EBA STAFF PAPER SERIES

N. 5 – 01/2020

APPLYING THE PRE-COMMITMENT APPROACH TO BOTTOM-UP STRESS TESTS: A NEW OLD STORY

by Simone Casellina, Giuseppe Pandolfo and Mario Quagliariello

EUROPEAN BANKING AUTHORITY

ABSTRACT

Stress tests have become a key tool for banks, supervisors and macro prudential authorities. An aspect of these exercises is the need for statistical models to obtain risk measurements under an adverse scenario and a fundamental question is who should develop such models. If models are developed by the authorities (top-down approach), homogeneity of treatment among banks and more control over results are achieved, but the authorities do not necessarily have all the information that individual banks have. Banks' models (bottom-up approach) may be more accurate. However, banks may have incentives to underestimate the impact of a shock, thus reducing the supervisory reaction. In this paper, we focus on bottom-up stress tests and suggest creating a system of monetary penalties (charges) proportional to the difference between the expected and the realised losses of a portfolio. The charges would aim to induce model developers to reveal their best forecasts. We show that this approach can be seen as an adaptation of the pre-commitment approach (PCA) developed and promoted by the US Federal Reserve in the 1990s but also as an application of the penalty criterion proposed by the Italian mathematician de Finetti as the foundation of the subjectivist definition of probability. We explain how the PCA could be adapted to bottom-up stress testing and provide a practical example of the application of our proposal to the banking book. What emerges is that the PCA can indeed mitigate banks' incentives to provide underestimated measures of risk under the adverse scenario and thus better align the incentives of banks and supervisors.

KEYWORDS

de Finetti penalty; principal-agent problem; pre-commitment approach; stress test; supervision

Introduction

Stress tests have become a key tool for banks, supervisors and macroprudential authorities. During the crisis, they were pivotal in identifying weak banks and recapitalising them accordingly. In normal times, they are a component of the toolkit available to supervisors for assessing banks' solvency and capital planning. A stress test aims to deliver measurements of risk related to banks' portfolios under an extreme but realistic macroeconomic scenario. An aspect of these exercises is the need for statistical models to obtain these measurements.

If models are developed by the authorities (top-down approach), homogeneity of treatment among banks and more control over results are achieved, since one single model is used. However, the authorities do not necessarily have as much information as individual banks, which could make models less accurate. Conversely, models developed by banks (bottom-up approach) are based on more granular data and better knowledge of the business. However, banks may have incentives to underestimate the impact of a shock in order to reduce the supervisory reaction. No approach is exempt from criticism and indeed there is debate on what stress tests would ideally look like.[1] The top-down model is criticised for being opaque and discouraging banks' own risk management. The bottom-up approach is considered not fully trustworthy, since it relies on banks' models, which may be less conservative given banks' incentives.

In particular, in bottom-up approaches, the bank–supervisor relationship can be seen as a principal–agent problem: the bank (the agent) develops a model and provides its outcome to the supervisor (the principal). The supervisor makes decisions relying on the bank's model under asymmetrical information. With banks becoming more and more complex, information asymmetries increase. Paradoxically, banks' complexity has also been the main motivation for moving towards the wider use of internal models in regulation. However, this has come at the cost of magnifying the principal–agent problem, particularly when models – as in the case of the stress test – are not validated by supervisors.

An alternative to a purely top-down or bottom-up approach is the hybrid approach applied in the EU-wide stress test. Following this approach, the authorities try to mitigate the agency problem by providing banks with a detailed methodology, including constraints, in the form of caps and floors, on the outputs of internal models. Along with a thorough quality assurance process, the 'constrained bottom-up approach' adopted by the European Banking Authority (EBA) is aimed at delivering conservative and comparable outcomes while maintaining reliance on banks' models. However, this approach has also been challenged, since the constraints are sometimes perceived as too strict or not realistic, with the risk of transforming a bottom-up exercise into a top-down one (Enria, 2018; Quagliariello, 2019). The question is whether there is a better way to allow the use of internal models and relax the constraints while still achieving accurate estimates of banks' capitalisation under an adverse scenario.

In this paper, we suggest an approach along the lines of what was proposed by the US Federal Reserve in the context of the use of internal models for market risk at the end of the 1990s. At that time, for the first time, the Basel Committee on Banking Supervision (BCBS) was opening the way for the use of banks' own models to quantify regulatory capital requirements, but there was debate on how to control their use. Under the pre-commitment approach (PCA) (Kupiec and O'Brien, 1995), banks could quantify their own requirements with the understanding that they would face penalties should losses exceed the pre-committed capital. We also show that the Federal Reserve proposal can be interpreted as an application of the penalty criterion suggested by de Finetti starting from the '30s as the foundation of the theory of probability following the subjectivist approach.

We review the debate on the original Federal Reserve proposal and try to extend it to credit risk. We then assess the feasibility of applying the PCA to bottom-up stress tests. In the Annex, we offer a practical example.

[1] See Quagliariello (2019) for a discussion.

A critical assumption we make is that the models used by the banks for internal purposes are the same as those used to provide the risk measures in the stress test. This implies that, for any scenario, banks will use the same model.

While it is true that banks may rather use satellite models for credit risk in a stress test, under the new accounting standards (IFRS9), they are required to compute the provisions that would be required under various future scenarios. We believe that, as a consequence, banks will increasingly use IFRS 9 models also for stress-testing purposes. In this sense, it is possible to define a bank's better forecast for the next year as that obtained through the internal model fed with the bank's own expected scenario. In this case, the comparison between the expectation and the error is acceptable. The same model is then assumed to be employed by the bank to provide the previsions conditional on different scenarios.

The pre-commitment approach

The PCA was developed by the Federal Reserve's economists Kupiec and O'Brien in 1995 (Kupiec and O'Brien, 1995). At that time, the BCBS had agreed to provide an alternative regime for setting capital requirements for the trading book based on internal value-at-risk (VaR) models. The debate during that period was on whether supervisors should set model parameters or allow banks to use their estimates. The latter approach would have produced more accurate results but could also have encouraged banks to underestimate risks. It was in this context that the PCA was proposed.

Following the PCA, a bank would specify the amount of capital needed to back its trading book over a period and would be penalised by a fine if cumulative losses exceeded this figure at any time during that period. In other words, rather than meeting the capital requirement set by the regulator, the bank itself sets the capital requirement, with the awareness that it will face penalties should its trading activities generate losses exceeding the pre-committed capital.

The proposal raised much interest. In one of his speeches, Alan Greenspan, the then Chair of the Federal Reserve (Greenspan, 1998) remarked that:

> As banks' internal risk measurement and management technologies improve, and as the depth and sophistication of financial markets increase, bank supervisors should continually find ways to incorporate market advances into their prudential policies, where appropriate. Two potentially promising applications of this principle have been discussed at this conference. One is the use of internal credit risk models as a possible substitute for, or complement to, the current structure of ratio-based capital regulations. Another approach goes one step further and uses market-like incentives to reward and encourage improvements in internal risk measurement and management practices. A primary example is the proposed pre-commitment approach to setting capital requirements for bank trading activities.

