BANK OF ENGLAND

Financial Stability Paper No. 44 – July 2018

Risk sensitivity and risk shifting in banking regulation

Marc Hinterschweiger, Tobias Neumann and Victoria Saporta

BANK OF ENGLAND

Financial Stability Paper No. 44 – July 2018

Risk sensitivity and risk shifting in banking regulation

Marc Hinterschweiger, Tobias Neumann and
Victoria Saporta

The authors are grateful to Marc Farag, Sam McPhilemy, Simon Pittaway and Rhiannon Sowerbutts for helpful conversations and comments.

We are also grateful to David Aikman and Vas Madouros for earlier conversations and contributions on this topic.

marc.hinterschweiger@bankofengland.co.uk
Bank of England, Threadneedle Street, London, EC2R 8AH

tobias.neumann@bankofengland.co.uk
Bank of England, Threadneedle Street, London, EC2R 8AH

victoria.saporta@bankofengland.co.uk
Bank of England, Threadneedle Street, London, EC2R 8AH

ISSN 1754–4262

Contents

Risk sensitivity and risk shifting in banking regulation

Marc Hinterschweiger, Tobias Neumann, and Victoria Saporta

The financial crisis exposed enormous failures of risk management by financial institutions and of the authorities' regulation and supervision of these institutions. Reforms introduced as part of Basel III have tackled some of the most important fault-lines. As the focus now shifts toward the implementation and evaluation of these reforms, it will be essential to assess where the balance has been struck between the robustness and the risk sensitivity of the capital framework. This paper contributes to this assessment by stepping back from the details of the recent reforms and instead taking a bird's eye view on the fundamental trade-offs that may exist between robustness, complexity, and risk sensitivity. We review the history of risk sensitivity in capital standards and assess whether a higher degree of risk sensitivity necessarily leads to a better measurement of risk. We also provide evidence that the more risk-sensitive Basel II framework may have reduced banks' incentives to engage in higher-risk mortgage lending in the UK. Our analysis suggests the need for a robust regulatory framework with several complementary standards interacting and reinforcing each other, even if, prima facie, subjecting banks to a number of regulatory constraints adds to complexity.

1 Introduction

The financial crisis exposed enormous failures of risk management by financial institutions and of the authorities' regulation and supervision of these institutions. Reforms introduced as part of Basel III have tackled some of the most important fault-lines: they raised the level and the quality of capital qualifying for the numerator of the capital ratio, introduced internationally consistent liquidity standards, a leverage constraint, and a floor requirement applied to risk-weighted assets (RWAs).

Having already fundamentally changed the frameworks for market risk, counterparty credit risk and securitisations, the Basel Committee recently finalised the reforms to the denominator of the risk-weighted capital ratio – RWAs. This followed an active debate on how to make the risk-weighting regime more robust without losing the benefits of a risk sensitive capital framework. As the focus now shifts toward the implementation and evaluation of these reforms, it will be essential to assess where the balance has been struck between the robustness and the risk sensitivity of the capital framework.

This paper contributes to this assessment by stepping back from the details of the recent reforms and instead taking a bird's eye view on the fundamental trade-offs that may exist between robustness, complexity, and risk sensitivity. The linchpin of our analysis is Basel II, as it represents the capital framework with the highest degree of risk sensitivity. Indeed, for its architects, Basel II had the core aim of increasing 'risk sensitivity', compared to Basel I. The idea was to align capital requirements more closely with the risks actually assumed by banks in order to improve the measurement of risk and reduce incentives for risk shifting and regulatory arbitrage. As the Basel Committee explains, this should lead to a capital standard that can distinguish 'with reasonable accuracy between sound banks and those that are likely to fail' (Basel Committee on Banking Supervision, 2013a). In this paper, we ask whether the goal has been achieved, and what the implications may be for the design of the capital framework.

The Basel Committee makes a useful distinction between ex-ante risk sensitivity, which mostly concerns the granularity of the approach taken, and ex-post risk sensitivity, which asks whether the framework distinguishes correctly between different risk profiles (Basel Committee on Banking Supervision, 2013a). We present evidence that Basel II may have achieved great ex-ante risk sensitivity. But Basel II assumed that ex-post risk sensitivity would automatically follow ex-ante risk sensitivity. We find that there are strong theoretical and empirical grounds for believing that this is a non sequitur: greater granularity and complexity does not necessarily lead to better risk differentiation.

