

B IDENTIFYING EXCESSIVE CREDIT GROWTH AND LEVERAGE¹

Excessive credit growth has often been associated with the build-up of systemic risks to financial stability. With the entry into force of a new macro-prudential policy framework in the EU on 1 January 2014, a set of policy instruments has been made available to regulators to address such risks by curbing excessive leverage and/or imposing capital buffers which increase the resilience of the system against potential future losses.

This special feature presents an early warning system designed to support macro-prudential policy decisions. Drawing on the historical experience of EU countries, the model aims to assess whether observed leverage dynamics might justify the activation of macro-prudential tools such as the counter-cyclical capital buffer proposed by the Basel Committee on Banking Supervision. The early warning indicators are based on aggregate credit-related, macroeconomic, market and real-estate variables, while the early warning thresholds are derived by considering conditional relationships between individual indicators in a unitary framework.

INTRODUCTION

Past financial crises, and in particular the global financial crisis, have shown that excessive credit growth often leads to the build-up of systemic risks to financial stability, which may materialise in the form of systemic banking crises. As mitigating systemic financial stability risks is the objective of macro-prudential policy, several macro-prudential tools have been designed to curb excessive leverage or to build up buffers against likely future losses. Such instruments include the counter-cyclical capital buffer proposed by the Basel Committee on Banking Supervision, as well as other capital instruments, such as the leverage ratio and the systemic risk buffer, and instruments directly targeting borrowers, such as loan-to-value and loan-to-income caps.

However, the application of macro-prudential policy is still at an early stage and much effort is currently being devoted to providing policy-makers with concrete advice on how to actually design macro-prudential instruments. Indeed, the macro-prudential policy strategy has been defined by the European Systemic Risk Board (ESRB) with reference to the *guided discretion* principle, whereby the exercise of judgement is complemented by quantitative information derived from a set of selected indicators and associated “early warning” thresholds. In particular, with respect to the counter-cyclical capital buffer, the Basel Committee on Banking Supervision identifies the aggregate private sector credit-to-GDP gap as a useful guide, as this variable would have performed well in signalling the build-up of excessive leverage in the past. However, policy-makers should supplement the signal coming from credit-to-GDP trend deviations with judgement based on a broader information set, as also suggested in the 2010 Basel guidance.²

Against this background, this special feature presents an early warning model to be used for identifying those periods in which the build-up of leverage can be defined as excessive and may warrant the activation of relevant macro-prudential instruments.³ As in any early warning exercise, the target event is defined first. In the present case, the model is designed to issue warning signals well ahead of systemic banking crises caused by excessive credit growth. The second step is the selection of the candidate early warning indicators: in this respect, the dataset used in this

An early warning model to identify excessive leverage

¹ Prepared by Lucia Alessi.

² See Basel Committee on Banking Supervision, “Guidance for national authorities operating the countercyclical capital buffer”, 2010.

³ For more details on the methodology, see Alessi, L. and Detken, C., “Identifying excessive credit growth and leverage”, *Working Paper Series*, ECB, forthcoming.


application comprises publicly available aggregate credit-related, macroeconomic, market and real-estate variables for the euro area countries together with the United Kingdom, Denmark and Sweden. The modelling technique is based on decision trees, in particular binary classification trees. One of the main advantages of this technology is that it takes into account the conditional relationships between indicators in setting the respective early warning thresholds. Finally, the in-sample and out-of-sample predictive performance of the model is evaluated.

