Applying the Pre-Commitment Approach to bottom up stress tests: a new old story

Simone Casellina[1] ; Giuseppe Pandolfo[2] ; Mario Quagliariello[1]

Abstract:

Stress tests have become a key tool for banks, supervisors and macro prudential authorities. An aspect of these exercises is the need of 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 the results are achieved but authorities do not necessarily have all the information that the individual banks have. Banks' models (bottom-up approach) could be more accurate. However, banks may have incentives to underestimate the impact of the shock so reducing the supervisory reaction. In this paper, we focus on bottom-up stress tests and suggest to create a system of monetary penalties (charges) proportional to the difference between the expected and the realised losses of a portfolio. The charge aims at inducing 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 Federal Reserve in the 90s 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 PCA can indeed mitigate banks' incentives to provide underestimated measures of risk under the adverse scenario and, thus, better align incentives between banks and supervisors.

Keywords: Stress test, Supervision, Moral Hazard, Pre-Commitment approach, de Finetti penalty criterion

JEL: C11, G21, G28

1. Introduction and motivation

Stress tests have become a key tool for banks, supervisors and macro prudential authorities. During the crisis, they have been pivotal for identifying weak banks and recapitalise 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 a measure of risk related to banks' portfolios that is conditional to an extreme but realistic macroeconomic scenario. An aspect of these exercises is the need of statistical models to obtain these measurements.

[1] European Banking Authority (EBA)
[2] Istituto Vigilanza sulle Assicurazioni (IVASS)
The opinion expressed are those of the authors and not reflect those of the authors belonging Institutions.

Acknowledgments. 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, Franco Varetto.

If models are developed by the authorities (top-down approach), homogeneity of treatment among banks and more control over the results are achieved since one single model is used. However, authorities do not necessarily have as much information as individual banks, which could make models less accurate. Conversely, models developed by the 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 the shock in order to reduce the supervisory reaction. No approach is exempt from criticism and indeed there is debate on how stress tests should ideally look like[3]. 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 could be less conservative given the banks' different 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 asymmetric information. With banks becoming more and more complex, information asymmetries increase. Paradoxically, banks' complexity has been also the main motivation for moving regulation towards a wider use of internal models. However, this came 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 pure top-down or bottom-up approaches is the hybrid approach applied in the EU-wide stress test. Under this approach, the authorities try to mitigate the agency problem providing banks with a detailed methodology, including constraints in the form of caps and floors to the output of internal models. Along with a thorough quality assurance process, the "constrained" bottom-up approach adopted by the European Banking Authority (EBA) aimed at delivering conservative and comparable outcomes, while maintaining the reliance on banks' models. However, this approach is also challenged since constraints are sometimes perceived as too strict or not realistic, with the risk of transforming a bottom-up into a top-down exercise (Enria, 2018; Quagliariello, 2019). The question is whether there is a better way for allowing the use of internal models, relaxing 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 Federal Reserve (FED) in the context of the use of internal models for market risks at the end of the 90s. That was the time when, for the first time, the Basel Committee on Banking Supervision (BCBS) was opening the way to the use of banks' own models for quantifying regulatory capital requirements, but there was debate on how to control their use. Under the Pre-Commitment Approach (Kupiec and O'Brien, 1995), banks could quantify their own requirement 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 already in the 30s as the foundation of the theory of probability under 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 using the Pre-Commitment Approach to bottom-up stress tests. In the Annex, we illustrate a practical example.

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

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

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 provisions under different future scenarios. In that respect, we believe that banks will increasingly use IFRS9 models also for stress testing purposes. In this sense, it is possible to define the bank's better forecast for next year as the one 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 to different scenarios.

A further assumption 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, which is typical in stress testing exercises.

2. The Pre-Commitment Approach

The Pre-Commitment Approach (PCA) was developed by FED's economists Kupiec and O'Brien in 1995[4]. In that period, the Basel Committee on Banking Supervision (BCBS) had agreed to provide an alternative regime for setting capital requirements for the trading book based on the internal Value at Risk (VaR) models. The debate at that time was on whether supervisors should set model parameters or rather allow banks to use their estimates. The latter approach would have produced more accurate results but could have also encouraged to underestimate risks. It was in this context that PCA was proposed.

