BANK OF ENGLAND

Staff Working Paper No. 636

Assessing vulnerabilities to financial shocks in some key global economies

Jack Fisher and Lukasz Rachel

December 2016

BANK OF ENGLAND

Staff Working Paper No. 636

Assessing vulnerabilities to financial shocks in some key global economies

Jack Fisher[(1)] and Lukasz Rachel[(2)]

Abstract

This paper describes a quantitative, data-driven method to assess vulnerabilities in a range of countries. We provide country-level vulnerability indices which can be used to gauge the level of fragility at any point in time. In particular, our results suggest that in the run-up to the Global Financial Crisis, vulnerabilities rose to extremely high levels in the United States, but were only a little above average in Europe and had actually receded across much of Asia. The picture has changed dramatically during the recovery, however, with vulnerabilities close to record-highs by end-2015 in some of the Asian economies. We document numerous practical challenges that arise when developing such a toolkit, the main one being to know the trend — the 'neutral' level — of a financial variable (for example credit-to-GDP). In that context, one important contribution of this paper is to document the robustness of vulnerability measures to different judgements about the trend level of financial variables. We find that for most countries results are fairly robust to different views of the underlying trend, but importantly that this robustness is not universal. In particular, at the moment differing views of what 'the new normal' is suggest dramatically different assessments of the level of fragility in the United States and South Korea.

Key words: Financial vulnerabilities, risks, crisis.

JEL classification: G10, G15, G17.

---

(1) Email: j.w.fisher@lse.ac.uk

(2) Bank of England and London School of Economics. Email: lukasz.rachel@bankofengland.co.uk and l.p.rachel@lse.ac.uk.

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees. We are grateful to colleagues from the Financial Stability Strategy and Risk Directorate at the Bank for helpful discussions and comments.

Information on the Bank's working paper series can be found at
www.bankofengland.co.uk/research/Pages/workingpapers/default.aspx

Publications Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email publications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

1. Introduction

International risk assessment is critical to the pursuit of monetary and financial stability, particularly in the UK. Modern economies are tightly interconnected. Events and shocks in one region spill over to other places through financial and trade linkages, while common shocks can reveal and propagate through vulnerabilities in many countries simultaneously. Policymakers, financial market participants and researchers have long been aware of these interdependencies, yet the global nature of the 2008 financial crisis has been a reminder of their importance. Naturally, the crisis has brought the international analysis of risks and vulnerabilities to the fore. The UK is a very open economy even by the standards of our interconnected world. Trade into and out of the UK is worth more than 50% of GDP, and London's role as a global financial centre is well known. As a result, the UK's prospects are particularly sensitive to events elsewhere (see by Chowla et al[1]). Consequently, it is crucial to keep a close eye on economic movements abroad so that policy may be set appropriately at home. The key to achieving this is a robust international risk assessment framework - a comprehensive set of tools and methods that permit systematic analysis of the global risk environment.

An evaluation of international risks should reflect the complex nature of cross-country linkages, the multifaceted structure of risks, and uncertainty surrounding the transmission of shocks. As a result, there should be several key elements to any international risk assessment framework. These include detailed country-specific analysis of prospects, risks and vulnerabilities; research and analysis of common global trends and shocks; and system-wide stress-testing exercises.

In this paper we document only one key element to an international risk assessment framework, namely the *country-specific, data-driven assessment of vulnerabilities*. These country-based assessments can be thought of as an initial input into the analysis of global risks. They consist of a set of easily-updateable indicators and a range of graphical presentation tools that deliver a purely quantitative assessment of the economic outlook. Such an approach allows for analysis of risks consistently over time, and aids the process of identifying the most pressing issues which require further investigation. In this paper, drawing on recent research, we describe a methodology that can be employed to implement such an assessment. We describe the various trade-offs that naturally arise when designing the toolkit, so as to provide a convenient resource for practitioners and to spur further debate and research in this area.

The risk indicators developed below are designed to incorporate a wide range of variables, and thus to capture a breadth of information. This broad approach reflects recognition of the complexity of modern economies, which means that no single indicator is capable of accurately and fully describing vulnerabilities that may exist in a macro-financial system. It is also driven by a desire to limit any hindsight bias, whereby indicators and variables that appear to have been important in previous crises tend to garner focus in risk assessment analysis, thus, reducing the chance of catching new risks that may emerge in the system.

We impose a loose structure on the data by hypothesising the existence of five vulnerabilities. Such a structure helps to pinpoint the areas of the macro-financial system where risks are elevated relative to their historic levels which is helpful when building a narrative around the evolution of risks over time.

The broad approach described above creates its own challenges. Data availability, especially for some of the emerging economies, is an important issue. Time series for some variables are short and others may not exist or may be of a different form for different countries. Furthermore, the indicators used are of different frequencies; some are more up-to-date than others; and their time-series properties (e.g. stationarity) vary enormously across different series. We discuss in some detail how one can go about addressing these challenges simply yet effectively.

Our approach is in part based on analysis by our colleagues in the Bank of England to analyse domestic financial stability risks. The techniques developed there are geared to aid the assessment of vulnerabilities within the UK financial system, but are naturally also useful in the international context. The analysis in this paper also draws heavily on the contribution of Aikman et al[2], who develop a framework based on numerous indicators for the US, and of a paper by Lee et al[3], who construct a comprehensive assessment of the past evolution of vulnerabilities across a wide range of countries.

More generally, our work is also related to at least two strands of literature. First, our drive to keep the analysis simple is motivated by work on uncertainty and complexity. Aikman et al[4] emphasise the attractive features of a simple modelling approach for complex systems, which are best characterised by 'Knightian' uncertainty (immeasurable and unknown risk) and 'unknown unknowns'. The second strand of related literature seeks to identify the predictors of crises and their severity in historical data. In this spirit, Schularick and Taylor[5] highlight the importance of credit growth in predicting crises across 14 countries between 1870 and 2008, and Krishnamurthy and Vissing-Jorgenson[6] highlight the importance of short-term debt. Our paper also relates to studies that aim to quantify the likelihood of future crises. Seminal pieces include Kaminsky and Reinhart[7] and Bussiere and Fratzscher[8].

The remainder of this paper is structured as follows: Section 2 provides an exposition of the methodology; Section 3 discusses the steps we take to ameliorate various data deficiencies; Section 4 discusses the results and their robustness. Finally, Section 5 concludes.

2. Methodology

The ultimate aim of the risk assessment framework is to identify risks to financial stability as early as possible, so that appropriate policy actions can be deployed to minimise stress and attenuate any impact. But how can one predict stress? The idea underpinning our analysis is that shocks hitting the economy will have a different impact depending on how vulnerable the economy is when the shock hits. The size and timing of the shocks themselves are inherently unpredictable – they are, by definition, beyond the limit of our understanding of the economic system. But if and when shocks materialise, their impact will vary according to exposures and leverage in the system that have been built up. Indeed, a shock may go unnoticed in the context of a resilient system, but can have severe consequences in an environment of heightened vulnerabilities (as illustrated by Lee and co-authors[4], Figure 1). Unlike shocks, these vulnerabilities can be assessed and monitored – and that is exactly what the toolkit presented here aims to do.

The shocks-and-vulnerabilities framework has become a standard approach in the recent risk assessment literature[2, 3, 9]. With a large enough sample of crisis events, it is in principle possible to check whether the data is consistent with this framework. Periods preceding crises should, on average, be associated with elevated vulnerabilities. Conversely, there should be fewer crises when vulnerabilities are contained. Lee et al[3] find that vulnerabilities are elevated prior to crises on average, lending support to the framework.

As outlined in the introduction, the approach here is to use many different data series to capture as much information as possible. Alessi et al[10] find evidence that multivariate approaches are better at explaining in- and out-of-sample crises than univariate signalling variables. In practice, however, the availability of data, especially in EMEs, often constrains the number of indicators included in an exercise such as ours.

Below we outline the key stages in producing the indices. The Appendix provides a more mechanical step-by-step summary.