DEFINING CREDIT-RELATED SYSTEMIC BANKING EVENTS

Standard banking crisis definitions include episodes in which: much or all of bank capital is exhausted; bank runs lead to the closure, merger, or takeover by the public sector of one or more financial institutions; there are significant signs of financial distress in the banking system; or significant banking policy intervention measures are required in response to significant losses in the banking system. However, macro-prudential tools such as counter-cyclical capital buffers and leverage ratios aim to avoid a broader array of circumstances than simply a banking crisis as defined in these terms alone. Therefore, the definition of a banking crisis used in this exercise, borrowed from the work of the European Systemic Risk Board Expert Group on Countercyclical Capital Buffers, is extended to include “near misses”, i.e. periods in which domestic developments related to the credit/financial cycle could well have caused a systemic banking crisis had it not been for policy action or an external event that dampened the credit cycle. Non-systemic banking crises and crises not related to the credit cycle are excluded.⁴

According to this definition, 25 episodes are identified in the countries under analysis over the period from the first quarter of 1970 to the end of 2013 (see Chart B.1). Owing to the time lag

Chart B.1 Pre-crisis and crisis periods



Sources: Detken, C. et al., “Operationalising the countercyclical capital buffer: indicator selection, threshold identification and calibration options”, *Occasional Paper Series*, ESRB, forthcoming and Alessi, L. and Detken, C., “Identifying excessive credit growth and leverage”, *Working Paper Series*, ECB, forthcoming.
 Notes: Credit-related systemic banking crises are marked in blue; early warning signals are expected in periods highlighted in red; periods marked in grey are excluded from the analysis.

4 For details on this banking crises dataset, see Detken, C. et al., “Operationalising the countercyclical capital buffer: indicator selection, threshold identification and calibration options”, *Occasional Paper Series*, ESRB, forthcoming.

between the adoption of a macro-prudential measure and its entry into force, the early warning model is designed to identify excessive leverage sufficiently early, namely at least one year prior to the start of the crisis, and up to five years ahead.⁵ The periods in which the model is expected to issue warning signals are highlighted in red in Chart B.1, while periods marked in grey are not included in the analysis because they are too close to the outbreak of a crisis, or are not classifiable as pre-crisis, given that we do not know whether a crisis will actually materialise in the next few years. The crisis periods themselves (in black in Chart B.1) are, of course, also excluded.

A BROAD SET OF INDICATORS

A battery of indicators which could contain valuable information is considered. This broad set includes mainly financial variables, in particular various transformations of credit aggregates. The key aggregate is broad credit, covering loans and debt securities provided by the domestic banking sector to non-financial corporations and households. This is entered into the model in the form of year-on-year rates of growth, as well as the ratio to GDP and deviations of this ratio from its trend (i.e. the “gap”). This latter transformation has been suggested by the Basel Committee on Banking Supervision and is therefore referred to as the “Basel gap”.⁶ The narrower bank credit aggregate and sectoral credit aggregates are also considered, as well as global liquidity measures. With respect to debt service costs, the aggregate debt service ratio and sectoral debt service ratios are included. Finally, public debt also features in the pool of credit-related indicators.

With respect to asset prices, the dataset comprises housing market indicators and equity price growth. Macroeconomic variables and interest rates are also considered, as they could be useful for conditioning the signals coming from credit-based indicators and asset prices.⁷

A MODEL FOR IDENTIFYING EXCESSIVE CREDIT GROWTH AND LEVERAGE

The model presented in this special feature aims to identify whether, in a given period, the European financial system is in a state of vulnerability owing to the build-up of excessive leverage, which in turn increases the likelihood and potential impact of a subsequent banking crisis. In such a situation, the activation of macro-prudential policy measures would be prudent. To this end, a purely statistical approach is adopted, based on decision trees.

A binary classification tree is a partitioning algorithm which recursively identifies the indicators and the respective thresholds that are able to best split the sample into the relevant classes, i.e. pre-crisis and tranquil periods. The output of the predictive model is a tree structure like the one shown in Chart B.4, with one *root* node, only two branches departing from each *parent* node (hence “binary” classification tree), each entering into a *child* node, and multiple terminal nodes (or “leaves”). Starting by considering all available indicators and threshold levels, the procedure selects the single indicator and threshold yielding the two purest sub-samples based on an impurity measure. A standard impurity measure is the Gini index:

$$GINI(f) = \sum_{i=1}^n f_i(1-f_i) = 1 - \sum_{i=1}^n f_i^2$$

5 For example, in the case of counter-cyclical capital buffers, banks are given one year to raise additional capital.

6 The Basel gap is computed with a recursive slowly adjusting (i.e. $\lambda=400,000$) HP filter, which implicitly assumes that the financial cycle is four times as long as the business cycle. As such a HP trend might be adjusting too slowly following a prolonged period of negative credit growth, an alternative gap computed with $\lambda=26,000$ is also considered, corresponding to a financial cycle which is twice as long as the business cycle.

