.54	197	**381**	**639**	38	**0.34**	0.84	0.94	**0.63**	**0.67**	0.15	0.81	0.36
SOFSM	Train	0.6	0.50	**214**	419	601	**21**	**0.34**	**0.91**	**0.97**	0.59	0.65	**0.18**	**0.83**	**0.39**
Logit	Test	0.4	0.72	**77**	57	249	**93**	0.57	**0.45**	**0.73**	0.81	**0.68**	**0.07**	0.72	**0.28**
SOFSM	Test	0.4	0.75	76	**56**	**250**	94	**0.58**	**0.45**	**0.73**	**0.82**	**0.68**	**0.07**	**0.75**	**0.28**
Logit	Test	0.5	0.72	77	**57**	**249**	93	**0.57**	0.45	0.73	**0.81**	0.68	0.13	0.72	0.28
SOFSM	Test	0.5	0.60	**112**	89	217	**58**	0.56	**0.66**	**0.79**	0.71	**0.69**	**0.18**	**0.75**	**0.36**
Logit	Test	0.6	0.54	110	**109**	**197**	60	0.50	0.65	0.77	**0.64**	0.64	0.05	0.72	0.28
SOFSM	Test	0.6	0.50	**134**	**109**	**197**	**36**	**0.55**	**0.79**	**0.85**	**0.64**	**0.70**	**0.13**	**0.75**	**0.41**

Notes: See the notes for Table 4.

Table 6: Robustness tests on in-sample and out-of-sample data for different horizons (μ=0.5)

								Positives		Negatives					
Model	**Data set**	**Horizon**	**Threshold**	**TP**	**FP**	**TN**	**FN**	**Precision**	**Recall**	**Precision**	**Recall**	**Accuracy**	**U**	**AUC**	**MCC**
Logit	Train	C6	0.72	70	**282**	**882**	21	**0.20**	0.77	0.98	**0.76**	**0.76**	**0.26**	0.81	**0.30**
SOFSM	Train	C6	0.51	**88**	530	634	**3**	0.14	**0.97**	**1.00**	0.54	0.58	**0.26**	**0.83**	0.27
Logit	Train	C12	0.72	117	**235**	**855**	48	**0.33**	0.71	**0.95**	**0.78**	**0.77**	**0.25**	0.80	**0.37**
SOFSM	Train	C12	0.69	**123**	267	823	**42**	0.32	**0.75**	**0.95**	0.76	0.75	**0.25**	**0.84**	**0.37**
Logit	Train	C18	0.72	162	**190**	**830**	73	**0.46**	0.69	0.92	**0.81**	**0.79**	**0.25**	0.81	**0.44**
SOFSM	Train	C18	0.60	**190**	314	706	**45**	0.38	**0.81**	**0.94**	0.69	0.71	**0.25**	**0.83**	0.40
Logit	Train	C24	0.58	**242**	286	673	**54**	0.46	**0.82**	**0.93**	0.70	0.73	0.26	0.81	0.45
SOFSM	Train	C24	0.63	233	**241**	**718**	63	**0.49**	0.79	0.92	**0.75**	**0.76**	**0.27**	**0.85**	**0.47**
Logit	Test	C6	0.72	18	**116**	**302**	40	0.13	0.31	0.88	**0.72**	**0.67**	0.02	0.57	0.02
SOFSM	Test	C6	0.51	**47**	205	213	**11**	**0.19**	**0.81**	**0.95**	0.51	0.55	**0.16**	**0.65**	**0.21**
Logit	Test	C12	0.72	49	**85**	**275**	67	**0.37**	0.42	**0.80**	**0.76**	**0.68**	**0.09**	0.64	**0.18**
SOFSM	Test	C12	0.69	**51**	102	258	**65**	0.33	**0.44**	**0.80**	0.72	0.65	0.08	**0.68**	0.14
Logit	Test	C18	0.72	77	**57**	**249**	93	**0.57**	**0.45**	0.73	**0.81**	0.68	0.13	0.72	0.28
SOFSM	Test	C18	0.60	**112**	89	217	**58**	0.56	**0.66**	**0.79**	0.71	**0.69**	**0.18**	**0.75**	**0.36**
Logit	Test	C24	0.58	132	68	185	91	0.66	0.59	0.67	0.73	0.67	0.16	0.76	0.33
SOFSM	Test	C24	0.63	**150**	**51**	**202**	**73**	**0.75**	**0.67**	**0.73**	**0.80**	**0.74**	**0.24**	**0.80**	**0.48**

Notes: See the notes for Table 4.