Figure 2 shows the (monthly) data for the 22 indicators for the 8 countries we consider. There are a few things to note. First, not all the series are available over the whole sample period, January 1980 – December 2015. This limited data availability poses practical and econometric challenges, which we discuss in more detail in the next section. Second, for each individual country, the different series have different units, scales, and time series properties. To make the series comparable and to enable the aggregation, we normalise the data, by subtracting the sample mean and dividing by the sample standard deviation of each variable.

The key methodological and economic challenge however is the appropriate de-trending of individual series. From an econometric and practical point of view, a series that is non-stationary will tend to push the overall vulnerability index ever higher (or ever lower), making it difficult to interpret. Moreover, a change over any period of time would have been largely predictable. That is why it is important to apply a transformation to make the data stationary – in other words, subtract the systematic trending part from each of the series. But that in itself poses a deeper, economic challenge: what is the underlying trend that one ought to use? Unfortunately this is a difficult question to which no fully satisfactory answer exists. Ideally, the systematic trend would reflect the 'neutral' level of a given variable, in the sense that a variable value equal to its trend is consistent with overall macro-financial stability and balanced growth. But these neutral counterparts are, of course, unobservable. Indeed, even conceptually there is a lot of uncertainty as to the equilibrium levels of these variables that would deliver stability. Our approach relies on earlier work by Borio and Drehman[11] and Borio and Lowe[12] who find that, despite these difficulties, exercises that use simple techniques to try to identify build-up of vulnerabilities above the safe levels are useful in that they provide a helpful guide to practitioners. And one of our contributions is precisely to investigate the robustness of vulnerability indicators to those different techniques.

While providing a definitive answer to the question of which particular technique is optimal is beyond the scope of this paper, it is useful to consider several simple solutions and their implications. This is best illustrated with an example. The household debt to income ratio for the US provides a suitable case study. It is clear that US household debt in relation to GDP exhibited a strong upward trajectory over the past 30 years (Figure 3). That ratio reached a peak of 98% in

2008Q3 and then declined sharply during the financial crisis and recent recovery, as credit availability tightened and households deleveraged.

This example raises some noteworthy issues. First, it is not clear whether some of the indicators that appear to be trending up should be non-stationary in a conceptual sense. In the case of debt-to-GDP example, in a growing economy both GDP and the level of debt do trend up, but it is less clear that the ratio of the two does. Even though we observe data for a relatively long period of time (35 years), it is still difficult to distinguish a long-run trend from low-frequency changes: if the ratio of debt to GDP continues to decline in the years to come, what looks like a trend today may turn out to be a long and persistent rise, followed by a sharp decline. In principle we can rely on formal econometrics to test for stationarity: in our example, all variants of the Dickey Fuller test and the Phillips-Perron test fail to reject the null of a unit root in the household debt-to-GDP ratio (for both the DF and the PP tests we carried out test with no constant and no trend, with constant only, and with both constant and trend; the smallest p-value of these 6 tests was 0.52). But particularly for the series which have a shorter sample these formal techniques are less useful; and in any case they would also struggle to tell apart persistent changes from a long-run trend in short sample.

Perhaps the key observation one can take away from this example is that different de-trending techniques can yield significantly different results. The left hand side panel of Figure 3 shows three simple alternatives: the deviation from a linear trend, the deviation from a 3-year moving average, and a Hodrick-Prescott filter with parameter λ=1,000,000. The right panel of the Figure illustrates the de-trended series obtained using these three methods (the series shown in the Figure have been demeaned and they all have standard deviation equal to unity). The three lines move similarly until recently, but the latest observations appear to constitute (at least a local) turning point. The moving average and the HP filter trend move close together with the raw data; conversely, the linear trend is charging ahead so that the de-trended series declines further and further.

This example illustrates disadvantages of these respective methods: the linear trend may struggle to pick out the turning points in the data; the moving average trend is driven by a relatively arbitrary choice of the length over which the average is calculated, and the HP filter method suffers from a well-documented end-point problem. And each of these techniques implicitly carries a rather different implication about the economic concept on the 'neutral' trend level of a variable. The bottom line is somewhat unsatisfactory – there is no perfect, universal solution. In our baseline results below we use the linear trend due to its simplicity – the linear trends we estimated for the US, for example, are illustrated in Figure 4. We discuss the robustness of our results to different de-trending techniques in Section 4.

There are three levels of aggregation, illustrated diagrammatically in Figure 5. The first stage – from individual indicators to components – is important because it allows us to control the weight which certain sectors of the economy and the financial market receive in the analysis. For instance, there is much greater data availability on the corporate side of the economy than the household side. As a result, we have more variables describing the state of the corporate sector. If we did not combine these variables into a corporate component, the corporate sector would effectively dominate the nonfinancial vulnerability. Combining the data into components first deals with this issue. The second stage of aggregation is from components to vulnerabilities. Vulnerabilities are vital in providing a more detailed insight into the possible drivers in the risk profile of a country. The third stage is the construction of overall vulnerability index (OVI) from the underlying vulnerabilities. This index is a focus of our interest and analysis in this paper. It provides a high level summary of fragility in a given economy. Finally, as we show below, we can introduce a fourth stage where we weight the OVIs for different countries to construct a global snapshot.

We allow for different ways of aggregating the series together; they are all nested by the following constant-elasticity aggregator:

$$y(t) = \left[\sum_{j=1}^{N} w(j)\left(x(j,t)\right)^{r}\right]^{1/r}$$

where $y(\cdot)$ is the aggregate, $x(j,\cdot)$ is the j[th] component, $w(j)$ is the weight associated with that component in the aggregator, and $1/(1-r)$ is the constant elasticity of substitution.

There are two parameter values to be selected at each stage of aggregation: the series' weight $w(j)$ and the parameter r.

A natural choice for the weights would be to set them according to how useful various indicators have been in predicting stress in the past. But that usefulness has to be estimated, which is a non-trivial exercise given data limitations and the small number of well-documented crises. Furthermore, a 'good' set of weights will be a victim of its own success: by predicting the build-up of risks and helping to mitigate crises such system would tend to provide many 'false positives' and hence would not look very attractive retrospectively. In what follows we use equal weights across components at each level of aggregation. With simplicity comes an added advantage of being very easy to present and understand.

One exception to the equal weights approach is at the final stage of aggregation, when we combine the country-specific indices into a global vulnerability index. Given that the transmission of shocks from overseas to the UK occurs in large part through the UK banking sector, UK banks' exposures to foreign countries is an appealing way to judge the relative importance of different jurisdictions. An alternative to this simple approach would be to view the UK as a node in a network of countries which acts to originate and transmit shocks. A measure of network centrality could then be used as the weighting, capturing indirect effects. We leave this possibility as an avenue for future work.

The r parameter governs how the aggregator resolves potentially conflicting signals from different components. Setting r equal to one implies an infinite elasticity of substitution and delivers an arithmetic average. The choice of r<1 brings the aggregator closer to a geometric mean, with the property that for a given arithmetic mean, the more spread out the individual components, the lower the aggregate. Conversely, for r>1, if the individual indicators are widely spread, the aggregate is relatively high. Therefore, it is clear that the choice of the r parameter should depend on whether the practitioner becomes concerned only when many of the constituent components are elevated, or if a single raised vulnerability is sufficient to cause alarm. In what follows we set r=1, but the broad brush conclusions of our analysis are little changed for reasonable values of this parameter.

3. Data and Data Availability

We build on the work in Lee et al[3] in choosing the 22 individual indicators. We compile the data for 8 countries using DataStream. The five vulnerabilities and the underlying data are represented in the diagram in Figure 6. The first layer of the diagram represents the overall vulnerability index (the OVI), with five vulnerabilities directly below. These are further split into their components, which are in turn formed of the 22 underlying indicators. As flagged above, the data are of monthly frequency and span 35 years, from January 1980 to December 2015.

The countries that we focus on are chosen on the basis of the strength of their links back to the UK. Given the UK's large financial services industry, we focus on UK banks' external exposures in order to select the countries covered (Table 1). The countries chosen based on that metric are: China, France, Germany, Hong Kong, India, Korea, and the US.

Table 1: UK-owned bank claims on an ultimate risk basis in 2015Q4 as a percentage of UK banks' tangible equity

US	Hong Kong	China	Germany	France	Japan	India	Korea
141%	66.70%	30.60%	30.20%	29.10%	21.50%	12.70%	12.60%

Source: BIS and the ECB.