7 All variables are entered into the model in real terms. Measures of funding liquidity (e.g. the LIBOR-OIS spread and the loan-to-deposit ratio) have not been included in the analysis owing to data availability issues.

Credit-related early warning indicators...

... asset prices and macroeconomic variables.

Binary classification trees

The Random Forest...

where f_i is the fraction of periods belonging to each category i , with $i=1,2$ in this case. The algorithm proceeds recursively, by finding the best split at each node, so that one of the child nodes contains mostly pre-crisis periods while the other contains mostly tranquil periods. Once the logical structure is constructed on the basis of historical data, the tool can be used in real time to map the current value of a set of indicators into a single prediction, expressed as the probability of being in each of the classes.

The main drawback of the tree technology is that, while it can be very good in-sample, it is known not to be particularly robust when additional predictors or observations are included. This problem is overcome by using the “Random Forest” algorithm.⁸ This framework is a state-of-the-art machine learning technique which involves *bagging*, i.e. bootstrapping and aggregating, a multitude of trees. Each of the trees in the Forest is grown on a randomly selected set of indicators and country quarters. Once a new quarter of data is available, the prediction of the Forest will be based on how many trees in the Forest classify it as a pre-crisis or tranquil period.


... and its out-of-sample predictive performance

Each of the trees in the Forest is in itself an out-of-sample exercise, as the observations that are not used to grow the tree (*out-of-bag* observations) can be put through the tree to get a classification. It is therefore possible to compute the total misclassification error of the Forest. Based on the out-of-sample error rate of a 100,000-tree forest grown on all of the considered indicators, the chance of misclassifying an incoming quarter of data is 6%. A more advanced measure of the performance of a classifier is the area under the receiver operating characteristic curve (AUROC). The receiver operating characteristic (ROC) curve plots the combinations of true positive rate (TPR) and false positive rate (FPR) attained by the model (see Chart B.2). The ROC curve of a random classifier will tend to coincide with a 45 degree line, corresponding to an AUROC of 0.5, while the AUROC of a good classifier will be closer to 1 than to 0.5. The ROC curve of the Random Forest presented in Chart B.2 corresponds to an AUROC above 0.9.

Measuring the indicators' importance

Finally, the Random Forest makes it possible to measure the importance of each of the input variables by evaluating the extent to which it contributes to improving the prediction. Chart B.3 shows the indicators' ranking derived from the Random Forest.⁹ Not surprisingly, since the model is designed to predict banking crises associated with a domestic credit boom, the most important indicator turns out to be bank credit, in the form of its ratio to GDP, followed closely by the gap derived with a very slowly adjusting trend. Global liquidity – in the form of both the global credit gap and the growth rate – turns out to be another key concept, ranking among the five most important indicators. The remaining two indicators among the top five are the level of household credit and the aggregate debt service ratio. In general, credit-to-GDP ratios appear helpful in assessing how vulnerable a country is because of excessive structural leverage rather than conjunctural developments, and are

Chart B.2 Receiver operating characteristic curve




Source: Alessi, L. and Detken, C., “Identifying excessive credit growth and leverage”, *Working Paper Series*, ECB, forthcoming.

8 See Breiman, L., “Random Forests”, *Machine Learning*, Volume 45, Issue 1, 2001, pp. 5-32.

9 Since an element of randomness is inherent in the Forest, the ranking may vary slightly from replication to replication.

