CISS – A Portfolio-Theoretic Framework for the Construction of Composite Financial Stress Indices[*]

Daniel Holló[†]   Manfred Kremer[‡]   Marco Lo Duca[§]

This version: August 3, 2012

Abstract

This paper introduces a new indicator of current stress in the financial system as a whole named Composite Indicator of Systemic Stress (CISS). Its specific statistical design is shaped in accordance with standard definitions of systemic risk. The main innovative feature of the CISS is the application of portfolio theory to the aggregation of individual stress indicators into the composite index. Along the lines of how portfolio risk is computed from the risks of individual assets, we propose to compute the level of stress in the system as a whole by aggregating five market-specific subindices of stress - comprising a total of 15 individual stress indicators - on the basis of a time-varying measure of the cross-correlations between them. The CISS thus puts relatively more weight on situations in which stress prevails in several

market segments at the same time, capturing the idea that financial stress is more systemic and hence more hazardous for the real economy if instability spreads more widely across the whole financial system. Applied to data for the euro area as a whole, we determine within a threshold VAR model an endogenous systemic crisis-level of the CISS at which financial stress tends to depress real economic activity materially.

KEY WORDS: Financial stress index; Systemic risk; Financial stability; Financial crisis; Macro-financial linkages; Threshold VAR; Portfolio theory

1   Introduction

The recent global financial crisis began when growing and increasingly visible strains in the US subprime mortgage market caused liqudity conditions to largly dry up in the markets for securities backed by pools of such mortgages. This eventually forced the European bank BNP Paribas to halt redemptions on three of its investment funds with large exposures to such asset-backed securities. That very moment in August 2007 turned local strains in certain US asset markets into an open systemic crisis affacting large parts of the financial system in particular in advanced economies. Financial stress further intensified in September 2008 when Lehman Brothers failed, an event which clearly shifted the crisis into a higher gear. Financial frictions now started to seriously damage the global economy which, in turn, further aggravated the level of strains in the financial system, and so forth. This vicious cycle also widened the scope of the turmoil, now spilling over into emerging markets and breeding the sovereign crisis in Europe in early 2010. In general, the crisis unfolded erratically, with catalytic events triggering new stress peaks followed by periods of gradual and partial recovery.

While it makes sense to associate financial crises with its main identifying events in such narrative accounts, a more profound characterisation of crisis episodes requires quantified information about the degree of financial instability prevailing at each point in time. One popular tool to tackle this measurement problem is what has become known as a financial stress index

(FSI). FSIs aim to quantify the current state of instability - i.e., the current strength of frictions and strains ("stress") - in the financial system as a whole by aggregating a certain number of individual stress indicators into a single (continuous) composite index. In their property as a coincident indicator of overall financial stability conditions, FSIs serve several main purposes. Since FSIs usually rely on input data which is recorded at relatively high frequency (e.g., daily or weekly) and available without much delay, they provide information in more or less real time and have thus become a standard tool for those in charge of regularly monitoring financial stability. FSIs also help to better describe, analyse and compare historical crisis episodes. For instance, FSIs provide a framework to delineate the start and end points of crises in a meaningful way and at higher frequencies of observation than the traditional annual classifications (see Reinhart and Rogoff, 2009, Chapters 1 and 16). They might also improve the information content and the statistical power of early warning models which typically apply binary crisis indicators as dependent variables (e.g., Illing and Liu, 2006; for recent applications of FSIs in early warning models see Misina and Tkacz, 2009, and Lo Duca and Peltonen, 2011). Last but not least, FSIs may offer a quick, but also rough and often biased gauge of the overall impact of policy measures aimed at alleviating financial instability.

In this paper we introduce an innovative financial stress index named Composite Indicator of Systemic Stress (CISS). The main innovative features of the CISS vis-à-vis alternative FSIs rest in its economic foundation on the notion of *systemic risk.* Systemic risk can be defined as the risk that instability becomes so widespread within the financial system that it impairs its functioning to the point where economic growth and welfare suffer materially (de Bandt and Hartmann, 2000). We interpret *systemic stress* - which is what the CISS aims to measure - as an ex post measure of systemic risk, i.e. systemic risk which has materialised.

Against this conceptual background, the CISS is designed in such a way that it operationalises both the idea of widespread financial instability and the importance of financial stress for the real

economy. At the level of individual financial stress indicators, the CISS selects 15 mostly market-based stress measures which are categorised into five market segments arguably representing the largest and systemically most important parts of a modern financial system. A separate financial stress subindex is computed for each of these five market segments after appropriate transformation of the individual stress measures. The resulting subindices are now aggregated into the composite indicator based on portfolio-theoretic principles. We see the application of portfolio theory to the essential aggregation problem in the construction of financial stress indices as the main "methodological" innovation of the CISS. The portfolio-theoretic framework offers two elementary avenues to incorporate systemic risk aspects. First, analogously to the computation of portfolio risk from the risk of individual assets, the five subindices of segment-specific stress are aggregated by taking into account a time-varying measure of the cross-correlations between them. In this way the CISS puts relatively more weight on situations in which stress prevails in several market segments at the same time, i.e. on situations in which financial instability *spreads widely* across the whole financial system. Second, the portfolio weights (shares) assigned to each subindex can be calibrated in proportion to their systemic importance which in turn may be approximated in different ways. For instance, the weights could reflect the relative size of the financial market segment covered by each subindex (*size weights*). Alternatively, the weights could also mirror the relative strength of the estimated impact of each subindex of financial stress on economic activity (*real-impact weights*). The latter route is taken in the construction of the CISS for the euro area as a whole which is the subject matter of the empirical part of the paper.

Independent of the portfolio-theoretic aggregation framework, the proposed euro area CISS possesses two further idiosyncratic features vis-à-vis most other FSIs, namely its recursive (real time) computation over expanding data samples, and its targeted robustness to the addition of new information achieved by transforming the raw stress indicators on the basis of order

4

statistics. Both features help to mitigate the risk of reclassifying crisis regimes/events ex post, a risk which may affect in particular those FSIs whose statistical design relies strongly on stable distribution properties of the raw input series in typically small samples. The empirical evaluation of the euro area CISS confirms the robustness of its information content. Furthermore, all peaks in the CISS can be associated to well-known periods of financial stress, and the recent financial crisis clearly stands out as a unique systemic event in the past two and a half decades.

The paper makes a further contribution to the literature on FSIs by proposing the use of econometric approaches to *endogenously* identify different stress regimes. We demonstrate it on the basis of a parsimoneous threshold vector autoregressive (TVAR) model that identifies a systemic crisis level of the CISS at or above which financial stress becomes very costly in terms of reduced real economic activity. The results from the TVAR suggest that while shocks in the CISS do not exert any statistically significant output reactions during low-stress regimes, industrial production truly collapses during high-stress regimes in response to a typical adverse shock in financial stress. Similarly, it is only in the high-stress regime that a negative output shock triggers increases in financial stress, supporting the idea that output and financial shocks might reinforce each other in a truly systemic crisis.

The remainder of this paper is organised as follows: Section 2 provides a very brief summary of the related literature. Section 3 motivates and describes the statistical design of the CISS and presents an empirical application to data for the euro area economy as a whole. The euro area CISS is evaluated in Section 4 in terms of its robustness properties and its ability to identify well-known periods of financial stress; in addition, it presents results from the TVAR model to determine endogenouosly different regimes in the CISS. Section 5 concludes.

2 Related literature

The paper relates mainly to two strands of literature, the first discussing different options to construct financial stress indices, and the second studying the impact of financial distress on aggregate economic activity.

As to the first strand, the development of FSIs has become a very active business in recent years, spurred by the analytical demands created by the crisis. For the sake of brevity the following literature review is neither very detailed nor exhaustive but tries to illustrate the broad range of existing methodologies. The ECB working paper version of this article (Hollo, Kremer and Lo Duca, 2012) provides a more detailed account of the recent literature. The seminal paper is Illing and Liu (2006). They develop a daily financial stress index for the Canadian financial system and propose several approaches to the aggregation of individual stress indicators into a composite stress index. The specification of their preferred FSI was chosen according to which variant performs best in capturing crisis events in the Canadian financial system identified on the basis of a survey among Bank of Canada policy-makers and staff. The preferred FSI comprises 11 financial market variables aggregated on the basis of weights determined by the relative size of the market to which each of the indicators pertains. Caldarelli, Elekdag and Lall (2011) present a monthly financial stress index for 17 advanced economies computed as the arithmetic average of twelve standardised market-based financial stress indicators, an aggregation method also known as *variance-equal weighting*. Nelson and Perli (2007) and Carlson, Lewis and Nelson (2012) present a weekly financial fragility indicator for the United States computed in two steps from twelve market-based financial stress measures. The standardised input series are first reduced to three summary indicators, namely a level factor, a rate-of-change factor and a correlation factor. In the second step, the financial fragility indicator is computed as the fitted probability from a logit model with the three summary indicators as explanatory variables and a binary pre-defined crisis indicator as the dependent variable. Refining the last step of the

6

approach by Nelson and Perli (2007), Blix Grimaldi (2010) computes a weekly FSI for the euro area, where the binary crisis indicator is systematically derived from crisis events identified on the basis of a keyword-search algorithm applied to relevant parts of the ECB Monthly Bulletin. Hakkio and Keeton (2009) construct a monthly FSI applying principal components analysis to US data. The idea is that financial stress is the factor most responsible for the observed correlation between the indicators, and this factor is identified by the first principal component of the sample correlation matrix computed for the standardised indicators. The weights of each input series is computed from its loading to the first principal component. The weekly financial conditions index developed by Brave and Butters (2011a, 2011b) also builds on factor analysis but is more complex and sophisticated than its competitors in terms of the number and the heterogeneity of the input data and the statistical indicator design. The computation of the FCI is cast into a dynamic factor model in state-space form which includes 100 indicators, where Kalman filtering takes account of the missing data problem resulting from the different sample lengths and frequencies of the input data. The FSI developed by Oet et al. (2011) integrates 11 daily financial market indicators grouped into four sectors. The raw indicators are normalised by transforming the values of each series into the corresponding value of their empirical CDF. The transformation method is similar to the one developed independently in the present paper. The transformed indicators are then aggregated into the composite indicator by applying time-varying credit weights which are proportional to the quarterly financing flows through the four markets concerned.

