Rethinking Operational Risk Capital Requirements

Peter Sands

Gordon Liao

Yueran Ma

Havard Kennedy School

Mossavar-Rahmani Center for Business and Government

Harvard Business School

Project on Behavioral Finance and Financial Stability

Rethinking Operational Risk Capital Requirements

Peter Sands

Gordon Liao

Yueran Ma

Rethinking Operational Risk Capital Requirements[1]

Peter Sands
peter_sands@hks.harvard.edu
Harvard University

Gordon Liao
liao@fas.harvard.edu
Harvard University

Yueran Ma
yueranma@g.harvard.edu
Harvard University

Abstract

Operational risk capital requirements represent a relative backwater of the Basel capital framework for banks. We examine both the existing Basel II framework and the latest Basel Committee proposals for reform and conclude that neither are effective in creating appropriate incentives and loss absorbency to minimize negative externalities from operational risk events. We evaluate an alternative approach that would appear to be much more effective in achieving the regulatory objectives. We do not offer a view on the amount of capital required, focusing instead on the methodology and structure of the capital requirement.

1 INTRODUCTION

Operational risk weighted assets ("RWA") are one of the three components of the denominator of any bank's risk-based capital ratio. Operational RWA represent 15.6% of the RWA of the 30 globally systemically important banks ("GSIBs").[2] With an average Common Equity Tier 1 ("CET1") ratio of 11.7%,[3] this means some USD$411bn of equity capital[4] is dedicated to protecting banks, their investors and ultimately, society from the consequences of operational risk events, as shown in Exhibit 1. Yet despite the amount of capital involved, the derivation and deployment of operational risk capital typically receives much less management, regulatory, investor or academic attention than the two other components of banks' RWA, credit and market risk.

Exhibit 1: Operational RWA: a component of the capital ratio denominator

$$\frac{\text{Adjusted Common Equity*}}{(\text{Credit RWA + Market RWA + Operational RWA})} = \frac{\text{Common Equity}}{\text{Tier 1 (CET1)}}$$

Represents 15.6% of total RWA for GSIBs

Average CET1 of 11.7% across GSIBs. So USD $411bn equity aligned against operational RWA

* Adjusted for deductions for expected loss, tax credits, etc

The Basel Committee on Banking Supervision ("BCBS") has proposed to reform the calculation of operational RWA, replacing an existing system which offers banks a choice between three approaches differing in complexity and reliance on internal models, with a single approach, named the Standardised Measurement Approach ("SMA"). This has prompted considerable debate, particularly around the overall impact on banks' capital requirements, the backward-looking nature of the proposed methodology, and the complexity and potential unintended consequences of the calculation, including so-called "cliff effects" arising from the thresholds.[5]

---

[2] Average across 30 GSIBs based on company filings and investor disclosures.

[3] BCBS Monitoring Report September 2016.

[4] CET1 of €2,424 billion as of 2015, BCBS Monitoring Report September 2016.

[5] Responses to the March 2016 BCBS consultative document have voiced a wide range of concerns including the overall impact, lack of transparency, disconnect with managerial actions, poor predictive power, and counter-intuitive consequences. The so-called "cliff effects" arise where losses slightly above or below the defined thresholds can have marked difference on capital requirements.

The objective of this paper is to answer the question of whether the BCBS proposals fix the widely recognized flaws in the existing system, or whether a more radical alternative would be better. We start by briefly summarizing the history of operational risk capital requirements within the Basel framework. We then set out what we think should be the objectives of an operational risk capital regime and evaluate the performance of both the existing system and new proposals in delivering against these objectives. Finally, we suggest and explore an alternative approach that might achieve the underlying regulatory objectives more effectively.

While the BCBS's desire to improve the current system is well-intentioned, our analyses corroborate many of the criticisms directed at the BCBS's proposals (hereafter described as the SMA). Yet our objection is more fundamental: our evaluation suggests that much of the current debate around the SMA largely misses the point, and that there should be a much more fundamental rethink of operational risk capital requirements. When one considers the underlying regulatory objectives, both the existing system and the SMA look deeply flawed, for several important reasons.

First, despite requiring significant levels of capital, in practice neither approach gives banks more ability to absorb the losses from operational loss events without negative externalities such as disruption to credit provision. In effect, the capital deployed against operational RWA to meet the required capital ratio is "dead" capital, largely incapable of being used. Without effective loss absorbency, banks are likely to react to operational losses by cutting lending and shrinking assets, precisely the actions regulators aim to avoid (Rajan and Stein, 2008; Hanson, Kashyap and Stein, 2011). In fact, the argument could be expressed more strongly. Since significantly increased operational losses would lead to increased operational RWA under the SMA and there is virtually no scope for management to reduce operational RWA in the way a credit book or trading positions can be managed to reduce credit or market RWA, the SMA approach arguably provides *negative* loss absorbency. This lack of loss absorbency is discussed in more depth in Section 5 (p24-25) using a highly simplified example to illustrate the point.

Second, both the existing and proposed approaches appear of little to no use in incentivizing bank management to improve the management of operational risk. Both approaches are almost entirely backward-looking, while operational risks are constantly evolving and the drivers of the biggest losses defy mechanistic prediction from historical data (see Ames, Schuermann and Scott, 2015).

Finally, the weaknesses in the determination of operational RWA across both existing and proposed approaches, and the sometimes counter-intuitive variability in outcomes, do little to contribute to the overall credibility and comparability of risk-based capital ratios. As Exhibit 2 shows, under the current regime, there is

significant variation in the percentage of a bank's total RWA contributed by operational RWA. Some of this variation can be explained by differences in strategy and business model (eg the fact that among the GSIBs, State Street and Bank of New York Mellon have the highest proportions of operational RWA reflects their focus on custody and settlement services rather than traditional lending), but many of the differences appear to reflect differences in the approach towards determining operational RWA across banks and regulatory jurisdictions, rather than differences in the underlying operational risk profile.

Exhibit 2. Operational RWA as a percentage of total RWA for GSIBs
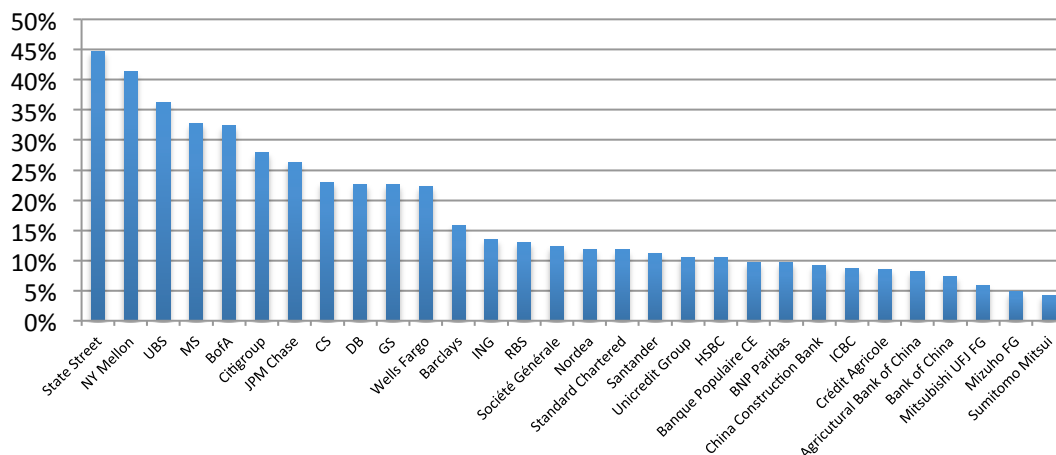
Note: This figure presents the fraction of operational risk weighted asset (RWA) as a percentage of total RWA for 30 GSIBs as disclosed in regulatory filings and investor reports.

The SMA will eliminate the comparability problems posed by having different banks use completely different approaches, but arguably, will provide the *appearance* of comparability rather than *real* comparability: under the SMA, two banks could have equivalent operational RWA whilst facing very different levels of risk, and vice versa.

These considerations lead us to propose an alternative approach that replaces the notion of operational RWA with a capital buffer approach, which could provide both loss absorbency and appropriate management incentives. Under this approach, the size of the buffer would be based on a combination of: 1) scale-based minima expressed in absolute dollar terms; 2) calculated capital to cover unexpected losses for those operational risk types that are amenable to statistical modelling, again expressed in absolute dollar terms; and 3) a judgment–based overlay determined by the regulator on the basis of a structured assessment of the specific risks faced by a bank and its capability to manage and mitigate these risks[6]. When significant loss

---

[6] A variant of this proposal would have only two components, the scale-based minimum and the discretionary element. The logic here is that the calculated component for operational risk types for which the definitions and data are sufficiently robust to enable robust and comparable predictions will be a relatively small contributor to the total and can probably be proxied by the scale-based calculation. This would have the advantage of increasing

events occurred, the bank could deploy part of this buffer to absorb the loss. How much of the buffer could be deployed would be agreed with the regulator, together with a plan to return the buffer to the required level.

This approach echoes what a number of regulators have already done. Recognizing the deficiencies of the existing Pillar 1[7] approach to operational risk in providing either loss absorbency or appropriate incentives, several regulators[8] have already included forward-looking assessments of specific operational losses (such as losses that arise from regulatory fines) in their Pillar 2 requirements or stress tests.

Our suggestion of shifting to a buffer approach does *not* imply a reduction in operational risk capital requirements. This approach could be implemented in a manner that is capital neutral, or even that requires more capital. However, for any given amount of capital deployed against operational risks, we believe this approach would achieve the underlying regulatory objectives more effectively than either the existing system or the BCBS proposals. Under our approach, banks would be able to use their operational risk buffers to help absorb significant loss events, thus protecting their core financial intermediation activities from unnecessary disruption.[9] Regulators could structure the discretionary component of the buffers to provide incentives, penalizing banks with control deficiencies and rewarding those that rectified problems or strengthened their defenses. Investors will see more clearly delineated the underlying CET1 ratio based on credit and market RWA, and the operational risk capital buffer as determined by the regulator, thus contributing to the Pillar III objective of reinforcing market discipline.

If the test of success is protecting society from the negative externalities of operational risks, then the SMA represents an attempt to reform a regulatory construct that has clearly failed. In the wake of the global financial crisis, banks have experienced operational risk losses of unprecedented scale as shown in Exhibit 3. The negative externalities from bank losses have been significant, including reduction in credit provision, disruption to specific markets, withdrawal of certain socially beneficial products, and financial exclusion as a consequence of "de-risking", plus a significant diminution of trust in the financial system (Ivashina and Scharfstein, 2011; Chodorow-Reich 2014; Zingales 2011). However, the SMA does not really address the the flaws in the current regime. It is more an exercise in methodological tinkering

---

simplicity.

[7] As introduced in Basel 2, Pillar 1 emphasizes bank capital adequacy by setting a minimal capital ratio of CET1 to risk-weighted asset (RWA) composed of market, credit and operational RWA; Pillar 2 focsus on supervisory review and enables local regulators to monitor and fine-tune risk management and supervision for individual banks; Pillar 3 enhances disclosure to the public to support market discipline.

[8] For example, the UK Prudential Regulation Authority, European Banking Authority, U.S. Federal Reserve, Monetary Authority of Singapore.

[9] The theoretical merits of a buffer approach to capital charges had been examined closely in Kashyap and Stein (2004).

that makes some aspects of the current system marginally better, others worse. The scale of recent losses from operational risk events, the amount of capital deployed (albeit ineffectively) against operational RWA, and the significance of the risks banks face in the future, all point to the need for a more radical rethink.