Table 7: Characteristics of the financial stability states

Variable	Pre crisis		Crisis		Post crisis		Tranquil	
	Centre	Range	Centre	Range	Centre	Range	Centre	Range
Inflation	0.49	[0.22,0.66]	0.55	[0.30,0.69]	0.59	[0.26,0.76]	0.37	[0.17,0.68]
Real GDP growth	0.67	[0.40,0.80]	0.48	[0.14,0.83]	0.34	[0.25,0.50]	0.53	[0.30,0.72]
Real credit growth	0.66	[0.28,0.85]	0.55	[0.35,0.82]	0.39	[0.18,0.68]	0.43	[0.21,0.75]
Real equity growth	0.68	[0.41,0.85]	0.28	[0.16,0.58]	0.39	[0.23,0.80]	0.61	[0.40,0.74]
Leverage	0.63	[0.31,0.80]	0.59	[0.37,0.81]	0.52	[0.23,0.83]	0.29	[0.18,0.51]
Equity valuation	0.73	[0.62,0.80]	0.55	[0.27,0.81]	0.33	[0.17,0.66]	0.45	[0.30,0.63]
CA deficit	0.58	[0.30,0.78]	0.54	[0.26,0.80]	0.48	[0.25,0.77]	0.41	[0.19,0.66]
Government deficit	0.38	[0.19,0.74]	0.45	[0.22,0.62]	0.53	[0.32,0.85]	0.61	[0.26,0.85]
Global inflation	0.33	[0.08,0.61]	0.61	[0.34,0.76]	0.46	[0.20,0.79]	0.63	[0.11,0.90]
Global real GDP growth	0.67	[0.54,0.74]	0.67	[0.30,0.86]	0.29	[0.13,0.69]	0.45	[0.13,0.71]
Global real credit growth	0.55	[0.28,0.77]	0.86	[0.61,0.92]	0.37	[0.16,0.67]	0.33	[0.15,0.52]
Global real equity growth	0.72	[0.47,0.80]	0.4	[0.23,0.63]	0.34	[0.11,0.79]	0.54	[0.20,0.73]
Global leverage	0.35	[0.18,0.60]	0.79	[0.57,0.91]	0.58	[0.17,0.77]	0.33	[0.16,0.73]
Global equity valuation	0.67	[0.48,0.82]	0.81	[0.54,0.91]	0.36	[0.14,0.76]	0.27	[0.19,0.55]

Notes: Columns represent characteristics (cluster centre and range) of the financial stability states on the SOFSM and rows represent indicators. Since data are transformed to country-specific percentiles, the summary statistics are comparable across indicators and clusters.