Chart B.3 Importance of indicators



Source: Alessi, L. and Detken, C., "Identifying excessive credit growth and leverage", *Working Paper Series*, ECB, forthcoming.

Note: The bars represent a normalised measure of the prediction error obtained by ignoring the information conveyed by a specific indicator.

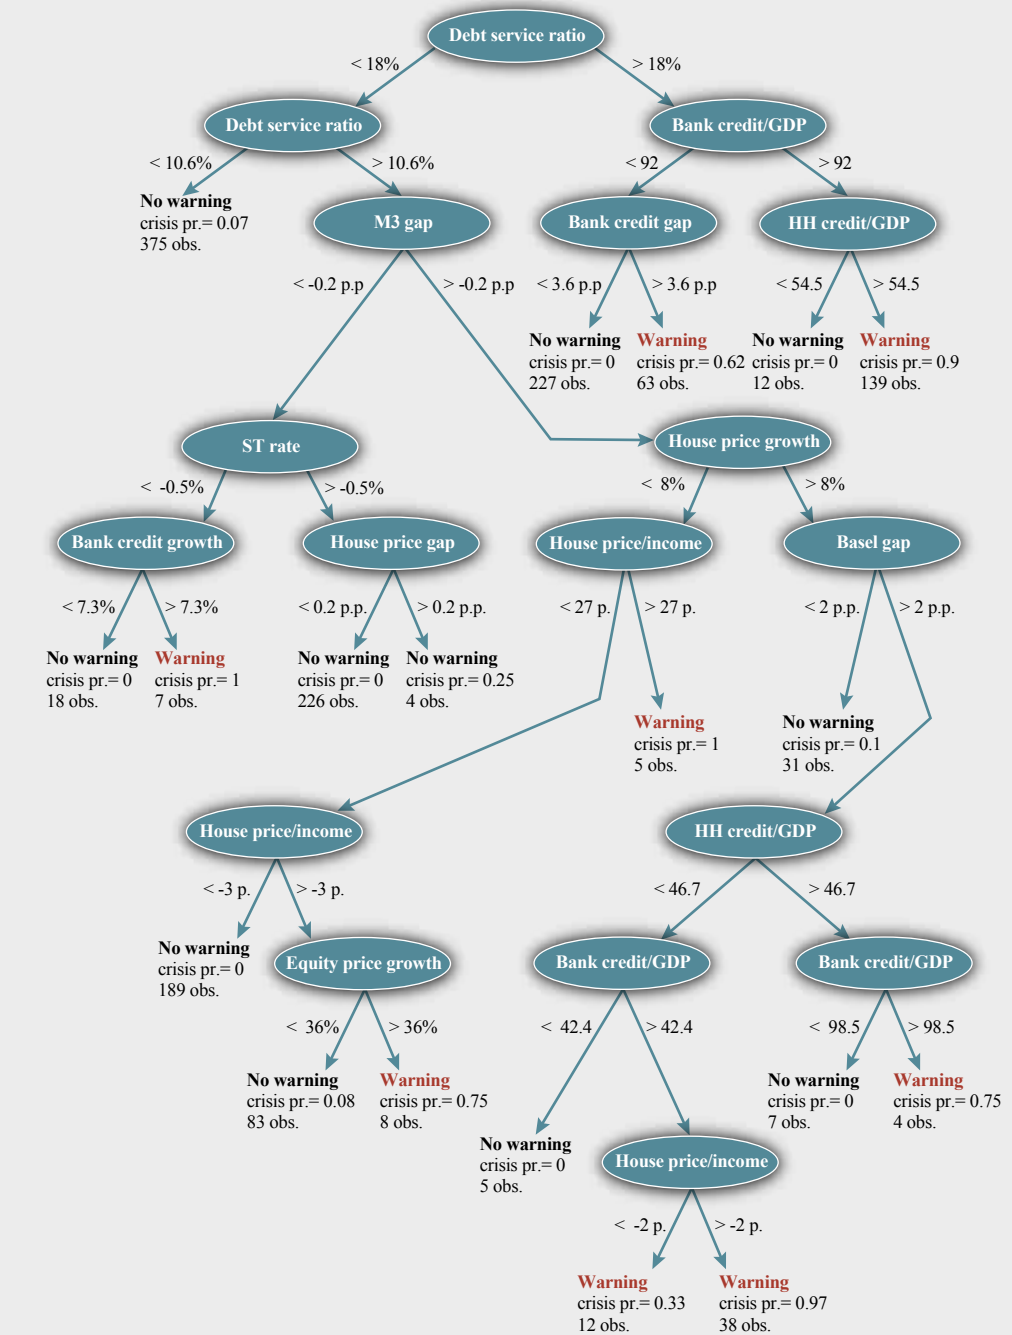
therefore useful in conditioning the information provided by gaps and rates of growth. As expected, the bank credit gap ranks relatively high, while the Basel gap, which considers broad credit, ranks lower, though still in the top half of all the indicators. Immediately following the top six indicators, there are some measures relating to house prices, namely the house price-to-income ratio, the house price gap and house price growth. Equity price growth ranks a little lower. Indeed, heated asset price growth might be associated with excessive credit growth fuelling a growing bubble. After considering the housing market, the Random Forest suggests that the real short-term rate should be looked at next, most likely because a low rate may encourage risk-taking in a search for yield. Also among the top half of all the indicators are the household debt service ratio, bank credit growth, the NFC credit-to-GDP ratio and M3 gaps.