Second, the present paper also relates to the general literature examining empirically the real impacts of financial stress (e.g., Hakkio and Keeton, 2009; Cardarelli, Elekdag and Lall, 2011; Hatzius et al., 2010; Li and St-Amant, 2010; Mallick and Sousa, 2011; Carlson, King and Lewis, 2011; and van Roye, 2011). The regime-dependence of the impact of financial stress on economic activity found in our study broadly corroborates the findings of Davig and Hakkio

(2010) from a bivariate Markov-switching model with the FSI developed by Hakkio and Keeton (2009) and a monthly measure of US economic activity as endogenous variables. Hubrich and Tetlow (2011) for the US and Hartmann et al. (2012) for the euro area provide qualitatively similar evidence on much stronger impacts of financial stress on economic activity in high-stress regimes within more richly specified small-scale macro-econometric Bayesian VAR models with Markov-switching in coefficients and residual variances, where the latter study uses the CISS to measure financial stress.

3 Statistical design of the CISS

The CISS aims to measure the current level of systemic stress in the financial system as a whole. Ideally, the indicator should capture strains in each part of the financial system, weighted by its systemic importance. However, a real-world financial system constitutes a highly complex and complicated network of a multitude of financial markets, financial intermediaries and financial infrastructures, and it is practically impossible - not least due to data limitations - to measure the level of stress in each and every of its elements. In order to reduce the level of complexity, it seems to make sense to limit attention to those parts of the financial system - subject to data availability - which can be regarded as both systemically important and sufficiently representative for the system as a whole.

Against this background, the design of the CISS is set up as a three-tier aggregation framework, with each tier featuring particular characteristics of systemic risk.

We start with the *intermediate level,* at which five highly aggregated market segments suppose to represent the main elements of a financial system. These segments capture in a stylised fashion the main flows of funds from ultimate lenders/savers to borrowers/spenders, channeling funds either indirectly through financial intermediaries or directly via short-term and long-term security markets. The five segments are: 1. The financial intermediaries sector (comprising

banks, insurance companies, pension funds and other financial services providers); 2. The bond market (only sovereign and non-financial corporate issuers); 3. The equity market (only non-financial corporations); 4. The money market (broadly defined as including in principle all forms of short-term wholesale financing in the economy, e.g., interbank and commercial paper markets); and 5. The foreign exchange market (capturing potential stresses affecting cross-border financing activities).

The choice of these five market segments can be justified on grounds of their systemic importance. Size, substitutability and interconnectedness are three of the main criteria usually applied to identify systemically important financial institutions and markets. According to the size criterion, it is probably fair to say that the five identified market segments collectively represent the core of any financial system. In addition, the markets and sectors included in the CISS are aggregated to such an extent that in case financial stress disrupts all of them at the same time, no major substitute forms of unimpaired finance presumably exist in the economy.

The interconnectedness criterion brings us to the *top tier* of our aggregation framework where the heart of the paper rests, namely the application of portfolio-theoretical principles to the aggregation of market segment-specific indices of financial stress into the CISS. The aggregation of the five subindices of stress by way of their time-varying cross-correlations operationalises the idea of *widespread financial instability* in a novel fashion. In addition, the variation in the cross-correlations may also capture state-dependent changes in the degree of interconnectedness between the market segments, which are likely to be relatively strongly interconnected in general but in particular so during times of stress. The calibration of segment-specific portfolio weights for each subindex of stress offers another route to bring in features of systemic risk.

Finally, the selection of individual indicators of financial stress takes place at the *lower tier* of the aggregation framework. Each selected indicator captures typical symptoms of financial stress in the market segment it is associated with.

Details on each of the three tiers of the statistical indicator design are provided in the subsequent subsections. The empirical implementation of the CISS concept is demonstrated on the basis of data for the aggregate euro area economy.

3.1   Raw stress indicators

Financial stress is a rather elusive concept. It is usually operationalised by drawing on the main features associated with financial crises defined as situations in which the normal functioning of a financial system is impaired. The list of typical crisis features includes (see, e.g., Hakkio and Keeton, 2009; Fostel and Geneakoplos, 2008): increased uncertainty (e.g., about asset valuations and the behaviour of other investors); increased differences of opinion among investors; increased asymmetry of information between borrowers and lenders (intensifying problems of adverse selection and moral hazard); and lower preferences for holding risky assets (flight-to-quality) or illiquid assets (flight-to-liquidity) resulting from stronger risk or uncertainty aversion, for instance (Caballero and Krishnamurthy, 2008).

Although the various stress features are not directly observable, they can be captured by observable stress symptoms like increased asset price volatility, large revaluations for risky assets, wider default and liquidity risk premia, as well as sharp reversals in financing flows linked to financial instruments or institutions perceived as being more risky. However, such symptoms measure the underlying stress characteristics only imperfectly, as the former typically also reflect the impact of other factors than the mentioned crisis features. The identification of individual stress features is further complicated by the fact they are often closely interrelated, with a tendency to reinforce each other as in the case of fire sales and liquidity spirals (Brunnermeier and Pedersen, 2009; Krishnamurthy, 2010). As a consequence, it is likely that certain financial market indicators - henceforth called "raw stress indicators" - capture several stress features at the same time.

The literature offers a vast variety of financial quantity and price variables reflecting characteristics of financial stress. Which ones to pick for the construction of a financial stress index appears to be a greatly arbitrary choice. For our purposes, we narrow down the list of candidate raw stress indicators to be included in the CISS by imposing several restrictions:

1. Each of the five segment-specific subindices of stress includes (not more than) three raw stress indicators. The composite indicator thus comprises a total of (at most) 15 individual indicators of financial stress. The same number of indicators per subindex ensures that the subindices do not possess different statistical properties by construction. In addition, the three raw indicators in each subindex should convey complementary information on the level of strains in the respective market segment; ideally, the information content of all three indicators should be perfectly correlated only under conditions of extreme stress.

2. To ensure representativeness, the raw stress indicators should cover market-wide developments. We therefore prefer indicators based on broad market indices, but sometimes revert also to certain assets with benchmark status for the pricing of a wider range of close substitutes (e.g., government bonds).

3. To make the CISS fit for real-time monitoring purposes, all raw stress indicators should be available at a daily/weekly frequency and with a publication lag of one day at most.

4. Raw stress indicators should carry sufficiently long data histories to comprise several episodes of financial stress.

These restrictions jointly imply that the CISS includes mainly fairly standard price-based financial market indicators available for many countries and over relatively long samples. We mostly rely on risk spreads and a measure of realised asset return volatility included in all five subindices. Table 1 provides details on the computation and the data sources of all individual stress indicators included in the euro area CISS.

As to their information content, asset return volatilities tend to increase with investors' uncertainty about future fundamentals and/or the behaviour and sentiment of other investors (Pastor and Veronesi, 2009; Veronesi, 2004). Chordia, Sarkar and Subrahmanyam (2005) present evidence that volatility shocks in bond and stock markets tend to predict shifts in liquidity condition in both markets. Stress in the money market is also captured by a euro area equivalent of the US TED spread, i.e. by the yield differential between a three-month unsecured inter-bank market rate and a comparable essentially risk-free Treasury bill rate. This spread reflects liquidity and counterparty risk in the inter-bank market (Heider, Hoerova and Holthausen, 2010; Acharya and Skeie, 2011) as well as the convenience premium on short-term Treasury paper, and thus captures stress features like flight-to-quality, flight-to-liquidity as well as the price impacts of enhanced adverse selection problems in times of stress in the banking system. Another variable measuring stress in the inter-bank money market is banks' recourse to the marginal lending facility at national central banks of the Eurosystem. The yield spread between long-term A-rated bonds of non-financial corporations and governments, respectively, measures bond market stress. Drawing on the empirical findings of Feldhütter and Lando (2009) for the US, the ten-year swap spread is arguably a relatively clean measure of the convenience premium embedded in the prices of German government bonds - the presumably safest and most liquid sovereign bonds in the euro area - which, in turn, captures flight-to-liquidity and flight-to-quality effects in this market segment (on the convenience yield in US Treasuries see Krishnamurthy and Vissing-Jorgensen, 2010, and Krishnamurthy, 2010). Stress in the equity market is captured by the so-called CMAX measuring the maximum cumulated loss in a stock price index over a moving two-year window. It was originally developed to identify crisis periods in international stock markets (Patel and Sarkar, 1998). Stress in the equity market is furthermore measured by a time-varying correlation coefficient between stock and government bond returns capturing, amongst others, flight-to-liquidity and flight-to-quality phenomena (Baele, Bekaert and Inghelbrecht, 2010). For

instance, in times of heightened systemic stress, investors try to shift funds out of more risky stocks into safer government bonds, thereby driving the return correlation between these two asset classes into negative territory. Since our stress factors ought to increase with higher levels of stress, we take the negative of the short-term stock-bond correlation (measured as the deviation from a longer trend-correlation). Stress in the financial intermediaries sector is measured by idiosyncratic stock return volatility of the banking sector and the yield differential between A-rated financial and non-financial corporations. A partly novel stress measure of the financial intermediaries segment is obtained by interacting the CMAX of this sector with its inverse price-book ratio. The idea behind this indicator is that any given large stock market loss puts financial intermediaries the more under stress the lower their current valuation levels as measured by the price-book ratio. Stress in the foreign exchange market is exclusively represented by the realised volatility of three bilateral euro exchange rates.