Bliss (1995) highlights that the PCA was not a model-based regulation, in that it was not intended to regulate how banks designed and developed their models, and nor did it require any supervisory validation. The PCA is all about managing incentives, while leaving modelling issues to banks.

Bliss also notes that the PCA is conceptually more appealing than the BCBS internal model approach, in which banks can use their estimates but supervisors meticulously regulate several modelling aspects. However, he also highlights the need to clarify a number of details of the practical implementation of the PCA and, in particular, of the design of the penalty system. He suggests that a penalty should be certain and hard for banks to game. In his opinion, capital-based penalties would not achieve the goal since – in the case of a breach of the pre-committed capital – the penalty would result only in an increase in the capital requirement in the future. The bank could, therefore, offset the penalty by reducing the reported pre-commitment capital for the next period. Fines, however, meet Bliss's criteria: once a pre-commitment level was breached, the bank would incur an immediate cost. The use of fines on a regular and automatic basis might appear strange, but it must be understood that these fines are not sanctions but, rather, a device for ensuring the incentive-compatible use of internal models. As well as reducing the incentive to game the rules, the PCA also allows banks to make their own decisions on risk-management investments: banks could invest resources in modelling so as to avoid fines.

Bliss's contribution is key in understanding not only the benefits of the PCA but also the inherent application challenges. Like an internal model approach, the PCA builds on banks' own risk assessment technologies. The PCA does not impose standardised parameters, and, at the same time, it facilitates – or makes redundant – the inherently difficult task of supervisory validation of internal models. Insofar as the penalty scheme creates the right incentives, banks will endeavour to make their models as accurate as possible. The most obvious difficulty is the calibration of the penalty. The penalty must be adequate to offset the incentive for a bank manager to set the capital level quite low, but it cannot undermine a bank's economic situation. Moreover, since there is no *ex ante* supervisory specification of capital levels, a bank that already finds itself in shaky circumstances may be both tempted and able to set a very low capital level and gamble on resurrection. The literature discussed in the following paragraphs focuses on some of these aspects.

Prescott (1997) notes that the PCA may be interpreted as a menu of contracts, a well-established economic concept already applied in several contexts. He suggests that a proper use of the PCA might minimise the distortions to capital holdings caused by private information and concludes that it could be beneficial, provided that fines were appropriately calibrated.

Santos (2000) examines the pros and cons of the PCA and compares it with other approaches. He observes that the PCA is more prone to causing a time consistency problem, because it applies penalties *ex post*. Regulators, for example, could be pressured to waive the penalty if it were to lead to bankruptcy. In addition, because of the limited liability of banks' managers, the approach does not protect against go-for-broke strategies. Similarly, Milne (2002) focuses on different categories of banks and concludes that, at least in the short term and for healthy banks, the threat of penalties would have a significant impact on portfolio decisions, while for 'dead walking' banks, the incentives would be ineffective.

Barth et al. (2006) observe that an attractive feature of the PCA is that it addresses concerns that the Basel approach is too rigid and arbitrary. Tarullo (2008) devotes ample space to the PCA, and an application of the PCA beyond market risk is also suggested. Colombini (2018) concludes that the use of the PCA in risk management is among the issues that need to be taken into consideration in order to gain a complete picture of the capital framework in Europe.

Interest in the PCA was so great and support from banks so vocal that the Federal Reserve decided to run an experiment involving 10 banks (Considine, 1998). According to those banks, the pilot demonstrated that the PCA was a viable alternative to the BCBS internal model approach, since it provides stronger incentives for prudent risk management and a more efficient allocation of capital than other existing capital standards. Interestingly, they also suggested that disclosure of the breach of pre-committed capital could be the appropriate penalty. In addition, they observed that, since the PCA goes directly to the basic question of whether a business holds adequate capital to absorb unanticipated losses, the PCA could cover any risk that has the potential to generate a loss.

The Federal Reserve Board formally requested comments on the PCA in parallel with the consultation on the market risk internal model approach. There was a mixed reaction: perhaps predictably, larger banks were favourably inclined to the idea, while some regional and smaller banks had reservations. However, the consultation seemed ill-timed from the outset, insofar as the other BCBS countries had just agreed to the VaR approach to market risk. The proposal was, therefore, abandoned at least as an official Federal Reserve stance.

de Finetti's penalty criterion

A key question with regard to the PCA is how the fine should actually be designed. In this respect, the seminal contribution of de Finetti on penalty criteria provides some useful guidance. According to de Finetti (1975), probability is never an objective measure, such as length and weight. Keynes (1921) had already noted that the quantification of probability depends on the information set available and that people with different information fragments attribute different probabilities to the occurrence of the same uncertain event. de Finetti takes a step further and thinks that people with the same information set can provide different probability evaluations. The probability then becomes a relationship between a particular subject (and the particular information set it has) and an uncertain event.[2]

For de Finetti, the problem is not finding the true probability of the occurrence of an uncertain event – because this measure simply does not exist – but, rather, getting the best evaluation of the probability that a subject can provide. The tricky aspect – which is similar to the agency problem in bottom-up stress tests – is that, in real life, the subject may have incentives to not reveal its true belief about the probability of the event. According to de Finetti, the revelation of true beliefs requires that the incentives be corrected through penalties.

The easiest example relates to the case of an uncertain event that can assume only two values: $E = \{0,1\}$. A subject is asked to provide an evaluation of the probability that case 1 occurs. Many practical situations can be traced back to this example. For example, a company requesting credit could be asked to provide its evaluation of the probability that a given event (the positive conclusion of a court case or the award of an important contract) will occur. Under the penalty criterion, the subject is informed that it will suffer a monetary fine proportional to $(1 - P)^2$ if the event occurs and P^2 if the event does not occur,[3] where P is the probability value provided by the subject.

de Finetti suggests fining the subject in any case, regardless of whether the event occurs or not. This appears to be necessary if the goal is to obtain the subject's best evaluation. It is easy to show that the expected cost of the fine is minimised by providing the true belief. Let us assume that the true evaluation of the probability is equal to P, but the subject decides to provide $P^* \neq P$. In this case, the subject can anticipate that the fine, f, he will pay is proportional to:

$$E(f|P) = (1 - P^*)^2 * \Pr(E = 1) + (P^*)^2 * \Pr(E = 0) = (1 - P^*)^2 * P + (P^*)^2 * (1 - P)$$

In this expression, the expected value of the fine is obtained conditional on P, which is the true evaluation of the probability. The rational subject can minimise the expected fine. The first-order condition for an extremum of $E(f|P)$ is:

$$\frac{\partial E(f|P)}{\partial P^*} = -2(1 - P^*) * P + 2(P^*) * (1 - P) = 0$$

[2] Nau (2001): 'The subjective theory of probability, which is now widely accepted as the modern view, is jointly attributed to de Finetti (1928/1937), Ramsey (1926/1931), and Savage (1954). Ramsey and de Finetti developed their theories independently and simultaneously, and Savage later synthesised their work and also incorporated features of von Neumann and Morgenstern's (1944/1947) expected utility theory. All three authors proposed essentially the same behavioural definition of probability, namely that it is a rate at which an individual is willing to bet on the occurrence of an event. Betting rates are the primitive measurements that reveal someone else's probabilities, which are the only probabilities that really exist.'