Under 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 exceed this figure at any time in that period. In other words, rather than meeting the capital requirements set by the regulator, the bank sets itself 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, the then chairman of the Federal Reserve Board (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 did not contemplate to regulate how banks should design and develop their models nor it required any supervisory validation. PCA is all about incentives, while leaving modelling issues to banks.

Bliss also notes that the PCA is conceptually more appealing than the BCBS internal model approach where banks can use their estimates but supervisors regulate meticulously several modelling aspects. However, he also highlights the need to clarify a number of details in the practical implementation of PCA and, in particular, in the design of the penalisation system. He suggests that a penalty should be certain and hard to game by the bank. In his opinion, capital-based penalties would

[4] Kupiec and O'Brien (1995)

not achieve the goal since – in case of 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 reducing next period reported pre-commitment capital. Fines instead meet Bliss' criteria: once a pre-commitment level is breached, the bank would incur an immediate cost. The use of fines on a regular and automatic basis could appear strange, but it must be understood that these fines are not sanctions but rather a device for creating an incentive-compatible use of internal models. Along with reducing the incentive to game the rules, PCA also allows banks to make their own decisions on risk management investments: banks would invest resources in modelling as to avoid fines.

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

Prescott (1997) notes that 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 PCA may minimise the distortions to capital holdings caused by private information and conclude that it can be beneficial provided that fines are appropriately calibrated.

Santos (2000) examines the pros and cons of the PCA also compared to other approaches. He observes that the PCA is more prone to a time consistency problem because it applies penalties *ex-post*. Regulators, for example, would be pressured to waive the penalty in case it were to lead to bankruptcy. In addition, because of 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, Caprio, Levine (2006) observe that an attractive feature of the PCA is that it addresses the concern that the Basel approach is too rigid and arbitrary. Tarullo (2008) devotes ample space to the PCA and an application of PCA beyond market risk is also suggested. Colombini (2018) concludes that a widespread Pre-Commitment Approach in risk management is among the issues that need to be taken into consideration in order to complete the broad picture of capital framework in Europe.

The interest for the PCA was so high and support from banks so vocal that the FED decided to run an experiment involving 10 banks (Considine, 1998). According to those banks, the pilot demonstrated that the PCA is a viable alternative to the BCBS internal models approach since it provides strong incentives for prudent risk management and a more efficient allocation of capital compared to 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 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 a comment on the PCA in parallel with the consultation on 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 Basel Committee countries had just agreed to the value-at-risk (VaR) approach to market risk. The proposal was, therefore, abandoned at least as an official FED's stance.

3. de Finetti's penalty criterion

A key question of the PCA is how the fine should be actually designed. In that respect, the seminal contribution of de Finetti on penalty criterion 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 further step 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[5].

For de Finetti, the problem is not to know the true probability of the realisation of an uncertain event, because this measure simply does not exist, but rather to get the best evaluation of such 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 could have incentives to not reveal its true belief about the probability of the event. According to de Finetti, the revelation of true beliefs requires to correct the incentives through penalties.

The easiest example is related 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 the case 1 occurs. Many practical situations can be traced back to this example. For example, a company requesting credit could be asked 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 he will suffer a monetary fine proportional to $(1-P)^2$ if the event occurs and P^2 if the event does not occur[6] where P is the probability value provided by the subject.

de Finetti suggests to fine the subject in any case, regardless of whether the event occurs or not. This appears 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's assume that the true evaluation of the probability is equal to P but the subject may decide to provide $P^* \neq P$. In this case, he can anticipate that the fine f he will pay is proportional to:

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

[5] "*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*.", Nau (2001).

[6] Piccinato (2009) Annex A

In this expression, the expected value of the fine is obtained conditional on P that 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:

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

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

It is important to highlight that, in de Finetti's approach, the penalty is only a tool for aligning incentives and does not imply any negative judgment. Since the use of terms like "fine" or "penalty" is likely to induce some misinterpretation in 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 they incur in the risk of not finding oil, the responsible of the projects (the decision maker) asks geologist for an advice. Grayson points out that the decision maker cannot ask the geologist for an exact prediction or prophecy but rather for an evaluation in terms of probability. The problem for the decision maker is how to induce the expert to provide an advice in line with his/her own genuine belief. If his/her remuneration is linked to potential profits, he/she could have incentives in the direction of over-optimism since he/she would take the upside but not the downside. On the contrary, the perspective to participate in potential economic losses in case the oil is not found, could lead him/her in the direction of over-pessimism. In this situation, the penalty criterion can mitigate the problem. Indeed, the geologist would know that a penalty would be applied in any case but he can minimise it by revealing its best evaluation.