3.2 Transformation of raw indicators by means of order statistics

The aggregation of individual stress indicators is arguably the most difficult aspect of constructing composite financial stress indicators (Illing and Liu, 2006). The literature offers several options, all coming with specific advantages and disadvantages. In most cases the first step consists in putting the individual raw stress indicators on a common scale by standardisation, i.e. by subtracting the sample mean from the raw score and dividing this difference by the sample standard deviation. The standardised indicators are then usually aggregated into a composite indicator by taking their arithmetic average ("variance-equal weighting"), a weighted average or by applying principal components analysis (PCA). Standardisation, however, implicitly assumes variables to be normally distributed. The fact that many standard stress indicators clearly violate this assumption (e.g., variances) enhances the risk that the results obtained from the use of standardised variables are sensitive to aberrant observations. In that case, for instance, both

Table 1: Individual financial stress indicators included in the CISS

Money market

1. Realised volatility of 3-month Euribor rate; weekly average of absolute daily rate changes; data start 8 Jan. 1999; source: Datastream.

2. Interest rate spread between 3-month Euribor and 3-month French T-bills; weekly average of daily data; data start 8 Jan. 1999; source: Datastream.

3. Monetary Financial Institution's (MFI) recourse to the marginal lending facility at Eurosystem central banks, divided by their total reserve requirements; MFIs can use the marginal lending facility to obtain overnight liquidity from the national central banks against eligible assets and, typically, at an interest rate which is higher than the prevailing overnight market interest rate; weekly average of daily data; data start 1 Jan. 1999; source: ECB.

Bond market

4. Realised volatility of German 10-year benchmark government bond index; weekly average of absolute daily yield changes; data start 5 Jan. 1990; source: Datastream.

5. Yield spread between A-rated non-financial corporations and government bonds (7-year maturity); weekly average of daily data; data start 3 Apr. 1998; source: Bloomberg.

6. 10-year interest rate swap spread; weekly average of daily data; data start 4 Mar. 1987; source: Datastream.

Equity market

7. Realised volatility of Datastream non-financial sector stock price index; weekly average of absolute daily log returns; data start 4 Jan. 1980; source: Datastream.

8. Maximum cumulated loss (CMAX) of Datastream non-financial sector stock price index (x_t) over a moving 2-year window: $CMAX_t = 1 - x_t / \max[x \in (x_{t-j} | j = 0, 1, ..., T)]$ with $T = 104$ for weekly data; data start 1 Jan. 1982; source: Datastream.

9. Stock-bond correlation; weekly average of the difference between the 4-year (1040 business days) and the 4-week (20 business days) correlation coefficients between daily log returns of Datastream total stock price index and the 10-year German government benchmark bond price index; final indicator is assigned a value of zero for negative differences; data start 8 Jan. 1982; source: Datastream.

Financial intermediaries

10. Realised volatility of idiosyncratic equity return of Datastream bank sector stock price index over the total market index; weekly average of absolute daily idiosyncratic returns; idiosyncratic return calculated as residual from OLS regression of daily log bank return on log market return over a moving 2-year window (522 business days); data start 1 Jan. 1982; source: Datastream.

11. Yield spread between A-rated financial and non-financial corporations (7-year maturity); weekly average of daily data; data start 3 Apr. 1998; source: Bloomberg.

12. CMAX of Datastream financial sector stock price index interacted with the sector's book-price ratio; both indicators transformed by their recursive sample CDF prior to multiplication; final indicator obtained by taking the square root of this product; data start 1 Jan. 1982; source: Datastream.

Foreign exchange market

13. Realised volatility of euro exchange rate vis-à-vis US dollar; weekly average of absolute daily log foreign exchange returns; data start 6 July 1990; source: Datastream.

14. Realised volatility of euro exchange rate vis-à-vis Japanese Yen; weekly average of absolute daily log foreign exchange returns; data start 6 July 1990; source: Datastream.

15. Realised volatility of euro exchange rate vis-à-vis British Pound; weekly average of absolute daily log foreign exchange returns; data start 6 July 1990; source: Datastream.

the conditional means and standard deviations - when calculated over expanding data samples - can be subject to large revisions if more and more outliers are added to the sample (Hakkio and Keeton, 2009) as it tends to happen during extended periods of severe financial stress. Such problems can distort the information content of financial stress indicators over time (see the general discussion of the event reclassification problem in Section 4.1). Applying PCA to the aggregation of standardised indicators may aggravate the problem of sub-sample robustness since PCA itself is sensitive to outliers (as it minimises squared distances from the multidimensional mean).

This stability problem can be mitigated by transforming the raw stress indicators based on more robust properties of their empirical distribution function (Stuart and Ord, 1994). For the CISS we propose to transform the raw indicators by means of their empirical cumulative distribution function (CDF) involving the computation of order statistics.

Let us denote a particular data set of a raw stress indicator x_t as $x = (x_1, x_2, ..., x_n)$ with n the total number of observations in the sample. The elements of x are now arranged in ascending order. The ordered sample is denoted $\left(x_{[1]}, x_{[2]}, ..., x_{[n]}\right)$ with $x_{[1]} \leq x_{[2]} \leq ... \leq x_{[n]}$ and r referred to as the *ranking number* assigned to each element of x. The order statistic $x_{[n]}$ accordingly represents the sample maximum and $x_{[1]}$ the sample minimum, respectively. The transformed stress indicator (the *stress factor*) z_t is now computed from the raw stress indicator x_t on the basis of its empirical CDF $F_n(x_t)$ as follows:

$$z_t = F_n(x_t) = \begin{cases} \frac{r}{n} \text{ for } x_{[r]} \leq x_t < x_{[r+1]}, & r = 1, 2, ..., n-1 \\ 1 \text{ for } x_t \geq x_{[n]} \end{cases} \tag{1}$$

for $t = 1, 2, \dots, n$. The empirical CDF $F_n(x^*)$ is simply computed as the number r^* of observations x_t not exceeding a particular value x^*, divided by the total number of observations in the sample (Spanos, 1999, p. 230f.). The value x^* thus represents the (r/n)-th sample quantile

(Koenker, 2005, p. 6). Equation 1 has to be modified in case of tied observations. When there are two or more identical values in the data, each of the observations involved is assigned the average of the ranks that they would have otherwise occupied. The empirical CDF is hence a function which is non-decreasing and piecewise constant with jumps being (multiples of) $1/n$ at the observed points. The transformation projects raw stress indicators into variables which are unit-free and measured on an ordinal scale with range $(0, 1]$. The transformation thus trades off gains in statistical robustness against losses of the marginal information contained in the original cardinal scale of the raw stress indicators.

Equation 1 does not yet feature the intended "real-time character" of the CISS. It is introduced by applying the transformation recursively over expanding samples. Precisely, the non-recursive transformation as defined in equation 1 applies to all observations from the pre-recursion period running from 8 January 1999 to 4 January 2002. All subsequent observations are transformed recursively on the basis of ordered samples recalculated with one new observation added at a time:

$$
z_{n+T} = F_{n+T}(x_{n+T}) = \begin{cases} \frac{r}{n*T} \text{ for } x_{[r]} \leq x_{n+T} < x_{[r+1]}, & r = 1, 2, ..., n-1, ..., n+T-1 \\ 1 \text{ for } x_{+T} \geq x_{[n+T]} \end{cases} \tag{2}
$$

for $T = 1, 2, \ldots, N$, with N indicting the end of the full data sample (here 24 June 2011). The total number of observations included in the ordered samples varies from indicator to indicator depending on the availability of historical data. The longest sample starts in 4 January 1980 (see Table 1) with the total number of observations included in the pre-recursion sample amounting to 1149, while the shortest pre-recursion sample starting in 8 January 1999 is left with 157 observations.

Transforming the raw stress indicators according to equations 1 and 2 on the basis of euro area data broadly confirms our presumption of robustness (see Figure A.1 in Hollo, Kremer

and Lo Duca, 2012, which displays the transformation of all 15 raw stress indicators computed both recursively and non-recursively). In most cases the differences between the empirical CDFs calculated in real-time and those computed from the full data sample are relatively small. While in a few cases the differences become somewhat more pronounced, they are still rather modest and thereby also contribute to the strong robustness of the composite indicator against variations of the sample length (see Section 4.1).