Exhibit 3: Operational losses by year of public disclosure
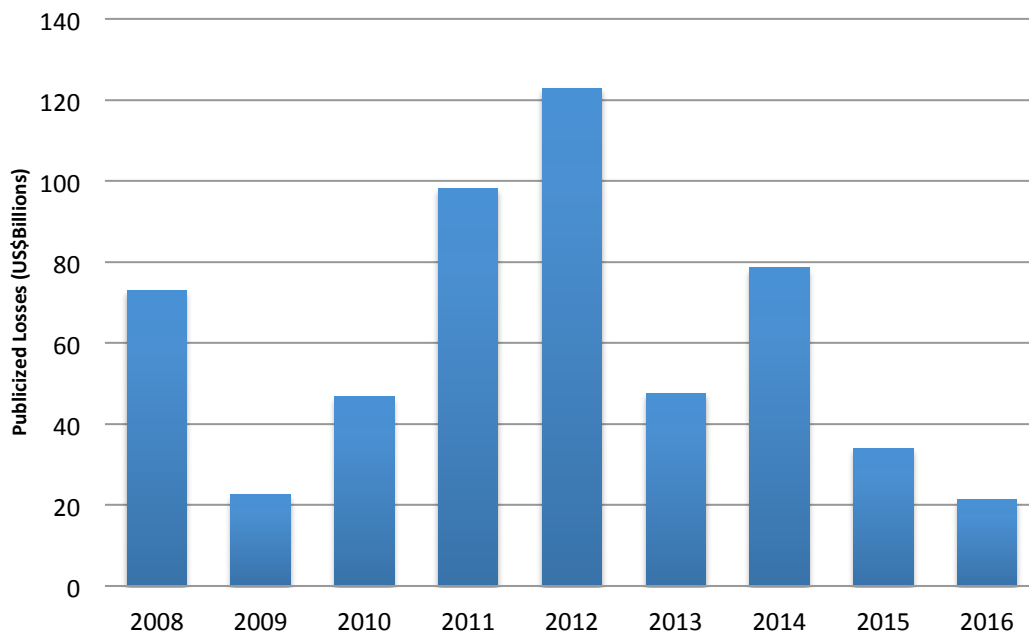
Note: This figure presents operational losses by the year of public disclosure in filings and media using the Public Source ORX News Data[10]. The date of public disclosure is generally around the announced fine and legal settlement and/or loss provision and accounting dates. The median time between event start (end) date and publication/settlement date is six (three) years.

In the rest of this paper we expand on this argument, setting out in turn:

- The background to today's operational risk capital requirements and the BCBS proposals (Section 2)
- The objectives of operational risk capital requirements (Section3)
- An assessment of the current approach to operational risk capital requirements against these objectives (Section 4)
- An assessment of the BCBS proposals against these objectives (Section 5)
- An alternative approach (Section 6)
- Potential arguments against the alternative approach (Section 7)
- Concluding comments (Section 8)

---

[10] Data on publicly announced loss events is provided by ORX (www.orx.org) and presented in Exhibits 3, 6, 7, 8 and 9. The ideas and conclusions do not represent any views of ORX or its members.

2 THE BACKGROUND TO TODAY'S OPERATIONAL RISK CAPITAL REQUIREMENTS AND THE BCBS PROPOSALS

As highly complex operational businesses dealing with money in all its forms, banks face a wide range of operational risks, from fraud and theft through regulatory penalties to system failures, as well as natural disasters, technology glitches and reputational issues. Failure to manage or mitigate these risks can lead to negative externalities. A bank that fails as a result of an operational risk event could bring down other banks, disrupt credit provision or disturb the functioning of markets. Alternatively, a bank that suffers a significant operational loss may cut back on lending or other activities to create the capacity to absorb the loss and ensure survival. Because of the resultant externalities, bank regulators have sought oversight of operational risks and how they are managed. However, it was only in 1998, following the collapse of Barings in 1995, that the BCBS formally introduced the notion of operational risk as a regulatory concern. The idea that banks should hold capital against operational risks is even more recent. Formal operational risk capital requirements were first introduced as part of Basel II in 2004.

2.1 The Introduction of Operational Risk

The term "operational risk" became prominent in banking and regulatory circles after a rogue trader caused the collapse of Barings Bank in 1995. This event highlighted the importance of internal controls and corporate governance in managing financial losses associated with fraud, human errors, and technical failures as well as other breakdowns in normal business processes and operations. In the early days, most banks and regulators defined operational risk simply as any risk not categorized as market or credit risk (BCBS, 1999). The consensus among banks was that "the primary responsibility for management of operational risk is the business unit" (BCBS 1998). The internal measurement, monitoring and control of operational risks at the overall bank level were still in their infancy. Furthermore, the prevailing view was that the role of regulatory supervisors in operational risk management should be limited to encouraging "qualitative improvements" through raising awareness and facilitating the sharing of best practices (BCBS, 1998).[11] The BCBS continued to develop the concept of strengthening supervisory oversight in subsequent documents such as "*Sound Practices for the Management and Supervision of Operational Risk*" (BCBS, 2003).

2.2 The Introduction of Operational Risk Capital Requirements

The concept of operational risk capital requirements was introduced with the Basel II

---

[11] The 1998 BCBS survey indicated that "few banks seem to have made considerable progress in developing more advanced techniques for allocating capital with regard to operational risk," although no regulatory discussion of operational risk capital were provided at the time.

Accord in 2004.[12] Operational risk was introduced as an additional risk category alongside the existing categories of market and credit risk in formulating the revised Pillar 1 framework of minimum regulatory requirements. The BCBS formally defined operational risk "as the risk of loss resulting from inadequate or failed internal processes, people and systems or from external events. This definition includes legal risk, but excludes strategic and reputational risk" (BCBS 2006, 2011). Seven broad categories of operational losses were delineated in Basel II: internal fraud, external fraud, employment practices and workplace safety, clients, products and business practices, damage to physical assets, business disruption and system failures, and execution, delivery and process management.

There is some debate as to the rationale for moving from a focus on supervisory oversight of operational risk management practices (known as Pillar 2 in the language of Basel II) to making operational risk an integral component of the minimum capital requirements regime (known as Pillar 1). Some have expressed the view that the BCBS included capital requirements for operational risk in Basel II primarily to offset the anticipated reduction in capital requirements for credit risk as a result of the newly introduced internal rating-based ("IRB") approach (Herring 2005, 2007). There was also considerable skepticism at the time about the wisdom of adding the heterogenous and vaguely-defined operational risk RWA, alongside the relatively well-defined market and credit RWA (Power, 2005).

Basel II stipulated three different approaches to assessing operational risk capital requirements depending on the complexity of a bank's business (see Appendix A). First, the Basic Indicator Approach ("BIA") calculates operational risk RWA based on a coefficient (set by the BCBS at 15%) multiplied by the rolling three year average revenue of the bank as a whole. This is the simplest approach and is widely used by small banks. Second, the Standardized Approach ("TSA") derives operational RWA by applying defined coefficients (set between 12-18%) to the rolling three year average revenue of eight broadly defined business lines. Under this approach a bank's overall operational risk capital requirement is simply the sum of these calculations.[13] Finally, large banks with more complex business lines were encouraged to adopt the Advanced Measurement Approach ("AMA") utilizing the bank's own internally developed risk measurement framework to develop a Value-at-Risk ("VaR") to a 99.9% confidence level. In adopting the AMA banks had to meet various criteria around the governance and management of operational risks and be able to demonstrate that they could take account of four variables in their calculation of their own capital requirements: i) internal data; ii) external data; iii) scenario analysis; and iv) the business environment and internal control factors. Whilst the BCBS provided guidance, ultimately it was for home regulators to determine how

---

[12] Before the introduction of an explicit operational risk capital charge, it was generally held that the credit RWA were designed and calibrated all relevant risks, thus implicitly covering operational risk (OCC, 2007).

[13] There is also a variant of the TSA called the alternative standardized approach ("ASA"), which is relatively rarely used.

much to encourage adoption of the AMA and to approve individual banks' methodologies. This inevitably led to differences between regions. For instance, while US-based GSIBs have mostly adopted the AMA, all but one UK GSIB are under the TSA.[14]

Under this framework, operational RWA has steadily increased and, given increased CET1 ratios, the implied levels of capital dedicated to operational risks have risen significantly. Exhibit 4 illustrates these developments for the five top European banks.

Exhibit 4: Operational RWA to total RWA ratio and Operational CET1 for 5 top European Banks
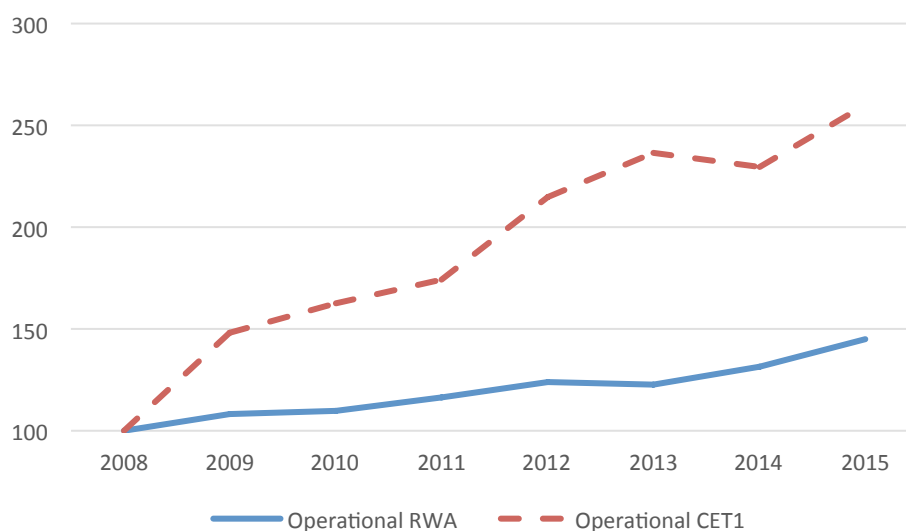
Note: This graph shows the time trend of operational risk RWA and the corresponding CET 1 capital (operational RWA*CET1 ratio) for the largest 5 European banks. Operational risk RWA was not consistently publicly disclosed by US banks before 2013.

2.3 BCBS Current Proposals – The SMA

Under the Basel II framework, banks were "encouraged to move along the spectrum of available approaches as they develop more sophisticated operational risk measurement systems and practices" (BCBS, 2006). However, the BCBS's latest proposals on operational risk, which were first put forward in October 2014, then subsequently updated in March 2016, envisage discontinuing all three current approaches, including the AMA, moving all banks to an updated standardized approach, the SMA, that incorporates both backward-looking and scale-based elements. The BCBS cited non-convergence of risk measurement methodologies as a key reason for withdrawing the option of internal modeling of operational risk under the AMA (BCBS, 2016). With a logic analogous to recent proposals for credit risk RWA, the BCBS is looking to achieve greater comparability of operational risk RWA across banks and regulatory jurisdictions by implementing a single standardized

---

[14] Among the four UK G-SIBS (Barclays, RBS, HSBC Standard Chartered), only Barclays utilizes AMA.

approach that is more risk-sensitive than current standardized approaches. As the BCBS states, "the combination of a simple standardized measure of operational risk and bank-specific loss data provides a sufficiently risk sensitive measure of operational risk…[and] meets its objectives of promoting comparability of risk-based capital measures and reducing model complexity."