Figure 1: The two-dimensional grid of the SOFSM

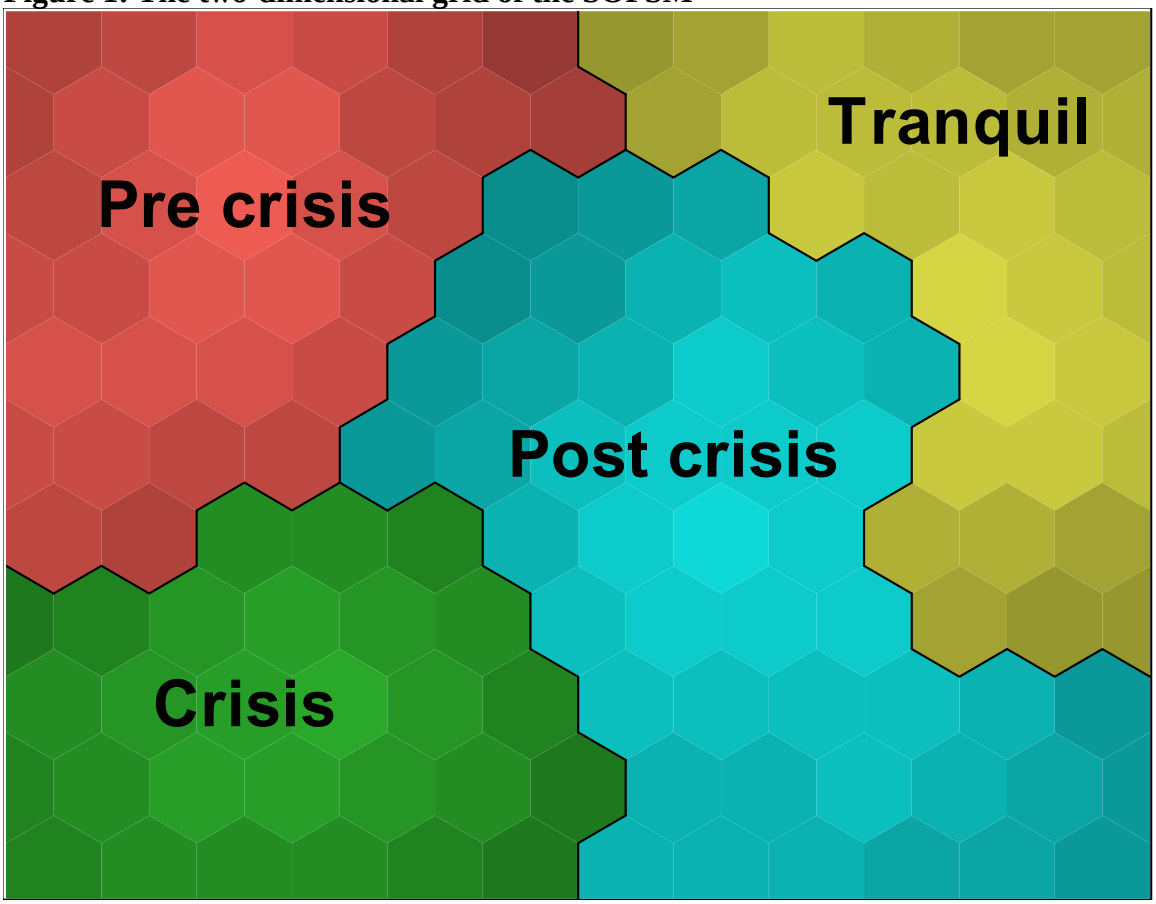

Notes: The figure displays the two-dimensional SOFSM that represents a high-dimensional financial stability space. The four clusters representing financial stability states, shown by lines and colours, are derived using the values of the class variables (C18, C0, P18, T0). Hence, the location on the SOFSM represents the state of financial stability, where the shades show distances of each unit to the centers of the financial stability states within a cluster. Distributions of the individual indicators and class variables are shown in Figures 2–3.