At the same time, there is no doubt that Basel I had become an inadequate framework by the turn of the millennium because of regulatory arbitrage and concerns about risk shifting. To our knowledge our paper is one of the first to provide evidence that the introduction of the more granular Basel II framework has indeed mitigated risk shifting incentives.

The crisis has also shown that Basel II on its own – that is to say a risk-weighted ratio on its own – would not be sufficient to provide a risk sensitive regulatory framework. Like any regulatory metric, its effectiveness is reduced due to leakages and arbitrage as soon as it becomes binding for banks; and, internal models can deliver highly variable capital requirements that may undermine system-wide risk sensitivity without additional constraints.

The conclusion of this paper is that one should neither demonise nor lionise the quest for greater risk sensitivity. Instead the thrust of our analysis suggests that neither of two ideals exist: there is no single risk-sensitive metric that captures all risks and cannot be arbitraged; and there is no simple metric that is sufficiently robust without incentivising risk shifting. The answer may then lie in using a suite of metrics, continuously examining the effectiveness of the framework, and being able and willing to make running repairs when necessary.

The paper is organised as follows. In Section 2, we provide a sketch of the history of risk-weighted capital requirements, a history that reached its pinnacle with the introduction of the Advanced Internal Ratings-Based approach of Basel II. In Section 3, we review whether Basel II has improved the way risk is measured, in particular why ex-ante risk sensitivity may not lead to greater ex-post risk sensitivity. In Section 4, we assess whether Basel II did lead to less regulatory arbitrage and less risk shifting as was the intention. Section 5 concludes.

2 The quest for greater risk sensitivity in capital standards

Assigning capital requirements that differ by perceived balance sheet risk has its origin in the German Banking Act of 1934 – a direct response to defaults on risky corporate loans contributing to the German banking crisis of 1931. The Banking Act effectively assigned a 0% weight to liquid assets (German sovereign debt) and a 100% weight to all other loans (Neumann & Haldane, 2016).

Though this is a very simple attempt at introducing greater risk sensitivity it was a step towards the goal of differentiating between sound banks and those more likely to fail. A comparison with the UK banking system at the time illustrates this. Whereas British banks were mostly active in the money market, and tended to only provide short-term corporate loans, German banks were tasked with providing long term capital to the real economy. This difference in business model resulted in much less liquid and more risky balance sheets for German banks than their British counterparts.

For example, Chart 1 shows that a capital to total assets ratio (i.e., a leverage ratio) would have suggested that British and German banks were similarly risky. Weighting liquid assets at 0%, on the other hand, reveals that British banks were almost a quarter better capitalised when adjusting for risk than German banks. Indeed, unlike Germany, Britain did not suffer a systemic banking crisis in the 1930s. This difference in the probability of failure would have been (and was) obscured by looking at a simple asset-to-equity ratio (Neumann, 2016).

Chart 1: Capital to total assets and capital to risk-weighted assets in 1930

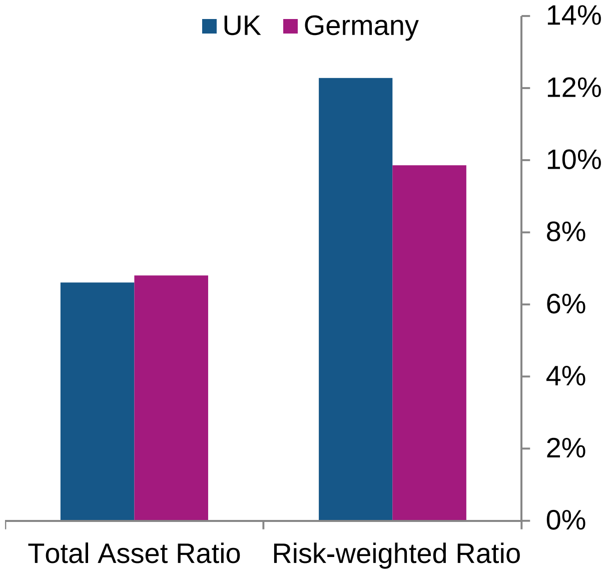

Sources: (Sheppard, 1971) (von Bissing, 1933)

Not that supervisors in the United Kingdom actively relied on risk weighting. Supervision in fact relied on the 'raised eyebrow' of the Governor of the Bank of England rather than statute, with a focus on liquidity rather than capital regulation (Norton, 1995).