Under the BCBS proposals, the SMA will be determined by two components (see Appendix B): first, a business indicator ("BI") component with coefficients increasing with scale (to reflect the BCBS's view that operational risks increase more than linearly with the size of banks); and second, an Internal Loss Multiplier ("ILM") based on a Loss Component ("LC") that factors in a bank's own loss history. The components are calibrated such that a "bank with a Loss Component equal to the BI Component is a bank with an exposure at the average of the industry." In this case the ILM is set at 1 and the SMA equals the BI component. If the LC is greater than the BI component, then the ILM increases the SMA. Whilst the principle of above (or below) average loss experience acting as a multiplier to a standardized calculation is relatively intuitive, the mathematics of the calculation turn out to be remarkably complex, given the issues around defining loss and different types of income, and the outcomes can be far from intuitive, not least due to the "cliff-effects" arising from the thresholds built into the methodology, and the persistent impact of large losses from many years ago.

There has been significant debate about the impact on banks' overall operational risk capital requirements. There appears to be some ambiguity about underlying intent. In its original consultation document in October 2014 the BCBS asserted that "the current standardised framework comprising the BIA, TSA and ASA is on average undercalibrated, especially for large and complex banks" and suggested that one objective of the proposed SMA is to address this concern. By contrast, in its March 2016 document the BCBS stated that "the objective of these proposals is to not significantly increase overall capital requirements." The consensus amongst analysts (and indeed many regulators) is that implementation of the proposals as currently calibrated would lead to significant increases across all jurisdictions, especially in Europe, where it is estimated that banks would face increases averaging 63% in operational RWA (ORX, 2016). Exhibit 5 shows the estimated potential impact of the SMA for a significant number of banks across different regions.

We note this debate because it is the aspect of the BCBS proposals that has arguably generated most heat. However, in this paper we are not offering an opinion as to whether operational risk capital risk requirements should stay at current levels or be increased (or decreased for that matter). Our concern is to suggest how any specificied quantum of capital can be made most effective in achieving the underlying regulatory objectives.

Exhibit 5: Impact of proposed SMA capital requirement by region
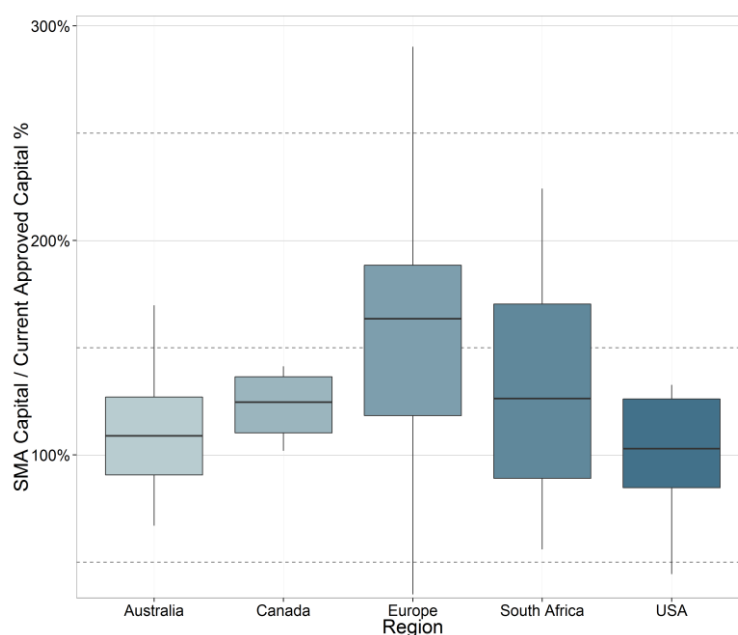
Note: The calculation and figure are sourced from Operational Riskdata eXchange Association (ORX) in its May 2016 response to the SMA proposal publicaly available.

3. THE OBJECTIVES OF OPERATIONAL RISK CAPITAL REQUIREMENTS

Here we offer some observations about prudential regulatory considerations concerning the operational risks that banks face and then distill some specific objectives for the design and functioning of an operational risk capital requirement regime. These objectives can then be used as the criteria by which to evaluate the existing system, the SMA, and any alternative.

3.1 High Level Observations on the Prudential Regulation of Operational Risks

We offer five high-level observations:

1) **Operational risk encompasses a wide variety of risk-types with very different loss characteristics and predictability**. At one end of the spectrum there are risks like individual credit card fraud, teller errors, employee expense fraud, or data entry errors, which are characterized by relatively large numbers of events and relatively small losses per event. These are amenable to the kinds of statistical analysis of historical data commonly used in the credit and market risk arenas. Here it is possible to build relatively robust models of expected and unexpected loss. However, as shown in Exhibits 6 and 7, banks financial losses from operational risks are dominated by relatively few large loss events from limited number of risk types, namely regulatory action, large-scale fraud and rogue traders.

Exhibit 6: Distribution of loss events by size
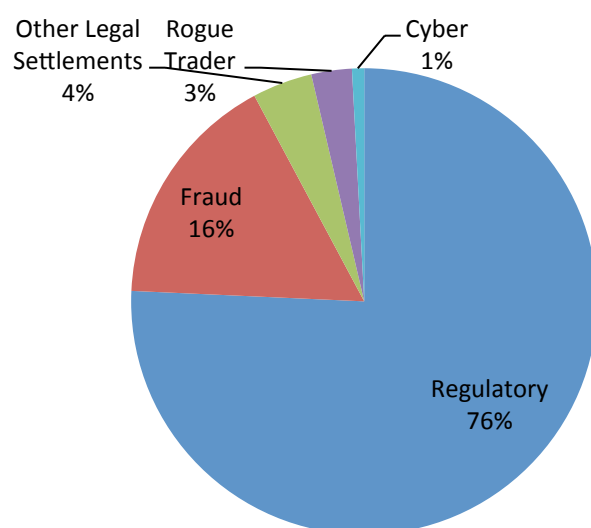
A. Aggregate loss amount
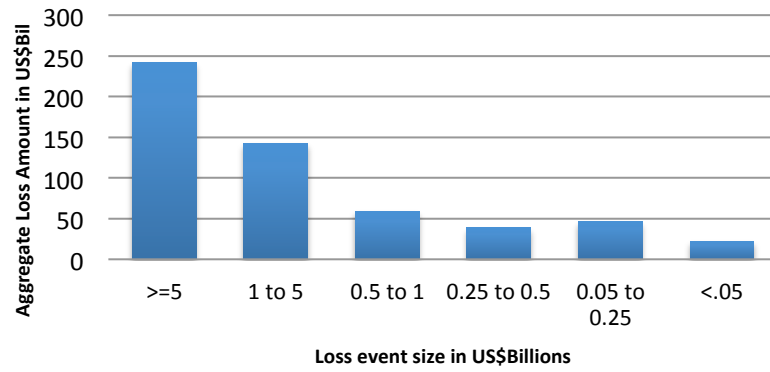
B. Number of events
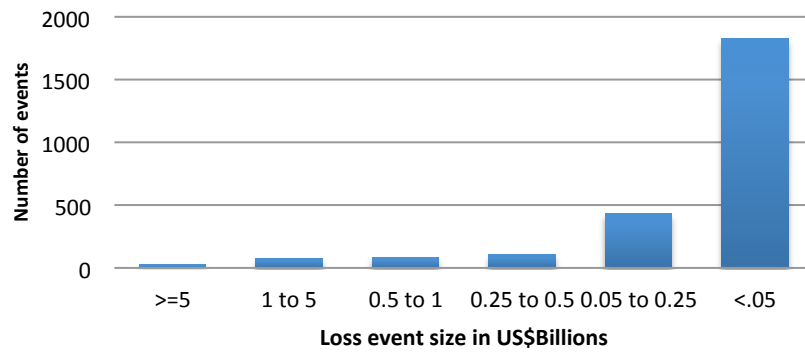
Note: The distribution of operational losses are compiled using the Public Source ORX News data. The data set consists of 3,250 publically-disclosed banking operational loss events.

**Exhibit 7: Categorization of top 250 operational losses**

Note: The top 250 losses in aggregate amount to US$468 billion, representing 85% of all losses in the Public Source ORX News data. The data consists of 3,250 publically-disclosed banking operational loss events.

The risk types that have led to the largest operational losses are much less amenable to prediction based on historical loss data than most credit and market risks for a number of reasons:

a. Individual banks are unlikely to experience enough loss events of the same type to enable robust prediction. Whilst many banks have experienced more than one significant regulatory penalty, very few have experienced multiple significant penalties for the same type of regulatory issue (and when they have it is normally been as a result of making separate settlements with different regulatory or law enforcement agencies for the same event). Likewise very banks have suffered more than one significant fraud or rogue trading instance.

b. Significant regulatory issues tend to result in major penalties (or customer recompense) for all or most of the banks participating in the specific market (e.g., auction-rate fixing, sanctions breaches) or selling the particular product (e.g., PPI, MBS). Exhibit 8 shows the composition of major regulatory penalties. Participation in a market in which one bank has been penalized appears a better predictor of a future loss of this kind than a bank's own loss history.

Exhibit 8: Composition of regulatory-related operational losses
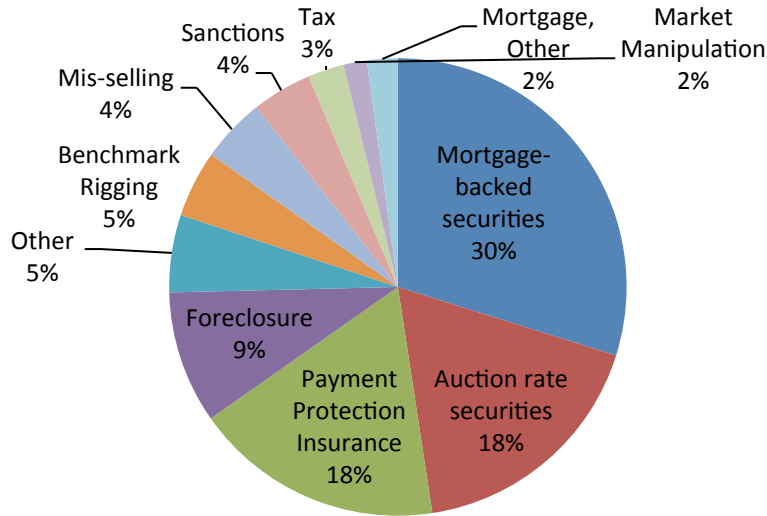

Note: This graph shows the composition of 151 regulatory-related settlements and fines categorized by misconduct type compiled using the Public Source ORX News Data. The aggregate total of these losses is US$347 billion, or 76% of all top 250 banking operational losses.

c. With regulatory penalties the lag between the event and recognition of the loss is usually several years, which means predictions based on losses would inevitably miss the intervening years. Moreover, the timing and definition of the losses can be somewhat opaque since they might involve multiple penalties, settlements, customer compensation, lookbacks, business restrictions and remedial action, and because accounting conventions differ on the timing and precise nature of loss recognition.

d. Significant regulatory penalties tend to result in major changes to market practices, ranging from the withdrawal of the offending products (e.g., PPI) to radical changes in management practices (e.g., salesforce incentivization) or market mechanisms (e.g., LIBOR). In these instances, significant penalties can be seen as indicating the closure of a problem rather than an indicator of future problems.

e. Regulatory penalties escalated enormously in the years following 2008, but have now fallen to some degree. Given that regulatory enforcement actions now dominate most banks' operational risk losses, a predictive model for operational risk based on a bank's historical losses would have massively underestimated the losses incurred by most major Western banks from 2009 onwards and may overestimate them now. Exhibit 9 shows regulatory penalties over time by types. The graph illustrates that

the majority of regulatory operational losses were relating to three categories – payment protection insurance ("PPI"), mortgage-backed securities ("MBS") and auction rate securities ("ARS") – all of which saw a wave of penalties and customer recompense before diminishing.