Figure 2: The feature planes for the 14 indicators and the main class variables

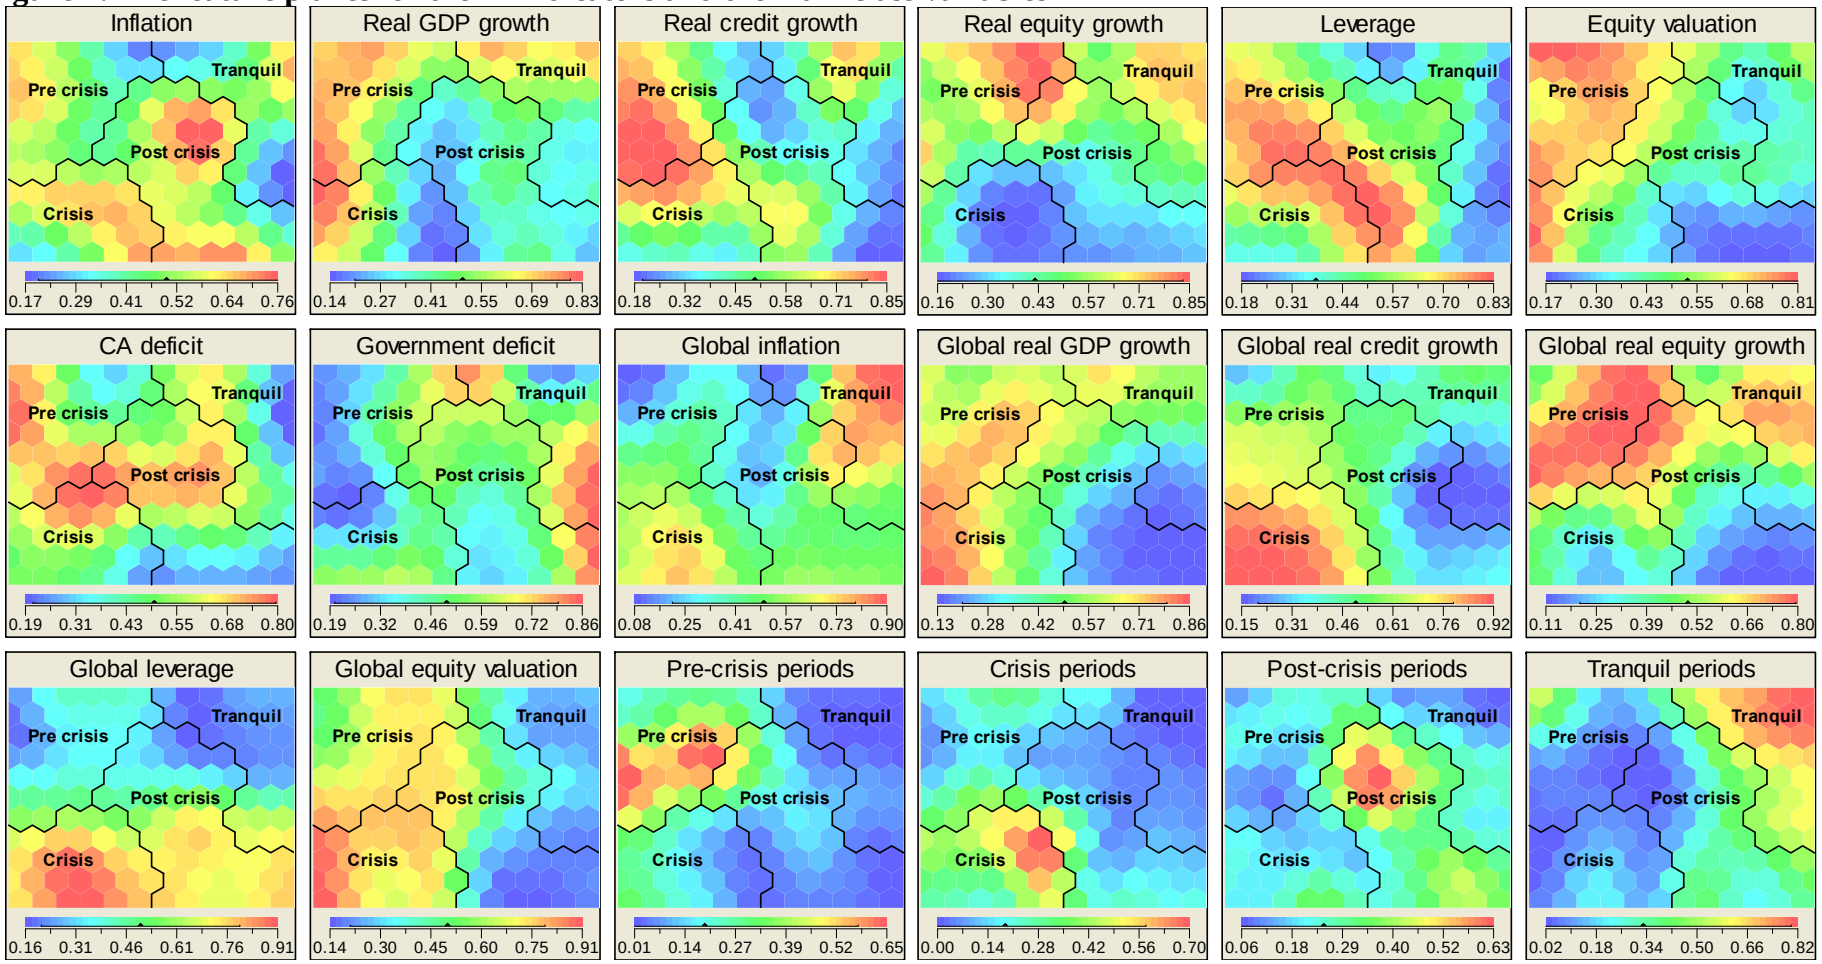

Notes: The feature planes are layers of the SOFSM in Figure 1. While the indicators are defined in Table 1, the four main class variables are Pre crisis (C18), Crisis (C0), Post crisis (P18) and Tranquil periods (T0). As each data vector consists of 14 indicators and 4 main class variables, these feature planes show the distribution of each data column on the SOFSM grid. In the case of binary class variables that take values 1 and 0, high values represent a high proportion of data in different periods (pre-crisis, crisis, post-crisis or tranquil periods). These views highlight the fact that location on the SOFSM represents the state of financial stability, where each location can be associated with variable values.