In order to make the most of the information that exists at each point in time, we introduce shorter series into the components as and when they become available. Therefore, a component is as long as its longest series, with weights adjusted dynamically. Other data issues are dealt with as simply as possible. Missing observations are linearly interpolated over, and variables which are not yet updated have their most recent values brought forward, although it would be relatively straightforward to incorporate nowcasts or 'best guesses' into our framework. This is left for future work. While most of the data we use are available at a monthly frequency, some indicators are only annual or quarterly. For those, we interpolate from annual into quarterly frequency then assume that the series takes the same value within a quarter.

4. Results

Figure 7 presents the baseline results. To aid the exposition, for each of the eight countries, we plot the overall vulnerability indices (the OVIs) using two different scales. The solid lines, shown on the left axis of each chart, show the indices which have been normalised to be distributed uniformly between 0 and 1. A value of 0.9, for instance, suggests that vulnerabilities have only been higher 10% of the time over the past 35 years. To perform this transformation, we first calculate the historical distribution of each of the indices using a non-parametric kernel estimator. The monthly observations are then transformed onto the [0,1] interval based on this historical distribution, in a way that ensures the distribution of the transformed series is uniform. The dotted lines, shown on the right axis, measure the vulnerabilities in standard deviations, relative to their historic mean.

It is interesting to note that the overall vulnerability indices in various countries follow quite different patterns over time: for example, the OVI for the US picks-up from the early 2000 and reaches its most elevated levels in 2008, just at the start of the financial crisis, and declines significantly after the crisis. A very different pattern is visible in China, Hong Kong and Japan, where vulnerability indices fell in the run up to the global financial crisis, but have picked-up strongly since then. The European countries fall somewhere in between these two cases. Our results thus provide an explanation for the severity of the financial crisis and the depth of the subsequent recession in the US. But for some other large economies, particularly in Asia, vulnerabilities were relatively low at the time when the global financial shock hit. This may help to explain a stronger performance of the Asian economies in the early phase of the crisis. Still, the sluggish global recovery took its toll particularly on the Asian economies, with vulnerabilities rising to elevated levels over the past few years. Our results put a new perspective on the recent macro-financial history and the financial crisis: while in absolute sense the global economy has been very badly hit, the damage could have been much larger had vulnerabilities outside the US been higher at the start of the financial crisis.

The set of charts in Figure 8 plots the overall vulnerability indices against the credit gaps, where both metrics have been transformed into measures distributed uniformly on the unit interval. This comparison is important because the credit-to-GDP gap often receives attention in work which tries to assess the stage of the credit cycle. Further, the Bank for International Settlements recommend a focus on the credit-to-GDP gap when local authorities set counter-cyclical capital buffers. The overall vulnerability indices developed here encompass a wider range of indicators and, therefore, often behave significantly differently to the credit-to-GDP ratios. Table 2 contains correlations between the credit gaps and the vulnerability indices for countries in our sample and across the lead/lag structure. For countries at a higher level of economic development: the US, France, Korea and to a lesser extent Germany (and with the exception of Japan), the correlation between the two measures is very high. On the other hand, the two measures send different signals in the case of Japan, India and China, highlighting how incorporating more information to private indebtedness changes the picture substantially. The upshot of this is that focusing solely on the credit gap may not be sufficient to pick-out all of the vulnerabilities emerging in the system, particularly in EMs. Overall the correlations between the two are highest contemporaneously, though for some countries (eg the US) the broader measures tend to lag credit gaps slightly.

Table 2: Correlations between the credit gap and the OVIs

Number of lags of OVI behind the Credit Gap	US	China	France	Germany	Japan	Hong Kong	India	Korea
-12	0.66	-0.31	0.49	0.22	-0.31	-0.18	-0.54	0.57
-6	0.77	-0.27	0.57	0.31	-0.20	-0.02	-0.54	0.69
0	0.85	-0.25	0.65	0.44	-0.08	0.14	-0.46	0.78
6	0.88	-0.25	0.67	0.44	0.00	0.21	-0.46	0.75
12	0.89	-0.27	0.65	0.47	0.08	0.24	-0.41	0.65

Given country-specific OVIs, it is possible to construct a global index of vulnerabilities. Figure 9 presents one such measure form the UK perspective. It is constructed by weighing the country-specific OVIs according to UK bank exposures (Table 1). The index is strongly driven by the US, given the large UK bank exposures there. Overall, the global index suggests that in the run up to the crisis, the UK-centric measure of global vulnerability netted out two opposing effects: on the one hand, vulnerabilities in the US were rising sharply; on the other hand, much of the rest of the world was becoming relatively more resilient. Since the crisis the roles have reversed: it is now emerging markets and Asia in particular that appear to be the most vulnerable.

Diving deeper into the country specific indices allows us to understand how vulnerabilities are developing and interacting and also focus our analytical resource on pockets of emerging risk. Figure 10 shows the contributions from each of the 5 vulnerability indicators to the total for each of the 8 countries. The contributions are distributed uniformly between 0 and 0.2 over the sample period (the 5 vulnerabilities themselves are distributed uniformly on the [0,1] interval, but we rescale the contributions to make sure that they add up to an index which lies with the [0,1] interval). The OVI in Figure 10 may therefore have a distribution that is close to the Bates distribution (average of the 5 uniform variables), which is bell shaped – meaning that the extreme values close to 0 and close to 1 are very unlikely indeed. Charts like these may be helpful in building a narrative about the evolution of risks, and can direct more detailed analysis of the particular sector in any given economy. For example, our index picks up the build-up of vulnerabilities in the financial and real sectors of the US economy in the run up to the crisis. It is also useful to spot any commonality of divergences across economies. For instance, it is interesting that the external vulnerability – reflecting the level of foreign exchange reserves, current account and external debt, appears to be very elevated across all of the Asian economies at the moment, perhaps in line with some concerns about global financial imbalances in that part of the world.

At an overall level, the outputs from our quantitative toolkit are broadly consistent with the assessment of risk performed elsewhere. For example, IMF's Global Financial Stability Report has in recent years highlighted the elevated level of vulnerabilities in some emerging markets, particularly China[13]. This broad picture of country risks has also been reflected in the scenarios underlying some of the recent stress testing exercises, e.g. those of Federal Reserve Board, the European Banking Authority or the Bank of England[14,15,16]. Specifically, the qualitative, judgement-based exercise on the assessment of country risk conducted as part of the Bank of England's 2016 Annual Cyclical Scenario concluded that risks were around standard in the US and Eurozone but elevated in Hong Kong and China due to high levels of credit growth and overall leverage, particularly in the corporate sector. Thus the newly available quantitative toolkit presented here broadly corroborates this assessment.

However, though the model indicates that vulnerabilities as of end-2015 were elevated in China and Hong Kong, this is driven more by the financial and external vulnerabilities, and to a lesser extent the non-financial vulnerabilities. In particular, a more detailed inspection reveals that the risks from the corporate component (within the non-financial vulnerability) are reported to be around average. This differing view stems from a shorter back run of data in the China and Hong Kong series and narrower data availability. This makes the estimation of a trend less reliable and places higher weight on fewer indicators, and highlights the need to further improve the data collection underlying the quantitative assessment.

As discussed in Section 2, the key judgement in the analysis of vulnerabilities concerns the estimation of the trend of the individual variables. So it is useful to consider what the implications are of different modelling choices.

First, it is interesting to see what happens when the trends in the variables are not corrected for at all. As discussed above, this choice could be conceptually consistent with at least some of the macroeconomic theory, which often predicts constancy of the key ratios such as debt-to-GDP over the long-run. Still, Figure 11 illustrates the practical difficulty such technique would run into: for many countries (and in particular for the US, France, India and Hong-Kong), the overall indices appear to be trending upwards. In practice, this means that the indices will on average indicate that the risks have increased compared to the past, and indeed that this increase could have been predicted. While these secular trends are interesting in their own right, they may make the overall vulnerability indicators of limited use to practitioners and policymakers, who are likely to focus on the genuine news and unpredictable changes. We therefore proceed to the analysis of the robustness to the choice of three basic de-trending techniques.

