Stress Testing: A Robust End-to-End Approach

There is one approach for stress testing corporate portfolios that can not only yield credible, transparent results but also lead to improved accuracy of probability of default and loss-given default forecasts.

By Bruno Eklund, Evguenia Iankova, Alessandra Mongiardino, Petronilla Nicoletti, Andrea Serafino and Zoran Stanisavljevic

Over the last few years, and in particular in the light of the financial crisis of 2007-2008, stress testing has become a key tool in the risk management "arsenal" of financial institutions and a crucial input into the decision making process of a bank's board and top management. It is also a core component of the dialogue between banks and regulators as part of the Supervisory Review Process (Pillar 2 of Basel II).

In order to be a truly effective risk management tool, stress and scenario analysis must make clear how a given scenario impacts different portfolios, highlighting the assumptions on which the analysis is based. Equipped with this information, a firm will then be in a position to manage its portfolios proactively, ensuring that the impact of a given stress is within the firm's tolerance to risk. Ultimately, robust stress testing is about preparing for anything that might happen.

In this article, we specifically focus on stress testing for banks' corporate portfolios and describe an approach that allows a firm to assess in a transparent fashion how a given scenario affects banks' capital requirements for corporate portfolios, via estimating stressed PDs and LGDs.

Our framework has three main advantages:

• First, it links high-level scenarios to a more granular description of the economy, by estimating values for virtually any variable that can be considered to be a risk driver for a sector. As such, it improves the accuracy of probability of default (PD) and loss-given default (LGD) forecasts.

• Second, by considering a large set of macroeconomic variables in a sound econometric model, it contributes to produce credible and transparent results.

• Third, it facilitates the identification of portfolios, or parts of portfolios, that are particularly vulnerable to stressed conditions. As such, it provides a sound basis for proactive risk and portfolio management.

The main aim of stress testing is to evaluate the impact of severe shocks on a bank and assess its ability to withstand shocks and maintain a robust capital or liquidity position. If properly conducted, it sheds light on vulnerabilities otherwise not identified, informs senior management in the decision-making process, and underpins risk-mitigating actions to ensure the long term viability of the firm.

When a scenario is set — for example, either by a bank's management or by the regulator — it is typically only articulated on the basis of a few variables, such as GDP and inflation. However, the determinants of the solvency of firms, which are likely to be sector specific, may not be included among the scenario variables. Trying to impose a relationship only between the given scenario variables and the PD and LGD may not be appropriate, and may lead to biased and spurious results that undermine the practical use of the stress analysis.

This article will describe a factor-model approach that allows us to obtain any variables considered to be determinants of the corporate PDs for each sector, based on broadly defined scenarios. On the basis of this model, it is then possible to estimate the impact of the scenario on portfolios' PDs and LGDs and, through these, on the bank's capital requirements.

This feature underpins a transparent end-to-end approach for stress testing of corporate portfolios that links the settings of high-level scenarios all the way to the estimation of capital requirements under stress. It also provides a proactive risk management tool as it helps to identify which portfolios are most affected by a given stress and what mitigating actions would be required, if any, to ensure a firm's financial strength under stress. The end-to-end process is summarized in the diagram below.

Diagram: From Scenarios to Mitigating Actions — A Robust End-to-End Approach

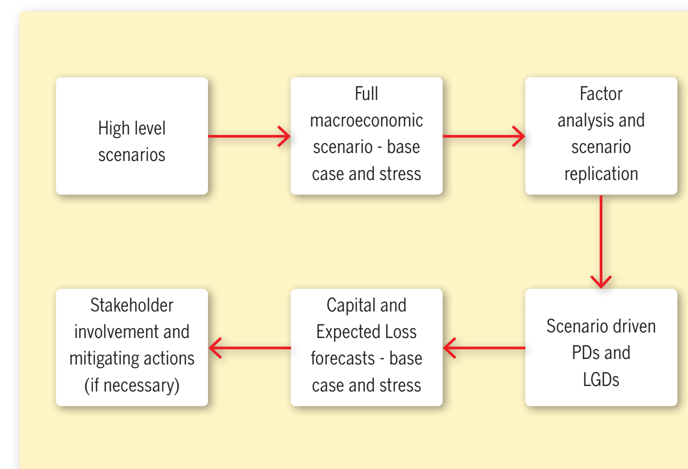

Our approach not only provides the tools to evaluate the impact of a scenario on sector-specific PDs, but is also less likely to suffer from under-specification, a problem shared by other macroeconomic models. In addition, unlike a purely judgmental approach to stress testing, it makes every step in the analysis transparent.