Figure 3: Feature planes for all classes

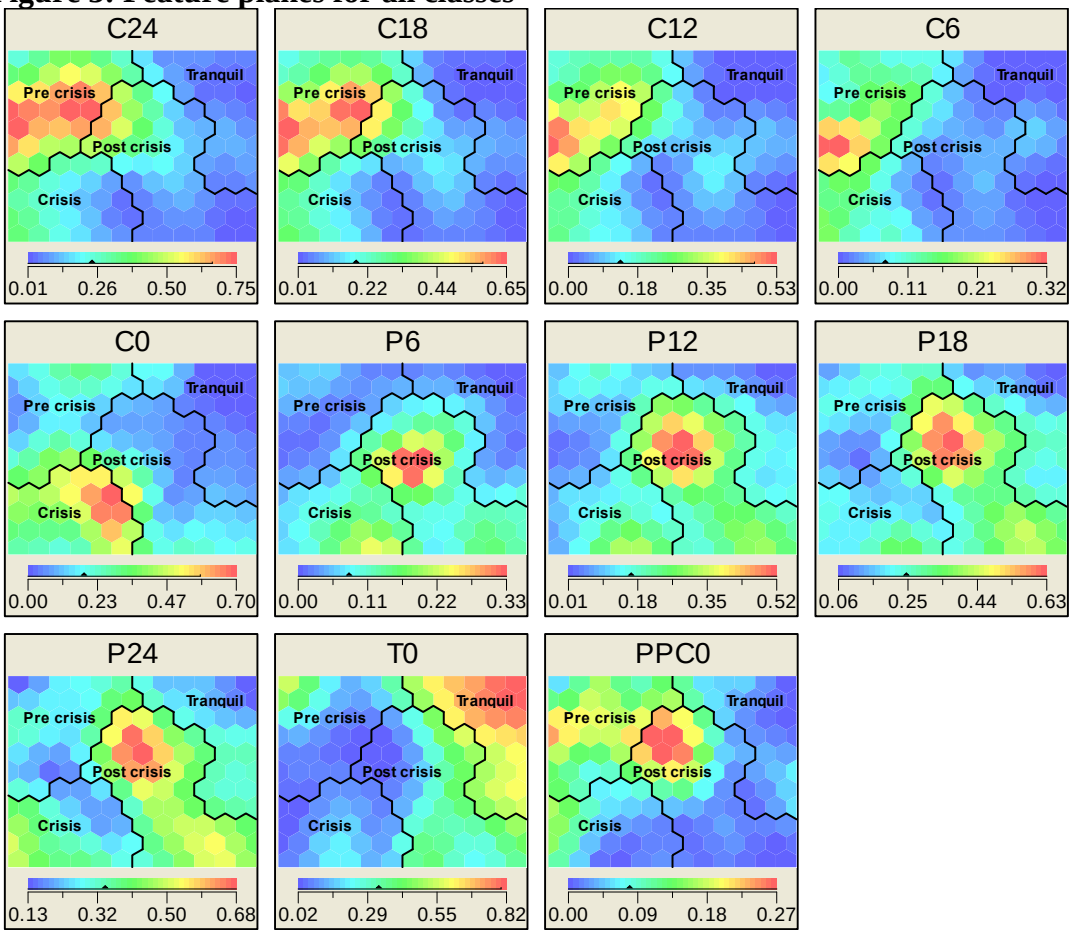

Notes: The figure shows the distributions of different pre- and post-crisis horizons. As in Figure 2, these are layers of the SOFSM in Figure 1. The feature planes C24, C18, C12, C6, P24, P18, P12 and P6 show the map distribution of class variables that represent 24, 18, 12 and 6 months before and after a crisis, respectively. While C0 and T0 show the distribution of crisis and tranquil periods, PPC0 represents the co-occurrence of pre- and post-crisis periods.

Figure 4: Early warning units for different policymakers' preferences

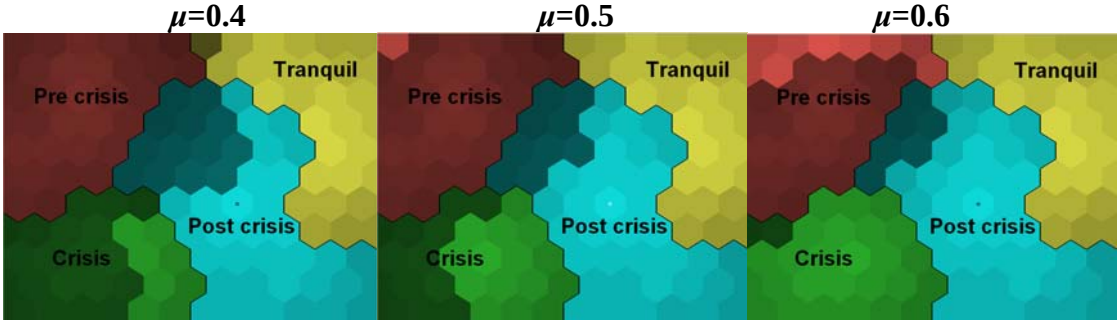

Notes: In the figure, the shaded area on the SOFSM (same map as in Figure 1) represents the part of the map that is classified as early warning units when maximizing the policymakers' preferences with three different parameter values (μ=0.4, μ=0.5 and μ=0.6) and a horizon of 18 months according to the evaluation framework.

Figure 5: In-sample and out-of-sample Receiver Operating Characteristics (ROC) curves for SOFSM and logit models (with μ=0.5 and horizon 18 months)

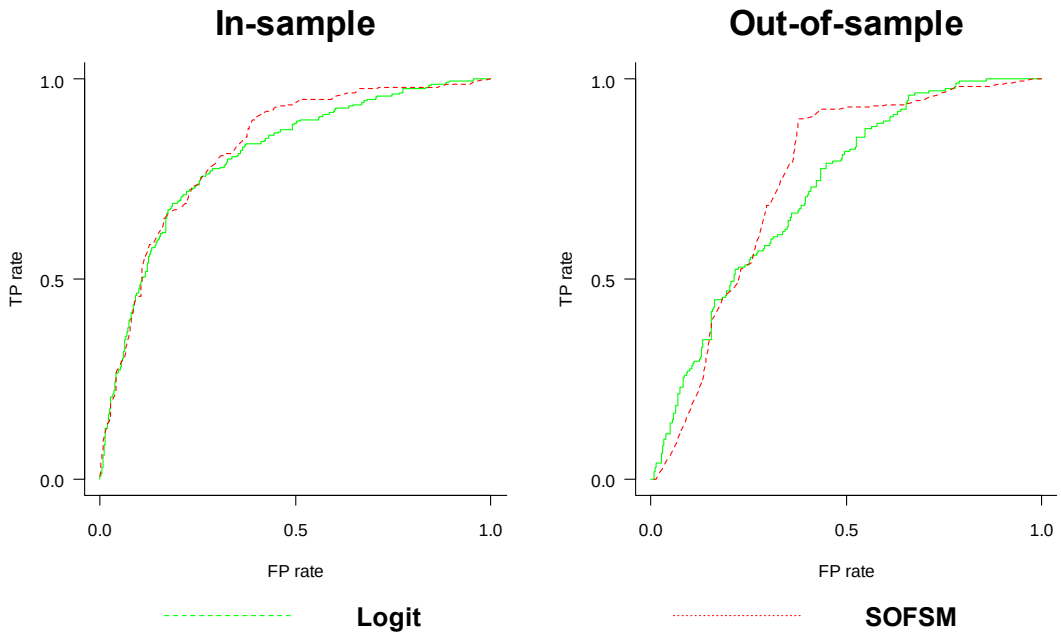

Notes: The vertical and horizontal axes represent True Positives (TP) rate (TP / (TP + FN)) and False Positives (FP) rate (FP/(FP+TN)). The area under the ROC curve (AUC), given in Tables 5–7, measures the area below these curves.