Figure 12 illustrates the robustness of the results to de-trending the data using three different methods: the linear trend (solid line – same as in the baseline results above); the 3 year moving average (the dashed line) and the Hodrick-Prescott filter (the dotted line). The interesting insight is that robustness varies greatly across countries. For some, such as Germany, France or Hong-Kong, the different measures paint a similar picture historically and over the recent past. For others – most notably the US, South Korea and perhaps Japan – the assessment of vulnerabilities can change dramatically depending on the de-trending method. For instance, the current assessment of vulnerabilities in the US or in South Korea depends crucially on the view of the sustainable path for the variables in question.

While the toolkit presented here does not present the answer to which of these techniques is best, it does motivate fascinating future research into the underlying secular trends and their sustainable and not-so-sustainable components. The finding that robustness varies across countries is an important one, especially given the early stages of the development of the largely empirical (indeed, mostly a-theoretical) research on this topic. Our paper strongly underscores the importance of digging deep into varying impacts of different assumptions. Given the high potential for fragile conclusions in this literature, it is vital that the authors of empirical analyses of financial crises or vulnerabilities provide readers with information to help in understanding the principal drivers of, and the robustness of, their specific results.

5. Conclusion

This paper has summarised a methodology that can be used to obtain quantitative metrics of vulnerabilities across several economies. There are many challenges when constructing such indices. A complete assessment of risks ahead will require in-depth, thematic analysis and judgement. Still, a simple, readily updateable quantitative metric of vulnerabilities for a set of the UK's economic partners is a useful starting point for any such analysis, and provides a way to make intertemporal comparisons and judgements that are consistent across time and jurisdictions.

Our attempts here are merely a first step; further work is needed to improve the toolkit and harness all of its potential. One such strand would carry out a more rigorous assessment of the vulnerability indices. For example, it may be interesting to investigate how useful they are in predicting crises, particularly when only the data available in real time are used. Of course, adding more countries and more variables into this framework would also increase its usefulness. Also, we have discussed in detail the robustness of the results to different de-trending techniques. Future work could usefully expand that analysis to include impacts of different elasticities and weights. All these strands of research would contribute to building increasingly reliable vulnerability metrics which can contribute to a better understanding of the risks in the global economic and financial system.

FIGURES:

Figure 1: Shocks and vulnerabilities framework

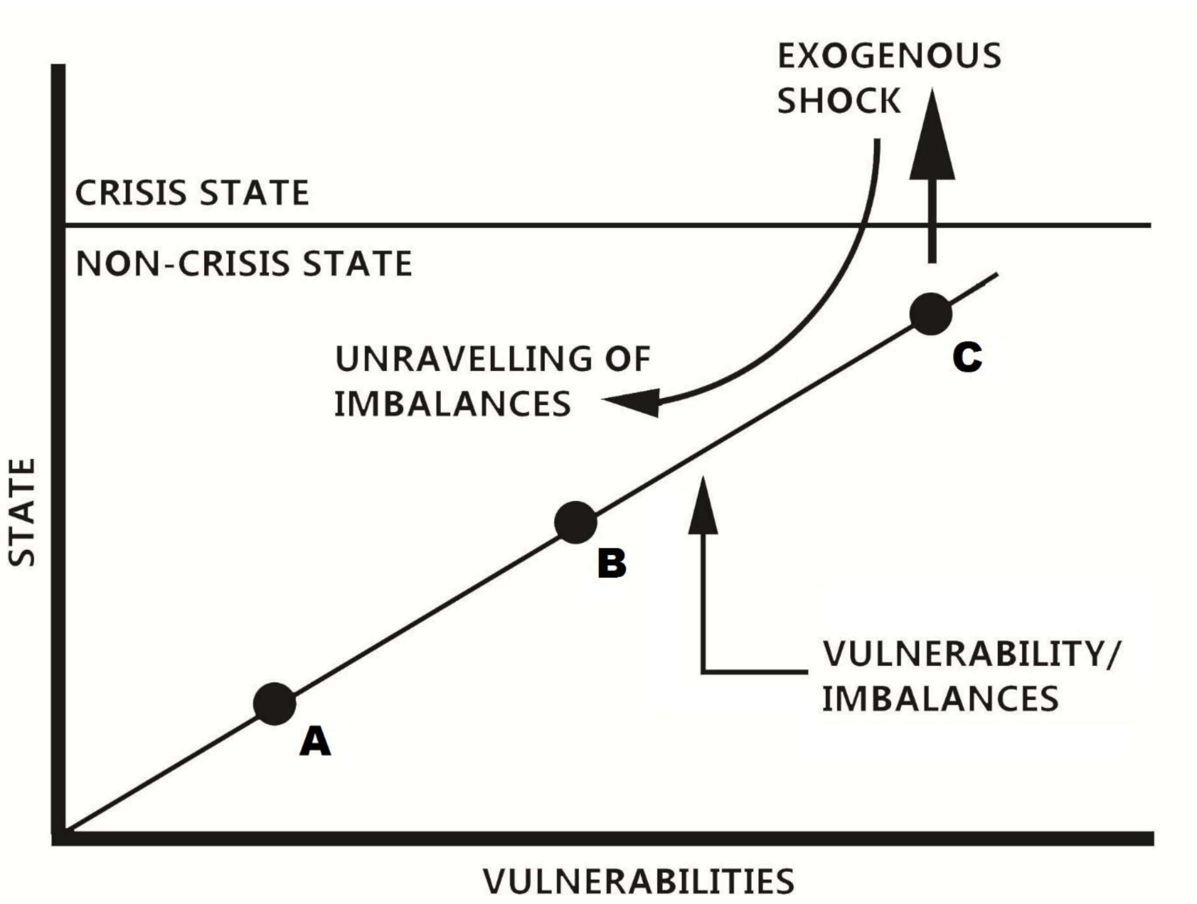

Source: Lee, S.J., Posenau, K.E., and Stebunovs, V., (2016), 'Anatomy of Financial Crises', unpublished, Board of Governors of the Federal Reserve System.

Figure 2: Raw data, all countries

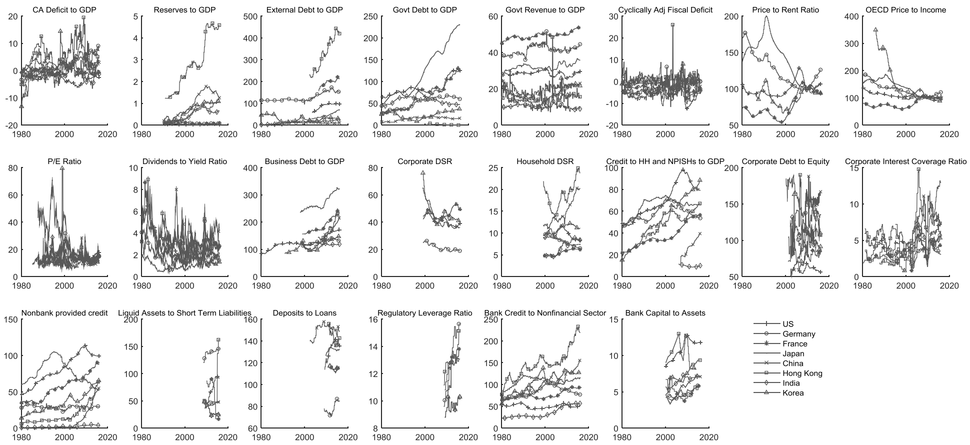

Figure 3: Example of different de-trending techniques: US household credit-to-GDP

