Operational Risk Stress Testing:

An Empirical Comparison of Machine Learning Algorithms and Time Series Forecasting Methods

Matteo Crimella

Goldman Sachs Global Investment Research

R/Finance 2017 - Chicago, May 19 2017

Outline

1. Motivation and Framework
2. Methodological Section
   - Multivariate Adaptive Regression Splines
   - Artificial Neural Networks
   - Error Correction Models
3. Empirical Applications
   - Data
   - Models Outcomes
4. Stress Testing Outcomes
   - Widespread Scenario
   - China Hard-Landing Scenario
5. Conclusions

Motivation

- The European Banking Authority launched **2016 EU-wide stress test** with a focus on Operational Risks and Conduct Risks.

- Models to measure the **response of the Operational Risk profile** of systemic financial institutions **to macroeconomic shocks** are needed.

General Framework

Definitions

Y = Historical Op.Loss. Series (e.g. frequency of operational loss events)
X = Historical Macroeconomic Variables (e.g. GDP growth, interest rates)
X^s = Macro Stress Scenarios

We look for a way to accurately predict Y given some values of X

1. Summarize X using PCA. Collect $\tilde{X}$ (principal components) and Γ(PCA weights).
2. Estimate a model $Y = f(\tilde{X}) + \varepsilon$, which maps $\tilde{X}$ into Y.
3. Construct the stressed principal components $\tilde{X}^s = \Gamma X^s$.
4. Compute $\mathbb{E}\left[Y|\tilde{X}^s\right] = f(\tilde{X}^s) = \hat{Y}^s$.

Modelling Approaches

- **MARS Model**: non-parametric regression that automatically models non-linearities and interactions among variables
- **Artificial Neural Network**: regression algorithms able to approximate highly complex functions, to learn and generalize to unknown data
- **Error Correction Model**: time series approach to simultaneously model long and short run dynamics of the variables

Multivariate Adaptive Regression Splines (MARS)

Standard MARS (Friedman 1991):

$$Y = \beta_0 + \sum_{m=1}^{M} \beta_m h_m(\tilde{X}) + \varepsilon,$$

Hinge functions $h_m(\tilde{X})$ are defined using a two-step algorithm that selects knots and relevant regressors for Y.

Bagged MARS (Breiman 1996)

Let $\hat{\theta}^*_M(x)$ be a bootstrap predictor; the bagged MARS estimate is defined as:

$$\hat{\theta}_{M;bag}(x) \approx J^{-1} \sum_{j=1}^{J} \hat{\theta}^*_{M,j}(x).$$

Artificial Neural Networks (ANN)

Model components:

$$Z_m = \sigma(\alpha_{0m} + \alpha_m^T \tilde{X}), \quad m = 1, ..., M,$$
$$T = \beta_0 + \beta^T Z,$$
$$Y = g(T) + \varepsilon.$$

The model is estimated using the **back-propagation algorithm**.

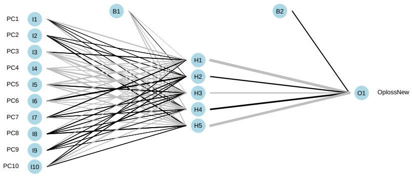

Error Correction Models (ECM)

Model Structure:

$$\Delta y_t = \psi_0 + \gamma_1 \hat{z}_{t-1} + \sum_{s=1}^{k^{diff}} \sum_{i=0}^{p} \psi_{1,s,i} \Delta \tilde{x}^{diff}_{s,t-i} + \\ + \sum_{j=1}^{p} \psi_{2,i} \Delta y_{t-i} + \sum_{s=1}^{k} \sum_{i=0}^{p} \psi_{3,s,i} \tilde{x}^{stat}_{s,t-i} + \nu_t .$$

$\hat{z}_t$ stands for the **error correction term** estimated via the Engle-Granger approach.

Stationarity has been tested using ADF test, KPSS test and PP test.

Empirical Applications: Data

- **Operational loss data** gathered from the Operational and Reputational Risk Department of a European systemic financial institution.

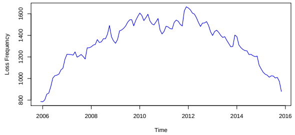

Figure: Deseasonalized Operational Loss Frequency Series: Monthly Observations

- **Macroeconomic data** span various regions and different classes. Data are summarized using PCA. Ten principal components are used in all the models.

Empirical Applications: MARS Outcomes (1/2)

Eight versions of MARS have been estimated based on:

1. Linear predictors - Nonlinear predictors
2. Degrees of interaction (one or two degrees)
3. Standard MARS - Bagged Mars

Example Predictors: hinge functions

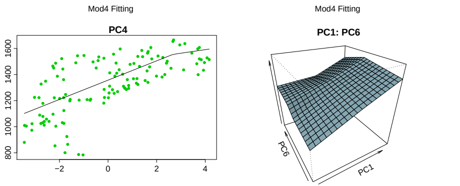

Empirical Applications: MARS Outcomes (2/2)

Cross validation has been used to detect the best performing model in terms of **out of sample forecast error**.