We are now equipped with a set of 15 homogenised stress factors $z_{i,j,t}$, with $i = 1, 2, \ldots 5$ indicating the respective market segment and $j = 1, 2, 3$ denoting the stress factors within each subindex i. The five subindices of financial stress are calculated as the arithmetic mean of the three constituent stress factors:

$$s_{i,t} = \frac{1}{3} \sum_{j=1}^{3} z_{i,j,t}. \tag{3}$$

We postpone the discussion of this presumably inconsistent choice of intra-subindex aggregation to the next section, as it requires an understanding of the portfolio-theoretic approach to the aggregation across subindices.

3.3 Aggregation of subindices into the composite indicator

The main innovative element of the CISS compared to alternative financial stress indicators is the application of standard portfolio theory to the aggregation of subindices. The portfolio-theoretic framework offers two elementary avenues to incorporate systemic risk aspects. First, analogously to the computation of portfolio risk from the risk of individual assets, the five subindices of segment-specific stress are aggregated by taking into account the cross-correlations between them. It is essential for our purpose that we allow for time-variation in the cross-correlations. In this way the CISS puts relatively more weight on situations in which stress prevails in several market segments at the same time, i.e. on situations in which financial instability spreads more widely across the financial system. The correlations thus focus on

capturing the systemic dimension of stress within the financial system (the "horizontal view" on systemic risk as defined by De Bandt and Hartmann, 2000). Second, the weights assigned to each subindex in the composite indicator can be calibrated in proportion to their systemic importance which in turn may be approximated in different ways. For instance, the weights may mirror the relative size of the financial market segment covered by each subindex (*size weights*) as in Illing and Liu (2006) and Oet et al. (2011). Alternatively, the weights may be determined on the basis of some reduced-form estimates of the relative importance of (stress in) each subindex for economic activity (*real-impact weights*), a route which has not yet been pursued in the literature. In both cases, the calibration of weights provides an opportunity to account for country differences in the structure of financial systems and the associated differences in the transmission of financial stress to the real economy, thereby capturing the "vertical view" on systemic risk which gives an idea about the potential real costs of a financial crisis (De Bandt and Hartmann, 2000). Since such structural features of an economy are not set in stone, the weights can in principle also vary over time.

Against this background, the CISS is computed in general terms according to equation 4:

$$CISS_t = (w_t \circ s_t)C_t(w_t \circ s_t)', \tag{4}$$

with $w_t = (w_{i,t})$ a 1×5 vector of subindex weights and $s_t = (s_{i,t})$ a 1×5 vector of subindices with $i = 1, ..., 5$; $w_t \circ s_t$ the Hadamard-product of both vectors; and C_t the symmetric 5×5 matrix collecting the time-varying cross-correlation coefficients $\rho_{ij,t}$ between subindices i and j. Just as its constituent stress factors, the CISS is a continuous, unit-free indicator bounded by the half-open interval $(0, 1]$.

Estimation of cross-correlations. In this paper, the time-varying cross-correlations $\rho_{ij,t}$ are recursively computed as exponentially-weighted moving averages (EWMA) of subindex co-

variances $\sigma_{ij,t}$ and variances $\sigma_{i,t}^2$ as approximated by the formulas collected in equation 5:

$$\sigma_{ij,t} = \lambda\sigma_{ij,t-1} + (1-\lambda)\widetilde{s}_{i,t}\widetilde{s}_{j,t}$$

$$\sigma_{i,t}^2 = \lambda\sigma_{i,t-1}^2 + (1-\lambda)\widetilde{s}_{i,t}^2 \tag{5}$$

$$\rho_{ij,t} = \sigma_{ij,t}/(\sigma_{i,t}\sigma_{j,t})$$

for $i = 1,\ldots,5$, $j = 1,\ldots,5$, $i \neq j$, $t = 1,\ldots,N$ with $\widetilde{s}_{i,t} = (s_{i,t} - 0.5)$ denoting de-meaned subindices (obtained by subtracting the theoretical mean of 0.5 rather than the sample-dependent empirical mean). EWMA is used by many practitioners for forecasting daily or weekly conditional asset price volatilities and correlations (see Cuthbertson and Nitzsche, 2004; González-Rivera, Lee and Yoldas 2007). In sufficiently large samples the above formulas approximate well the true infinite exponentially weighted moving averages. The decay factor or smoothing parameter is held constant through time at a value of 0.93. (This value roughly equals the average smoothing parameter estimated over expanding samples within a five-dimensional IGARCH model for the demeaned subindices.) The covariances and variances are initialised (at $t = 0$, i.e. 1 January 1999) at their average values over the pre-recursion period 8 January 1999 to 4 January 2002. Figure 1 displays the EWMA-estimates of all the cross-correlations between the five subindices of the euro area CISS.

Since the raw stress indicators are transformed by means of order statistics, the cross-correlations can be broadly interpreted as a time-varying variant of Spearman's rank correlation coefficient. The cross-correlations thus simply indicate whether the historical ranking of the level of stress in two market segments is relatively similar or dissimilar in any point in time.

Calibration of the subindex weights. For the present purpose, the weights attached to each stress subindex are calibrated in proportion to their average differential impact on real economic activity in the euro area. The real impacts are estimated using two different econometric approaches, where the results from in each case four (slightly) different model

Figure 1: Cross-correlations between subindices of the CISS. Correlation pairs are computed as exponentially-weighted moving averages with smoothing parameter $\lambda = 0.93$. The cross-correlations are labelled as follows: 1 – money market, 2 – bond market, 3 – equity market, 4 – financial intermediaries, 5 – foreign exchange market. Weekly euro area data from 8 Jan. 1999 to 24 June 2011.

specification are combined to ensure some degree of robustness. We also hold the weights constant over time, implicitly assuming that the structural features of the euro area financial system which determine the way financial stress is transmitted to the real economy have not undergone dramatic changes over the relatively short sample considered.

We first run conventional bivariate VARs with industrial production and one of the subindices of stress as endogenous variables. The models are estimated on the basis of monthly data (monthly averages for stress indices) with a uniform optimal lag order of two as suggested by standard selection criteria. Two model variants differ only in their respective sample length, with one starting in January 1987 (i.e., including pre-EMU data) and the other one in January 1999 (when the euro was introduced). The two remaining VAR specifications differ in the transformation of the industrial production data (log level and its 12th difference, respectively). We compute cumulated 24-month structural impulse responses of industrial production to a unit shock in each stress subindex. Structural identification of shocks is obtained by applying

20

the usual Cholesky decomposition to the variance-covariance matrix of reduced-form residuals, with the stress subindex ordered first and industrial production second (for a justification of this ordering see Section 4.3). The subindex weights associated with each model variant are then determined as each subindex's share in the sum of cumulated impulse responses across the five subindices.

Linear VARs, however, only measure the mean impact of financial stress on industrial production, since the impulse response functions are computed from the models' estimated (via least squares) conditional mean functions. While this can make perfect sense under many circumstances, it might be less suitable in the present context. For instance, financial crises are rare events often associated with unusually severe output losses. This may recommend to focus on the dependence structure in the lower tails of the conditional distributions when calibrating the subindex weights. Least squares regressions may also provide biased coefficient estimates because of the influence from extreme values in the data brought about by episodes of severe financial stress. Against this background, we also perform single-equation quantile regressions as introduced by Koenker and Bassett (1978). Quantile regressions are based on minimizing asymmetrically weighted absolute residuals and are more robust to extreme values and other forms of non-normality in the residual distributions than least squares. Resembling the set up of the VARs, we regress the annual growth in industrial production ($y_t = \Delta_{12} \log IP_t$) on each one of the subindices of financial stress (with lags 0 to 2) along with the lagged endogeneous variable (lags 1 and 2). This gives rise to the following linear conditional quantile functions estimated for all stress subindices $s_{i,t}$, $i = 1, \ldots, 5$, and for the full range of regression quintiles $\tau = 0.05, 0.10, ..., 0.95$:

$$Q_{y_t}(\tau|\Omega_t) = \beta_0(\tau) + \beta_1(\tau)y_{t-1} + \beta_2(\tau)y_{t-2} + \beta_3(\tau)s_{i,t} + \beta_4(\tau)s_{i,t-1} + \beta_5(\tau)s_{i,t-2} \quad (6)$$

with $\Omega_t = (y_{t-1}, y_{t-2}, s_{i,t}, s_{i,t-1}, s_{i,t-2})$ the conditioning information set available at time t, and

$t = 3, ..., T$. Estimating equation 6 for all τ yields a set of coeffients characterising the entire distribution of industrial production conditional on each subindex of financial stress. Figure 2 plots the coefficient sums $\overline{\beta}(\tau) = \beta_3(\tau) + \beta_4(\tau) + \beta_5(\tau)$ for the five subindices against the whole range of τ-values. The coefficient sums summarise the long-term impact of subindex stress on economic activity. Some notable features emerge from the coefficient plots: i) In a textbook-style fashion, all estimated impact functions are upward sloping, i.e. the coefficient sums tend to increase for higher quantiles. The strongest negative impacts - in all cases statistically significant - accordingly materialise in the lowest quantiles, in line with what one would expect from a financial crisis point of view. ii) Around median quantiles the estimated coefficient sums become uniformly very small in absolute terms and lose their statistical significance. This indicates that economic activity becomes unrelated to our measures of segment-specific financial stress during periods of "normal" growth. iii) The coefficient sums turn positive - but in only one case statistically signfiant at the 95% confidence level - at the highest quantiles. This may suggest that economic boom periods tend to be associated with somewhat higher uncertainty and risk aversion among financial market participants, possibly reflecting investors' growing concerns about the nature and duration of the boom and the eventual responses of (monetary) policy makers. iv) The impact functions look rather similar across subindices. This notwithstanding, each subindex still possesses some independent predictive power for economic activity. For this purpose we also run quantile regressions pooling all five subindices (with lags 0 to 2) as regressors. It results that all of them retain independent and statistically significant exlanatory power in particular at lower quantiles (results not shown).