4. The application of the PCA to bottom-up tress tests

In bottom-up stress tests, banks are requested to provide supervisors with their true beliefs of the losses that can 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 light of the de PCA/penalty criterion: through a system of charges, supervisors could induce banks to reveal the true outcome of their models because this would reduce the ex-post monetary cost.

The application of PCA to stress testing models seems also less controversial than to models for setting minimum capital requirements. In particular, there were two concerns on PCA: i) the capital level autonomously set by the banks could eventually be too low, bringing to an insolvency; ii) an application beyond market risk, where back-testing is relatively straightforward, could go too far.

As regard the first issue, stress tests are not meant for setting minimum capital requirements but rather for assessing capital planning and, possibly, require banks to hold additional buffers or reach some capital targets to meet supervisory expectations. In the EU, for instance, competent authorities use the findings of the stress test as an input to the supervisory review and evaluation process (EBA, 2019). Therefore, leaving more freedom to banks *ex-ante* is less likely to bring to the breach of minimum capital requirements.

On the application to other risks, the risk of losses can always be split into frequency of the event and severity of the losses. In credit risk, for instance, this partition corresponds to the probability of default (frequency) and loss given default (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 the severity since recovery procedures can take many years, while the frequency side is less problematic.

In practice, we propose to make 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 back-tested. This could also contribute to the discussion on which risks could be modelled by banks – under the PCA incentives – and which would be rather 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 like oil or the level of the 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 her Ω_t^i, the model she decided to be most appropriate M^i and her choice about the 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 will be informed that the value $E_t^i(L_{t+h})$ will be compared with the realised L_{t+h} and that the bank will 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 is provided directly by the authority, which means $E_t^i(y_{t+h})$ is replaced with $E_t^{AUT}(y_{t+h})$. Banks are thus required to produce the expected figures conditional to different scenarios – their own and the supervisory one – and using always the same model and to provide the outcomes through standardised templates.

Going back to the comparison between the previsions or forecasts and realizations, let's 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 $E_t^i(y_{t+h})$ that 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. Indeed, it would be harder to define an appropriate comparison term for the expectation of the loss amount conditioned to the adverse scenario.

The charge would be proportional to the quantity ϵ_{t+h}. This aspect is central and deserves some comments. It is clear that the forecast error can be due to several unpredictable events like mergers and acquisitions, or sales of significant portions of 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 so easy to define in practice which events were truly unforeseeable and it is even more difficult to do it *ex-ante,* i.e. before these events are observed.

However, in order for the 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 the bank 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 the 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 the revealed one.

However, it could be possible to decompose the forecast error. Being $f\left(\Omega_t^i, M^i, E_t^i(y_{t+h})\right)$ 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 $f(\Omega_t^i, M^i, y_{t+h})$ that is the forecast that the model would have produced given perfect foresight relatively to the scenario. The charge could be a proportion of the difference between the realization L_{t+h} and the output of the model given the realised scenario i.e. $f(\Omega_t^i, M^i, y_{t+h})$. This way, banks would be only penalised for shortcomings in the model and not also for not having correctly forecasted the scenario, which is irrelevant in a stress test where the scenario is provided by the authorities themselves.

As regards the concrete definition of the quantity L_{t+h} , it would be preferable that the charges were 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 rare i.e. the expected transition rate from non-defaulted to defaulted status[7] and $E_t^i(LGD_{t+h})$ is the expected loss rate. Though different models can be developed for different sub-portfolios, the results of these models, once expressed in terms of expected loss, can be always summed.

As we have already discussed, the loss given default 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 long. Keeping this aspect in mind, our proposal is to compare the forecasted 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 EAD and DR are substituted with the observed values while the loss rate is kept constant. Alternatively, it would be possible to substitute $E_t^i(LGD_{t+h})$ with an updated forecast

[7] It could represents the transition rates from the Stage 1 to Stage 2 and 3 using IFRS9 concepts.

produced in t+h, i.e. $E^{i}_{t+h}(LGD_{t+h})$ but this would not be fair given that here we would compare two forecasts based on two different information sets.