We should point out, though, that the model does not take into consideration firm-specific or idiosyncratic shocks, and is used only to evaluate the impact of macroeconomic stresses on PDs, LGDs and capital requirements. Furthermore, our discussion focuses specifically on assessing the impact of stress and scenario analysis for banks' corporate portfolios. The extension of this approach to other types of portfolios — for example, sovereign — is an avenue for further work.

Our model does not take into consideration firm-specific or idiosyncratic shocks, and is used only to evaluate the impact of macroeconomic stresses on PDs, LGDs and capital requirements.

The article is divided in sections, loosely following the diagram above. Section I explores the link between base case, stress scenarios and losses. Section II discusses our analytical approach (the macroeconomic model, the construction of factors and their use to replicate scenarios), while Section III concentrates on modelling risk drivers to obtain scenario-driven PDs and LGDs. Section IV explains the usefulness of the approach as a basis not only for forecasting capital and expected loss (EL), but also for productive discussions involving different stakeholders and risk mitigating actions. Conclusions are found in Section V.

Section I: Why do Scenarios Matter? Base Case vs. Stress

Banks' profitability, liquidity and solvency depend on many factors, including, crucially, the economic environment they face. In the assessment of extreme scenarios, there are two pressing questions for risk management: (1) What are the consequences of the extreme scenario on a particular bank? (2) Under that scenario, what can be done to improve the bank's resilience to the shock?

To address the first question, the future impact of the economy under normal conditions (base case) is compared with that of the extreme scenario (stress). By affecting the risk drivers, the stress shifts the loss distribution (see Figure 1, pg. 21) and raises expected and unexpected losses and capital requirements. To understand the extent of the impact on those variables, it is important to consider the link between risk drivers and economic dynamics.

Figure 1: Scenarios and Losses

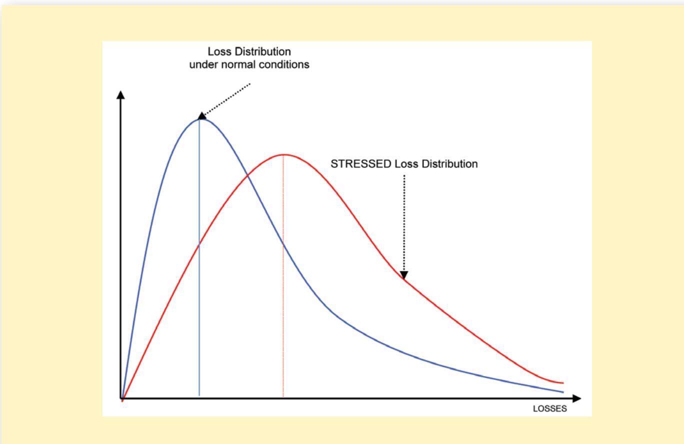

For corporate credit risk, the approach described allows for forecasting virtually any variable considered to be a risk driver given any scenario, and thus provides the basis for a robust answer to the first question about assessing extreme scenarios. The results are also useful in discussions on risk mitigating actions.

Section II: The Analytical Framework

The first building block of our approach is a UK macroeconomic model that uses 500+ UK and US quarterly macro and financial time series from 1980-Q1. The large number of variables considered over a fairly long sample period, which covers a few economic cycles, limits the risk of missing important drivers.

This model provides the background for stress and scenario analysis. To use an analogy, suppose that we want to evaluate the impact of a stone being dropped unexpectedly in the middle of a lake. The task is to predict the number of boats sinking and the number of fatalities. In this example, the probability of a boat sinking can be viewed as a "PD" and the mortality rate as a "LGD."

The "PDs" and "LGDs" are affected by several factors – including, for example, the size of the stone; the size of the boat; the distance between the boat and the stone; the experience of the captain and crew; the availability of lifeboats; the strength of the wind; and the proximity to the shore. Assume now that we only use the first two variables, which, in this analogy, represent the set of scenario variables provided. On the basis of those two variables only, we may draw substantially wrong conclusions about the actual values of the "PDs" and "LGDs."