The real-impact weights are determined from two sets of quantile regressions which differ in the specification of the dependent variable as in the case of the VARs, namely industrial production in log levels and in annual log growth. Moreover, in line with the notion of systemic stress, we focus attention on the lower regression quantiles. More precisely, we compute the

Figure 2: Quantile regressions for annual log growth of industrial production. The lines represent the sum of slope coefficients (for lags 0 to 2) of the financial stress subindices estimated for 19 equally spaced quantile functions.

real-impact weights for both specifications in two ways: we first calculate the average coefficient sums from the 5th to the 30th regression quantiles and determine the relative share of each subindex in the overall sum; the second method computes the weights from each subindex' maximum absolute impact within the same range of lower quantiles.

The VARs and the quantile regressions thus provide eight different measures of the subindex weights. (Detailed results from the VARs and the quantile regressions are available upon request.) Averaging across this set of weights leads to the following (rounded) subindex weights applied in the empirical part of this paper: 19 % money market, 22% bond market, 14% equity market, 25% financial intermediaries, and 20% foreign exchange market ($w_t = \overline{w} = (0.19, 0.22, 0.14, 0.25, 0.20)$ in equation 4) . However, it turns out that the differences in the CISS computed for different sets of real-impact weights are generally minor (see Figure A.2 in Hollo, Kremer and Lo Duca, 2012).

We have now compiled all the ingredients necessary to compute the CISS for the euro area

Figure 3: CISS versus the squared simple weighted-average of subindices (perfect-correlation case). Weekly euro area data from 8 Jan. 1999 to 24 June 2011.

economy according to equation 4. The resulting time series of weekly data from January 1999 is plotted in Figure 3 as the black line. Within the portfolio-theoretic aggregation framework, the square of the simple weighted average of the five subindices, i.e. $\left(\sum_{i=1}^{5} w_i s_{i,t}\right)^2$, emerges as a special case. If all subindices were perfectly correlated all the times, the CISS and the squared weighted average would coincide. The weighted average (the grey line in Figure 3) thus actually serves as an upper bound of the CISS. The CISS and its perfect-correlation counterpart indeed almost overlap when correlations are generally very high. This happened, for instance, in the run-up to the crisis around 2005 at very low levels of the CISS, as well as in the aftermath of the Lehman bankruptcy at very high levels of financial stress (see Figure 1). Most of the time, however, correlations are quite diverse and relatively moderate such that the CISS assumes much lower levels in "normal times" than the simple-average composite indicator. This, in turn, suggests that the CISS concept particularly helps to better identify periods of moderate levels of financial stress compared to conventional financial stress indices which are akin to our squared weighted average of subindices.

The difference between the weighted average of subindices and the CISS can also be used to

derive a decomposition of the CISS into the contributions coming from each of the subindices and the overall contribution from all the cross-correlations. Such a decomposition may appear particularly attractive for regular monitoring exercises as part of the financial stability surveillance functions performed by macro-prudential authorities (see ECB, 2011).

We still owe a discussion of the choice to take arithmetic means of three stress factors to compute the subindices of financial stress. It can be argued that within the portfolio-theoretic framework of the CISS, the arithmetic means imply perfect correlation between all three subindex components and thus run counter to our idea of stress factors providing complementary information. This inconsistency notwithstanding, the arithmetic mean also has its respective merits for our purposes. For instance, within our three-tier aggregation framework, applying EWMA-based correlation-weights also *within* subindices would smooth out the CISS excessively because of the double-smoothing entailed by applying correlation-weights also *between* subindices at the subsequent final tier. In addition, data limitations would in many cases obviate the application of portfolio-weighting or PCA within subindices; it often happens, for example, that one only finds one or two constituent stress factors to populate a certain subindex when composing a CISS for economies with less developed financial systems, but also for advanced countries in the past when data coverage was thinner.

4   Evaluation of the euro area CISS

Evaluating the performance of FSIs is an inherently complicated task. First of all, the CISS, just as any other existing FSI, is far from being an "ideal" composite indicator in the sense that neither the selection of raw stress indicators, their transformation, nor their weighting are determined on the basis of an underlying structural model embedding the concept of systemic risk. The measurement problem is aggravated by the fuzziness of the concepts of systemic risk and financial stability, the complexity of modern-world financial systems, and the difficulties

in empirically identifying certain features of financial stress. The construction of composite stress indicators thus involves many arbitrary and subjective choices. Any FSI therefore limits attention to only a few segments of the financial system, and draws on a broad array of largely imperfect measures of financial stress. In addition, reflecting the fact that financial crises are "rare events" (according to Reinhart and Rogoff, 2009, crises occur on average about once every five years or so world wide), the data samples of FSIs are typically rather short, covering merely a few crisis episodes, a fact which severely impairs the statistical reliability of empirical analyses. The discrepancy between the degrees of freedom available in constructing and in testing FSIs, respectively, makes it extremely difficult to assess whether a particular indicator performs well both in absolute terms (What is a good indicator?) and in relative terms (Which indicator is better?). Against the background of these caveats, this section evaluates the performance of the CISS on the basis of economic plausibility checks as well as a few statistical/econometric criteria. In order to enlarge the set of historical crisis-like episodes, we base the analysis on a version of the euro area CISS extended backward until January 1987, i.e. including 12 years of data from the pre-EMU period (for details see Hollo, Kremer and Lo Duca, 2012).

4.1 Robustness

The signals issued by any FSI should be stable over time in order to avoid the so-called event reclassification problem. For instance, assume that in a particular point in time an indicator suggests that the prevailing level of stress is unusually high by historical standards. It is then desirable that the indicator still classifies this period as a particularly stressful episode say ten years hence, i.e. when ten years of data are added to the sample. Otherwise no robust historical comparison can be made, and the calculation of certain threshold levels for the indicator would not make sense either.

In order to limit the event reclassification problem from the outset, we opt for a procedure

Figure 4: It shows two variants of the (backward-extended proxy) CISS. The first one is the proper CISS computed recursively as described in the body of the text; the recursion starts in Jan. 1990. The second series shown uses full-sample information in the transformation of raw stress indicators. Reflecting the robustness of the CISS against expanding samples, the differences between both series are very small in general. Weekly euro area data from 2 Jan. 1987 to 24 June 2011.

that transforms the raw indicators based on order statistics. Figure 5 illustrates the robustness of the (backward-extended) CISS when computed recursively (black line; recursion starting in January 1990!) and non-recursively (grey line) based on the full sample information, respectively. The two time series track each other remarkably closely. The average absolute difference amounts to only 0.015 (standard deviation: 0.022) with a mean error of 0.010. The largest deviation between the two differently computed indicators occurs in February 2008 with a value of 0.076. We therefore conclude that the CISS is a markedly robust statistic in the time dimension, implying that it is hardly affected by the event reclassification problem.

As a second statistical robustness check, we compute the CISS for a range of values of the smoothing parameter λ that governs the speed at which the cross-correlations adjust to latest information. In Hollo, Kremer and Lo Duca (2012) we compare the time series for three λ-values, namely 0.89, 0.93, and 0.97 (see their Figure 7). As expected, the CISS with the lowest smoothing parameter displays wider swings, and it spikes somewhat more pronouncedly in response to large

27

stress shocks than our preferred CISS with an intermediate λ-value of 0.93. Conversely, setting the smoothing parameter to a higher level produces a CISS with dampened swings and spikes. All in all, however, the differences produced by different smoothing parameters are relatively low and, importantly, they do not alter the general pattern of behaviour of the CISS. Its basic information content, namely the broad classification of financial stress events or regimes, thus also remains robust.

4.2   Identification of stress events

The most widely adopted evaluation criterion for financial stress indicators is their performance in identifying well-known past episodes of financial stress. Illing and Liu (2006) developed the event-based evaluation criterion into a probabilistic evaluation framework employed to decide which financial stress indicator performs best among a broader set of candidates. The evaluation framework rests on a survey of experts to define the most "critical" stress events for the Canadian economy out of 40 pre-selected potentially stressful events since the early 1980s. Their preferred financial stress indicator is the one which matches best the survey results balancing Type I errors (failure to report a high-stress event) against Type II errors (falsely reporting a high-stress event).

While the event-based criterion appears rather obvious and straight-forward, it suffers from conceptual and measurement problems. First, in a certain sense it relies on knowing a priori what the indicator is supposed to identify in the first place, namely systemic financial stress. Second, in particular when the data of the stress index is available at a higher frequency (daily or weekly), the criterion also requires knowing when the stress began and, even more difficult, when it subsided. Third, mere focus on popular incidents a priori excludes stressful periods which cannot be associated with specific triggering events but which rather build up gradually over time as a result of cumulated smaller pieces of bad news. The "dot-com" boom and bust episode

around the turn of the millennium may exemplify such a case. Hence, evaluation approaches relying on "crises defined by events" are likely to miss such more hidden periods of systemic stress, while "crises defined by quantitative thresholds" determined on the basis of financial stress indicators are less prone to such Type I errors (on these two crises definitions see Reinhart and Rogoff, 2009). In the light of these problems, we argue against relying too strongly on a formalised version of the event-criterion when studying the performance of FSIs.