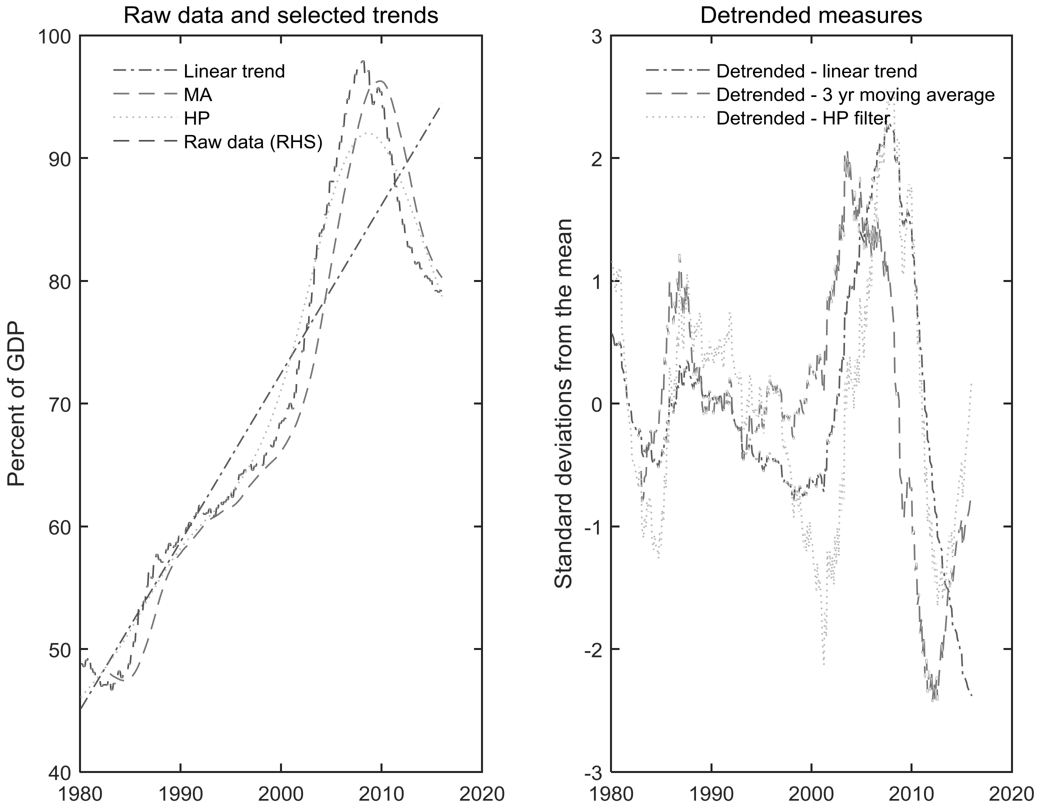

Figure 4: US individual series and trends

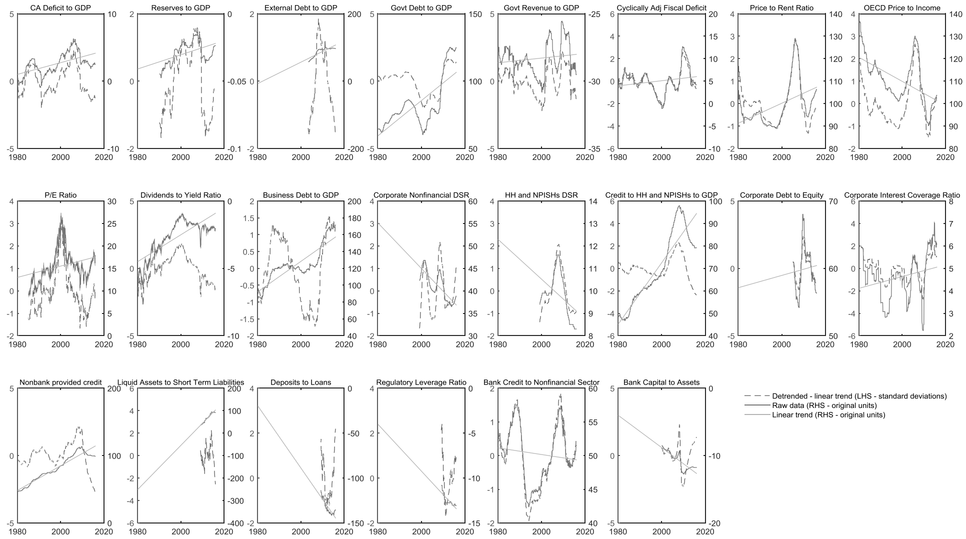

Figure 5: Levels of aggregation

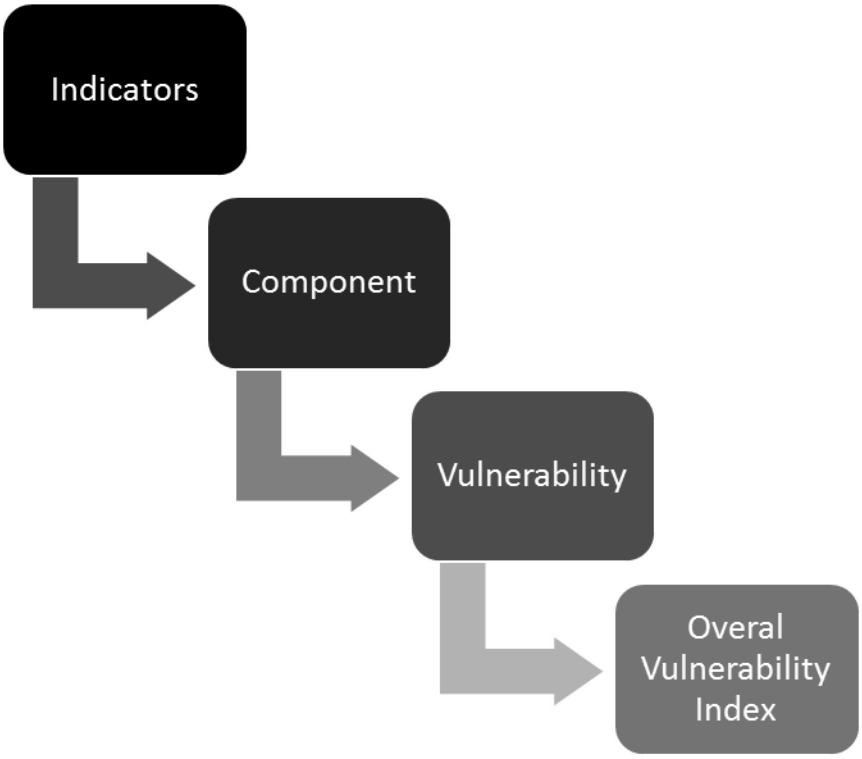

Figure 6: The individual indicators, the components, vulnerabilities, and the overall index