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

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

The above quantity is expressed in monetary terms and could be summed across different portfolio, status (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 of a symmetric penalty. This implies that the bank[8] is charged when estimates are not accurate, regardless of whether they are more or less conservative than the realised losses. This can 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 symmetric approach also appears consistent with the general principle regarding the application of IFRS9 accounting set out by the Basel Committee. BCBS (2015) for example highlights that "[..] *the Committee recognises that 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 the paper we suggest how the penalty system could be supplemented with additional interventions for penalising more those banks that systematically underestimate losses.

Consider now the possibility that the bank decides to provide to the authorities a distorted version of its previsions: the bank has a genuine prevision of the losses obtained with its best model $E_t(L_{t+h})$ but may decide to not reveal it to the authority.

Instead of providing the expected losses stemming from its true model $E_t(L_{t+h})$, the bank may decide to provide systematically lower figures. For simplicity, we can think that the bank employs a model like $\widehat{E_t(L_{t+h})} = E_t(L_{t+h}) - d$ where its true expectations about losses conditioned to any scenario are reduced by a constant quantity. We then have two models, the bank true-best model $E_t(L_{t+h})$ and the model used to provide risk measures to authority $\widehat{E_t(L_{t+h})}$. Clearly, the amount of the charge will be computed on the ground of the risk measures officially provided, that is: $\epsilon_{t+h} = \widetilde{L_{t+h}} - \widehat{E_t(L_{t+h})}$.

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

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

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

[8] For now we refer generically to the bank, we refer to section 6 for a broader discussion on the calibration of fines and also on who should be paying them.

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 prevision $E_t(L_{t+h})$ and the measure provided to the authority $\widehat{E_t(L_{t+h})}$, this should lead to provide an incentive to reduce the difference between the two quantities.

It is interesting to compare our proposal with the back testing approach (BTA) proposed by BCBS (2005) for the validation of models for the probability of default in the context of the IRB system. Under the BTA, it is recognised that a bank's Probability of Default (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 PDs merits a critical assessment. Four approaches are typically used: the binomial test, the chi-squared test, the normal test and the traffic lights approach. Following Bliss (1995), also the BTA is a non-model based 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 about *if* the models work and not on *how* they work. The PCA and BTA may appear quite similar but they are not. Three main differences between the BTA and our version of the PCA exist:

(1) the BTA requires to define one or more thresholds. 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 pre-defined thresholds. Setting these boundaries implies both the necessity to rely on a model and to set arbitrarily the confidence level. The simplest approach to set such thresholds is to refer to the binomial model under the hypothesis of independence of the defaults and infinite granularity but this model can be considered too unrealistic[9].

(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 performances in the 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 the new model would deliver lower measures of the risk. In practice, it would not be possible to include *ex-ante* such evaluations in the decision process. On the contrary, 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 prevision $E_t(L_{t+h})$ and the measure provided to the authority $\widehat{E_t(L_{t+h})}$.

(3) The BTA implies a judgment about the bank model. If ϵ_{t+h} have been higher than the upper threshold for a period, the authorities will say that the model is "wrong". This kind of judgments somehow support the simplistic idea that a better model, maybe the perfect one, was possible while the forecast was the best possible model given the information available. With our proposal, no judgment is formulated, as it is left to the banks to try and minimise the charges the best they can.

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

5. 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 the banks' incentives but this, in turn, would require to know precisely the banks' cost function. In the absence of such information, a simple function like the one above (in practice the one originally suggested by de Finetti) could represent a valid solution.

We already discussed our preference of a symmetric mechanism behind the quantification of the charges. Without the symmetry, the bank would have incentives to reveal a prudent evaluation and not its best 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 on models' outcome but it is a second best compared to the best forecast. The point here is to provide banks with incentives to develop accurate models and to reveal them. Producing a prudent evaluation is a simpler task compared to producing an unbiased measure. A prudent evaluation can be obtained starting with a quite 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 to take actions if a 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 substituting the bank's models with external ones.