We rather prefer a narrative approach like in Hakkio and Keeton (2009) to find out whether peaks in the CISS can be plausibly associated with well-known crisis events. Figure 5 illustrates that the sharpest spikes in the CISS indeed tend to occur around very popular events which caused, at least temporarily, severe stress in the global financial system (for a full account of the most important stress events identified by the CISS see Hollo, Kremer and Lo Duca, 2012). The first major stress event in the sample is the stock market crash in October 1987. On October 19, the US stock market experienced its largest one-day loss in market valuations ever, causing extreme stress in the financial industry worldwide. However, stress subsided relatively quickly when market participants realised that financial firms had been able to remain financially sound (Cardarelli, Elekdag and Lall, 2011). About five years later the European financial system was shaken by the collapse of the European Exchange Rate Mechanism (ERM). Tensions in the currency markets culminated in the British Pound and the Italian Lira eventually withdrawing from the ERM on September 16 and 17, 1992, respectively. But the financial turmoil caused by the ERM crisis again turned out rather short-lived with the CISS reverting quickly back to pre-crisis levels. It took another six years for financial stress to return to Europe in the context of the global market reactions to the Russian debt moratorium in August 1998 and the subsequent collapse of the hedge fund Long-Term Capital Management (LTCM) in September 1998.

The next period of elevated stress appears to be closely related to the above-mentioned downturn in high-tech stocks in early 2000. More widespread tensions occurred in the wake of

Figure 5: CISS and major financial stress events in the global and European financial markets. Weekly euro area data from 8 Jan. 1999 to 24 June 2011.

the strong initial losses in the high-tech segment. The CISS remained relatively high in general over the subsequent two years fed by the continued "crash in instalments" in technology stocks (by October 2002, the NASDAQ had lost about 75% of its peak level in early March 2000) and recessions in core parts of the global economy. The terrorist attacks in the US on September 11, 2001, caused a sharp abrupt increase in the CISS in between. Investors soon realised, though, that their initial fears about the potential financial and real economic impacts of the attacks were exaggerated such that the global financial system recovered relatively quickly from this severe shock.

However, none of those previous events pushed the CISS towards similarly high levels reached during the most recent financial and economic crisis. The CISS first signalled an extreme level of stress in August 2007, when BNP Paribas suspended three investment funds that invested in asset backed securities linked to subprime mortgage debt which had become virtually illiquid. Spreading announcements of severe losses incurred by banks, mortgage lenders and other financial institutions lifted the CISS further up, and it peaked again in response to the collapse of Bear Stearns in March 2008. The CISS experienced its largest jump in September 2008 when

Lehman Brothers filed for bankruptcy protection and AIG was rescued to avoid bankruptcy. The index reached its historical maximum in November 2008 when the US plan to buy toxic assets under the Troubled Asset Relief Program (TARP) was abandoned, which undermined global market confidence. After November 2008 the CISS signalled a steady decline in financial stress until mid-April 2010 when serious concerns about sovereign credit risk in the euro area emerged. (Most recent updates of the CISS can be downloaded from the ECB homepage as an electronic annex to Hollo, Kremer and Lo Duca, 2012).

To sum up, it appears that all extreme peaks in the CISS can be associated with specific financial stress events, suggesting that it does not suffer from type II errors. It is harder to judge whether it also performs well on the dimension of type I errors, i.e. whether there are severe crises which it failed to indicate. Potential candidates in this regard are the global bond market crisis and the Mexican peso crisis both in 1994 and the Asian crisis in 1997, for instance. The CISS suggests that these events did not trigger significant systemic stress in the euro area financial system as a whole, but rather represented more isolated tensions in specific market segments and other parts of the world economy. This view is broadly consistent with findings from the international contagion literature (e.g., Bekaert, Harvey and Ng, 2005). Overall, developments in the CISS appear in general rather plausible, not least because it singles out very clearly the recent financial and economic crisis as the by far most stressful period over the past quarter of a century of available data for the euro area, comparable probably only to the Great Depression.

4.3 A threshold VAR to identify systemic crises

One of the main objectives of FSIs is to help policymakers identifying stress levels in the financial system that may be of serious concern. However, truly systemic stress levels which might eventually disrupt the process of financial intermediation and thereby economic activity, can not be easily identified. The literature suggests several ways to tackle this problem. One approach

is to benchmark the current level of stress against levels observed during historical crises known to have caused such serious disruptions.

An alternative is to identify quantitative thresholds or regimes for the level of the financial stress indicator based on statistical or econometric methods. The most widely used approach is to classify financial stress as severe if the index exceeds its historical mean by one or more standard deviations (e.g., Illing and Liu, 2006; Cardarelli, Elekdag and Lall, 2011). This approach, however, manifests several shortcomings. First, it implicitly assumes that the stress indicator is normally distributed, a presupposition which is clearly violated in the case of the CISS (a histogram of the CISS is shown as Figure 9 in Hollo, Kremer and Lo Duca, 2012). Violating this assumption might exacerbate the issue of temporal instability in the conditional means and standard deviations of the financial stress index in smaller data samples which, in turn, might give rise to the event reclassification problem discussed above. This problem might be particular pronounced in the present case since in times of crisis, the new data added to the sample usually take on extreme values. Furthermore, this approach also suffers from the ad hoc nature of the identified threshold, in the sense that it is not obvious how many standard deviations the index should exceed its mean in order to signal serious stress.

To overcome some of these shortcomings, we propose applying econometric regime-switching models which endogenously identify periods of extreme financial stress. The basic idea behind such approaches is that the dynamics of the financial system and its interactions with the real sector may be subject to multiple equilibria depending on whether the economy is in a state of financial crises and non-crises, respectively (Hansen, 2000). This may reflect the fact that the interaction between externalities (e.g., contagion), information problems (e.g., adverse selection) and certain special features of the financial sector (e.g., the existence of maturity mismatches and high leverage) can lead to powerful feedback and amplification mechanisms driving the system from a state of relative tranquillity to a state of turmoil, also altering the system's normal laws

of motion. In order to identify such regime changes, Hollo, Kremer and Lo Duca (2012) identify three different level-regimes of the euro area CISS based on an autoregressive Markov-switching model.

In this paper we focus on a threshold VAR (TVAR) model that captures in a stylised fashion regime-dependent dynamic interactions between financial stress and the real economy. The regimes are identified on the basis of an estimated threshold of the CISS based on the idea that financial stress becomes a cause of major concern when it adversely impacts on the real economy, thereby integrating the "vertical view" of our favoured definition of systemic risk. According to this viewpoint, we would expect that economic activity slows down or drops sharply whenever the CISS reaches a certain critical level. One major advantage of such an estimated threshold level of the CISS and the corresponding regime classification consists in its direct economic interpretation.

In general, threshold regression models represent a class of regime-switching which assumes that state transitions are triggered any time an observable variable crosses a certain threshold level (Franses and van Dijk, 2000). We assume a priori that the CISS is the relevant threshold variable and that at most two regimes and therefore one single threshold exists. We follow Tsay (1998) and identify potential threshold effects within a bivariate TVAR with the CISS (C_t) and annual growth in industrial production (y_t) as the endogenous variables. Anticipating a shortage of degrees of freedom in the high-stress regime recommends a specification of the TVAR as parsimonious as possible. Hence, we also opt for the shortest lag-order suggested by standard specification tests. While information criteria (weakly) prefer a higher lag order (four lags), an exclusion F-Test suggests that a VAR with two lags may suffice. The basic regression

setup is as follows:

$$x_t = \alpha^H + \Phi_1^H x_{t-1} + \Phi_2^H x_{t-2} + \epsilon_t^H \quad \text{if} \quad z_{t-d} > \tau \tag{7}$$

$$x_t = \alpha^L + \Phi_1^L x_{t-1} + \Phi_2^L x_{t-2} + \epsilon_t^L \quad \text{if} \quad z_{t-d} \leq \tau \tag{8}$$

with $x_t = (C_t, y_t)'$ a two-dimensional vector; α^s and Φ_j^s the vector of intercepts and the two matrices collecting the slope coefficients, respectively, for states $s = \{H, L\}$ (with H and L standing for high-stress and low-stress regimes, respectively) and lags $j = \{1, 2\}$. The threshold variable is denoted z_{t-d} with $d \in \{1, ..., d_0\}$ and $d_0 = 2$ the maximum threshold lag or *delay* foreseen. The threshold parameter is labelled τ and the vector ϵ_t^s contains the state-dependent regression errors with variance-covariance matrices Σ^s. As mentioned above, the once or twice lagged CISS plays the role of the threshold variable exciting the switches in regimes any time it crosses the threshold τ.

Tsay (1998) proposed a two-step conditional least squares procedure to estimate this TVAR under the assumption that the lag order, the number of states and the threshold variable are all known. It is furthermore assumed that z_{t-d} is stationary and continuous with a positive density function on a bounded subset of the real line. As the first step, for given d and τ, the model parameters α^s, Φ_j^s and Σ^s can be estimated by ordinary least squares. Given the parameter estimates, Tsay (1998) developed test procedures to determine d and τ simultaneously. The main criterion of the selection procedure is Tsay's $C(d)$-Statistic testing for statistically significant threshold effects in the VAR. The $C(d)$-Statistic is asymptotically chi-squared distributed, and results for $d = 1$ and $d = 2$ (i.e., the once and twice lagged CISS as the threshold variable) are shown in Table 2. In both cases the $C(d)$-Statistic clearly rejects the null hypothesis of no-threshold effects (linear VAR against TVAR) with p-values below the 5%-confidence level. The optimal threshold value for each d is determined by a grid search procedure (over a range of

34

Table 2: Testing for threshold delay and threshold values. d denotes the threshold delay and the threshold value. AIC is the Akaike information criterion. The C(d)-Statistic (p-value shown in the next column) tests for a statistically significant threshold effect in a bivariate VAR with two lags and the CISS and annual growth in industrial production for the euro area as endogenous variables. The F-Statistic tests for the presence of a single threshold in a regression of output growth on a constant, two of its own lags and the CISS with same lag length. Monthly data from Jan. 1987 to June 2011.