On the other hand, the inclusion of too many variables can make a model unmanageable. Our preferred solution is to build a relatively small model, while still retaining most of the information contained in the dataset by constructing principal components.

The basic idea behind this method is that many economic variables co-move, as if a small set of common underlying and unobservable elements — the principal components (also called factors) — is driving their dynamics. Under this assumption, a small number of factors explain most of the variation in a large data set, and thus those few factors can be used to predict the variables in the dataset quite well.[1]

We believe that building a vector autoregressive (VAR) model based on the factors is a far better solution than the standard approach of specifying a small VAR model using a subset of the macro variables. A major drawback of a small VAR model is the high chance of under-specification, which may lead to unrealistic estimates and therefore significantly limits its practical use. For example, early VAR applications exhibited a price puzzle where a positive monetary shock was followed by a counterintuitive increase of the price level.[2,3] Furthermore, it would be difficult to construct a meaningful VAR model that also would include the relevant determinants for the corporate sectors' PDs.

Constructing the Factors

After collecting the economic and financial data, we use them to derive factors, which are a parsimonious and manageable "description" of the state of the economy. All data are collected into a single matrix, X_t, where each variable has been transformed to be stationary and standardized to have a zero mean and unit variance. The matrix X_t corresponds to the "macro variables" box of historical data in Figure 2 below.

Figure 2: From Macro Variables to Factors, and Vice Versa

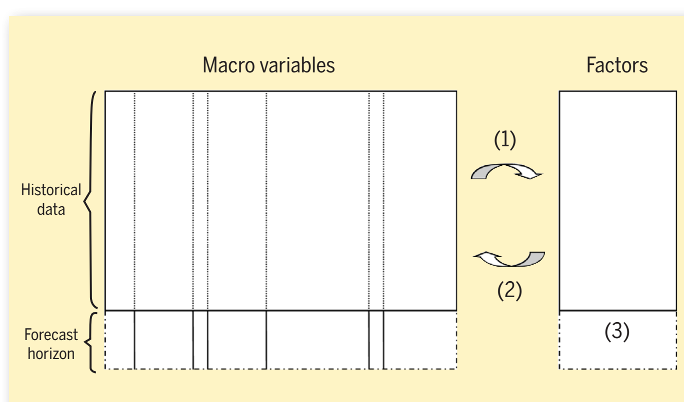

CREDIT RISK

The principal components or factors can then be constructed using the so-called singular value decomposition, which expresses the matrix X_t as a product of three separate matrices:

$$X_t = U_t L_t A_t^T , \tag{1}$$

where N is the number of observations and K is the number of variables in the standardized data set X.[4] Using this decomposition of the matrix X, the factors are constructed as

$$F_t = U_t L_t . \tag{2}$$

Note that in this step, we obtain as many factors as variables included in the data set. Equation (2) corresponds to arrow (1) in Figure 2 (see pg. 21), linking the macro data to the factors. The data set can be retrieved by post-multiplying F_t with A_t^T, arrow (2) in Figure 2, and thus mapping the factors back to the macro data, as follows:

$$F_t A_t^T = U_t L_t A_t^T = X_t . \tag{3}$$

Since the general idea with using principal component factors is to reduce the number of variables that needs to be included in the analysis, we only use the first $r < K$ factors. This limited number of factors summarizes the information contained in the underlying macroeconomic variables efficiently.[5]

When neglecting the remaining columns of F_t, a small error is normally made when transforming the factors back to the macro data set. The re-constructed matrix X_t can thus be expressed as a linear combination of the factors plus an error, as follows:

$$X_t = \tilde{F}_t \tilde{A}_t^T + \hat{\varepsilon}_t , \tag{4}$$

where $\tilde{F}_t$ and $\tilde{A}_t$ contain the first five columns of F_t and A_t, respectively, and $\hat{\varepsilon}_t$ contains the estimated error made by not using all factors. This relationship is used to retrieve forecasts for the data matrix X_t from the factor forecasts over a given horizon.