[3] Piccinato (2009), Annex A.

The above equality is satisfied for $P^* = P$ and since $\partial^2 E(f|P)/\partial P^{*2} = 2 - P > 0$ because $0 \leq P \leq 1$ then the extremum is a minimum.

Less formally, the penalty needs to be symmetrical, since, otherwise, a distortion will be introduced and the provided evaluation, P^*, will not be equal to the true belief, P. More importantly, the difference will depend on the level of risk aversion of the subject, which is idiosyncratic and not predictable.

It is important to highlight that, following de Finetti's approach, the penalty is only a tool for aligning incentives and does not imply any negative judgement. Since the use of terms such as 'fine' or 'penalty' is likely to induce some misinterpretation of our proposal, we will use the terms 'incentive mechanism' and 'charge'.

A practical application of the penalty criterion to drilling projects for oil exploration can be found in Grayson (1960). Since such projects are quite costly and incur a risk of not finding oil, the person responsible for the project (the decision-maker) asks a geologist for advice. Grayson points out that the decision-maker cannot ask the geologist for an exact prediction or prophecy; rather, the decision-maker asks for an evaluation in terms of probability. The problem for the decision-maker is how to induce the expert to provide advice in line with his or her own genuine belief. If his or her remuneration is linked to potential profits, he or she could have an incentive for over-optimism, since he or she would take the upside but not the downside. On the other hand, the idea of participating in potential economic losses if oil were not found could lead him or her in the direction of over-pessimism. In this situation, the penalty criterion can mitigate the problem. The geologist knows that a penalty will be applied in either case, but he or she can minimise it by giving his or her best evaluation.

The application of the pre-commitment approach to bottom-up stress tests

In bottom-up stress tests, banks are requested to provide supervisors with their true beliefs about the losses that would occur under a baseline and an adverse scenario. However, their incentives are such that they may tend to provide supervisors with an underestimation of those losses (Quagliariello, 2019). The setting is thus perfectly suited to be interpreted in the light of the PCA and the de Finetti penalty criterion: through a system of charges, supervisors could induce banks to reveal the true outcomes of their best models, because this would reduce the *ex post* monetary cost.

The application of the PCA to stress-testing models also seems less controversial than its application to models for setting minimum capital requirements. In particular, there were two concerns about the PCA: (i) the capital levels autonomously set by banks could be too low, resulting in insolvency; and (ii) application to an area other than market risk, where backtesting is relatively straightforward, could go too far.

As regards the first issue, stress tests are intended to be used not to set minimum capital requirements but, rather, to assess capital planning and, possibly, require banks to hold additional buffers or reach certain capital targets to meet supervisory expectations. In the EU, for instance, competent authorities use the findings of stress tests as input into the supervisory review and evaluation process (EBA, 2019). Therefore, providing more freedom to banks *ex ante* is less likely to result in a breach of minimum capital requirements.

With regard to its application to other risks, the risk of losses can always be split into the frequency of the event and the severity of the losses. In credit risk, for instance, this division corresponds to probability of default (PD, frequency) and loss given default (LGD, severity). On closer inspection, the 'opaqueness of credit losses' mentioned by Tarullo (2008) as a possible obstacle to the application of the PCA to credit risk is more connected with severity, since recovery procedures can take many years, while the frequency side is less problematic.

In practice, we propose making use of the PCA to induce banks to reveal their best forecast models for the probabilities of default. While we focus on credit risk, the approach can be generalised to those components of the stress test that can be more easily backtested. This could also contribute to a discussion on which risks should be modelled by banks – under the PCA incentives – and which would be better left to top-down modelling.

Let L_{t+h} be a measurement of losses between t and h. At time t, L_{t+h} is not known and only a forecast is available. This forecast is based on three elements. First, the information set available at t, Ω_t. Second, a model that we identify as *M*. For example, *M* can be a linear combination of some variables, i.e. $\sum_i \beta_i x_i$ with $x_i \in \Omega_t$, and the set of weights (the betas) are attributed through some estimation process. The third element is the definition of a scenario. If, for example, y_t is the value of a stock index, the price of a commodity such as oil or the level of gross domestic product (GDP) of a country, $E_t(y_{t+h})$ is the baseline scenario defined in t for the period $t+h$.

For a given subject i (the bank or, better, the risk manager), the best forecast of L_{t+h} will be a function of the information set available to him or her, Ω_t^i, the model he or she has decided is most appropriate, M^i, and his or her choice of scenario, $E_t^i(y_{t+h})$:

$$E_t^i(L_{t+h}) = f\left(\Omega_t^i, M^i, E_t^i(y_{t+h})\right)$$

Under a penalty criterion, the bank would be informed that the value $E_t^i(L_{t+h})$ was to be compared with the realised L_{t+h} and that the bank would be subject to a charge proportional to the difference between $E_t^i(L_{t+h})$ and L_{t+h}.

In practice, in a supervisory stress test, the scenario would be provided directly by the authority, which would mean that $E_t^i(y_{t+h})$ would be replaced with $E_t^{AUT}(y_{t+h})$. Banks would thus be required to produce the expected figures conditional on different scenarios – their own and the supervisory one – always using the same model, and to report on the outcomes using standardised templates.

Going back to the comparison between the previsions or forecasts and realised losses, let us define the forecast error as:

$$\epsilon_{t+h} = L_{t+h} - E_t^i(L_{t+h})$$

$E_t^i(L_{t+h})$ depends on $E_t^i(y_{t+h})$, which is the expected scenario for $t+h$. So $E_t^i(L_{t+h})$ is the expectation about the loss amount under the baseline scenario provided by the bank. It would be harder to define an appropriate comparison term for the expectation of the loss amount conditional on the adverse scenario.

The charge would be proportional to the quantity ϵ_{t+h}. This aspect is central and deserves some comment. It is clear that the forecast error could be due to several unpredictable events, such as mergers and acquisitions, the sale of significant portions of the portfolio, etc. This could lead to think it is necessary to disentangle the genuine error from the error induced by unforeseeable events. The problem is that it is not easy to define, in practice, which events are truly unforeseeable, and it is even more difficult to do it *ex ante,* i.e. before these events are observed.

However, in order for a bank to anticipate the amount of the charge today and therefore take it into account in its decisions, it is necessary that there are no doubts about how this will be calculated. If the computation of the charge requires to wait and observe the forecast error and disentangle its components, it would not possible to evaluate in advance the charge. Furthermore, disputes between banks and the authorities would open up about what was predictable and what was not, and this would lead to some confusion.

The charge is a mechanism aimed at providing incentives to reveal a true opinion (in this case a forecast). What we want to eliminate is not the forecast error but the difference between the true opinion and that revealed.