THE EARLY WARNING TREE

Notwithstanding the remarkably good predictive performance of the Random Forest, this model is a black box and its predictions would be hard to link with a convincing narrative describing an identified risk, in particular if they would support the activation of macro-prudential tools. Therefore, this special feature also presents a benchmark tree, constructed on the set of key

*A benchmark
early warning tree
based on the key
indicators...*

Chart B.4 Early warning tree



Sources: Alessi, L. and Detken, C., "Identifying excessive credit growth and leverage", *Working Paper Series*, ECB, forthcoming.
 Note: In each terminal node (leaf) of the tree, the crisis probability corresponds to the in-sample crisis frequency associated with that particular leaf, while the number of observations indicates the number of country quarters ending up in that particular leaf, considering the historical data on which the tree has been grown.

indicators identified by the Random Forest and discussed above.¹⁰ The underlying preferences of the policy-maker are such that twice as much weight is attached to failing to identify excessive leverage compared with issuing a false alarm. To avoid overfitting, the tree shown in Chart B.4 has been grown by imposing a minimum number of eight country quarters per parent node and four country quarters per terminal node, while less relevant branches have been *pruned*.

The indicator appearing in the root node is the debt service ratio, associated with a threshold of 18%. According to end-2012 data, this threshold splits the sample equally, with around half of the countries ending up in the right branch and the other half in the left branch. The next node along the right branch of the tree corresponds to the bank credit-to-GDP ratio, with a threshold of 92. If this threshold is breached, the next relevant indicator is household credit as a percentage of GDP, with a threshold of 54.5%. At the end of 2012 a relatively large number of countries breached all of these thresholds, ending up in the “warning” leaf associated with a 90% in-sample crisis frequency. As cyclical developments might be less relevant along this branch of the tree, one could consider employing macro-prudential instruments like the systemic risk buffer to increase resilience in the system, given the elevated leverage identified by the model. However, this estimate of the probability of a crisis should be interpreted with caution for the following two reasons. The first is that the better the tree is at fitting in-sample data, the purer the leaves it will yield, with associated in-sample frequencies close to 1 or 0. However, in assessing a country’s situation, one should consider whether the relevant indicators only marginally exceed (or not) the respective thresholds. The second caveat relates to country specificities, which cannot be captured by the model. With respect to this leaf, for example, the inclusion of the debt service ratio could be misleading for specific countries that, for reasons not harmful for financial stability, have structurally high private sector debt. In such a case, a net debt concept taking into account accumulated private sector wealth would be more suitable.