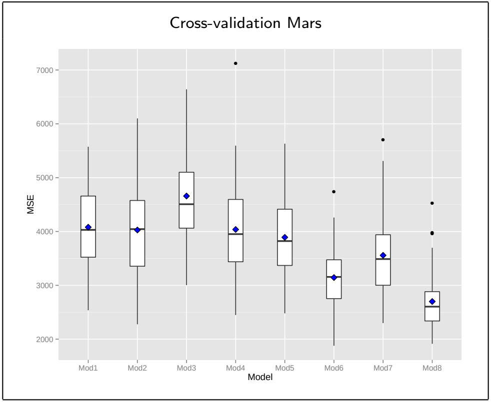

Empirical Applications: ANN Outcomes (1/2)

Twenty different networks have been fitted according to various combinations of the following features:

1. Starting values of the network weights
2. Number of hidden units and layers
3. Linear and non-linear networks

All variables have been scaled (to ease convergence of the backprop. algorithm) according to:

$$X_t^* = \frac{\tilde{X}_t - \tilde{X}_{min}}{(\tilde{X}_{max} - \tilde{X}_{min})}$$

Empirical Applications: ANN Outcomes (2/2)

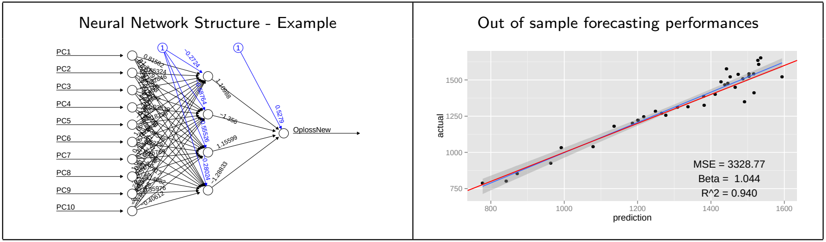

Figure: Neural Network Outcomes

Empirical Applications: ECM Outcomes

The graph shows in-sample fitting performances of the ECM model

ECM: Scatter plot actual and fitted values

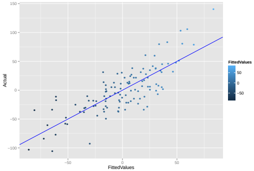

ECM: Ts-plot actual and fitted values

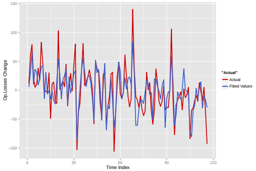

ECM: Residuals q-q plot

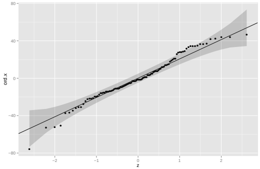

ECM: Residuals autocorrelation function

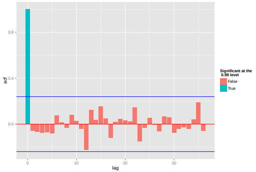

Applications to Macroeconomic Scenarios

1. Once the models are fitted, they can be used to **forecast the Operational loss dynamics, conditional to some macroeconomic scenarios.**
2. Stress scenarios are passed through the models as stressed principal components:

$$PC1_{T+s} = \alpha_{11}x_{1,T+s} + \alpha_{12}x_{2,T+s} + ... + \alpha_{1k}x_{k,T+s} \quad \textit{for } s = 1, ...S$$

3. Stress on OpLoss frequency has been studied using the following models:

- Non-linear bagged MARS
- Linear and non-linear ANNs
- ECM
- Forecast combination (average of forecasts from previous models)

Widespread Scenario - response of loss series

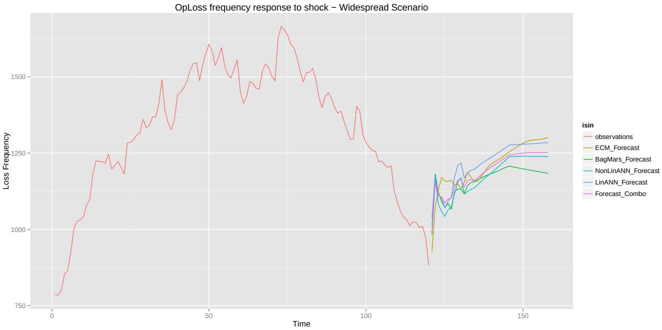

Figure: Scenario Widespread Contagion - Forecasts

China Hard-Landing Scenario - response of loss series

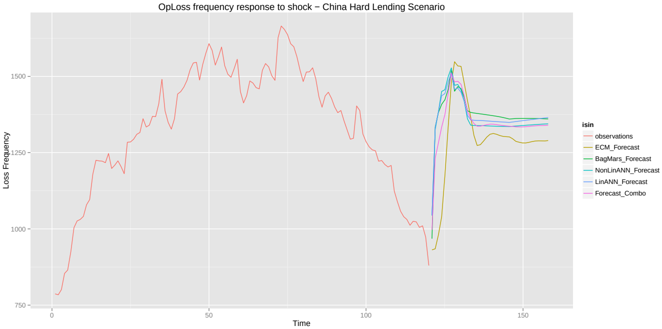

Figure: Scenario China Hard Landing - Forecasts

Conclusions

- The macro economy and its effects on Operational losses are complex processes: sophisticated methods are needed.
- Empirical testing indicates that ANN, MARS and ECM provide very similar predictions to the same scenarios. Results appear to be robust.
- In terms of forecasting performances, bagged MARS models seem to be the best choice to approach EBA stress tests on Operational Risks.
- Operational losses tend to increase as a consequence of negative shocks from the macro economy.

Future enhancements

1. Estimate time-lag between macro shocks and losses response
2. Estimate non-linear ECM using MARS algorithm
3. Test performance of ANN allowing for adaptive learning rates

Thank you for your attention.

Matteo Crimella

Goldman Sachs Global Investment Research
matteo.crimella@gs.com
+44(0)7719568357

Disclaimer: *This presentation is made in the author's individual capacity. The views expressed are those of the author alone and not of Goldman Sachs & Co. or any of its affiliate companies.*