However, it is possible to decompose the forecast error. With $f\left(\Omega_t^i, M^i, E_t^i(y_{t+h})\right)$ being the forecast at t given the information set available at that time, Ω_t^i, at $t+h$ it is possible to compute the quantity $f(\Omega_t^i, M^i, y_{t+h})$, which is the forecast that the model would have produced given perfect foresight relative to the scenario. Then, it could be possible to define the charges as a proportion of the difference between the realised losses, L_{t+h}, and the output of the model, given the realised scenario, i.e. $f(\Omega_t^i, M^i, y_{t+h})$. In this way, banks could be penalised only for shortcomings in the model and not for not having correctly forecast the scenario, which is irrelevant in a stress test, as the scenario is provided by the authorities themselves.

Although we believe that this decomposition could be useful in understanding the sources of the error, we still think that the charge should be based on a simple comparison between forecast and realised losses. This is because, even though we could easily agree that there are many unpredictable factors that could justify a forecast error, once we tried to draw up operational/practical definitions of things such as unexpected events or exceptional situations, we could not reach agreement. Consider the case of mergers and acquisitions: in what circumstances would we accept that these were unpredictable events?

The problem is that the system we are suggesting works if and only if all the rules are perfectly known at the beginning. Saying the amount of the charge would be determined only *ex post* based on necessarily qualitative considerations about the 'predictability' of the events would shift the efforts of the banks from trying to provide *ex ante* their best forecasts to trying to convince the authorities *ex post* that some event was 'exceptional'.

As regards the concrete definition of the quantity L_{t+h}, it would be preferable for the charges to be proportional to a unique, final, monetary quantity than to granular model parameters. A good candidate is the amount of new provisions. In broad terms, the provisions are meant to cover the expected loss amount, and, in turn, this quantity can be expressed as:

$$E_t^i(L_{t+h}) = E_t^i(EAD_{t+h}) * E_t^i(DR_{t+h}) * E_t^i(LGD_{t+h})$$

Where $E_t^i(EAD_{t+h})$ is the expected level of exposure entered in default, $E_t^i(DR_{t+h})$ is the expected default rate, i.e. the expected transition rate from non-defaulted to defaulted status,[4] and $E_t^i(LGD_{t+h})$ is the expected loss rate. Although different models can be developed for different sub-portfolios, the results of these models, once expressed in terms of expected loss, can always be summed.

As we have already discussed, the LGD is more difficult to back-test, because the effective loss rate can be known only at the end of the recovery process, which can take a long time. Keeping this aspect in mind, our proposal is to compare the forecast expected loss amount with the following *mixed* quantity:

$$\widetilde{L_{t+h}} = EAD_{t+h} * DR_{t+h} * E_t^i(LGD_{t+h})$$

This quantity represents a proxy of the true loss amount realised in $t+h$, where the expectations for the exposure at default (EAD) and the default rate (DR) are replaced with the observed values, while the loss rate is kept constant. Alternatively, it would be possible to replace $E_t^i(LGD_{t+h})$ with an updated forecast produced in t+h, i.e. $E_{t+h}^i(LGD_{t+h})$, but this would not be fair, given that in that case we would be comparing two forecasts based on two different information sets.

To sum up, the forecast error would be defined as:

$$\epsilon_{t+h} = \widetilde{L_{t+h}} - E_t(L_{t+h}) = E_t(LGD_{t+h})[EAD_{t+h} * DR_{t+h} - E_t(EAD_{t+h}) * E_t(DR_{t+h})]$$

[4] It could represent the transition rates from stage 1 to stage 2 and stage 3, using IFRS 9 concepts.

The above quantity is expressed in monetary terms and could be summed across different portfolios and statuses (or stages). The charge would be proportional to this quantity:

$$f \approx \delta * (\epsilon_{t+h})^2$$

where δ is a parameter set by the authority. The quadratic transformation of the error is in line with the principles set by de Finetti and the preference for a symmetrical penalty. This implies that the bank[5] is charged when estimates are not accurate, regardless of whether they are more or less conservative than the realised losses. This may be counterintuitive for stress-test practitioners, banks and supervisors, who expect conservatism to be rewarded. The point, however, is to incentivise banks to provide accurate, rather than conservative, estimates. Adopting a symmetrical approach is also consistent with the general principles regarding the application of IFRS 9 accounting set out by the BCBS. The BCBS (2015), for example, highlights that 'the Committee recognises that [expected credit loss] ECL accounting frameworks are symmetrical in the way that subsequent changes (both deteriorations and reversals of those deteriorations) in the credit risk profile of a debtor should be considered in the measurement of the allowances'. We recognise that this aspect is open to debate and in this paper we suggest how the penalty system could be supplemented with additional interventions to penalise more those banks that systematically underestimate losses.

Now consider the possibility that a bank decides to provide the authorities with a distorted version of its predictions: the bank has a genuine prediction of the probable losses obtained using its best model, $E_t(L_{t+h})$, but decides not to reveal it to the authority.

Instead of providing the expected losses stemming from its true model, $E_t(L_{t+h})$, the bank decides to provide systematically lower figures. For simplicity, we can imagine that the bank employs a model such as $E_t\widehat{(L_{t+h})} = E_t(L_{t+h}) - d$, where its true expectations about losses conditional on any scenario are reduced by a constant quantity. We then have two models, the bank's true-best model, $E_t(L_{t+h})$, and the model used to provide risk measures to the authority, $E_t\widehat{(L_{t+h})}$. Clearly, the amount of the charge will be computed based on the risk measures officially provided, that is: $\epsilon_{t+h} = \widetilde{L_{t+h}} - E_t\widehat{(L_{t+h})}$.

Knowing that it will be charged, the bank will try to figure out how much the charge could be. The expected amount of the charge is $E_t[f] \approx \delta E_t[(\epsilon_{t+h})^2] = \delta E_t\left[\left(\widetilde{L_{t+h}} - E_t\widehat{(L_{t+h})}\right)^2\right]$, where, obviously, the realised losses $\widetilde{L_{t+h}}$ cannot be known in advance. However, the bank can rationally employ its best forecast to evaluate *ex ante* the expected amount of the charge.

$$E_t[f] \approx \delta E_t\left[\left(E_t(L_{t+h}) - E_t\widehat{(L_{t+h})}\right)^2\right]$$

Under the simplified example, $E_t\widehat{(L_{t+h})} = E_t(L_{t+h}) - d$, $E_t[f]$ is proportional to d only, and this quantity represents the level of gaming that the bank has decided to employ.

It is assumed that the bank will try to minimise the charge, and given that, *ex ante*, the expected amount of the charge is proportional to the difference between the bank's true prediction, $E_t(L_{t+h})$, and the measurement provided to the authority, $E_t\widehat{(L_{t+h})}$, this should create an incentive to reduce the difference between the two quantities.

It is worth noticing that the proposed system of charges in practice provides an incentive to use models that are able to minimise the expected mean square error – that is, a classical model selection approach.