Figure 6: A mapping of the financial stability states of the United States and the euro area in 2002–2011

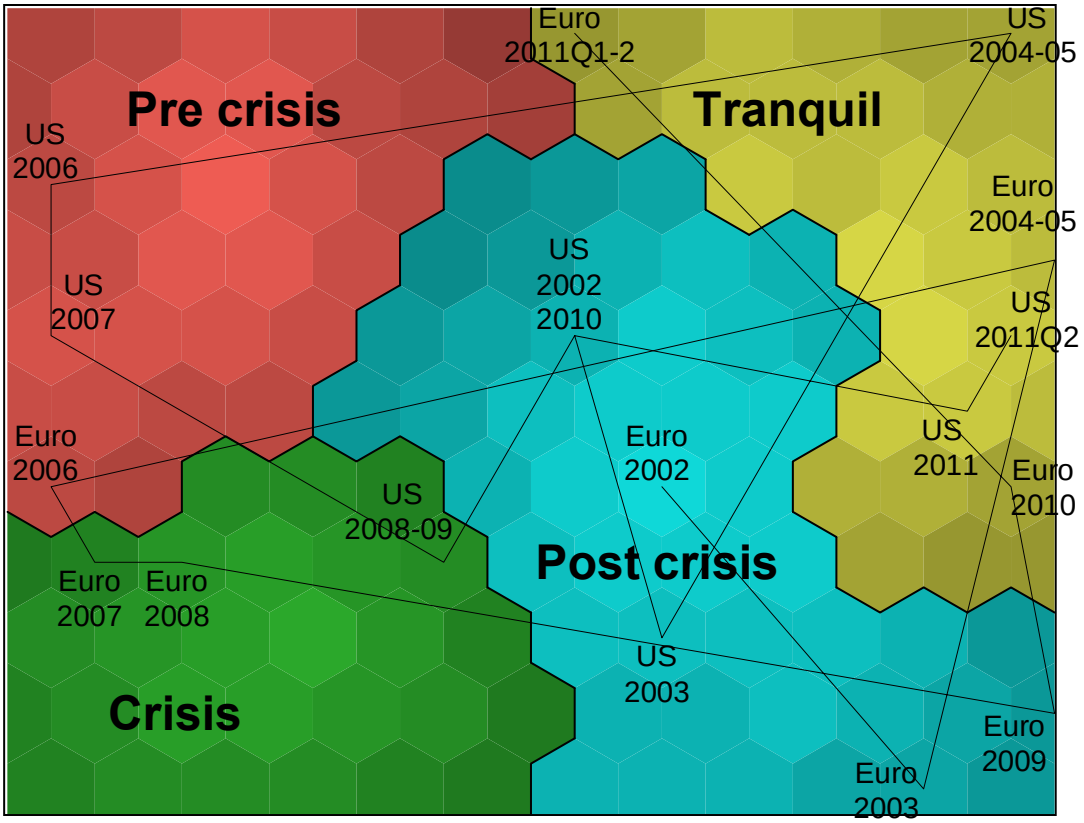

Notes: The figure displays the two-dimensional SOFSM that represents a high-dimensional financial stability space (same as in Figure 1). The lines that separate the map into four parts are based on the distribution of the four underlying financial stability states. The shades on the SOFSM show within a cluster the distance of each unit to the centers of the financial stability states. Data points are mapped onto the grid by projecting them to their best-matching units (BMUs) using only macro-financial indicators. Consecutive time-series data are linked with lines. The data for both United States and the euro area represent the first quarters of 2002–2011 as well as the second quarter of 2011.

Figure 7: A cross-sectional mapping of financial stability states for countries in the sample in 2010Q3 and 2011Q2