Another issue is the amount of the charges: 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 provide some criteria.

We believe that the penalisation should hit directly the person directly responsible for the models' development. Here we follow Varotto and Daripa (1997), who argue that Kupiec and O'Brien disregard the incentives of the bank vis-à-vis those of the managers. A key feature of large modern banks is the separation of owners from day-to-day decision-making and they conclude that in case of agency problems, the incentives could not reach the banks' decision makers unless they are hit directly by the charge. Therefore, the charge could be proportional to a share of the Chief Risk Officer annual bonus 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 to know L_{t+1} 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 process. An alternative is to set a rough upper-level for $\epsilon_{t+h} = L_{t+h} - E_t(L_{t+h})$ like 1.5*ϵ_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, devolve the amount to charity, and not translated into an increase of capital requirements or provisioning level. The PCA is based on the idea that a charge should be paid at all time, even for minor forecast error. However, in practice, a materiality threshold under which the charge should 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 judgment regarding the staff involved in models' development, but it is rather a device for making the bottom-up stress test incentive-compatible. Indeed, if the charges were correctly calibrated and then the incentives aligned, the charge would only reflect an unavoidable forecasting error.

6. Conclusions

In this paper we suggest to apply the PCA to bottom-up stress test in order to mitigate banks' incentives to game the exercise. Focusing on credit risk, we propose to introduce 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 made is that risk measures for the stress test are built using the same models adopted by the bank for internal/accounting purposes. Under the new accounting framework set out by the IFRS9 principles, banks are required to base the level of provisioning on the Expected Credit Loss (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 prevision/forecast and it make sense to compare it with realised figures (back-testing).

A further assumption 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, which 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 to scenarios that are less extreme. This would also make our linearity assumption more realistic.

The charge should be symmetric – i.e. it should be linked to any estimation error, regardless of whether the realised losses are higher or lower than the estimated ones – and as long as they do not jeopardise the stability of the bank if systematic problems with the models are identified. Given agency problems in the bank, the charge should affect directly the risk manager in charge of developing the model. This monetary charge should not be interpreted as a judgment on the model itself but simply as a device to better align the incentives and ensure that banks provide their best estimate 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 a 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 the deregulation. On the contrary, it can be used as a complement to existing tools, including benchmarking and quality assurance.

The design and calibration of the system of charge would clearly require additional work and discussion. An important aspect, in particular, is related to the degree of disclosure of the approach, the possible publication of the charges paid by each risk manager, banks' responses to models' shortcomings. The publication of the difference between forecasts 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.

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 ensure the reliability of stress test outcomes. In that respect, the PCA applied to bottom-up stress tests can be the way forward to aligning supervisors and banks' incentives, and to 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 think to a FIRB bank) so that there is no need to have a model for it and that the exposure at default is set equal to the value of the exposure at the beginning of the period. In this context, the only uncertainty is about the probability of default.

Let's assume that the default rate generated each year by a bank credit portfolio obeys to a well-defined model like the following.

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

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.

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

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

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

We assume now the model (1) is the private information of the 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 stressed 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 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 known also to the authorities, the forecasted 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 we adopted to keep the example simple.

Now the bank evaluates the possibility to game. For example, it can provide a default rate that is systematically lower than its forecast. The bank decides then to provide the result of 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 would not appear as a credible if it produces zero default rate in case of zero growth rate of the GDP.

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

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

$E_t^{bank,\widehat{AUT}}(DR_{t+1})$ is used to stress that this forecast is obtained with the model provided by the bank ($bank$) but it is the model ($\widehat{\ldots}$) corrected with the parameter b and it is alimented 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:

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

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

Given the portfolio value, the LGD and the bank provided forecast of the default rate under the baseline scenario, the Credit Risk Adjustments 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:

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

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

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 bonus of the Risk Manager is connected somehow with the return on capital, then 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 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, also the return increase. Without the credibility lien the bank could increase the return indefinitely.

The Lagrangean associated to the problem is:

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

and the 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. not 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's introduce a charge proportional to the forecast error and to 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.

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

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 in advance its value and to 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 increase the return is affected by a positive component LGD/B proportional to b but also to a negative element C/B proportional to b squared.