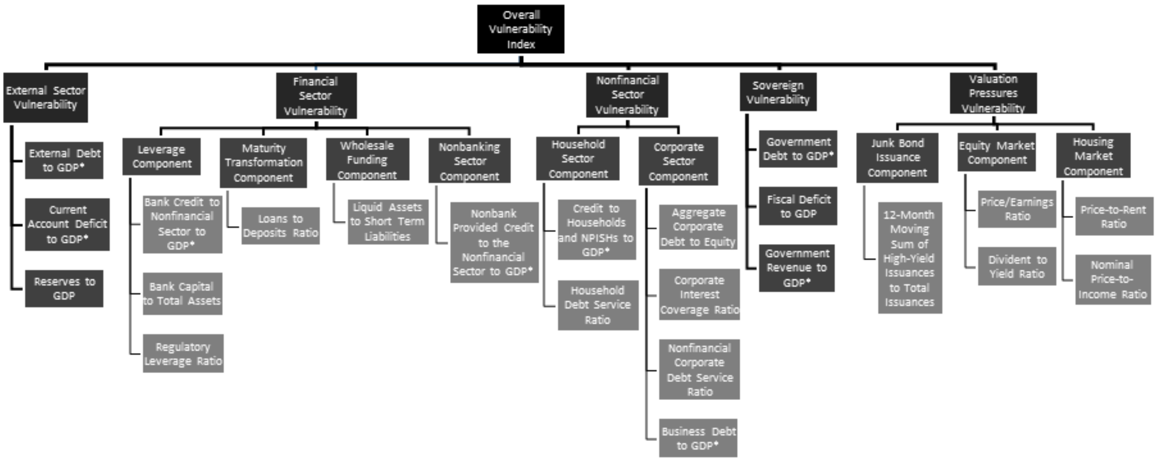

Figure 7: Overall vulnerability indices: uniform on [0,1] and standardised

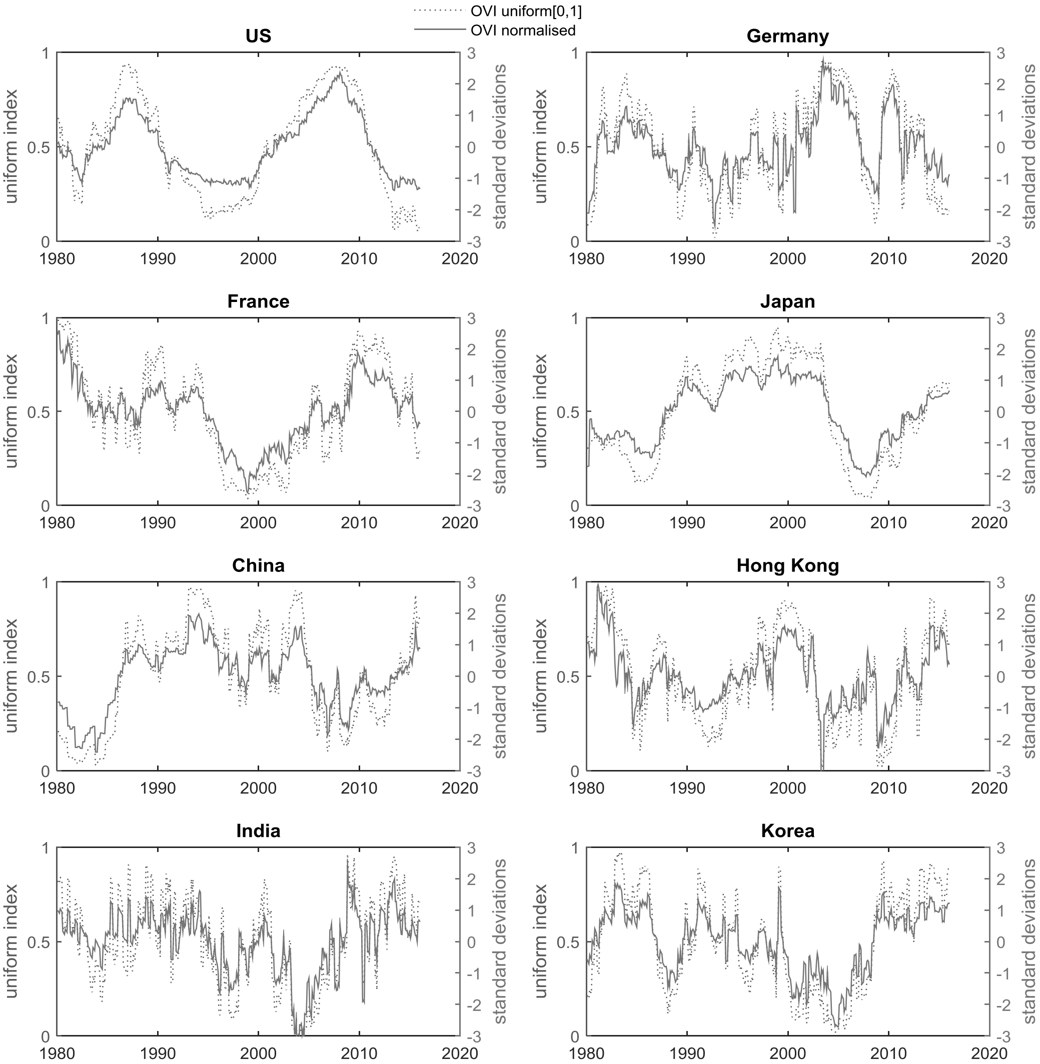

Figure 8: Overall vulnerability indices and the credit gap

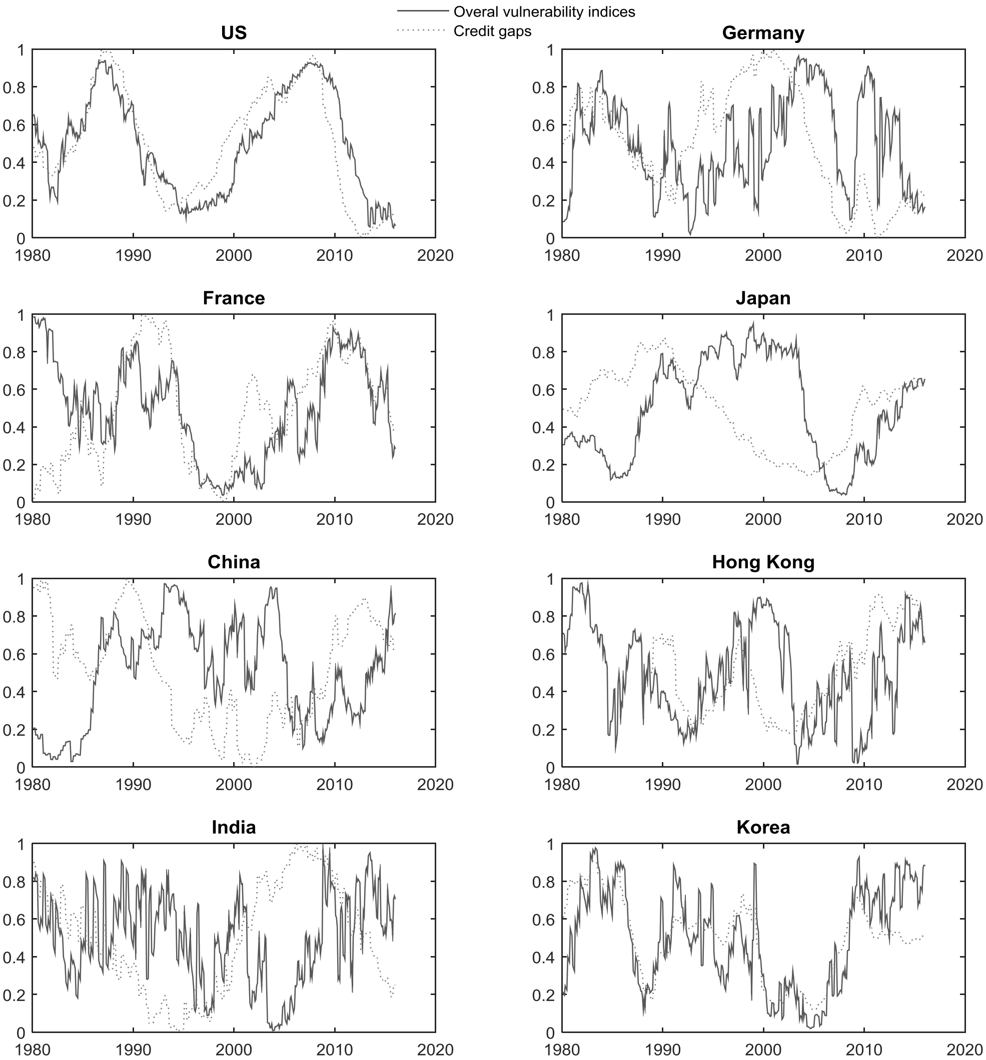

Figure 9: Global vulnerability index (weighted using the UK banks' exposure)