These characteristics make us skeptical about the predictive quality of historical loss data for events such as regulatory penalties, major frauds and rogue traders – and given that these types of operational losses dominate overall operational risk losses, skeptical about the value of historical loss data in predicting future operational risk losses. Some (eg Curti and Miguies, 2016) have claimed to demonstrate that "past losses are predictive of future exposure". However, we think this analysis is flawed, at least when applied to determining operational RWA. While past losses are predictive for the more frequently recurring operational risk types, this does not appear to hold for the more extreme events that dominate overall losses and thus drive the determination of operational RWA. As Ames et al (2015) point out: "as operational risk is characterized by highly skewed and extreme outcomes, the assumed severity distribution must also be higly skewed and allow for extreme outcomes. The problem is that such distributions also tend to be highly sensitive to the sparse data used for parameter estimation" (Ames, Schuermann and Scott, 2015). Moreover, spurious time-series correlations can arise where there are multiple legal settlements/fines over time for a singular regulatory failing, as Exhibit 10 illustrates using the LIBOR example.

Exhibit 9: Regulatory-related operational losses by category and year
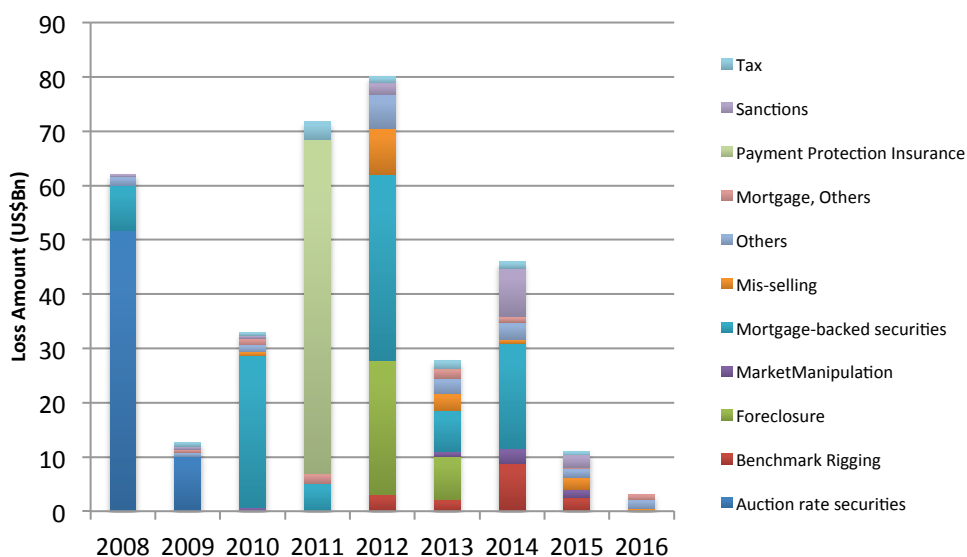
Note: This graph shows aggregate amount of regulatory-related settlements and fines over time totaling US$347 billion, compiled using the Public Source ORX News Data.

Exhibit 10: Libor-rigging regulatory timeline
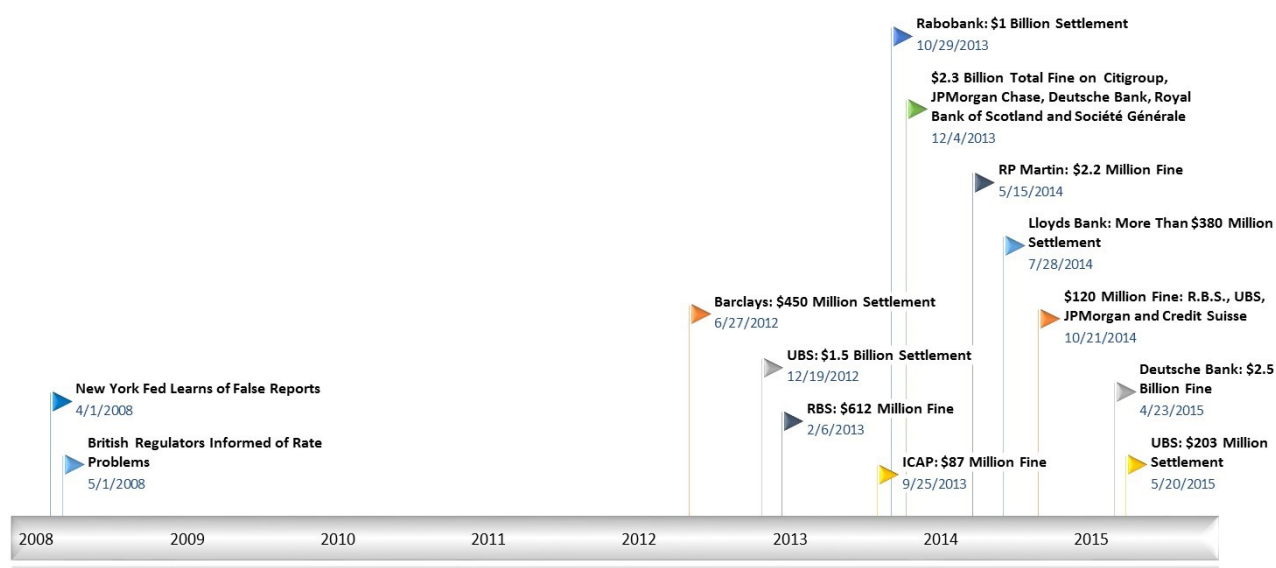
Note: This graph plots the timeline of major events and regulatory fines related to libor-rigger issues as recorded in the press.

Finally, it is worth considering cyber risk, which most bank CEOs would put at or near the top of their operational risk management priorities. Thus far, banks' financial losses directly attributable to cyber risk have been relatively low.[15] An empirical model based on past losses would almost certainly underestimate future risks. Moreover, cyber risks would appear to have the potential to generate disproportionate negative externalities relative to the direct losses incurred by the bank suffering the risk event (eg leakage of customer data, or compromising the integrity of the payment system). An operational risk capital requirement based on historical loss patterns may incentivize too much "closing the stable door" rather than more forward-looking assessment and prevention.

2) **The regulation of operational risks should focus on avoiding negative externalities.** As with most aspects of the prudential regulation of banks, operational risk capital requirements should be designed with the objective of curbing negative externalities. Significant operational risk events can disrupt the provision of credit, the functioning of markets especially if the banks affected are key players in certain specialized markets, and could even threaten financial stability. Therefore, capital requirements should create incentives for banks to internalize and mitigate the social costs associated with these risks.

---

[15] For example, a partial review indicates that financial institutions' average annual costs due to cyber attacks are around $28 million in 2015 (Ponemon Institute 2015).

Since the term "operational risk" covers an enormously wide (arguably, unbounded) variety of risk types, it is possible to postulate an extremely broad range of operational risk events that could have negative externalities. Two types of event where operational risk capital might appear immediately relevant are: 1) the event that leads to a financial loss so catastrophic that the bank fails with direct implications for its customers and potential implications for financial stability; and 2) the event that leads to the bank suffering a financial loss so significant that it needs to constrain its lending activities in order to survive. In both cases it is easy to see how having substantial operational risk capital resources in place might help the bank minimize the negative externalities.

However, it is also possible to envisage operational risk events that might not pose significant direct financial losses to the bank suffering the event, but could create negative externalities of far greater magnitude. For example, a bank with a relatively small share in a particular market for traded instruments might fall victim to a cyber attack that propagates to an extent that it cripples the entire market, imposing far greater costs on other market participants than on the bank itself. Ultimately, legal redress and regulatory penalties might lead to the bank bearing a fuller share, but at the first instance, the direct cost to the bank might be only a tiny fraction of the societal cost. The point here is that the scale of the loss to the bank might not be the best measure of societal impact, and therefore of regulatory concern.

3) **Operational risk capital requirements can reduce negative externalities by enhancing loss absorbency.** Where the operational risk event does cause significant losses to the bank suffering the event, which in turn cause significant negative externalities, capital can obviously provide loss absorbency and thus mitigate the knock-on consequences. Where an operational loss event is of such catastrophic scale as to cause the bank to fail, the regulator wants to ensure sufficient capital is in place to enable enough "gone concern" loss absorbency to facilitate orderly resolution, dismemberment, or sale without recourse to the taxpayer. Where an operational loss event does not lead a bank to collapse, but is of a magnitude to put a dent in capital ratios, the regulator wants to ensure sufficient capital is in place to provide enough "going concern" loss absorbency for the bank to take the loss in its stride without undue disruption to the bank's core financial intermediation activities. In both cases, there is a clear rationale for requiring banks to hold capital for operational risk events, as long as that capital can be used when the event happens.

One may argue that since operational events tend to be idiosyncratic, the negative externalities from operational losses should be limited, as firms can

switch to lenders that are not affected. We make several observations in light of this view:

- First, it is well documented that switching to different lenders is costly and substitution is limited, especially for smaller firms (Petersen and Rajan, 1994; Ashcraft, 2005; Chodorow-Reich, 2014).
- Second, operational losses due to regulatory actions often affect most banks in a given market (eg ARS, MBS, PPI), which has systematic as opposed to idiosyncratic impact on financial intermediation activities.
- Finally, if it were true that operational losses pose little threat of negative externalities, then it seems hard to explain why there are operational risk capital requirements in the first place.

4) **Operational risk capital requirements are only one component of the prudential regulation of operational risks.** As the BCBS has repeatedly stressed (BCBS 2011), supervisory oversight and other regulations relating to specific risk-types (eg on data privacy) or governance arrangements (eg senior management accountability) are also necessary, since operational risk capital requirements are a relatively blunt and imperfect tool for internalizing the social costs of operational risk failures and mitigating their consequences. So the effectiveness of an operational risk capital requirements regime has to be assessed in the light of how it complements such non-capital based mechanisms. In this context it is perhaps worth observing that the regulatory regimes for other industries in which operational risk events can significant negative externalities – such as aviation, shipping, pharmaceuticals or nuclear – tend not to use capital requirements as regulatory instrument, but instead put more reliance on standards, reporting, inspection and accountability.

5) **The regulatory approach to operational risk must recognize that banks do not approach operational risk in the same way as credit and market risk.** Banks willingly take on credit and market risks to generate a return. This is how they make money and perform their role in the broader economy. The allocation of capital across different types of market and credit risks to generate returns is central to how banks are run. Measures of return on credit and market risk RWA are typically embedded at every level of decision-making from client relationship management and product pricing, to overall portfolio management and strategy.

This is not how banks approach operational risk. Banks do not take on operational risk to generate a return in the way they do credit or market risk, deliberately accepting a quantifiable risk exposure for a prospective return. Depending on the type of operational risk, their approach will typically vary on a spectrum from reducing the risk to acceptable parameters within a cost constraint (eg with credit card fraud), or simply minimizing the risk as far as possible without any constraint (eg certain types of regulatory or cyber risk).

Even in businesses where the operational risks are large relative to the credit or market risks, such as custody or clearing, the operational risks are seldom measured and managed through operational RWA. Indeed, the operational RWA associated with different types of operational risk rarely gets mentioned or considered in managing operational risks, often because it is impossible to identify separately, frequently because it bears no resemblance to management's judgement of the prospective risk, and also because operational RWA is seen more as a tax than a controllable variable.

The implication of this is that whereas credit and market risk RWA are powerful influences on management behavior, operational risk RWA has very limited, if any, influence.