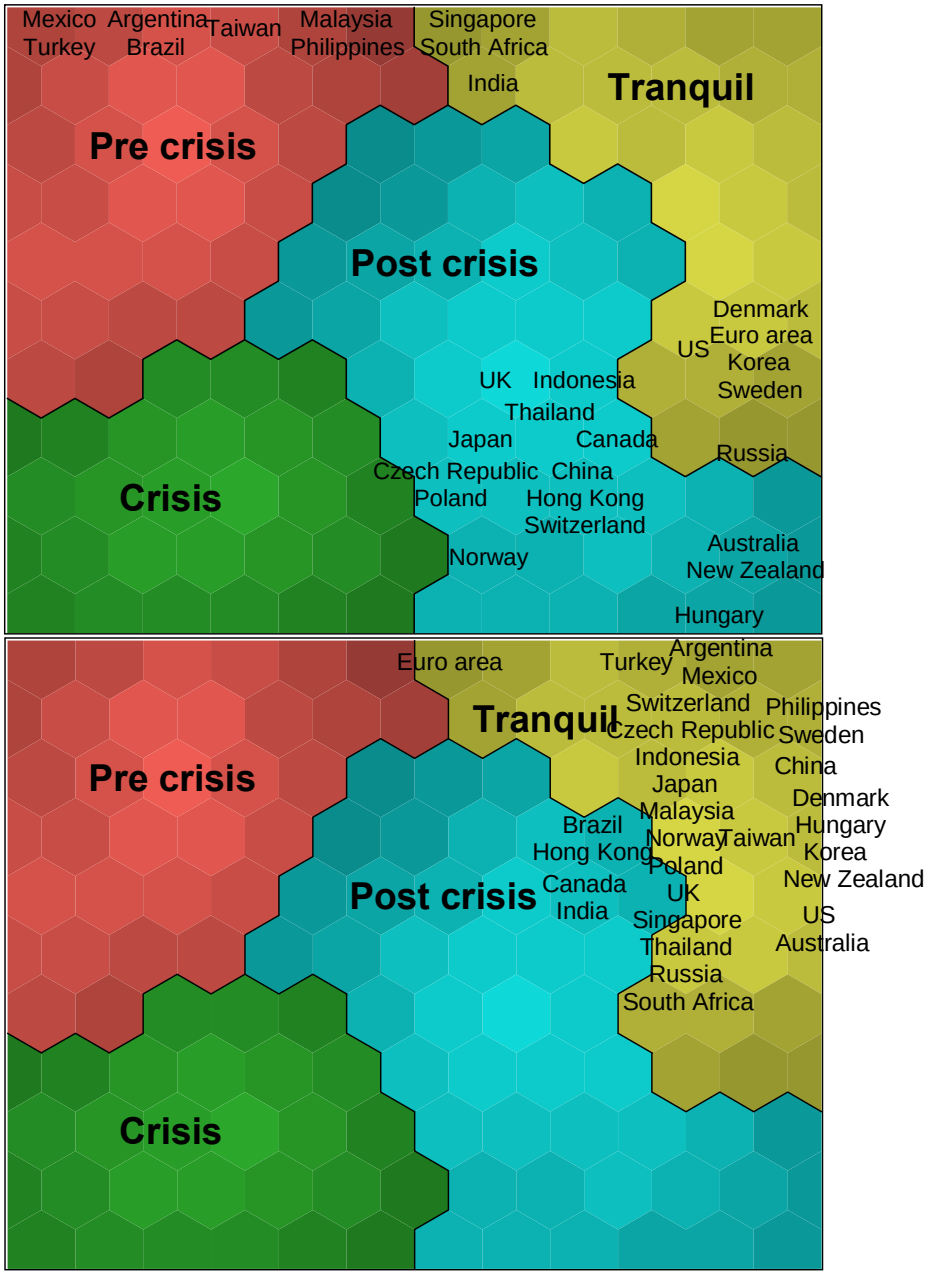

Notes: The figure displays the two-dimensional SOFSM that represents a high-dimensional financial stability space (same as in Figure 1). The lines that separate the map into four parts are based on the distribution of the four underlying financial stability states. The shades on the SOFSM show within a cluster the distance of each unit to the centers of the financial stability states. Data points are mapped onto the grid by projecting them to their best-matching units (BMUs) using only macro-financial indicators. The data for all economies represent the third quarter of 2010 and the second quarter of 2011.

Figure 8: A mapping of financial stability states for the aggregated world economy in 2002–2011