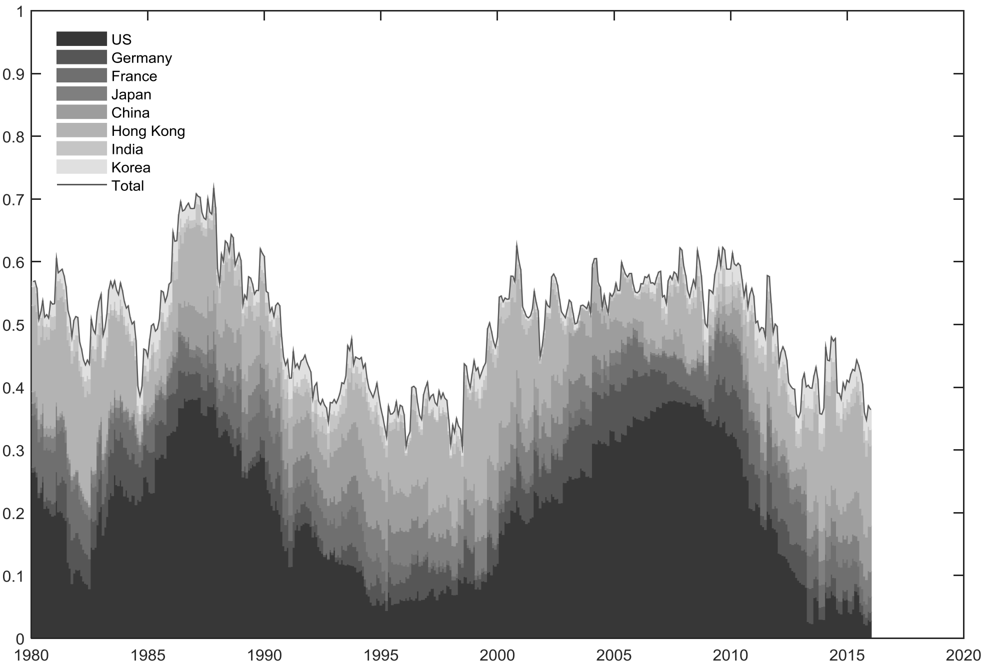

Figure 10: Overall vulnerability indices and their components

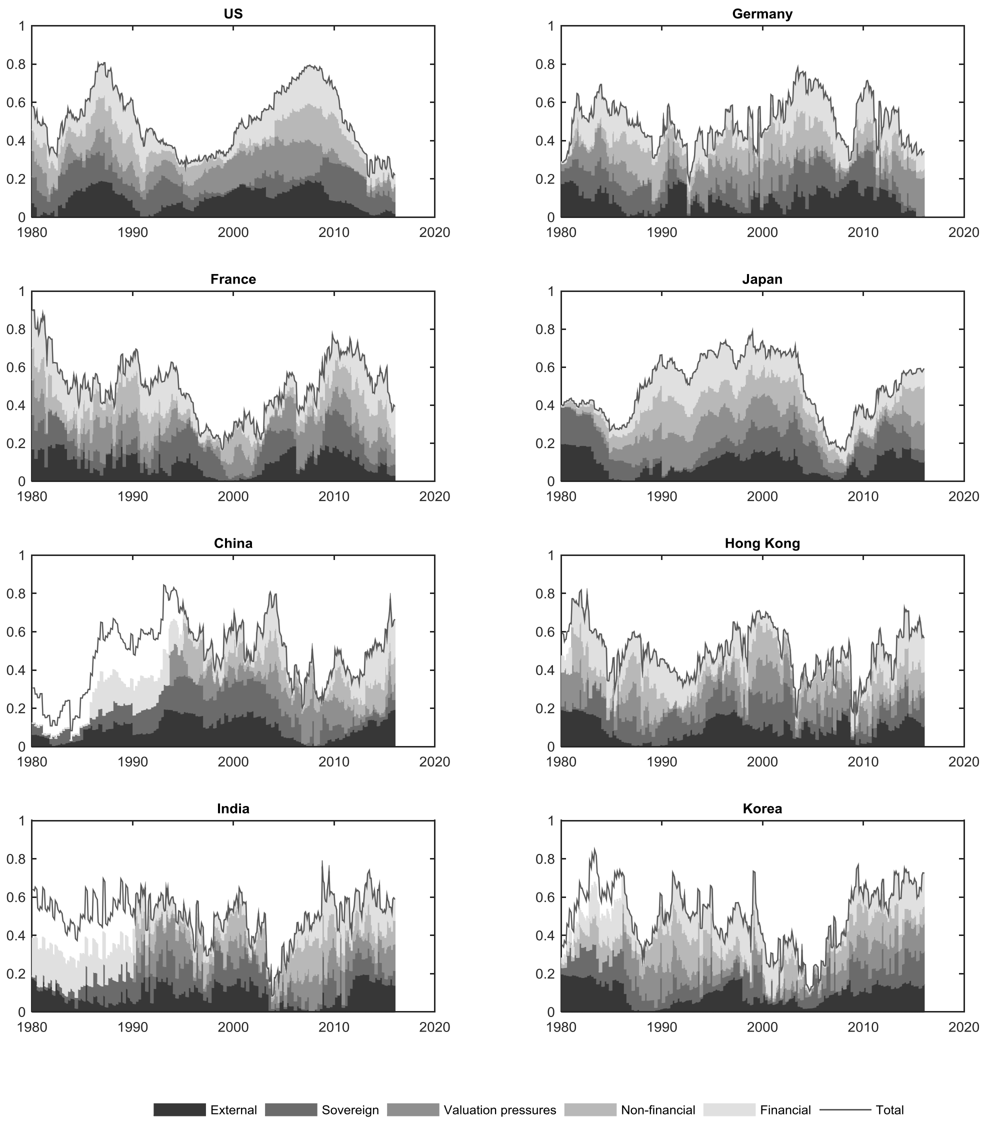

Figure 11: The overall vulnerability indices when no detrending is applied to the underlying data