	Tsay (1998)-Test				Hansen (2000)-Test		
d	$C(d)$-Stat	p-value	τ	AIC	F-Stat	p-value	τ
1	20.03	0.0166	0.2960	-2741	13.46	0.0000	0.2957
2	19.24	0.0402	0.3233	-2766	12.59	0.0000	0.2747

CISS values) which minimises the Akaike information criterion (AIC). The optimal specification is found to be a TVAR(2) model with the twice lagged CISS ($d = 2$) as the threshold variable and an estimated threshold value of 0.3233 (see Figure 12 in Hollo, Kremer and Lo Duca, 2012, plotting the AIC against different potential threshold values of the CISS). This is suggested by the fact that the AIC is lower for $d = 2$ than for $d = 1$ (see the fifth column in Table 2).

As a robustness check we also perform Hansen's (2000) test for thresholds in a single-equation regression of output growth on a constant, two of its own lags and the CISS with the same lag length. This regression can thus be regarded as one equation of the bivariate TVAR model. Hansen developed an F-Test for the existence of threshold effects. The test results are shown in the last three columns of Table 2, clearly suggesting the existence of statistically significant threshold effects with threshold values very similar to those of the Tsay-procedure.

Equipped with a fully specified and estimated TVAR model we are now in a position to assess whether the effects of the identified threshold of the CISS are both qualitatively and quantitatively consistent with our expectation that particularly high levels of financial stress tend to depress economic activity. Visual inspection of a scatter plot relating output growth to the twice lagged CISS seems to vindicate this expectation (see Figure A.4 in Hollo, Kremer and Lo Duca 2012). While at lower levels of the CISS (non-crisis times) the scatter plot appears purely random, at higher levels of the CISS a clear negative relationship emerges between industrial

black solid line: IRF of high-regime; grey solid line: IRF of low-regime. IRFs with one std. dev. analytical error bands

Figure 6: Regime-dependent impulse response functions (IRF) of output growth to shocks in the CISS from the TVAR(2) model. TVAR(2) denotes the bivariate threshold-VAR model with 2 lags, one threshold (two regimes) and the CISS and annual growth in industrial production for the euro area as endogenous variables. High-stress regime occurs when the CISS (twice lagged) stands at or above a threshold value of 0.3233, and the low-stress regime prevails below the threshold. IRFs for orthogonalised innovations (Choleski factorisation) with shocks in the CISS allowed to have a contemporaneous impact on industrial production growth. IRFs of the low-stress regime rescaled to the typical size (one standard deviation) of CISS innovations from the high-stress regime. Monthly data from Jan. 1987 to June 2011.

production and financial stress.

In order to substantiate this claim further we compute the impulse response functions (IRFs) from the estimated TVAR-coefficients separately for the high-stress and the low-stress regimes. Of course, computing conventional impulse response functions in non-linear VARs ignores their history- and shock-dependence in such setups and are therefore valid only under certain assumptions (Koop, Pesaran and Potter, 1996). Figure 6 displays the two state-dependent IRFs of industrial production growth for a uniform one-standard deviation structural shock in the CISS from the high-stress regime. The dotted lines around the IRFs represent analytical one-standard-deviation error bands (Lütkepohl, 1990). The structural innovations are obtained from the triangular Choleski-factorisation of the variance-covariance matrix of residuals. The endogenous variables are ordered in such a way (CISS first, output second) that shocks in the CISS can

have a contemporaneous impact on economic output but not conversely. This structural shock identification can be justified from an information perspective, for instance. Owing to the lag in the publication of the euro area industrial production index (released in the second third of the second month following the reference month), one may argue that the output innovation of a given month cannot be perfectly predicted by financial market participants, in turn implying that they cannot be fully reflected in contemporaneous asset prices either. In addition, it may appear plausible to assume that CISS shocks tend to originate mainly from within the financial sector particularly during crisis times, and that producers react quickly to increased uncertainty with a rapid drop in aggregate output reflecting a (temporary) pause in their investment and labour hiring decisions (as in Bloom, 2009). However, since our favoured structural identification scheme may not always properly describe the true causal ordering, the IRFs may be better interpreted as an upper bound (in absolute terms) of the output reactions to shocks in the CISS. This notwithstanding, the qualitative results from the impulse-response analysis remain robust to a reverse ordering of variables.

Figure 6 indeed confirms our expectations that the real economic impacts of financial stress are in fact dramatically different across the two regimes. While shocks in the CISS do not exert any statistically and economically significant reactions in output over whatever horizon during low-stress regimes, industrial production virtually collapses in response to a large positive CISS shock in the high-stress regimes. The maximum impact is reached after four months, when annual output has declined by about 2.7% in response to an initial shock in the CISS of 0.06. It takes about a year for the marginal effects to taper off. Similarly, it is only during high stress regimes that, for instance, a negative output shock leads to a subsequent increase in financial stress (see Table 3 and Figure A.5 in Hollo, Kremer and Lo Duca 2012, for the full set of IRFs in the high-stress regime). Taken together, these mutual reaction patterns seem to suggest that when hit by a sufficiently large (financial or real) shock an economy faces the risk of entering a

Table 3: Parameter estimates of the TVAR(2) model. TVAR(2) denotes the bivariate threshold-VAR model with 2 lags, one threshold (two regimes) and the CISS (C_t) and annual growth in industrial production (y_t) for the euro area as endogenous variables. High-stress regime occurs when the CISS (twice lagged) stands at or above the estimated threshold. Estimations based on monthly averages of weekly data from Jan. 1987 to June 2011.

	High-stress regime		Low-stress regime	
	C_t	y_t	C_t	y_t
Constant	0.2171	0.0650	0.0161	0.0019
	(2.8863)	(3.5790)	(3.5123)	(1.1221)
C_{t-1}	0.9296	-0.0938	0.9022	-0.0401
	(5.3965)	(2.2556)	(14.7494)	(1.7711)
C_{t-2}	-0.4106	-0.0749	-0.0652	0.0386
	(1.9639)	(1.4846)	(0.9948)	(1.5936)
y_{t-1}	-2.2133	0.8239	0.1273	0.5827
	(2.6935)	(4.1551)	(0.8282)	(10.2482)
y_{t-2}	1.8947	-0.0932	-0.1073	0.3510
	(2.7918)	(0.5693)	(0.7051)	(6.2343)
σ	0.0594	0.0126	0.0337	0.0125
Exclusion F-tests (p-value)				
lagged C	15.63	8.49	336.69	1.59
	(0.0000)	(0.0014)	(0.0000)	(0.2069)
lagged y	3.90	82.89	0.34	553.33
	(0.0330)	(0.0000)	(0.7093)	(0.0000)

vicious downward spiral with financial and economic stress reinforcing each other over time, a finding which could be explained theoretically by some financial accelerator mechanism (e.g., as in Bernanke, Gertler and Gilchrist, 1999).

In contrast, during normal times with low financial stress the CISS tends to become a negligible quantity as evidenced by the absence of any statistically significant cross-equation relationships in this regime according to standard exclusion F-Tests. Accordingly, the IRFs quantifying the impacts of CISS shocks on output, and of output shocks on the CISS, are basically flat in the low-stress regime such that the bivariate VAR degenerates into a set of two independent autoregressions.

We have to conclude this section with adding some words of caution. Any econometric analysis of financial stress indicators in the time series dimension must suffer from the low number of crisis events and the resulting lack of statistical degrees of freedom. Financial crises

are rare events, and even more so are the truly systemic ones with effects as devastating as in the case of the present crisis. Hence, the results obtained from the threshold VAR are clearly dominated by the dynamics observed during the past four years or so and might therefore have only limited power in predicting the dynamics of any similar future crisis.

5   Conclusions

The recent financial and economic crisis revealed considerable gaps in the theoretical underpinning and the empirical toolkits available to analyse and monitor financial stability in general and systemic risk in particular. Academics and financial authorities all around the globe have been stepping up their efforts to improve the suit of tools and models in this field accordingly. This paper contributes to this branch of literature by proposing a new composite indicator of systemic financial stress called CISS which aims to measure the contemporaneous state of instability in the financial system as a whole; it can therefore be interpreted as a measure of systemic risk which has materialised already. The main distinguishing features of the CISS are its explicit foundation on standard definitions of systemic risk and, as its main methodological innovation, the application of portfolio-theoretic principles to the aggregation of individual financial stress indicators into the composite indicator. We have also proposed a parsimoneous econometric approach to estimate a critical level of the indicator as the endogenous outcome of a threshold VAR. Its statistical robustness to computation over expanding samples ensures that past signals issued by the CISS remain valid also at later points in time. The CISS can be updated quickly on a weekly basis and is thus particularly suitable for real-time surveillance tasks as typically conducted in central banks and other macroprudential authorities.