The factors can be compared to indices — i.e., they can be viewed as weighted averages constructed using a number of different variables. For our UK macroeconomic model, four out of the five estimated factors have direct interpretations that allow us to gain valuable insights into the final results. The first factor is highly correlated with GDP growth and related variables, and can be thought of as a proxy for economic activity. Similarly, the second factor is related to asset prices, the third to real interest rates and the fourth to productivity and employment costs.

An important advantage of our framework is that it can be employed to forecast a wide variety of economic variables that can be used to derive PDs and LGDs.

Replicating Economic Scenarios

The framework previously described can be used to obtain estimates of the risk drivers, even if they are not included in the set of scenario variables provided, both under a base case and under a stress scenario. Good replication of a scenario is important for the overall performance of the stress testing procedure. Its purpose is to analyze how all other macro and financial variables — the potential risk drivers — are affected by the given stress scenario and, at a later stage, to analyze how this in turn affects the estimates of the PDs and LGDs.

As mentioned in the introduction, an important advantage of our framework is that it can be employed to forecast a wide variety of economic variables that can be used to derive PDs and LGDs. In our model, we can replicate a scenario by imposing the dynamics of the scenario on factors and on all other variables in the data set. The scenario usually consists of explicit values of the most common macro and financial variables over a given forecast horizon.

The macroeconomic scenario can be represented in Figure 2 (pg. 21, see the vertical lines in the macro data set box). Each column in this matrix is a separate variable. The lines are dotted over the known sample period and solid over the forecast horizon.

To be able to estimate the effect of the stressed macro variables on all other macro variables, we first need to obtain good estimates of the factor forecasts in box (3) of Figure 2. Once these forecasts have been estimated, we can use the back transformation, arrow (2) or equation 4, to re-construct the forecasts of the remaining macro variables. This final step will fill in the forecast horizon of all macro variables (not just for the variables that were included in the scenario), and will thus allow us to use any macro variable we desire in later steps of our stress testing process.

Because the replication of the scenario is subject to an approximation error, it is important that this process not only can replicate the given scenario variables but also gives rise to realistic forecasts of the other variables — especially the possible risk drivers. Below, we describe two alternative methods that can be used to replicate scenarios, and the conditions under which one is preferable to the other.

The first method calls for the inclusion of the variables in the stress scenario as exogenous variables in a VAR model (with factors as endogenous variables), specifying a so called VARX model. This can be depicted as follows:

$$F_t = \Phi(L)F_{t-1} + \Pi(L)H_t + \eta_t , \tag{5}$$

where $\Phi(L)$ represents a lag polynomial vector of finite order d, $\Pi(L)$ is a lag polynomial of finite order s, η_t is an error term and the matrix H_t contains the variables whose forecasts are given in the scenario — such as, for example, GDP growth and CPI inflation. The model should be specified to produce the best possible forecast performance and/or the best fit. Since the stressed variables have known paths over the forecast horizon, the factor forecasts (box (3) in Figure 2) can easily be estimated.[6]

The second method, an alternative to the VARX model, is to model the factors individually, which corresponds to setting the lag polynomial $\Phi(L)$ equal to zero in equation (5).

In some cases, the second method is better at replicating the scenario than the VARX, possibly because it is more parsimonious. Including lagged dependent variables, as in the VARX model, might put a lot of weight on the history of the factors and could thus be worse at capturing the given stress dynamics. Individual factor models will also guarantee that any given shock to a macro variable will feed into the factor forecasts directly.

However, this second method might miss key interactions between the factors. The VARX also performs better at forecasting factors with low correlation to the set of scenario variables given, and when the set of scenario variables available is very limited.

One final point to remember is that a successful scenario replication requires an economic scenario derived using an analytical approach, so that the included variables are economically consistent with each other. This is especially important for stressed scenarios where variables might exhibit unexpected relationships between each other. Any ad hoc choices can severely affect any step of the stress testing process, introducing spurious results and unrealistic effects on macro and financial variables, and in turn on the forecasts for PDs and LGDs.

Section III: Scenario-driven PDs and LGDs

Once a scenario has been replicated, forecasts of the associated risk drivers can be obtained. These estimates are then used when modeling and deriving forecasts of the PDs and LGDs, which will then drive the capital and EL estimation.[7] Our framework has two distinct advantages: (1) it produces estimates of the risk drivers by linking them to the macroeconomic environment; and (2) it takes into account the heterogeneity across business sectors.