[5] For now, we refer generically to the bank; we refer to Section 6 for a broader discussion on the calibration of fines and who should pay them.

It is interesting to compare our proposal with the backtesting approach (BTA) proposed by the BCBS (2005) for the validation of models for PD, in the context of the internal ratings-based system. Under the BTA, it is recognised that a bank's PD estimates will usually differ from the default rates that are afterwards observed, and the key question is whether the deviations are purely random or occur systematically. A systematic underestimation of PD merits a critical assessment. Four approaches are typically used: the binomial test, the chi-squared test, the normal test and the traffic lights approach. As stated by Bliss (1995), the BTA is also a non-model-based system of regulation. Indeed, both the PCA and the BTA are more interested in models' performance than in understanding/regulating how models are built. In other words, the focus is more on *if* the models work and not on *how* they work. The PCA and BTA may appear quite similar, but they are not. There are three main differences between the BTA and our version of the PCA.

(1) The BTA requires that one or more thresholds be defined. Indeed, with this approach the difference, ϵ_{t+h}, between the expected and the realised default rate must be classified as *normal* or *abnormal*, and this classification is obtained by comparing ϵ_{t+h} with predefined thresholds. Setting these boundaries implies the necessity of both relying on a model and arbitrarily setting the confidence level. The simplest approach to setting such thresholds is to refer to the binomial model under the hypothesis of independence of the defaults and infinite granularity, but this model may be considered too unrealistic.[6]

(2) With the BTA, it is not clear what the consequences of breaking the thresholds should be. The obvious remedy in this situation is to develop a new model in the hope of obtaining better performance in future years, perhaps requiring additional conservatism in the estimates as a temporary solution. However, developing and validating a new model is a long process, and there is no guarantee that the new model will deliver lower measures of risk. In practice, it would not be possible to include *ex ante* such evaluations in the decision process. By contrast, under our proposal, the impact of the forecast error ϵ_{t+h} is clearly and easily quantifiable *ex ante,* even if realised figures are known only *ex post*. Indeed, the bank is supposed to try to minimise the expected charge, a quantity that, *ex ante*, depends on the difference between the bank's true prediction, $E_t(L_{t+h})$, and the measurement provided to the authority, $\widehat{E_t(L_{t+h})}$.

(3) The BTA implies a judgement about the bank's model. If ϵ_{t+h} has been higher than the upper threshold for a period, the authorities will say that the model is 'wrong'. This kind of judgement somehow supports the simplistic idea that a better model, maybe a perfect one, is possible, although the forecast was the best possible model given the information available. With our proposal, no judgement is formulated, as it is left to the banks to try and minimise the charges as best they can.

The design and calibration of the charges

The design of the charges can be very simple, as shown in the previous section:

$$f = \delta * \left(L_{t+h} - E_t(L_{t+h})\right)^2$$

More complex approaches could be developed, with the aim of correctly aligning banks' incentives, but this, in turn, would require knowing precisely banks' cost function. In the absence of such information, a simple function, such as that above (essentially the one originally suggested by de Finetti), could represent a valid solution.

We have already discussed our preference for a symmetrical mechanism behind the quantification of the charges. Without symmetry, a bank would have incentives to reveal a prudent evaluation and not its best

[6] Tasche (2003) discusses how to set the thresholds under different models.

forecast, i.e. its true opinion. In particular, without a fair/unbiased measure of risk, the comparison with realised figures would not be reliable. The quantity ϵ_{t+h} (the difference between realised and expected losses) would depend partly on the risk aversion of the models' builders, but this effect cannot be quantified. In addition, conservatism is not good per se. It can be a safeguard if there is uncertainty about models' outcomes, but it is second best compared with the best forecast. The point here is to provide banks with incentives to develop accurate models and reveal them. Producing a prudent evaluation is a simpler task than producing an unbiased measure. A prudent evaluation can be obtained by starting with quite a rough model and then increasing its results by some quantity.

However, this approach does not prevent supervisory authorities from monitoring, over time, the gap between realised and expected losses and taking action if systematic underestimation is observed (i.e. if the hypothesis $E(\epsilon_{t+h}) > 0$ cannot be rejected). These could include reviewing the calibration of the charges or even replacing a bank's models with external ones.

Another issue is the amount of the charge: it must be adequate to offset banks' incentives to provide underestimates of the risk, but it should not cause banks to hold high levels of reserves inefficiently that could otherwise be used in a more productive way. On the precise calibration of δ, we believe that this should be left open to discussion at this stage. Nevertheless, we would like to suggest some criteria.

We believe that the penalty should hit the person directly responsible for the model's development. Here, we follow Varotto and Daripa (1997), who argue that Kupiec and O'Brien disregard the issue of incentives of the banks vis-à-vis those of their managers. A key feature of large modern banks is the separation of the owners from day-to-day decision-making, and they conclude that, in the case of agency problems, the incentives may not reach the banks' decision-makers unless they are hit directly by the charge. Therefore, the charge could be proportional to a share of annual bonus of the chief risk officer (CRO), for example $f \leq$ 5%*bonus, so that:

$$\delta \leq 5\%\ BONUS/\left(L_{t+h} - E_t(L_{t+h})\right)^2$$

The last expression would require L_{t+1} to be known, but this is not possible in t, and waiting to observe this value would mean not providing *ex ante* to the banks all the information needed to evaluate the charge. This, in turn, would prevent them from considering it in their decision-making process. An alternative would be to set a rough upper level for $\epsilon_{t+h} = L_{t+h} - E_t(L_{t+h})$, such as 1.5*$\epsilon_t$, where ϵ_t is the forecast error observed the year before. In this case, $\delta \leq 2{,}2\% * bonus/[\epsilon_t]^2$.

The charges should be paid by the risk manager, who could, for example, donate the amount to charity, and should not be translated into an increase in capital requirements or provisioning level. The PCA is based on the idea that the charge should be paid at all times, even for minor forecast errors. However, in practice, a materiality threshold under which the charge need not be paid could be set.

A last point is whether the amount of the charge paid should be published along with the outcomes of the stress test. In case of publication, it should be clear that the charge itself does not represent a judgement regarding the staff involved in the models' development but, rather, is a device for making the bottom-up stress test incentive-compatible. Indeed, if the charges are correctly calibrated and the incentives aligned, the charge will reflect only an unavoidable forecast error.

Clearly, it is possible that, even if the system of charges were able to align the incentives of the CRO with those of the supervisors, decision layers above the CRO (e.g. the chief executive officer, the board and shareholders) could be in a position to override the CRO. However, the problem is not if this system is able to completely offset the 'wrong' incentives of banks' stakeholders. The problem is if the system reduces these incentives. If so, a second-best solution has been achieved. Moreover, we suggest that the results of a comparison between predicted and realised losses should be published. This would generate an additional, non-monetary incentive that could be substantial, as it would have an impact on not the remuneration but the reputation of the CRO and consequently that of the bank.[7]

Conclusions

In this paper, we suggest applying the PCA to bottom-up stress tests to mitigate banks' incentives to game the exercise. Focusing on credit risk, we propose introducing a system of monetary charges proportionate to the difference between the expected losses and the realised losses of a credit portfolio, with the aim of inducing model-builders to reveal their best forecasts.