The Lagrangean associated to the problem is:

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

and the Karush-Kuhn-Tucker (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 $\lambda = 0$ (not 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 depends also on the parameter δ that 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.

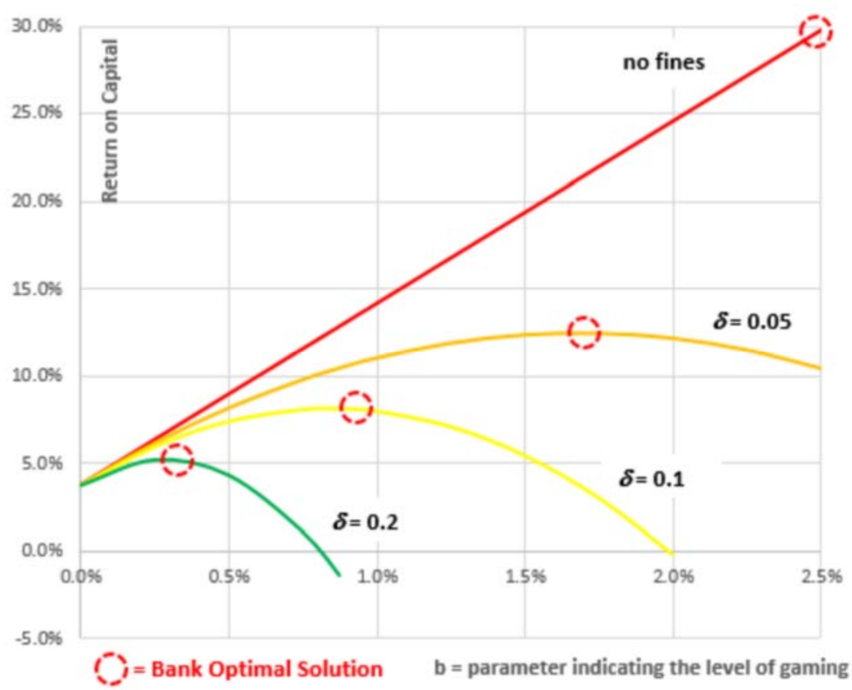

The penalty reduces the economic convenience to game i.e. to underestimate the losses. If we add to this that the charge could be accompanied by the publication of the results (such as the ranking of how much the banks have paid), then the incentive to game could become null.

References:

AIFIRM (2016), "Position Paper 7 -Validation of rating model calibration", http://www.aifirm.it/wp-content/uploads/2016/12/2016-Position-Paper-7-Validation-of-rating-model-calibration.pdf

Barth J., Caprio G., Levine R. (2006) "Rethinking Bank Regulation Till Angels Govern", Cambridge University Press

BCBS (2015), “Guidance on credit risk and accounting for expected credit losses”

BCBS (2018), “Stress testing principles”

Bliss R. (1995), `Risk-based bank capital: issues and solutions', Federal Reserve Bank of Atlanta, Economic Review, September/October, Vol. 80, pages 32-40.

Considine J. (1998), "Pilot Exercise—Pre-Commitment Approach to Market Risk", Federal Reserve Bank of New York, Economic
Policy Review October 1998 Volume 4 Number 3

de Finetti B. (1975), "Theory of Probability: A critical introductory treatment", Wiley

EBA (2019), "2020 EU-wide stress test Draft Methodological Note"

Enria A. (2018), “What we have learnt from EU-wide Stress Tests”, speech delivered at the National Bank of Romania, November

F. Colombini (2018), "Raising Capital or Improving Risk Management and Efficiency", Palgrave Macmillan

Grayson, C.J., 1960. Decisions under Uncertainty, Drilling Decisions by Oil and Gas Operators, Harvard Business School, Division of Research, Boston.

Greenspan A., 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 26/2/98

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, July, No. 95-34

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

Nau F. R. (2001), "De Finetti was Right: Probability Does Not Exist", Theory and Decision vol. 51

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

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

Quagliariello M. (2019), "Are stress test beauty contests", EBA STAFF PAPER SERIES, N. 4 – JULY/2019

Santos J. (2000), "Bank capital regulation in contemporary banking theory:a review of the literature" BIS WP No. 90 Sept 2000

Tarullo D. (2008), "Banking on Basel" Peterson Institute for International Economics

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

Varotto S., 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