Regarding the first aspect, our approach is original because it models the relationship between systematic factors, such as GDP growth and interest rates, and PDs and LGDs, while incorporating as much information as possible from the economic environment. Several papers estimate links between PDs and economic variables; however, our framework can use both macro variables and factors as explanatory variables for the PDs and LGDs.

As for the second aspect, our approach can differentiate between sectors in order to identify the specific vulnerabilities of a particular scenario. Different sectors are sensitive to sector-specific drivers and respond differently to the systematic risk. Therefore, we model each sector PD separately and use variables related to that particular sector — e.g., house prices for real estate — in order to capture the sector-specific dynamics.[8]

Section IV: Estimation of Stressed Expected Loss and Capital, and Risk Mitigating Actions

After obtaining forecasts of PDs and LGDs, we can revert to our original question of assessing the impact of the stress on losses and capital requirements. The results of the stress test, usually compared to the results under the base case, will normally show an increase in EL and capital requirements. Figure 3 (next page) presents some stress test results in which the blue lines show the response to the base case and the red lines to the stress.

The approach presented in this article shows how to establish a clear link between a broadly defined scenario, often defined only for a few macro variables, and a fully defined macroeconomic scenario.

Figure 3: An Example of Stress Test Results

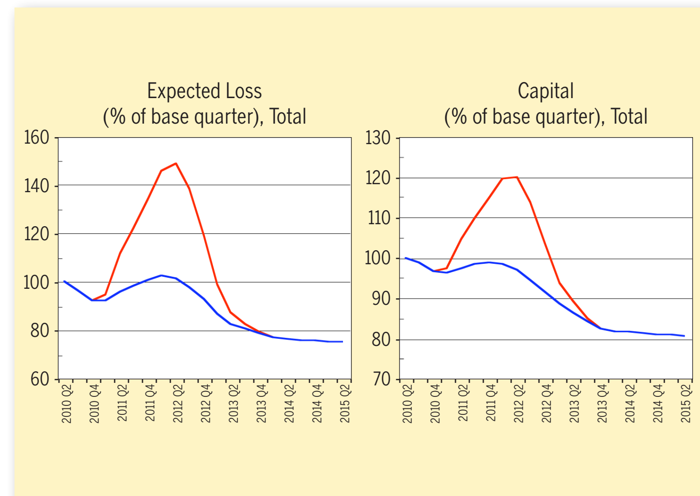

The analysis is then followed by a comprehensive discussion that challenges both assumptions and results by stakeholders in different areas of the bank: e.g., risk and finance. The results emerging from this debate can then be used to inform risk- mitigating actions.

In particular, our framework allows for the evaluation of corporate sector contributions to capital requirements and expected losses. If a sector under a certain stress, for example, drives a rise in capital requirements and/or impairments above the levels compatible with the bank's risk appetite, actions can be taken to limit the bank's exposure to that sector.

In this process, the involvement of senior management and the board of directors ensures that any decision taken is aligned with the bank's risk appetite and is effectively incorporated into the wider portfolio and risk strategy.

Section V: Conclusions

The approach presented in this article shows how to establish a clear link between a broadly defined scenario, often defined only for a few macro variables, and a fully defined macroeconomic scenario. It also demonstrates how to assess the credit capital requirements for corporate portfolios via stressed PDs and LGDs.

The merits of our framework are multiple. First, by estimating values for virtually any variable that can be considered to be a risk driver for each corporate sector, it raises the accuracy of PD and LGD forecasts. Second, by efficiently considering the information contained in a large set of macroeconomic variables in a sound econometric model, it produces transparent results, which form the basis of discussion for proactive risk management. Third, it can derive estimates of PD and LGD determinants consistent with scenario variables produced by a bank's top management and regulators.

The approach is practical and transparent, and can be used as a key input to assess the bank's capital position at times of stress, with respect to its own risk appetite as well as regulatory requirements. Whenever mitigating actions need to be considered, the framework allows one to identify the specific portfolios toward which these actions should be targeted.

Of course, the output of the analytical framework should not be used in a mechanistic way. Rather, it has to be subject to a critical review based on sound judgment. The combination of robust modeling and sound management is, we believe, the basis for good risk management.