The working assumption that we have made is that risk measures for the stress test are built using the same models adopted by a bank for internal/accounting purposes. Under the new accounting framework set out by the IFRS 9 principles, banks are required to base the level of provisioning on the ECL. The ECL is a point-in-time measure of risk that should reflect the expected scenario. In this sense, the ECL is a genuine prediction/forecast, and it makes sense to compare it with realised figures (backtesting).

A further assumption that we make is that models are relatively linear and that a good/bad model in the scenario selected by banks is equally good/bad under the extreme scenarios that are typical in stress testing. On this point, we would like to be a bit provocative. Once we know banks have the right incentives not to game the stress test, we could possibly move towards scenarios that are less extreme. This would also make our linearity assumption more realistic.

The charge should be symmetrical – i.e. it should be linked to any estimation error, regardless of whether the realised losses are higher or lower than the estimated ones – as long as it does not jeopardise the stability of the bank if systematic problems with the models are identified. Given agency problems in banks, the charge should affect directly the risk manager in charge of developing the model. This monetary charge should not be interpreted as a judgement on the model itself but simply as a device for better aligning incentives and ensuring that banks provide their best estimates of losses. Bad models will remain bad, but banks with good models will be incentivised to fully reflect their outcomes in the results submitted to the authorities. The charge would most likely not work on its own, but it should be part of a toolkit. Authorities could – and should – still intervene if there is systematic underestimation of risks, possibly taking measures beyond the monetary fee.

This is why this proposal should not be interpreted as a step in the direction of deregulation. On the contrary, it can be used to complement existing tools, including benchmarking and quality assurance.

The design and calibration of the system of charges would clearly require additional work and discussion. An important aspect is the degree of disclosure of the approach, the possible publication of the charges paid by each risk manager and banks' responses to models' shortcomings. The publication of the difference between forecast and actual losses could even be enough, without any need to impose a monetary charge. We expect a lively debate on if and how to communicate to external stakeholders the level of the charges.

[7] This proposal also seems in line with the Ramsey–Savage version of the penalty criterion (see Nau, 2001, for a comparison between de Finetti and Ramsey and Savage). All three authors proposed essentially the same behaviouristic definition of probability, namely that it is a rate at which an individual is willing to bet on the occurrence of an event. While in de Finetti's theory bets are for money – your estimate of the probability of an event is effectively the price that you are willing to pay – Ramsey and Savage, recognising that the marginal utility of money could vary, introduced the notion of a 'consequence'.

We realise that there are still several implementation questions to be addressed before making the proposal operational, but we see some merit in discussing a completely different approach to ensuring the reliability of stress-test outcomes. In this respect, the PCA applied to bottom-up stress tests can be seen as the way forward to aligning supervisors’ and banks’ incentives and increasing the reliability of the results.

Annex – a practical example

In the following example, for simplicity, we assume that the LGD parameter is constant (we can imagine that the bank in question is one that uses the foundation internal ratings-based approach) so that there is no need to have a model for it and so that the EAD is set equal to the value of the exposure at the beginning of the period. In this context, the only uncertainty is about the PD.

Let us assume that the default rate generated each year by a bank's credit portfolio obeys a well-defined model such as the following:

$$DR_t = d + \rho DR_{t-1} + \theta y_t + \varepsilon_{1,t} \qquad (1)$$

where:

DR_t default rate of the portfolio at time t

y_t macroeconomic variable, for example the GDP growth rate

$|\rho| < 1$ … no unit roots (a simplifying assumption)

$\theta < 0$ if the GDP increases (i.e. $y_t > 0$) then the default rate diminishes.

The macroeconomic variable y_t follows a simple AR(1) stochastic process.

$$y_t = \beta_0 + \beta_1 y_{t-1} + \varepsilon_{2,t} \qquad (2)$$

$|\beta_1| < 1$ again to avoid unit roots

$$\varepsilon_{2,t} \sim N(0, \sigma_i^2)$$

We now assume that model (1) is the private information of a bank, while model (2) is widely known. In this way, we are depicting a situation characterised by the presence of asymmetrical information.

From model (2), we obtain:

$$E[y_t] = \frac{\beta_0}{1-\beta_1}$$

i.e. the long-run growth rate.

$$V[y_t] = \frac{\sigma_2^2}{1-\beta_1^2}$$

The stress scenario is set by the authorities as the lower GDP growth rate observable with a given confidence level.

$$E_t^{AUT}(y_{t+1}) = q_\alpha(y_{t+1}) = y^* | P(y_{t+1} < y^*) = 1 - \alpha$$

Then being:

$$y_t \sim N\left(\frac{\beta_0}{1-\beta_1}, \frac{\sigma_2^2}{1-\beta_1^2}\right)$$

We have:

$$E_t^{AUT}(y_{t+1}) = y^* = \frac{\beta_0}{1-\beta_1} + \Phi^{-1}(\alpha)\sqrt{\frac{\sigma_2^2}{1-\beta_1^2}}$$

The authorities ask the bank to provide a forecast for next year's default rate and to do it using its own model, including its own scenario. The bank knows the true model given by (1) so that its best forecast is:

$$E_t^{bank}(DR_{t+1}) = d + \rho DR_t + \theta E_t^{bank}(y_{t+1}) = d + \rho DR_t + \theta(\beta_0 + \beta_1 y_t)$$

Notice that, since we assumed that (2) is also known to the authorities, the forecast scenario for y_{t+1}, i.e. $E_t^{bank}(y_{t+1})$, is obtained from (2). In other words, we are saying that the bank and authorities agree on the baseline scenario, but this is only a simplifying assumption that we have adopted to keep the example simple.

Now the bank evaluates the possibility of gaming. For example, it can provide a default rate that is systematically lower than its forecast. The bank then decides to provide the results from the following model, where the internally estimated default rate is systematically reduced by a quantity equal to b:

$$E_t^{\widehat{bank}}(DR_{t+1}) = E_t^{bank}(DR_{t+1}) - b = d - b + \rho DR_t + \theta(\beta_0 + \beta_1 y_t)$$

Here, the only limit is given by a principle of minimum credibility from which the bank derives the lien $d - b \geq 0$. In practice, the model will not appear credible if it produces zero default rate in the event of zero growth rate of GDP.

Given the stress scenario set by the authorities, i.e. $E_t^{AUT}(y_{t+1}) = y^*$, the model provided by the bank produces the following stressed default rate:

$$E_t^{\widehat{bank,AUT}}(DR_{t+1}) = d - b + \rho DR_t + \theta y^*$$

$E_t^{\widehat{bank,AUT}}(DR_{t+1})$ is used to emphasise that this forecast is obtained with the model provided by the bank ($bank$) but that it is the model ($\widehat{...}$) corrected with the parameter b and fed with the scenario provided by the authorities (AUT). The bank forecast would have been $d + \rho DR_t + \theta y^*$ had the bank provided its best model.