Pillar 3 disclosures provide an indication of how differently banks approach operational RWA versus credit and market risk RWA. Pillar 3 disclosure obligations vary significantly by jurisdiction, but two examples illustrate the point. Barclays' annual Pillar 3 disclosure is 170 pages long and 58 pages are devoted to a discussion of RWA. Of these 58 pages, three are dedicated to the £56.7bn of operational RWA that constitutes 16% of Barclays' total. For credit and market risk Barclays provides a detailed disaggregation by business line, credit grade and geography, and explains the derivation of the RWA figure, showing the performance of their IRB models. By contrast, the disclosures on operational RWA are much more cursory (Barclays, 2016). Bank of America's Q3 Pillar 3 disclosure amounts to 27 pages, of which roughly half a page is devoted to the precisely $500bn of operational RWA they carry, representing 32% of their total RWA (Bank of America, 2016).

3.2 Objectives of Operational Risk Capital Requirements

Where do these observations take us in thinking about what should be the objectives of an operational risk capital requirement approach? Whilst we recognize that this is a matter of judgement, we would offer for consideration the following five objectives:

1) The approach should be designed to *provide incentives to minimize and mitigate negative externalities from operational risk events* by internalizing potential societal costs. This implies focusing on the risk types with most potential to cause negative externalities.

2) The approach should *provide both "going concern" and "gone concern" loss absorbency* to minimize the negative externalities arising from operational risk events creating significant financial losses. This implies making operational risk capital "usable"

3) The approach should be designed to *complement supervisory interventions on operational risk* , reflecting the fact that capital requirements will not be the most effective tool for all risk types. This implies making some element of the capital requirements subject to regulatory discretion

4) The approach should *be as simple as possible avoiding over-complex modelling* to facilitate implementation, maximize transparency and minimise the potential for gaming. This is consistent with the BCBS' emphasis on reducing complexity.

5) The approach should seek to *enable meaningful comparability across institutions and jurisdictions*, so that to the extent possible, equivalent risks result in equivalent capital requirements. This is also consistent with the BCBS' desire to enhance comparability.

4. ASSESSMENT OF THE CURRENT APPROACH TO OPERATIONAL RISK CAPITAL REQUIREMENTS

The Basel II approach to defining operational risk capital requirements, combining the BIA, TSA and AMA methods arguably fails to meet any of the five objectives described above. Specifically:

1) *Provide incentives to minimize and mitigate negative externalities from operational risk events.* The BIA and TSA approaches create no incentives to minimize and mitigate the societal downsides of operational risk, because there is no connection between the effectiveness of a bank's operational risk management and its operational risk RWA. Under AMA, there is in theory more scope for creating such incentives. However, in practice, most banks that use the AMA see signifcant disconnects between how they manage their risks and how they determine their operational RWA (Ames et al, 2016).

2) *Provide both "going concern" and "gone concern" loss absorbency*. Here, for reasons discussed in more detail in the next section (because the same applies to the SMA), none of the three methods appears satisfactory, particularly in terms of "going concern" loss absorbency.

3) *Complement supervisory interventions on operational risk.* The BIA and TSA approaches contribute little or nothing to a more comprehensive approach towards regulating operational risk. Moreover, as the BCBS have noted, the fact that they are based on revenue means that the operational risk capital requirements can fall even when underlying operational risks are rising: "The most common situation involved banks experiencing a decline in their GI [Gross Income] due to systemic or bank-specific events, including those

involving operational risk losses, and seeing a commensurate decline in operational risk capital when intuitively this should have either stayed at the same level or increased" (BCBS, 2014). By contrast, the AMA approach can be used as part of a broader array of tools, but this depends on how it is implemented by individual regulators.

4) *Be as simple as possible, avoiding over-complex modelling.* The BIA and TSA approaches are commendably simple, but precisely because they are so simple, fail to give any insight about the operational risks faced by an institution. The AMA, on the other hand, is intrinsically complex. As the BCBS put it in justifying their decision to withdraw the approach "The inherent complexity of the AMA and the lack of comparability arising from a wide range of internal modelling practices have exacerbated variability in risk-weighted asset calculations" (BCBS, 2016).

5) *Enable comparability across institutions and jurisdictions.* The BIA and TSA approaches are directly comparable across institutions because they are formulaically driven from historical revenue. However, it is a consistency of an unhelpful kind, in that it is equally uninformative about the operational risks faced by any institution. The AMA, on the other hand, produces RWA results that are very difficult to compare across institutions and even more difficult to compare across jurisdictions.

5. ASSESSMENT OF THE BCBS PROPOSALS AGAINST THE OBJECTIVES

The BCBS proposals offer one significant advantage over the current system in that they envisage one method of calculating RWA rather than three, which certainly aids simplicity and comparability. However, in general they still score poorly against the objectives as we have described them, since these proposals fail to address the lack of loss absorbency in the current system, provide limited incentives and remain mechanistically backward-looking.

1) *Provide incentives to minimize and mitigate negative externalities from operational risk events.* Because they combine a size-based component and a multiplier related to past loss history, it is hard to see how the SMA provides anything but very weak incentives to minimize and mitigate operational risk events with negative externalities. The principal channel through which such incentives can be created is through the impact of an institution's loss history on the ILM. Yet given that the ILM is derived from a 10 year loss history, actions taken today will have limited impact. A bank that takes decisive action to remedy control weaknesses following a particular loss event will see no RWA benefit relative to a bank that is much more dilatory in its rectification unless and until the less diligent bank suffers another loss event. In addition,

past losses from business lines that are no longer relevant can have a major impact on capital requirement. For instance, analysts have estimated that the CET1 ratio of Lloyds will decline by 390bps under the SMA, of which 210bps is due to PPI-related historical losses (Deutsche Bank, 2016). Since Lloyds no longer sells PPI and has revamped its sales processes, the logic of PPI being a powerful driver of its operational RWA for the next decade could be questioned. Some of the secondary incentives created by the SMA seem somewhat perverse. The methodology is strongly skewed to advantage smaller banks, thus disincentivising scale. We are not convinced that this make sense. The exclusion of risk mitigants like insurance from the SMA methodology effectively disincentivizes their use (see Ames et al, 2016, for a discussion of insurance).

2) *Provide "going concern" and "gone concern" loss absorbency.* The SMA does not address one of the biggest flaws in the current system, the lack of loss absorbency. Whilst banks are required to hold significant amounts of capital, ostensibly to enable them to absorb operational risk losses, in practice the capital is unusable, in particular on a "going concern" basis. This is best illustrated with a stylized example. Suppose a bank has RWA of $100, comprising credit RWA of $70, market RWA of $15 and operational RWA of $15. Suppose also that the bank has $10 of equity, so a CET1 ratio of 10%. If the bank suffers an operational loss of $1 and wants to maintain its CET1 ratio of 10% its only immediate option (ignoring the flow of profits and retained earnings for the moment) is to reduce credit and market risk RWA by $10. In fact, because the operational RWA will increase as a result of the loss, the bank will have to reduce credit and market RWA by more than $10. In this highly stylized example, the result will be exactly what a regulator would not want: the operational risk loss results in a sharp reduction in credit provision and market facilitation.

There are of course a number of potential objections to this stylized example. For instance:

- If the bank is operating at 11% CET1 – i.e., it has equity of $11 – it can absorb the loss without breaching the 10% CET1 limit or reducing credit or market RWA. This is true, but begs the question of what the $1.5 of equity represented by CET1 of 10% on operational risk RWA of $15 is for, because it will not contribute to absorbing this loss. Moreover, given investor pressure for returns on equity, banks do not typically carry significantly more equity than they believe they need (which is typically substantially above the prescribed regulatory minimum, given the consequences of breaching the minimum)

- If the bank's retained profits in the next period amount to at least $1 then it can absorb this loss without denting its capital ratio or reducing credit or market RWA. Again, this is true, and in practice this is how much of the most significant operational risk losses are absorbed, but once more it begs the question of the purpose of the operational RWA (It also misses the point that it is through retaining earnings that the bank can support credit growth: in essence the deployment of retained earnings to absorb the loss is a sacrifice of future credit and market RWA)

A more subtle objection points to the fact that if our stylised bank lost $10 on credit risk (or market risk) rather than operational risk, the same problem might arise. From this perspective, the problem is less about operational risk than the construct of requiring banks to hold minimum capital ratios: capital below the required ratio cannot be used, at least not on a "going concern" basis because breaching the minimum would compromise the bank's status as a "going concern". The higher the perceived minimum capital ratio, the more capital is available for "gone concern" loss absorbency, but with potentially no gain (or even a negative impact) on "going concern" loss absorbency. There is merit in this argument, but this broader discussion is beyond the remit of this paper. Here we would simply observe that there may be reasons to believe the problem of "going concern" loss absorbency is even more acute for operational risk than it is for credit or market risk:

- First, when a bank writes off a credit or closes out a loss-making position then the RWA associated with the loan or position is extinguished. In this way some capital is released, although typically only a small fraction of the loss. This will not happen with operational RWA under the BCBS proposals – indeed a significant loss will lead to an increase in RWA through the ILM[16].

- Second, a bank will typically react to significant credit or market losses by reducing the related portfolio with a consequent reduction in RWA (although this may be offset by negative credit migration). The point is that a bank will be actively managing its credit and market RWA both to maintain target capital ratios and in response to the risk environment. The SMA does not provide any scope for operational RWA to be managed in an equivalent way.

---

[16] As an illustrative example, suppose a bank holds $100 of risk-weighted assets composed of $70 of credit RWA, $15 of market RWA and $15 of operational RWA, and furthermore, suppose the bank holds $10 of equity, yielding a risk-weighted capital ratio of 10%. A loss of $2 in the credit portfolio reduces equity by 2 and (assuming 100% risk-weghting for simplicity) credit RWA by $2. To maintain a CET1 ratio of 10%, the bank would need to cut $18 of credit or market RWA. On the other hand, if the loss of $2 is operational in nature, under the SMA proposal operational risk RWA might increase to $16 . To maintain a 10% CET1 the bank would need to cut $21 of credit or market RWA.

- Third, it could be argued that whilst both are relevant, the relative need for "going concern" and "gone concern" loss absorbency is different for operational risk versus credit and market risk. Whilst bank failures typically involve a complex combination of loss drivers, it is extremely rare for the primary driver to be an operational risk event. Barings is probably the clearest example in recent history. By contrast we have seen numerous failures of banks driven by credit losses and losses on market portfolios. If this is correct, we might want to strike a different balance between "going concern" loss absorbency and "gone concern" loss absorbency" for operational risk.[17]

The SMA does not provide "going concern" loss absorbency, but neither does it automatically help in a "gone concern" scenario. Suppose our stylised bank suffers an operational loss of $10 wiping out all the equity. In this case the $1.5 of equity devoted to operational risk will have contributed in part to protecting the taxpayer. However, this is only true to the extent that the bank is completely dissolved. To the extent that some aspects of the bank are rescued, folded into another bank or in some other way kept operational, as is likely to be the case, then these component parts will have operational risk RWA, which will likely be elevated given the reason for the failure. So even in a "gone concern" scenario, there may be limits to the loss absorbency provided by the bank's operational risk capital.

3) *Complement supervisory interventions on operational risk.* As with the current system, it is difficult to see how the SMA can be used to complement or reinforce non-capital based regulatory interventions on operational risk, since the RWA is mechanistically determined on the basis of scale and past loss history. It is a "top-down" approach that does not link to identified risk drivers (Peters et al, 2016). Regulators who want to give banks a capital incentive to rectify particular issues or in anticipation of new threats will need to use a Pillar 2 buffer or the stress tests, as some already do.