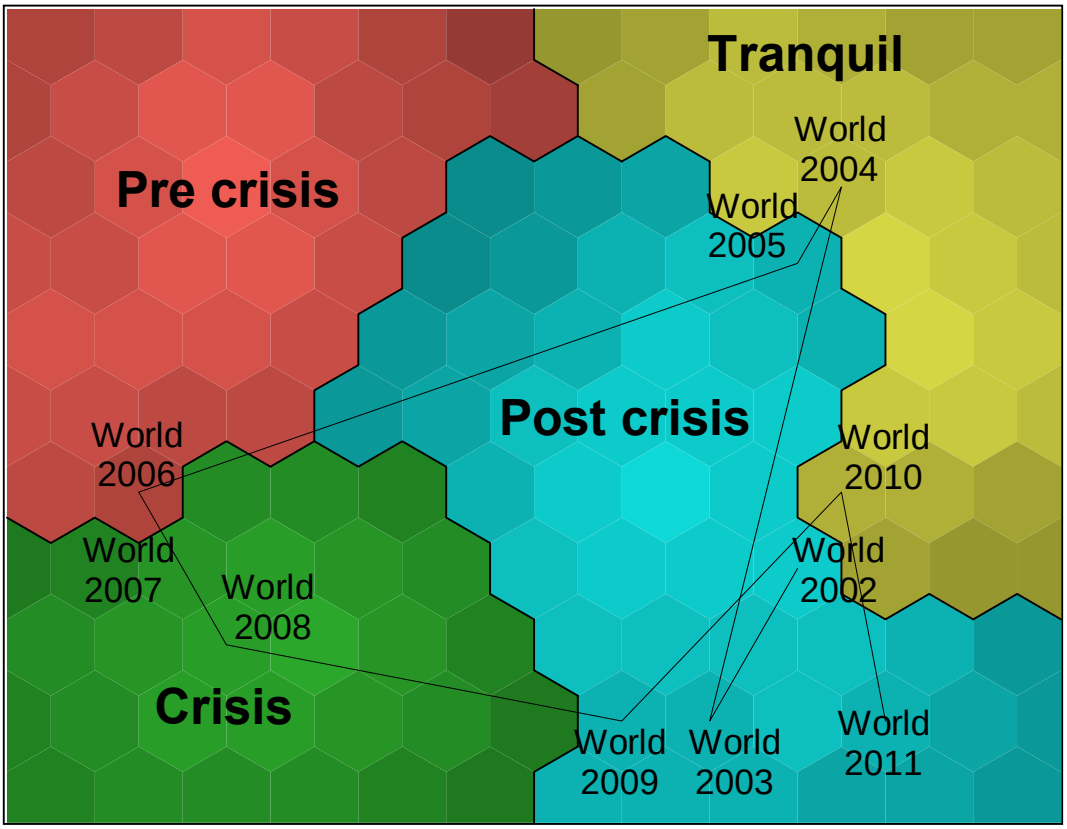

Notes: The figure displays the two-dimensional SOFSM that represents a high-dimensional financial stability space (same as in Figure 1). The lines that separate the map into four parts are based on the distribution of the four underlying financial stability states. The shades on the SOFSM show within a cluster the distance of each unit to the centers of the financial stability states. Data points are mapped onto the grid by projecting them to their best-matching units (BMUs) using only macro-financial indicators. Consecutive time-series data are linked with lines. The data for the aggregated world economy represents the first quarters of 2002–2011.

Figure 9: A mapping of financial stability states for the advanced and emerging market economies in 2002–2011

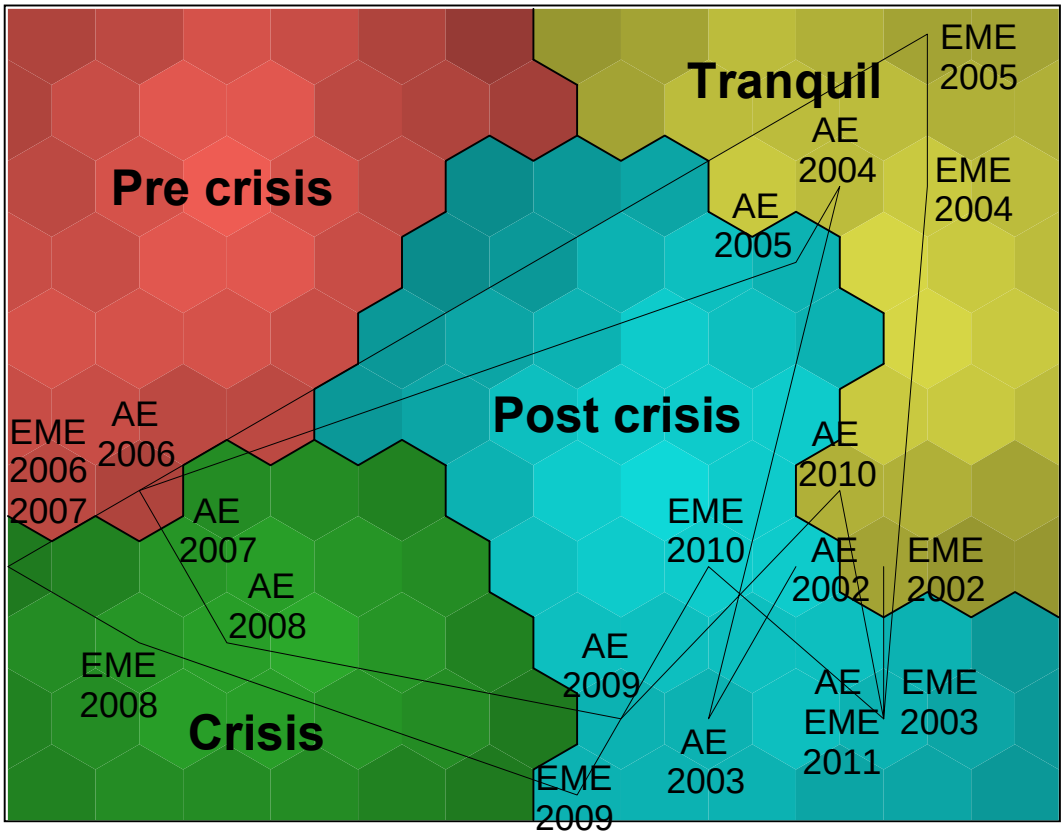

Notes: The figure displays the two-dimensional SOFSM that represents a high-dimensional financial stability space (same as in Figure 1). The lines that separate the map into four parts are based on the distribution of the four underlying financial stability states. The shades on the SOFSM show within a cluster the distance of each unit to the centers of the financial stability states. Data points are mapped onto the grid by projecting them to their best-matching units (BMUs) using only macro-financial indicators. Consecutive time-series data are linked with lines. The data for both advanced and emerging market economies (AEs and EMEs) represent the first quarters of 2002–2011.