The following quantity is the loss amount estimated given the stress scenario:

$$L_{stress} = P_t * E_t^{\widehat{bank,AUT}}(DR_{t+1}) * LGD$$

where P_t is the monetary value of the portfolio and LGD is the estimated loss rate given default (severity).

Given the portfolio value, the LGD and the forecast of the default rate under the baseline scenario provided by the bank, the credit risk adjustments (CRA) amount to:

$$L_{baseline} = CRA_t = P_t * E_t^{\widehat{bank}}(DR_{t+1}) * LGD$$

Comparing the loss under the scenario with the CRA and capital buffer, it is possible to define a shortfall of resources as:

$$k = L_{stress} - (CRA_t + C - C_{min}) = L_{stress} - CRA_t - (C - C_{min})$$

where the minimum capital C_{min} includes the minimum required capital and regulatory/supervisory buffers, while the capital is $C = C_{min} + c * P_t$ so that $c * P_t$ is the additional capital buffer voluntarily held by the bank.

The quantity k is meant to ensure that provisions and the additional capital buffer are able to cover losses under the stress scenario. If k is positive, the bank could be required to increase its capital or the level of provisioning by this amount. Rearranging, we obtain:

$$k = P_t * LGD * \left(E_t^{\widehat{bank,AUT}}(DR_{t+1}) - E_t^{\widehat{bank}}(DR_{t+1}) \right) - c * P_t =$$

$$= P_t * [\theta(y^* - \beta_0 - \beta_1 y_t) * LGD - c]$$

Now we need to define the objective function. We assume that the shareholders are also the decision-makers. In this case, it is safe to imagine that the objective function is the return on capital. In a more complex scenario, the decision-maker, for example the risk manager, could have a different objective function. However, assuming that the risk manager's bonus is connected in some way to the return on capital, the following analysis is still valid.

We define the return on capital as follows:

$$R_t = \frac{P_t r - CRA_t}{C + k} = \frac{P_t \left(r - E_t^{\widehat{bank}}(DR_{t+1}) * LGD \right)}{P_t(8\% * rw + c) + P_t \, [\theta(y^* - \beta_0 - \beta_1 y_t) * LGD - c]}$$

and simplifying it a bit:

$$= \frac{r - [d - b + \rho DR_t + \theta(\beta_0 + \beta_1 y_t)] * LGD}{8\% * rw + \theta(y^* - \beta_0 - \beta_1 y_t) * LGD}$$

by placing as:

$$A = r - \left(d + \rho DR_t + \theta(\beta_0 + \beta_1 y_t) \right) * LGD$$

$$B = 8\% * rw + \theta(y^* - \beta_0 - \beta_1 y_t) * LGD$$

We arrive at:

$$R_t = \frac{A + b * LGD}{B}$$

If the bank maximises the return on capital, we have the following constrained optimisation problem:

$$max_b R_t = \frac{A}{B} + \frac{LGD}{B} b$$

$$-b \geq -d$$

$$LGD > 0$$

Notice that without the constraint $LGD > 0$, the estimated losses would be null for any level of the default rate.

The target function $R_t(b)$ is a line upward shaped (see figure below). The quantity A/B is the return the bank may obtain without cheating at all (i.e. with $b = 0$). As long as b increases, the return also increases. Without the credibility lien, the bank could increase the return indefinitely.

The Lagrangian associated with the problem is:

$$L = \frac{A}{B} + \frac{LGD}{B} b - \lambda(b - d)$$

and the Karush–Kuhn–Tucker (KKT) conditions are:

$$\frac{\partial L}{\partial b} = \frac{LGD}{B} - \lambda = 0; \lambda \frac{\partial L}{\partial \lambda} = \lambda(b - d) = 0$$

$$\lambda \geq 0; b \leq d; LGD > 0$$

The case $\lambda = 0$ (i.e. a non-binding constraint) is not acceptable, since it implies $LGD = 0$; the case $\lambda > 0$ is acceptable and it implies $b = d$ e $\lambda = LGD \backslash B$. In brief, the only solution is $b = d$, that is, the bank optimises its position by pushing the level of cheating to the maximum.

Now let us introduce a charge proportional to the forecast error and the portfolio value:

$$f_{t+1} = \delta * \left(L_{t+1} - E_t(L_{t+1})\right)^2 = \delta * (\epsilon_{t+1})^2$$

where δ is a proportionate parameter set by the authority and $E_t(L_{t+1})$ is the forecast provided by the bank.

$$f_{t+1} = \delta * (P_t * LGD)^2 * \left(DR_{t+1} - \widehat{E_t^{bank}(DR_{t+1})}\right)^2$$

In this expression, LGD and P_t are fixed and known.

The charge level will be known only in t + 1 when DR_{t+1} is observable. However, the bank will try to estimate its value in advance and include it in its decisional process. For the bank, the best forecast for DR_{t+1} is its true model, i.e. $E_t^{bank}(DR_{t+1})$, so that the expected charge depends on how much it decides to game.

$$E_t\left(f_{t+1} | E_t^{bank}(DR_{t+1})\right) = \delta * (P_t * LGD)^2 * (b)^2$$

Assume the bank discounts the charge so that the return on capital becomes:

$$R_t^f = \frac{A + b * LGD - E_t\left(f_{t+1} | E_t^{bank}(DR_{t+1})\right)(1 + i)^{-1}}{B} = \frac{A}{B} + \frac{LGD}{B} b - \frac{C}{B} b^2$$

where $\delta * (P_t * LGD)^2 * (1 + i)^{-1} = C$

Now the target function is a parabola in respect to b. This means that, as b increases, the return is affected by a positive component LGD/B proportional to b but also by a negative element C/B proportional to b squared.

The Lagrangian associated with the problem is:

$$L = \frac{A}{B} + \frac{LGD}{B} b - \frac{C}{B} b^2 - \lambda(b - d)$$

and the KKT conditions are:

$$\frac{\partial L}{\partial b} = \frac{LGD}{B} - 2\frac{C}{B}b - \lambda = 0$$

$$\lambda\frac{\partial L}{\partial \lambda} = \lambda(b - d) = 0$$

$$\lambda \geq 0; b \leq d; LGD > 0$$

The case $\lambda = 0$ (non-binding constraint) is now admissible, and indeed it implies $b < d$, that is, the bank will set the level of gaming below the maximum admissible level:

$$\frac{LGD}{B} = 2\frac{C}{B}b$$

i.e.

$$b = \frac{LGD}{2C} = \frac{LGD}{2\delta * (P_t * LGD)^2 * (1+i)^{-1}} = \frac{1+i}{2\delta * LGD * P_t^2}$$

Notice that now the level of gaming also depends on the parameter δ, which is under the control of the authorities. In other words, the presence of the charges makes the bank control variable dependent on the control variable of the authorities. This puts the authorities in a position to reduce b to an acceptable value.