4) *Be as simple as possible, avoiding overly-complex modelling.* The SMA replaces three methodologies with one.[18] However, the methodology is not

---

[17] The importance of "going concern" loss absorbency also depends on whether firms can find alternative lenders when the original lender cuts credit availability due to a loss event. It has been well documented that lending relationships are sticky; switching is costly and substitution is limited (Petersen and Rajan, 1994; Ashcraft, 2005; Chodorow-Reich, 2014). Moreover, major operational losses (e.g. regulatory fines) tend to be correlated across banks, thus aggregating the potential impact on credit availability.

[18] Banks that currently use the AMA tend to mention one advantage of the SMA: it will enable them to determine their operational RWA figure, however arbitrary, without devoting nearly as much analytical effort and time as is currently involved in calculating the (almost equally arbitrary) AMA figure. Yet this is hardly a ringing endorsement.

particularly simple, since it combines a scale-based element that increases in a non-linear manner with a multiplier based on loss history. The interaction of these two components can produce capital requirements that are unstable and excessively sensitive to large historical losses: two banks with the same risk distribution can end up with drastically different capital requirements depending on the random realization of past losses (Peters et al., 2016). A large GSIB that suffers a one-off significant operational loss (say the settlement of a regulatory issue related to a business that is now closed) could face a significant capital penalty for years thereafter, given the interaction of the loss multiplier with the BI component, exacerbated by the GSIB buffer. Conversely, a smaller bank that has under invested in cyber security, but has been lucky enough so far to escape significant operational loss might enjoy a significant capital benefit. The consultation responses also highlighted a range of methodological issues in the derivation of the ILM, including pronounced "cliff effects" in the banding of loss events by size, the determination of the timing of events, and the interaction with accounting practices.

5) *Provide comparability across institutions and jurisdictions.* One of the strongest arguments made by the BCBS for its proposals is increased comparability of operational risk capital requirements across banks and regulatory jurisdictions, because there would no longer be the differences between banks adopting BI, TSA or AMA, nor the differences in definitions and methodology between banks using AMA. Indeed the proposals are framed as part of the broader effort of "[e]stablishing consistency in the implementation of post-crisis regulatory reforms" (BCBS, 2016). It is true that implementation of the SMA would lead to consistent application of more comparable methodology. However, the comparability would be largely meaningless, since it would not translate into comparable treatment of operational risks. Banks of the same size facing very different prospective risks might have the same operational risk RWA. Banks facing equivalent risks could have very different operational risk RWA due to size, random historical events, or the number of subsidiaries (Peters et al, 2016).

One aspect of the proposals that seems to demonstrate historical "path dependency" rather than any underlying logic is the translation into RWA. The SMA methodology generates a required capital amount which is then multiplied by 12.5 (a throwback to the old Basel I 8% capital weighting) to arrive at a figure for operational RWA. This creates an entirely illusory appearance of equivalence to market and credit RWA. Moreover, in including this contrived operational RWA in the calculation of a bank's CET1 ratio we are arguably eroding the credibility of this ratio (which of course is already challenged by ongoing arguments about discrepancies in credit and market risk methodologies).

This is a point worth emphasizing. Whatever the flaws of the methodologies to derive credit and market RWA, they at least represent relatively coherent approaches to defining the risks attaching to loans and market positions. The risks are well defined, and in most cases, bounded, in that it is possible to define a maximum exposure or value- at-risk (this is not the case when writing certain types of derivatives, but even here it is possible to plot a meaningful loss distribution curve). Definitions of default, exposure, collateral, mitigants, etc, are well understood and being made increasingly consistent, and there are an array of well understood and demonstrably effective analytical tools to translate vast amounts of data on historical losses into assessments of future risk. Of course there are still many areas of debate, such as with low default portfolios, or asset classes with low idiosyncratic risk and high systemic risk, but the fact is that our analytical tools for identifying and measuring credit and market risks are far more advanced than they are for operational risks.

Moreover, while it is important to create a common metric across credit and market risks, since banks can often create the synthetic equivalent of a credit exposure through a market risk approach and vice versa, this is simply not true of operational risks. Few bankers would consider a dollar of operational RWA as conceptually representing an equivalent risk as a dollar of credit or market RWA. Indeed, most bankers would consider a dollar of operational RWA as simply a somewhat arbitrary regulatory construct that has very little to do with anything, including the management of operational risk.

6. AN ALTERNATIVE APPROACH

We believe the BCBS proposals fail to achieve the regulatory objectives, largely because they are focused on fixing perceived problems about comparability and complexity with the existing system, applying an approach not dissimilar to the logic used to address the issues with IRB and the Standardized methodologies for credit risk RWA, rather than going back to first principles to examine the underlying objectives underpinning the imposition of regulatory requirements for operational risk capital.

We would suggest that, rather than just debating the finer points of the ILM methodology, the non-linearity of the BI component, or the potential impact on overall capital requirements, serious consideration be given to a radically different approach designed from first principles to achieve the regulatory objectives.

For example, we would see considerable advantages to an alternative approach comprising the following core components:

1) *Replace operational RWA with an incremental operational risk capital buffer requirement expressed in absolute dollar terms*. The operational risk capital

requirement would be switched from being a determinant of the denominator of the CET1 ratio to being an increment to the required numerator. Specifically, the total RWA would comprise only market and credit RWA, but the CET1 ratio would have a new component, the operational risk capital buffer.[19] This change could be made capital neutral, with the required CET1 ratio increasing to reflect the operational risk capital buffer.[20] Making this switch would, when combined with the other proposed changes, facilitate loss absorbency, and help align incentives appropriately. This will also enhance the transparency and credibility of the underlying capital ratio

2) *Determine the quantum of this operational risk capital buffer using a combination of scale, history and judgement about prospective risks.* Whilst there are various ways of constructing the buffer and we do not attempt to flesh out a detailed approach in this paper, we see merit in combining three separately identified components: one related to scale; a second based on proven models for those types of high-frequency operational risk for which a bank has sufficient data to provide statistically robust predictions; and a third based on a process of structured judgement by the regulator.[21]

The first component would be deliberately crude, like the SMA BI component (although we are skeptical about the non-linear scalar in the BI), but would ensure a basic minimum level of operational risk capital.

The second component would only cover those risk types for which there is sufficient clarity of definition and data to provide confidence that past experience represents a reasonable guide to future losses. This will be most relevant for frequently recurring losses like credit card fraud, data entry errors and petty theft. Banks that have the data and models to do this sufficiently well (typically those that have been using AMA) would get an appropriate discount off the scale-based minimum. This second component would not contribute

---

[19] For example, Bank of America has total RWA of around $1,550 bn, of which $500 bn is operational RWA. Its common equity tier 1 ratio is 11.6%. This implies that about $58 bn CET1 capital (out of $180 bn total CET1 capital) is accounted for by operational risks. To be capital neutral the $58 billion will become the buffer. The new CET1 ratio under our approach will be calculated as 180/(1550-500) = 17.1%. The difference between this ratio and the current CET1 ratio is 17.1%-11.6%=5.5%. This is equal to the buffer divided by our new total RWA, namely the sum of market and credit RWA.

[20] In practice, banks' models calculate the equity capital required based on 8% capital requirement, and multiply it by 12.5 to calculate the RWA. In the Bank of America example, the model-based equity capital would be $40 billion. To be conservative, we calculate the buffer using the actual CET 1 ratio (as opposed to 8%), which is $58 bn.

[21] A simpler variant would exclude the second component on the grounds that it will represent only a small part of the total capital requirement and can be reflected in the calibtration of the scale-based component. While this is true, and a two component model would be simpler, we see benefit in encouraging banks to build their data gathering and modelling capabilities for these kinds of risks.

much to the overall total and could be omitted on grounds of simplicity, but has the merit of encouraging the development of modelling for those risk types for which this is appropriate.

The third component will be explicitly judgmental, determined by the regulator informed by their knowledge of prospective risks facing the bank, the quality of the bank's operational risk management capabilities, infrastructure and governance, and broader industry developments and challenges. To derive this component of the buffer, regulators would need to take account of: i) specific known risks, such as pending legal settlements for which it is possible to estimate potential losses (such as those that US banks disclose under Reasonable and Probable Losses); ii) risk types where it is possible to describe the loss mechanisms and loss drivers, even if the amounts at stake are uncertain, such as rogue trading, or emerging regulatory issues; iii) risk types where the dynamics of loss and negative externalities remain extremely uncertain, such as cyber risk; iv) assessment of an individual bank's relative strengths and weakenesses in infrastructure, capabilities and governance. The process could involve banks submitting structured self-assessments along these lines to their supervisor, as input into the regulatory judgement. To minimize differences across jurisdictions, it would make sense for BCBS to establish some principles for determining this component of the buffer, but because it will be separately identifiable, differences in regulatory approach will be relatively transparent to the market.

3) *Provide a mechanism by which banks can seek regulatory approval to use at least part of this buffer* on a "going concern" basis when they suffer a significant operational risk loss, subject to agreeing whatever remedial action is appropriate and a time frame and plan to return the buffer to whatever level the regulator determines. Here again, it would make sense for the BCBS to establish some principles.

In this paper we do not attempt to provide the details of how to determine or calibrate the components of the buffer, nor a detailed description of the mechanics by which it could be used and rebuilt. Our focus here is on discussing the potential pros and cons of such a radically different approach versus the SMA. How would this be different?

First, by configuring the operational risk capital requirement as an absolute buffer with a mechanism to allow it to be used, this approach facilitates "going concern" loss absorbency. This in turn reduces the need for a bank to respond to an operational risk loss by reducing lending or conducting a fire-sale of marketable securities.

Second, the judgmental component of the buffer enables the regulator to make a forward-looking assessment of operational risks (which can of course include an appraisal of the extent to which past losses indicate future risks) and to connect

capital requirements to identified risk drivers. The advantages of this can be seen by considering a situation where one bank is penalized by the regulators for miss-selling a product, but it is recognized that other banks sold the same product in a similar way and are therefore likely to get penalized as well. The SMA (unlike the AMA) gives no scope for including this risk in the operational risk capital requirement, while under the alternative approach the regulator can increase the required capital buffer of the affected banks. This approach also allows a regulator to judge the difference between a loss that is a harbinger of future problems and a loss, such as a final settlement, that closes off a problem. This cannot be done under the SMA.

Third, the judgmental component of the buffer allows the regulator to create powerful incentives for banks' management to reinforce their operational risk controls. Banks that fall behind on enhancement programs can be penalized with incremental capital requirements. Banks that demonstrate material improvements can be rewarded with reductions in the required buffer. No such flexibility exists in the SMA.

Fourth, including historical data for those risk types for which sufficient loss data exists to provide robust predictions, leverages the experience of those banks that adopted AMA and encourages such practices amongst other banks. It seems perverse to throw away the benefits of the considerable investments in data collection and modelling where these are of demonstrable benefit. We anticipate that this element of the buffer would be relatively small, since while these types of risks account for a very proportion of the operational risk events by number, they account for a very small proportion of the total by value. For example, Barclays notes that 83% of its operational losses by frequency were of amounts of less than £50,000.