FOOTNOTES

1. See Jolliffe (2004) for details.
2. See Sims (1972) and Sims (1980).
3. For more technical information, please see Hamilton (1994), Lütkepohl (2005), Sims (1972, 1980), Stock and Watson (2002), and Bernanke, Boivin and Eliasz (2005).
4. U_t and A_t are (N x N) and (K x K) orthonormal matrices; $U_t^T U_t = I_N$, $A_t^T A_t = I_K$, L_t is a (N x K) diagonal matrix with nonnegative elements in decreasing order.
5. The methodology developed by Bai and Ng (2002) has been adopted to determine the optimal number of factors to use (in our case, five), which explains a major proportion of the variance in X_t.
6. A method to replicate scenarios exactly has been developed. However, when applying it, forecasts of many of the non-scenario variables are unrealistic and meaningless. There is an apparent trade-off

between realistic results and the degree of accuracy of the scenario replication.

7. Here we focus on modeling probabilities of default. However, a way forward would be to apply a similar methodology to LGDs.

8. A common problem is that internal time series of PDs normally are too short, which means, in most cases, that the sample period does not cover a full credit or business cycle. This is the case for us, which is why we use MKMV Expected Default Frequencies (EDF) as a proxy for the PDs.

REFERENCES

Bai, J. and S. Ng (2002). "Determining the number of factors in approximate factor models," Econometrica, 70, 1, 191-221.

Bernanke, B.S., Boivin, J. and P. Eliasz (2005). "Measuring the effects of monetary policy: A factor-augmented vector autoregressive (FAVAR) approach," Quarterly Journal of Economics, 120, 1, 387-422.

Forni, M., Giannone, D., Lippi, M. and L. Reichlin (2007). "Opening the black box: Structural factor models with large cross-sections," European Central Bank Working Paper, n.712.

Hamilton, J.D. (1994). Time Series Analysis, Princeton University Press.

Jolliffe, I.T. (2004). Principal Component Analysis, Springer.

Lütkepohl, H. (2005). New Introduction to Multiple Time Series Analysis, Springer.

Sims, C.A. (1972). "Money, income and causality,"American Economic Review, 62, 540-552.

ibid. (1980). "Macroeconomics and reality," Econometrica, 48, 1-48.

Stock, J.H. and M.W. Watson (2002). "Macroeconomic forecasting using diffusion indexes," Journal of Business & Economic Statistics, 147-162.

Alessandra Mongiardino (FRM) is the head of risk strategy and a member of the risk management executive team for HSBC Bank (HBEU). She is responsible for developing and implementing the risk strategic framework, to ensure that risk management is effectively embedded throughout the bank and across all types of risk.

Zoran Stanisavljevic is the head of wholesale risk analytics at HSBC. After working at Barclays Capital as a credit risk data manager and a quantitative credit analyst, he joined HSBC as a senior quantitative risk analyst in 2005. Shortly thereafter, he took over the leadership on the bank's stress testing and economic capital projects, and was promoted to his current role in 2010.

Evguenia Iankova is a senior quantitative economist at HSBC, where she is currently involved in several projects, such as scenario building and stressing risk drivers. Following a stint as a quantitative economist in the economic research department at Natixis in Paris, she joined HSBC in 2008 to work on macroeconomic stress testing

Bruno Eklund is a senior quantitative analyst at HSBC. Before joining HSBC, he worked at Bank of England, where he built models for stress testing for the UK banking system. He also worked previously as a researcher at the Central Bank of Iceland, developing econometric models for forecasting the Icelandic business cycle.

Petronilla Nicoletti is a senior quantitative economist at HSBC, where she focuses on macroeconomic scenario development, replication for stress testing and the impact of shocks to risk drivers. Before joining HSBC, she worked as a senior risk specialist at the Financial Services Authority, the UK's financial regulator.

Andrea Serafino is a senior econometrician at the FSA. Prior to joining the regulator in 2010, he worked at HSBC, specializing in macroeconomic stress testing. He has also served as a consultant for the SAS Institute in Milan.

This article reflects the work of HSBC employees at the time of their employment at HSBC, but does not represent the views of either HSBC or the FSA.