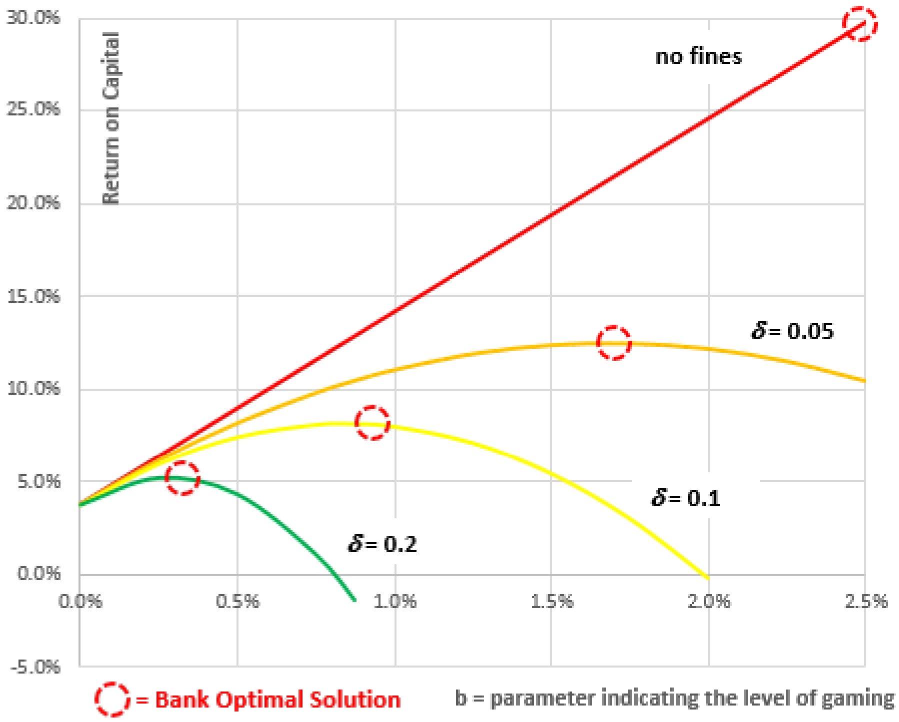

The penalty reduces the economic incentive to game, i.e. to underestimate the losses. If we add to this that the charge could be accompanied by the publication of the results (e.g. a ranking of how much the banks have paid), then the incentive to game could become null.

References

AIFIRM (2016), *Position paper 7: validation of rating model calibration*, Associazione Italiana Financial Industry Risk Managers, Milan (http://www.aifirm.it/wp-content/uploads/2016/12/2016-Position-Paper-7-Validation-of-rating-model-calibration.pdf).

Barth, J., Caprio, G. and Levine, R. (2006), *Rethinking bank regulation till angels govern*, Cambridge University Press, Cambridge, UK.

BCBS (2015), *Guidance on credit risk and accounting for expected credit losses*, Bank for International Settlements, Basel, Switzerland.

BCBS (2018), *Stress testing principles*, Bank for International Settlements, Basel, Switzerland.

Bliss, R. (1995), 'Risk-based bank capital: issues and solutions', *Economic Review* 80, Federal Reserve Bank of Atlanta, pp. 32-40.

Colombini, F. (ed.) (2018), *Raising capital or improving risk management and efficiency?* Palgrave Macmillan, London, UK.

Considine, J. (1998), 'Pilot exercise: pre-commitment approach to market risk', *Economic Policy Review* 4(3), Federal Reserve Bank of New York.

de Finetti, B. (1975), 'Theory of probability: a critical introductory treatment', John Wiley & Sons Ltd, Chichester, UK; Hoboken, NJ.

EBA (2019), *2020 EU-wide stress test draft methodological note*, European Banking Authority, Paris.

Enria, A. (2018), 'What we have learnt from EU-wide stress tests', speech delivered at the National Bank of Romania, November.

Grayson, C. J. (1960), *Decisions under uncertainty: drilling decisions by oil and gas operators*, Harvard Business School, Division of Research, Boston, MA.

Greenspan, A. (1998), Remarks by the Chairman of the US Federal Reserve System, Mr Alan Greenspan, before the Conference on Capital Regulation in the 21st Century, held at the Federal Reserve Bank of New York on 26 February 1998.

Keynes, J. M. (1921), *A treatise on probability*, Macmillian and Co.

Kupiec, P. H. and O'Brien, J. M. (1995), 'A pre-commitment approach to capital requirements for market risk', Federal Reserve Board, Finance and Economics Discussion Series No 95-34.

Milne, A. (2002), 'Bank capital regulation as an incentive mechanism: implications for portfolio choice', *Journal of Banking and Finance* 26, pp. 1-23.

Nau, F. R. (2001), 'De Finetti was right: probability does not exist', *Theory and Decision* 51.

Piccinato, L. (2009), *Metodi per le decisioni statistiche*, Springer

Prescott, E. (1997), 'The precommitment approach in a model of regulatory banking capital', *Economic Quarterly*, Federal Reserve Bank of Richmond.

Quagliariello, M. (2019), 'Are stress tests beauty contests?', EBA Staff Paper Series No 4, European Banking Authority, Paris.

Santos, J. (2000), 'Bank capital regulation in contemporary banking theory: a review of the literature', BIS Working Papers No 90, Bank for International Settlements, Basel, Switzerland.

Tarullo, D. (2008), *Banking on Basel*, Peterson Institute for International Economics, Washington DC.

Tasche, D. (2003), *A traffic lights approach to PD validation*, Deutsche Bundesbank, Frankfurt.

Varotto S. and Daripa A. (1997), 'Agency incentives and reputational distortions: a comparison of the effectiveness of value-at-risk and pre-commitment in regulating market risk', Bank of England Working Papers No 69, Bank of England, London, UK.

ACKNOWLEDGEMENTS

We are grateful for the discussion with Silvio Cuneo, Piero Franchini, Jan Hemmers, Paul Kupiec, Roberto Mosca, Andrea Resti, Anton Resuli, Giovanni Rinna, Fabio Salis, Carmelo Salleo and Franco Varetto.

The opinions expressed are those of the authors and do not reflect those of the institutions to which the authors belong.

Simone Casellina:
Bank sector analyst, European Banking Authority (EBA)

Giuseppe Pandolfo:
Advisor, Istituto Vigilanza sulle Assicurazioni (IVASS)

Mario Quagliariello:
Director of Economic Analysis and Statistics Department (EBA)

EUROPEAN BANKING AUTHORITY

20 avenue André Prothin CS 30154
92927 Paris La Défense CEDEX, France

Tel. +33 1 86 52 70 00
E-mail: info@eba.europa.eu

https://eba.europa.eu/

ISBN 978-92-9245-651-1 doi: 10.2853/709447
ISSN 2599-7831 DZ-AH-20-001-EN-N

This paper exists in English only and it can be downloaded without charge from https://eba.europa.eu/about-us/staff-papers, where information on the EBA Staff Paper Series can also be found.