Finally, removing operational risk from the calculation of a bank's RWA should improve the credibility and comparability of banks' RWA and capital ratios in the eyes of investors and counterparties. Investors have typically paid relatively little attention to the determinants of banks' operational RWA, rightly considering them rather arbitrary and unrelated to the prospective operational risks the banks are facing. However, neither do investors usually strip out operational risk RWA from comparison of capital ratios. One indication this lack of interest in is the fact that bank's operational RWA figures are not readily available on the most common industry data sources, as Exhibit 11 demonstrates. Under both the current system and the BCBS proposals, operationl RWA serve to muddy the assessment of banks' capital strength. Under the alternative approach, investors and other market participants will be able to assess a bank's operational risk capital requirements separately from its capital position relative to the credit and market risks it has taken on.

Exhibit 11: Coverage of Operational RWA in commercially available databases

Source	Coverage of Operational RWA
Factset	Not available
CapitalIQ	Not available
Bloomberg	Limited (missing data for many GSIBs)
Bankscope	Limited (missing/inconsistent data quality)
SNL Interactive	Limited (missing data for many GSIBs)
Compustat	Not available

Note: This table presents an assessment of the data availability across major commercial financial databases that provide information on other relevant data including T1 capital ratio and total RWA.

How would our alternative approach stack up against the five objectives discussed earlier?

1) *Provide incentives to minimize and mitigate negative externalities from operational risk events.* The inclusion of a component of the capital requirements that is determined by regulatory judgement would enable to regulator to incentivize banks' management teams as they see fit.

2) *Provide "going concern" and "gone concern" loss absorbency.* Constructing the operational risk capital requirement as an incremental buffer rather than as RWA and with an explicit mechanism for drawdown would make the capital much more "usable". It could be used to absorb significant losses on "going concern" basis and it would be available for "gone concern" loss absorbency.

3) *Complement supervisory interventions on operational risk.* The judgement-based component of the buffer could be used in conjunction with other initiatives – either industry-wide or institution specific – to reinforce other actions to address operational risk concerns.

4) *Be as simple as possible, avoiding over-complex modelling.* At its simplest the alternative approach would comprise just two elements: a fixed requirement based on scale, and a judgemental component at the regulator's discretion. Where banks are able to demonstrate to regulators' satisfaction that historical loss modelling for specific risk types can provide robust forward-looking risk assessment this can be included. However, BCBS should determine the risk types for which this option is available and define methodological and data standards to be used. The greatest potential complexity would lie in the determination of the judgmental component of the buffer. Here it would be helpful for the BCBS to agree some overarching principles.

5) *Enable comparability across institutions and jurisdictions*. Whilst on the face of it, there is more room for inconsistency in this alternative approach, given the inclusion of a component determined by regulatory judgement in our view, the alternative approach would provide more meaningful comparability than the BCBS proposals. Within a jurisdiction, the regulator would presumably ensure consistency of approach towards the banks under its jurisdiction. Inconsistencies would mainly arise where regulators took different approaches towards evaluating the judgmental component. This could be constrained in part by agreeing some principles on how to exercise this judgement for different risk types and perhaps reinforced by periodic peer review of processes. Yet even if such differences persisted, the fact that this component would be separately identifiable would enable market participants to form their own view of comparability.

7. ARGUMENTS AGAINST THE ALTERNATIVE APPROACH

In this section we discuss potential counter arguments against the alternative approach. Through our discussions thus far on this topic we have encountered five principal arguments:

7.1 Is this an attempt to reduce the capital requirements disguised as an alternative methodology?

Some will argue that moving from an RWA approach to a buffer approach will reduce discipline and enable banks to achieve a reduction in operational risk capital requirements through "regulatory capture". This is not our intent. In fact we would suggest that the starting point for defining the scale of the buffer should be the same amount of capital as is currently implied by a bank's current operational risk RWA and CET1 ratio (e.g. if a bank has operational risk RWA of $10 and a CET1 ratio of 10%, and thus implied operational risk capital of $1, then the starting point for defining the scale of the buffer should be $1). Exhibit 12 shows, for GSIB banks, what the new capital ratio would be for the buffer to be capital neutral. The blue part is the current CET1 ratio, and the red bar is the part of the new ratio that is accounted for by operational risk capital buffer. Since part of the buffer will be mechanistically defined as a function of scale, part empirically derived for predictable risk types, the only scope for gaming or reduction is the component determined by regulatory judgement. Our view is that regulators should be asked to give a rationale for reducing the capital required below the current number, or below the average for the banks in that jurisdiction.

Exhibit 12: Operational risk capital buffer and CET 1 ratio under proposed capital-neutral change
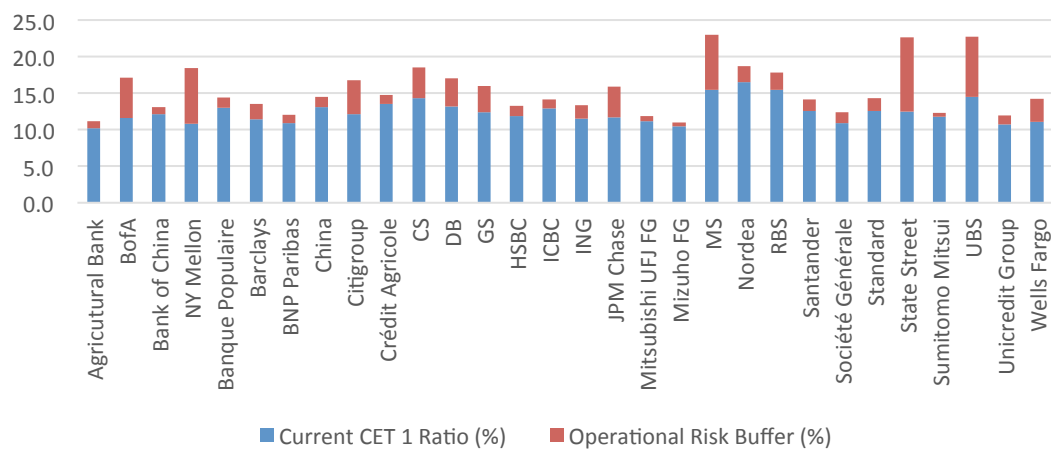
Note: This figure shows the current core tier 1 ratio for GSIB banks (blue bar), and the new ratio according to our approach (blue plus red). The red bar reflects the part of the new ratio that is accounted by operational risks.

In our view, the alternative approach would have significant benefits beyond operational risk, in that it would improve the meaning and coherence of the CET1 ratio. The overall CET1 ratio could now be thought of as expressing the amount of equity capital a bank has against its financial assets as described by its market and credit RWA, with a defined component of this equity identified as being held against operational risks. Market particpants could then form their own views of the adequacy the overall CET1 ratio, the component dedicated to operational risk and the balance deployed against the financial risks.

7.2 Why not achieve the benefits of the alternative approach by using Pillar 2 buffers or stress tests in addition to the BCBS proposals?

The argument here is that some regulators are already using Pillar 2 and stress tests to provide loss absorbency and incentivize management as a complement to the existing system of operational risk capital requirements. Why not encourage all regulators to do this *in addition* to implementing the SMA? We have three responses to this argument:

- First, the fact that regulators are building significant operational risk buffers into Pillar 2 and incorporating them into stress test scenarios is a powerful indication that the current system is failing to provide the right incentives and loss absorbency. The SMA will not change this.

- Second, creating Pillar 2 buffers or using tests in addition to the SMA means that banks' overall operational risk capital requirements would make more

acute the question of what purpose the Pillar 1 SMA capital requirement serves. If equity were a free resource that might not matter. But since the Modigliani-Miller theorem does not apply very well to banks (Baker and Wurgler, 2013; Stein, 2012), an arbitrary capital requirement that sefves limited regulatory purpose is arguably a misuse of society's resources.

- Third, increasing reliance on Pillar 2 and stress tests would seem an odd outcome given the desire to achieve greater comparability. The operational risk buffer we propose would be implemented according to consistent principles, and the judgemental component would be readily identifiable. By contrast, regulators' approaches to incorporating operational risks in Pillar 2 and stress tests are far more opaque.

- Finally, there is real benefit in removing operational risk from the determination of RWA, since its inclusion only erodes the credibility of this metric.

7.3 Won't differences between regulators' use of discretion in determining operational RWA under the alternative approach lead to regulatory arbitrage and investor confusion?

We recognise that some regulators see the elimination of differences in the regulatory treatment of operational risk across jurisdictions as a key benefit of the SMA. From this perspective, our suggestion of a buffer determined in part by regulatory judgment may not strike a chord. We are sympathetic to the view that overreliance on regulatory judgement can lead to a corrosion of standards, particularly where the exercise of such judgment is not transparent (as is typically the case in the determination of RWA, Pillar 2 add-ons, or stress tests). Yet we believe that a component of operational risk capital requirements should be determined by forward-looking judgement, rather than relying entirely on entirely consistent, but also entirely backward-looking formulae. By making the judgemental component of the operational risk capital buffer transparent, our approach will expose differences in regulatory approach. Moreover, the scope for differences in the exercise of judgement could be further constrained by agreement of principles through the BCBS and potentially, by peer review. Even where the characteristics of the risk types preclude robust predictive modelling from past losses, it is possible to take a structured approach to evaluating the risk exposure and drivers (see, for example, JPMorgan, 2016)

7.4 If the real driver here is to enhance "going concern" loss absorbency should we not be doing the same thing for credit risk?

Whilst a reasonable argument can be made about whether the current calibration of capital ratios and credit risk can create a pro-cyclical dynamic with credit losses,

forcing credit contraction as losses are reflected in RWA reduction rather than a reduction in CET1 (Kashyap and Stein, 2004), this is not our focus here. In any case, we think operational risk merits different treatment for three reasons:

- First, because banks' credit RWA plays a central role in the way banks are run (alongside market RWA). It is the basic measure of the risk capital the bank is deploying to generate a return. Operational RWA does not play much role, if any, in the way banks manage themselves.

- Second, whilst far from perfect, historical modelling of credit losses has proved to be a powerful tool in assessing future credit risks and plays a central role in bank's risk management. Historical modelling of operational risk losses is far less effective, particularly for the types of risk that cause the greatest losses and have most scope to generate negative externalities. With operational risk it makes more sense to place greater reliance on regulatory judgement.

- Third, when banks suffer significant credit losses, it is often sensible to for the bank to reduce exposure (and the RWA) to correlated portfolios from both a shareholder and micro-prudential perspective. If the regulator is concerned about the impact on the broader economy they already have the counter-cyclical buffer as a tool they can flex in response.

7.5 Won't the buffer be quickly exhausted, and what happens then?

We do not anticipate that banks will be allowed to draw down on their operational risk capital buffers as a regular occurrence. Most operational risk losses will be absorbed, as now, through depletion of current year earnings. We envisage that regulators would only allow use of the buffer when the alternative is an undesirable reduction in financial intermediation activities, and where the remaining buffer still appears not unreasonable given other prospective risks. Thus it could make sense to allow partial deployment of the buffer in the face of a particularly large one-off loss or to absorb the loss from the final settlement of a regulatory issue relating to a business that has since been closed or where the regulatory deficiencies have been rectified. When an operational incident is indicative of a broader vulnerability or control weaknesses, regulators should be much more cautious about allowing use of the buffer.

Of course one can imagine scenarios in which there is an operational loss of such magnitude that absorbing it purely through the buffer would exhaust the buffer entirely. In these circumstances, we would envisage the regulator deciding how much of the loss could be absorbed through the buffer versus through depletion of retained earnings, reduction in credit or market RWA, or reduction in the overall capital ratio. The regulator's decision on use of the buffer should be made transparent and should reflect a judgement on the minimum acceptable level of the residual operational risk buffer. In these circumstances, the advantage of the buffer approach is that it enables

at least part of the loss to be absorbed through operational risk capital, whereas under the current system or the SMA, none of it would be.