As to the way forward, several companion projects are ongoing or can be envisaged. For instance, an expansion of the geographical coverage of the CISS promises to lead to a better understanding and assessment of its indicator properties, for instance through econometric

analysis that also exploits the cross-country dimension. In a single-country context, the dynamic interactions between financial stress and the real economy should be more thoroughly investigated within richer non-linear econometric model setups as in Hubrich and Tetlow (2011) and Hartmann et al. (2012). In addition, the development of adequate evaluation criteria for running "horse races" between different financial stress indices would be highly welcome by the profession.

References

[1] Acharya, V. V. and D. Skeie (2011): "A Model of Liquidity Hoarding and Term Premia in Inter-bank Markets", CEPR Discussion Paper No. 8705, December.

[2] Baele, L., G. Bekaert and K. Inghelbrecht (2010): "The Determinants of Stock and Bond Return Comovements", Review of Financial Studies, Vol. 23, No. 6, pp. 2374-2428.

[3] Bekaert, G., C. A. Harvey and A. Ng (2005): "Market Integration and Contagion", Journal of Business, Vol. 78, No. 1, pp. 39-69.

[4] Bernanke, B. S., M. Gertler and S. Gilchrist (1999): "The Financial Accelerator in a Quantitative Business Cycle Framework", in J. Taylor and M. Woodford (eds.), Handbook of Macroeconomics, Vol. 1, Part 3, Elsevier, pp. 1341-1393.

[5] Blix Grimaldi, M. (2010): "Detecting and Interpreting Financial Stress in the Euro Area", ECB Working Paper Series No. 1214, June.

[6] Bloom, N. (2009): "The Impact of Uncertainty Shocks", Econometrica, Vol. 77, No. 3, pp. 623-685.

[7] Brave, S. and R. A. Butters (2011a): "Monitoring Financial Stability: A Financial Conditions Index Approach", Federal Reserve Bank of Chicago, Economic Perspectives, First Quarter, pp. 22-43.

[8] Brave, S. and R. A. Butters (2011b): "Gathering Insights on the Forest from the Trees: A New Metric for Financial Conditions", Mimeo, May.

[9] Brunnermeier, M. and L. Pedersen (2009): "Market Liquidity and Funding Liquidity", Review of Financial Studies, Vol. 22, No. 6, pp. 2201–2238.

[10] Caballero, R. J. and A. Krishnamurthy (2008): "Collective Risk Management in a Flight to Quality Episode", Journal of Finance, Vol. 63, No. 5, pp. 2195-2230.

[11] Caldarelli, R., S. A. Elekdag and S. Lall (2011): "Financial Stress and Economic Contractions", Journal of Financial Stability, Vol. 7, pp. 78-97.

[12] Carlson, M. A., T. King and K. Lewis (2011): "Distress in the Financial Sector and Economic Activity", B.E. Journal of Economic Analysis & Policy, Vol. 11, No. 1, Art. 35.

[13] Chordia, T., A. Sarkar and A. Subrahmanyam (2005): "An Empirical Analysis of Stock and Bond Market Liquidity", Review of Financial Studies, Vol. 18, No. 1, pp. 85-129.

[14] Cuthbertson, K. and D. Nitzsche (2004), "Quantitative Financial Economics: Stocks, Bonds and Foreign Exchange", 2nd edition, Wiley.

[15] Davig, T. and C. Hakkio (2010): "What is the Effect of Financial Stress on Economic Activity?", Federal Reserve Bank of Kansas City, Economic Review, Second Quarter, pp. 35-62.

[16] De Bandt, O. and P. Hartmann (2000), "Systemic Risk: A Survey", ECB Working Paper Series No. 35, November.

[17] European Central Bank (2011), "Special Feature C: "Systemic risk methodologies", Financial Stability Review, June, pp. 141-148.

[18] Feldhütter, P. and D. Lando (2008): "Decomposing Swap Spreads", Journal of Financial Economics, Vol. 88, No. 2, pp. 375-405.

[19] Fostel, A. and J. Geneakoplos (2008): "Leverage Cycles and the Anxious Economy", American Economic Review, Vol. 98, No. 4, pp. 1211-1244.

[20] Franses, P. H. and D. van Dijk (2000): "Non-Linear Time Series Models in Empirical Finance", Cambridge University Press, Cambridge.

[21] González-Rivera, Lee and Yoldas (2007): "Optimality of the RiskMetrics VaR Model", Mimeo.

[22] Hakkio, S. C. and W. R. Keeton (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, pp. 5-50.

[23] Hansen, B. E. (2000): "Sample Splitting and Threshold Estimation", Econometrica, Vol. 68, No. 3, pp. 575-603.

[24] Hartmann, P., K. Hubrich, M. Kremer and R. J. Tetlow (2012): "Widespread Instabilities and the Macroeconomy - Regime Switching in the Euro Area", European Central Bank and Federal Reserve Board, Mimeo.

[25] Hatzius, J., P. Hooper, F. Mishkin, K. L. Schoenholtz and M. W. Watson (2010): "Financial Conditions Indexes: A New Look after the Financial Crisis", NBER Working Paper 16150, July.

[26] Heider, F., M. Hoerova and C. Holthausen (2009): "Liquidity Hoarding and Interbank Market Spreads: The Role of Counterparty Risk, European Banking Center Discussion Paper No. 2009-11S.

[27] Holló, D., M. Kremer and M. Lo Duca (2012): "CISS - A Composite Indicator of Systemic Stress in the Financial System," European Central Bank Working Paper Series, No. 1426, March.

[28] Hubrich, K. and R. J. Tetlow (2010): "Financial Stress and Economic Dynamics: The Transmission of Crises", European Central Bank, Federal Reserve Board and IMF, Mimeo.

[29] Illing, M. and Y. Liu (2006): "Measuring Financial Stress in a Developed Country: an Application to Canada", Journal of Financial Stability, Vol. 2, No. 4, pp. 243-265.

[30] Koenker, R. (2005): "Quantile Regression", Cambridge University Press, New York.

[31] Koenker, R. and G. Bassett (1978): "Regression Quantiles", Econometrica, Vol. 46, No. 1, pp. 33-50.

[32] Koop, G., M. H. Pesaran and S. M. Potter (1996): "Impulse Response Analysis in Nonlinear Multivariate Models", Journal of Econometrics, Vol. 74, pp. 119-147.

[33] Krishnamurthy, A. (2010): "How Debt Markets Have Malfunctioned in the Crisis", Journal of Economic Perspectives, Vol. 24, No. 1, pp. 3–28.

[34] Krishnamurthy, A. and A. Vissing-Jorgensen (2010): "The Aggregate Demand for Treasury Debt", Mimeo, April.

[35] Li, F. and P. St-Amant (2010): "Financial Stress, Monetary Policy, and Economic Activity", Bank of Canada, Working Paper 2010-12, May.

[36] Lo Duca, M. and T.A. Peltonen (2011): "Macro-Financial Vulnerabilities and Future Financial Stress – Assessing Systemic Risks and Predicting Systemic Events", ECB Working Paper Series No. 1311, March.

[37] Lütkepohl, H. (1990): "Asymptotic Distributions of Impulse Response Functions and Forecast Error Variance Decompositions of Vector Autoregressive Models", Review of Economics and Statistics, Vol. 72, No. 1, pp. 116-125.

[38] Mallick, S. K. and R. M. Sousa (2011): "The Real Effects of Financial Stress in the Eurozone", Mimeo, March.

[39] Misina, M. and G. Tkacz (2009): "Credit, Asset Prices, and Financial Stress", International Journal of Central Banking, Vol. 5, No. 5, pp. 95-122.

[40] Nelson, W. R. and R. Perli (2007), "Selected Indicators of Financial Stability", Irving Fisher Committee's Bulletin on Central Bank Statistics, Vol. 23, pp. 92-105.

[41] Oet, M. V., R. Eiben, T. Bianco, D. Gramlich and S. J. Ong (2011): "The Financial Stress Index: Identification of Systemic Risk Conditions", Federal Reserve Bank of Cleveland, Working Paper 11-30, November.

[42] Pastor, L. and P. Veronesi (2009): "Learning in Financial Markets", NBER Working Paper No. 14646, January.

[43] Patel, S. and A. Sarkar (1998), "Stock Market Crises in Developed and Emerging Markets", Financial Analysts Journal, Vol. 54, No. 6, pp. 50-59.

[44] Reinhart, C.M. and K.S. Rogoff (2009), "This Time is Different. Eight Centuries of Financial Folly", Princeton University Press, Princeton and Oxford.

[45] Stuart, A. and J. K. Ord (1994), "Distribution Theory", Kendall's Advanced Theory of Statistics, Vol. 1, 6th edition, London and New York.

[46] Spanos, A. (1999): "Probability Theory and Statistical Inference. Econometric Modeling with Observational Data", Cambridge University Press, Cambridge.

[47] Tsay, R. S. (1998): "Testing and Modeling Multivariate Threshold Models", Journal of the American Statistical Assocation, Vol. 93, No. 443, pp. 1188-1202.

[48] Van Roye, B. (2011): "Financial Stress and Economic Activity in Germany and the Euro Area", Kiel Working Paper No. 1743, November.

[49] Veronesi, P. (2004): "The Peso Problem Hypothesis and Stock Market Returns", Journal of Economic Dynamics and Control, Vol. 28, pp. 707–725.