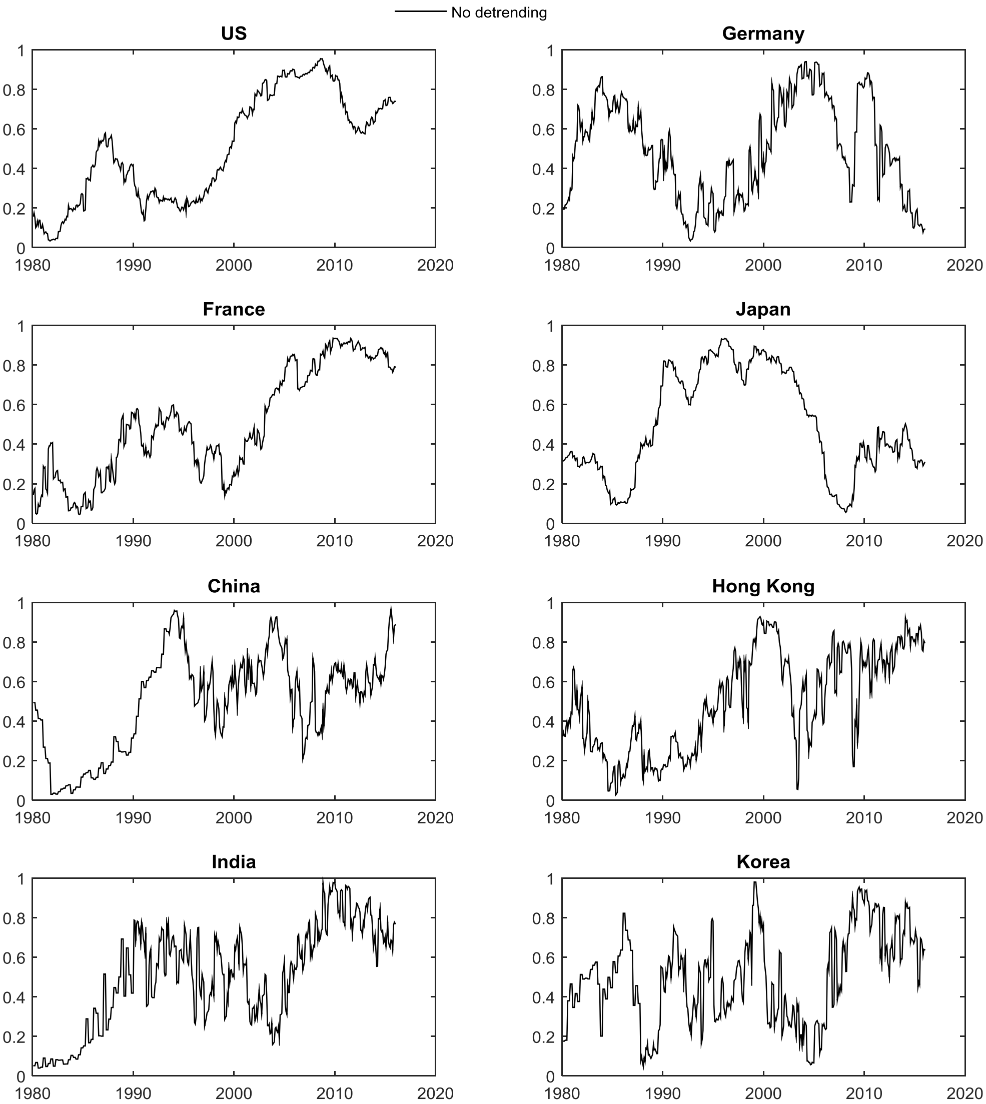

Figure 12: Robustness of the results to using different detrending techniques

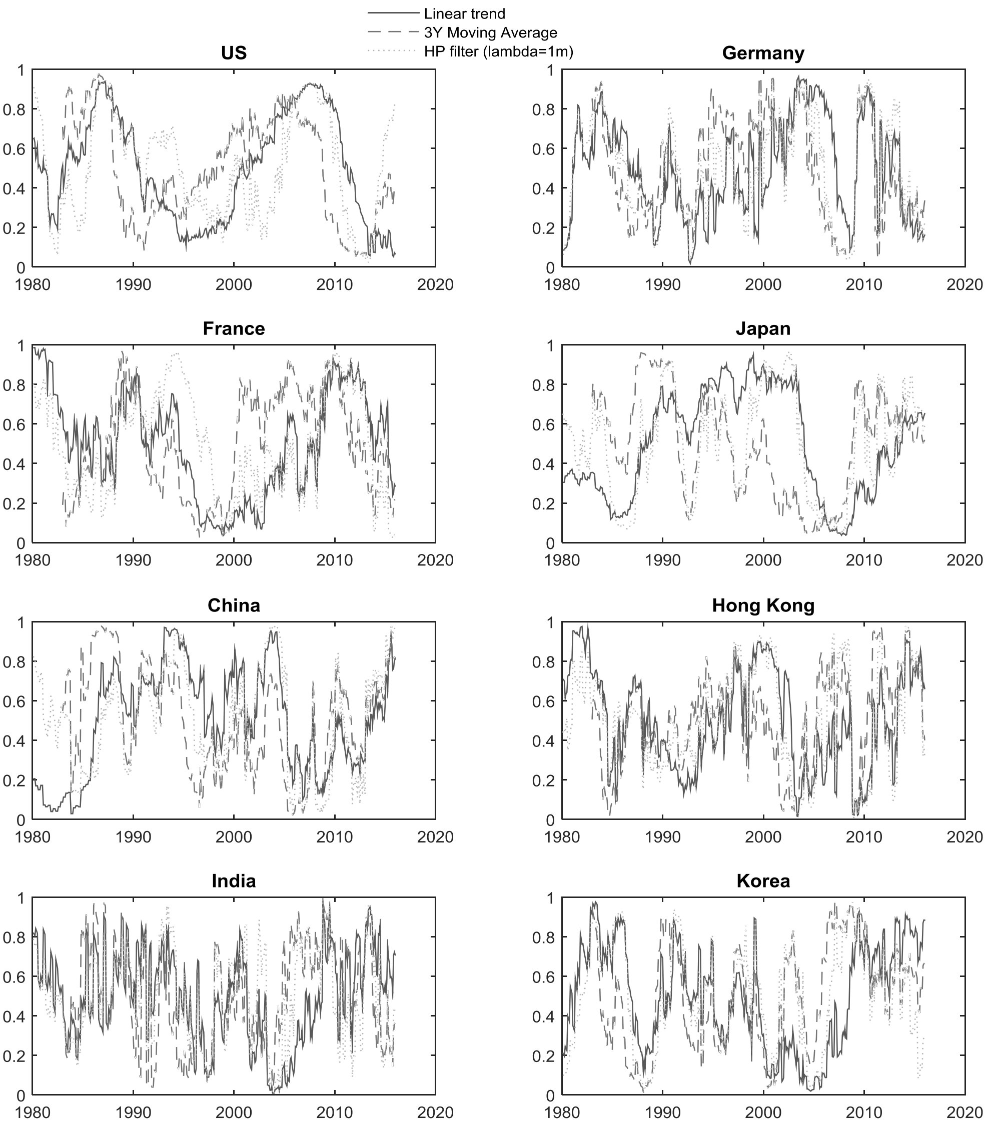

Appendix: steps taken to calculate the overall vulnerability index in a given country

1. Compile the raw data; interpolate lower frequency data into monthly frequency
2. Adjust ragged edges by extrapolating missing observations forward
3. Invert data series for which negative value indicates vulnerability
4. Estimate the trends and de-trend the variables
5. Standardise by subtracting the sample mean and dividing by standard deviation of each de-trended series
6. Weight the individual standardised and de-trended series into components (as shown in Figure 6). Use dynamic weights which sum to 1 in every period, but take account of data that become available over time
7. Weight the components into vulnerabilities (we use equal weights, but as discussed in the text other plausible weights may be used)
8. Standardise vulnerabilities by subtracting the sample mean and dividing by standard deviation of each vulnerability series
9. Perform kernel density estimation on each of the normalised vulnerabilities and rescale the vulnerabilities so that they are distributed uniformly on the [0,1] interval
10. Weight normalised vulnerabilities (we use arithmetic average)
11. Perform kernel density estimation on the weighted index to obtain the OVI that is uniformly distributed on the interval [0,1]

References

[1] Chowla, S., Quaglietti, L., and Rachel, Ł., (2014), 'How have world shocks affected the UK economy?', Bank of England Quarterly Bulletin, Vol. 54, No. 2, pg. 167-179.

[2] Aikman, D., Kiley, M.T., Lee, S.J., Palumbo, M.G., and Warusawitharana, M.N., (2015), 'Mapping Heat in the US Financial System', Finance and Economics Discussion Series 2015-059, Board of Governors of the Federal Reserve System.

[3] Lee, S.J., Posenau, K.E., and Stebunovs, V., (2016), 'Anatomy of Financial Crises', unpublished, Board of Governors of the Federal Reserve System.

[4] Aikman, D., Galesic, M., Gigerenzer, G., Kapadia, S., Katsikopoulos, K., Kothiyal, A., Murphy, E., and Neumann, T., (2014), 'Taking uncertainty seriously: simplicity versus complexity in financial regulation', Bank of England Financial Stability Paper, No. 28.

[5] Schularick, M. and Taylor, A.M., (2012), 'Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870-2008', American Economic Review, Vol. 102, No. 2, pg. 1029-1061.

[6] Krishnamurthy, A., and Vissing-Jorgensen, A., (2013), 'Short-term Debt and Financial Crises: What we can learn from U.S. Treasury Supply', unpublished, Northwestern University.

[7] Kaminsky, G.L., and Reinhart, C.M., (1996), 'The Twin Crises: The Causes of Banking and Balance-of-Payments Problems', American Economic Review, Vol. 89, No. 3, pg. 473-500.

[8] Bussiere, M., and Fratzscher, M., (2006), 'Towards a new early warning system of financial crises', Journal of International Money and Finance, Vol. 25, No. 6, pg. 953-973.

[9] Adrian, T., and Brunnermeier, M., (2008), 'CoVar', Federal Reserve Bank of New York Staff Reports, No. 348.

[10] Alessi, L., et al., (2013), 'Comparing different early warning systems: Results from a horse race competition among members of the Macro-prudential Research Network', MPRA Paper, No. 62194.

[11] Borio, C., and Drehmann, M. (2009), 'Assessing the risk of banking crises – revisited', BIS Quarterly Review, Bank for International Settlements, March, pp 29-46.

[12] Borio, C., and Lowe, P. (2002), 'Assessing the risk of banking crises', BIS Quarterly Review, Bank for International Settlements ,December, pp 43–54.

[13] International Monetary Fund (2016), 'Global Financial stability Report', available at https://www.imf.org/external/pubs/ft/gfsr/2016/02/pdf/ExecSummary.pdf

[14] Federal Reserve Board (2016), 'Stress Tests Required under the Dodd-Frank Act Stress Testing Rules and the Capital Plan Rule', available at: https://www.federalreserve.gov/newsevents/press/bcreg/bcreg20160128a2.pdf

[15] European Banking Authority (2016), 'Adverse macro-financial scenario for the EBA 2016 EU-wide bank stress testing exercise', available at https://www.eba.europa.eu/documents/10180/1383302/2016+EU-wide+stress+test-Adverse+macro-financial+scenario.pdf

[16] Bank of England, (2016), 'Stress testing the UK banking system: key elements of the 2016 stress test', available at: http://www.bankofengland.co.uk/financialstability/Documents/stresstesting/2016/keyelements.pdf