7.6 Will abandoning operational RWA and switching to a buffer approach upset the calibration of other RWA-based metrics such as the counter-cyclical buffer?

The short answer is that other RWA-based buffers, such as the counter-cyclical and GSIB buffers, and other ratios, such as the Total Loss Absorbing Capital ("TLAC') ratio, will need to be looked at to see whether they need recalibration. However, we would turn this argument on its head and suggest that rather than being a problem, this should be seen as another advantage of this alternative approach. Most of these buffers have been designed with market and credit RWA in mind, without much considerations of their impact and relevance on operational risk capital requirements Consider how adjustments to the counter-cyclical buffer apply to the capital supporting operational RWA. Increasing the counter-cyclical buffer to "lean against the wind" in buoyant markets would lead to banks holding more operational risk capital, while easing requirements to respond to economic stress would lead to a reduction. Given that operational risk events appear more likely in times of economic stress (Hess, 2011; Chernobai, Jorion and Yu, 2011; Moosa, 2011), this seems perverse.

The GSIB buffer and TLAC ratio could be recalibrated for neutrality at an industry level, although there would be some wnners and losers amongst individual banks. However, we would suggest that the BCBS avoid automatically recalibrating for neutrality and consider carefully the purpose of these other ratios as they relate to operational risk. If the primary objective of these of the GSIB buffer and the TLAC ratio is to protect the taxpayer against the risks inherent in the financial activites of a bank then it is not clear whether these buffers and ratios should be recalibrated pro-rata.

8. CONCLUSION

Operational risk capital requirements drive a substantial proportion of banks' capital requirements. The BCBS has recognized the current system works poorly and has developed proposals for reform. Yet the SMA proposal does not address the most fundamental problems with the current approach and neither would the suggestions for refinement that have emerged thorugh the consultation process.

We believe serious consideration should be given to a radically different approach, along the lines of the alternative capital buffer approach outlined in this paper. Replacing operational RWA with an operational risk buffer approach would:

- Enable regulators to use operational risk capital far more effectively to incentivize banks to manage and mitigate prospective operational risks with potential negative externalities.

- Create a mechanism to provide "going concern" loss absorbency to minimize the negative externalities from significant operational risk loss events.

- Enhance the clarity and coherence of the CET1 ratio as a measure of capital against risk-adjusted financial assets as represented by credit and market RWA.

Reference

Ames, M., Schuermann, T. and Scott, H.S., 2015. Bank Capital for Operational Risk: A Tale of Fragility and Instability. *Journal of Risk Management in Financial Institutions*, 8(3), pp.227-243.

Ashcraft, A.B., 2005. Are Banks Really Special? New Evidence from the FDIC-Induced Failure of Healthy Banks. *American Economic Review*, *95*(5), pp.1712-1730.

Baker, M. and Wurgler, J., 2013. Do Strict Capital Requirements Raise the Cost of Capital? Banking Regulation and the Low Risk Anomaly. Working Paper.

Bank of America, 2016. Pillar 3 Regulatory Capital Disclosures. Availabel at: http://investor.bankofamerica.com/phoenix.zhtml?c=71595&p=irol-basel#fbid=RZu_npO-xC2.

Barclays, 2016. 2015 Annual Report. Available at: https://www.home.barclays/annual-report-2015.html.

Basel Committee on Banking Supervision, 1998. Operational Risk Management. Availabel at: http://www.bis.org/publ/bcbs42.pdf.

Basel Committee on Banking Supervision, 1999. A New Capital Adequacy Framework. Available at: http://www.bis.org/publ/bcbs50.pdf.

Basel Committee on Banking Supervision, 2003. Sound Practices for the Management and Supervision of Operational Risk. Available at: http://www.bis.org/publ/bcbs183.pdf.

Basel Committee on Banking Supervision, 2006. International Convergence of Capital Measurement and Capital Standards. Available at: http://www.bis.org/publ/bcbs128.pdf.

Basel Committee on Banking Supervision, 2011. Principles for the Sound Management of Operational Risk. Availale at: http://www.bis.org/publ/bcbs195.pdf.

Basel Committee on Banking Supervision, 2014. Operational Risk – Revisions to the Simpler Approaches Issued. Availabel at: http://www.bis.org/publ/bcbs291.pdf.

Basel Committee on Banking Supervision, 2016. Standardised Measurement Approach for Operational Risk. Available at: http://www.bis.org/bcbs/publ/d355.pdf.

Chernobai, A., Jorion, P.,and Yu, F., 2011. The Determinants of Operational Risk in US Financial institutions. *Journal of Financial and Quantitative Analysis*, 46:6, pp.1683-1725

Chodorow-reich, G., 2014. The Employment Effects of Credit Market Disruptions: Firm-Level Evidence From the 2008-09 Financial Crisis. *Quarterly Journal of Economics*, 129 (1): 1–59.

Curti, F. & Migueis, M., 2016. Predicting Operational Loss Exposure Using Past Losses, Working Paper.

Deutsche Bank, 2016. Operational Risk: Sense Check, Research Report.

Hanson, Samuel G., Anil K. Kashyap, and Jeremy C. Stein, 2011. A Macroprudential Approach to Financial Regulation." *Journal of Economic Perspectives* 25 (1):

3-28.

Herring, R., 2005. Implementing Basel II: Is the Game Worth the Candle?, *Financial Markets, Institutions & Instruments,* 14(5), pp.267-287.

Herring, R.J., 2007. The Rocky Road to Implementation of Basel II in the United States. *Atlantic Economic Journal* 35(4):411-429

Hess, C., 2011. The Impact of the Financial Crisis on Operational Risk in the Financial Services Industry: Empirical Evidence. *Journal of Operational Risk*, 6(1), pp.23–35.

Ivashina, V. and Scharfstein, D., 2010. Bank Lending During the Financial Crisis of 2008. *Journal of Financial Economics*, 97(3), pp.319-338.

JPMorgan Chase, 2016. Bridging The Gap Between Operational Risk Measurement And Management. Conference Presentation.

Kashyap, A.K., Rajan, R.G. and Stein, J.C., 2008. Rethinking Capital Regulation. In *Proceedings of the Jackson Hole Economic Policy Symposium,* pp. 431-471. Federal Reserve Bank of Kansas City.

Kashyap, A.K. and Stein, J.C., 2004. Cyclical Implications of the Basel II Capital Standards. *Federal Reserve Bank Of Chicago Economic Perspectives*, 28(1), pp.18-33.

Office of the Comptroller of the Currency, 2007. Risk-Based Capital Standards: Advanced Capital Adequacy Framework -- Basel II. Available at: https://www.regulations.gov/document?D=OCC-2007-0018-0001.

Operational Risk eXchange Association, 2016. ORX Response: The Standardised Measurement Approach. pp.1–28. http://www.orx.org

Moosa, I., 2011, Operational Risk as a function of the state of the economy. *Economic Modelling*, 28, pp.2137-2142

Peters, G.W., Shevchenko, P.V., Hassani, B. and Chapelle, A., 2016. Should the Advanced Measurement Approach be Replaced with the Standardized Measurement Approach for Operational Risk? Working Paper.

Petersen, M.A. and Rajan, R.G., 1994. The Benefits of Lending Relationships: Evidence from Small Business Data. *Journal of Finance*, *49*(1), pp.3-37.

Ponemon Institute, 2015. 2015 Cost of Cyber Crime Study: Global. Ponemon Institute Research Report, p.30. Available at: http://www8.hp.com/us/en/software-solutions/ponemon-cyber-security-report/index.html?jumpid=va_fwvpqe387s.

Power, M., 2005. The Invention of Operational Risk. *Review of International Political Economy*, 12(4), pp.577–599. Available at: http://www2.lse.ac.uk/researchAndExpertise/units/CARR/home.aspx.

Stein, J.C., 2012. Monetary Policy as Financial Stability Regulation. *Quarterly Journal of Economics*, *127*(1), pp.57-95.

Zingales, L., 2011. Trust and Finance. The NBER Digest.

Appendix A: Derivation of Operational Risk RWA under Basel II

Basel II employs three different approaches increasing in sophistication and the complexity of the bank's operations.

- Basic Indicator Approach (BIA) The simplest approach of the three. Bank must hold operational risk capital equal to the average of past three years' of gross income, where positive, multiplied by 15%.

- Standardized Approach (TSA) Bank's activities are divided into eight business lines. Gross income is calculated within each business line, and capital charge is levied on each business line as a fixed fraction of the gross income ranging from 12% to 18% (see tale below). The total capital charge is calculated as the three-year average of the simple summation of the capital charges across each of the business lines in each year.

Business Line	Factor
Corporate finance	18%
Trading and sales	18%
Payment and settlement	18%
Commercial banking	15%
Agency services	15%
Retail Brokerage	12%
Retail banking	12%
Asset Management	12%

- Advanced Measurement Approaches (AMA) The most risk-sensitive of the three approaches. Capital requirement is determined by internally developed risk models of banks. The AMA requires banks to take into acount four variables in their calculation of capital requirements: i) internal data; ii) external data; iii) scenario analysis; and iv) the business environment and internal control factors. The required amount of capital is derived through simulations that estimate the value-at-risk at the 99.9[th] percentile.The BCBS provides general guidance for AMA implementation, and home regulators determine qualification to use the AMA and approve individual banks' methodologies. Banks can adopt AMA for some business lines and TSA for other less-complex business lines at the same time.

Appendix B: Derivation of Operational Risk RWA under the proposed Standardised Measurement Approach (SMA)

Under the proposed SMA in the March 2016 BCBS Consultation Paper, operational risk capital charges is determined by multiplying a Business Indicator (BI) component that assess the scale of the bank's operation with an Internal Loss Multiplier (ILM) that reflects historical operational losses.

- The BI is comprised of three components: financial component, interest, lease and dividend component and services component. The BI aims to capture P&L items that are found in the composition of gross income. It is determined based on the past three-year average of the individual components. The main conceptual difference relative to the BIA under Basel II is that the newly proposed BI emphasizes non-linearity that disproportionally raises the capital charges for banks with larger operations. Banks are categorized into five buckets based on the size of their operation. The BI component increases non-linearly across buckets.

$$BI\ Component = \begin{cases} 0.11 \cdot BI, if\ Bucket\ 1 \\ 110Mln + 0.15(BI - 1Bln), if\ Bucket\ 2 \\ 410Mln + 0.19(BI - 3Bln), if\ Bucket\ 3 \\ 1.74Bln + 0.23(BI - 10Bln), if\ Bucket\ 4 \\ 6.34Bln + 0.29(BI - 30Bln), if\ Bucket\ 5 \end{cases}$$

- The BI is multiplied by the ILM to arrive at the SMA capital charge. The loss multiplier is formulaically determined by average historical losses over the past 10 years calculated as follows:

$$Internal\ Loss\ Multiplier = Ln\left(\exp(1) - 1 + \frac{Loss\ Component}{BI\ Component}\right)$$

Where:

$$\begin{aligned} Loss\ Component = \ &7 * Average\ Total\ Annual\ Loss \\ &+ 7 * Average\ Total\ Annual\ Loss\ only\ including\ loss\ events\ above\ €10\ million \\ &+ 5 * Average\ Total\ Annual\ Loss\ only\ including\ loss\ events\ above\ €100\ million \end{aligned}$$

- The constants in the formula for BI and ILM are determined through calibration in BIS's Quantitative Impact Study conducted in 2015 with the intention of preserving capital neutrality relative to existing charges.