If the bank credit-to-GDP threshold of 92 is not breached, the next relevant indicator is the bank credit gap with a threshold of 3.6 percentage points. If this threshold is breached, the crisis probability increases to above 60%. In this case, there would be a role for macro-prudential tools such as the counter-cyclical capital buffer.

Looking at the left branches of the tree, the main messages are as follows. If the debt service ratio is below 10.6%, the crisis probability is negligible. A relatively large number of countries, however, are in the middle range, with a debt service ratio between 10.6% and 18%. For these countries, essentially depending on the sign of the M3 gap, different variables become relevant. These indicators relate to the following: (i) house prices, in the form of house price growth and gap and in relation to income; (ii) equity prices; (iii) the Basel gap; (iv) the short-term real interest rate; (v) bank credit level and growth; and (vi) household credit. As an example, a country falling in the “warning” leaf associated with a house price-to-income ratio 27 points above its long-term average might consider adopting measures such as caps to loan-to-value and loan-to-income ratios.

With respect to the in-sample predictive performance of this benchmark tree, the true positive rate and the false positive rate (or share of type 2 errors) are equal to 85% and 4% respectively, while the share of type 1 errors is 15%. The noise-to-signal ratio is 5%. A more sophisticated measure of the usefulness of the model, taking into account the policy-maker’s greater aversion to type 1

... its in-sample
predictive ability...

¹⁰ Global liquidity variables are not included in the decision tree as they are not suited for such a model, given that they take the same value for all of the countries.

... and an out-of-sample exercise.

errors, indicates that a policy-maker using this tree increases his/her utility by 65% compared with ignoring it.¹¹

The out-of-sample performance of the Random Forest/early warning tree approach has been evaluated by using data up to the first quarter of 2006 only and ignoring whether the period starting in mid-2001 would later be classified as a pre-crisis period. Six of the eight countries for which the model would have issued a warning in mid-2006 actually experienced a crisis in the five subsequent years (France, Ireland, Spain, Sweden, Denmark and the United Kingdom). Overall, the crisis would have been correctly predicted for all of the large EU economies that did indeed later undergo one. A prompt policy reaction, assuming the current macro-prudential legislation had already been in place, would have meant, for example, that counter-cyclical capital buffers could have been raised in these countries one year before the Lehman collapse. Notably, no warning signal would have been issued for Germany, which indeed did not experience a crisis afterwards.

CONCLUDING REMARKS

Policy-makers at the national authorities responsible for macro-prudential policies in the EU as well as at the European level, i.e. at the ECB and ESRB, will have to use their judgement in setting the macro-prudential policy stance for the respective countries. This special feature describes an early warning model that can be used to support policy decisions on whether to activate macro-prudential tools targeting excessive leverage. Based on the experience of EU countries over the last 40 years, decision trees can be effective at identifying excessive credit growth and leverage with a sufficient lead time to allow for policy reactions.

The analytical models presented can serve several purposes in the policy process. First, their good out-of-sample performance should help to overcome the possible inaction bias on the part of policy-makers. In case risks are emerging which have in the past led to systemic banking crises, the onus is on those who wish to rely on judgement alone to justify why macro-prudential policy tools are not activated. Second, the intuitive nature of a decision tree model and its easy visualisation is likely to increase acceptance of an analytical approach as a starting point for policy discussions. In particular, the tool can be used to trigger discussions on country specificities affecting the risk assessment. Third, a further advantage of such a model is that, depending on the characteristics of the leaf associated with a certain crisis probability, the nature of the vulnerability can also be identified, which in many cases would then suggest the use of one specific policy instrument over another.

¹¹ See Alessi, L. and Detken, C., “Quasi real time early warning indicators for costly asset price boom/bust cycles: A role for global liquidity”, *European Journal of Political Economy*, 2011, and Sarlin, P., “On policymakers’ loss function and the evaluation of early warning systems”, *Economics Letters*, 2013.