In the United States, shortly after the war, the Federal Reserve Board (FRB) and the Office of the Comptroller of the Currency (OCC) devised a ratio of capital to "risk assets", defined as total assets minus cash, balances due from other banks, and government securities – essentially the same as the earlier German standard (Alfriend, 1988). These were only used as examination tools as the US regulators did not have the legal powers to issue capital requirements until 1983 (Neumann & Haldane, 2016).

Capital standards in the US evolved in a fragmented manner across the supervisory agencies. In 1952, the FRB revised its standards by making them more granular: assets were categorised according to risk, with separate capital requirements assigned to each category. In 1962, the OCC abandoned the risk asset standard on the grounds that it was arbitrary and did not consider factors such as management, liquidity, asset quality, or earnings trends. The Federal Deposit Insurance Corporation (FDIC), meanwhile, used variations on the simple total capital-to-assets ratio throughout this period. This fragmented approach continued until the early 1980s, when, in response to a steady downward drift in banks' capital ratios, the FRB and the OCC (later followed by the FDIC) jointly established new guidelines based on simple ratios of capital to total assets, temporarily reversing the push towards increased risk weighting (Alfriend, 1988).

In contrast, European countries developed risk weighted regulations throughout the 1970s. In Germany, for example, more differentiated risk weights of 0%, 20% and 50% were assigned to different exposures in the 1970s (Deutsche Bundesbank, 1973). In the late 1970s, the majority of member states of the European Economic Community used risk-weighted capital regulation (Inter-Bank Research Organisation, 1978).

In the UK, too, risk weighting increased in importance. In 1975, in the wake of the "fringe" banking crisis that threatened financial stability in the United Kingdom, the Bank of England published the conclusions of a Joint Working Party it had formed with the London and Scottish clearing banks to develop proposals for assessing capital and liquidity adequacy. Two metrics were introduced. First, a simple "free resources" ratio based on the ratio of current liabilities to capital resources; like the leverage ratio, this took no account of the riskiness of a bank's assets. Second, a "risk asset" ratio, which under the proposals consulted on in 1979, was based on the ratio of capital resources to the risk weighted value of a bank's assets. Risk weights were assigned to reflect the relative potential for losses arising from "credit", "investment" and "forced sale" risks in each class of asset. Commercial loans were weighted at unity; market loans to listed banks at 20%; and property exposures at 200% (Bank of England, 1980).

By the early 1980s, the United States was an outlier in its continued use of simple capital-to-total asset ratios that did not attempt to discriminate assets by risk. Over this period, US banks' investments in low-risk liquid assets declined and regulators became concerned that the trend could be partly explained by the lack of risk-sensitivity in the capital standard (Bardos, 1987). Partly in response to this, and partly with a view to bringing the US standards more closely in line with other industrial countries, the US authorities began in 1986 to negotiate with the Bank of England on a common system of risk-weighting capital standards.

The outcome of the negotiations between the Bank of England and the US authorities was critical in reaching agreement across G10 countries on the original Basel Accord, published in 1988 (Goodhart, 2011). This established a common risk-weighted capital standard across eleven countries. A bank's credit risk assets were allotted to one of five broad risk categories, each with a fixed risk weighting that ranged from 0-100%. A portfolio of corporate loans, for instance, received a risk weight of 100%, while retail mortgages – perceived to be a safer bet – received a more favourable weighting of 50%. Minimum capital was set in proportion to the weighted sum of these assets (Basel Committee on Banking Supervision, 1988).

In 1996, the Basel Committee amended Basel I to require banks to have sufficient capital to cover their market risk exposures. In an attempt to codify best practice risk management as practiced by 'sophisticated' banks, the Committee permitted the use of banks' own internal models as the basis for these capital requirements. The internal model the Basel Committee required was the 'value-at-risk' (VaR) methodology, which had been introduced two years earlier by JP Morgan (Basel Committee on Banking Supervision, 1996).

In light of the introduction of internal models in 1996, the approach for credit risk was criticised for being insufficiently risk sensitive over the next few years. It is worth quoting the then-President of the Federal Reserve Bank of New York, who summarised the proceedings of a joint conference between his institution, the Bank of England, the Bank of Japan and the Federal Reserve Board.[2]

"Conference participants suggested that in the future, supervisory practice and capital regulation will be based less on specific rules and prescriptions and more on a system of general principles for sound and prudent management. For supervisors, the most important challenge involves developing an approach to capital regulation that works in a world of diversity and near-constant change. […] Whatever the approaches eventually adopted, the next generation of supervisory capital rules must take into account the vital role of incentives in determining the behavior of financial institutions." (McDonough, 1998)

The reason for the concerns raised at the conference was that the simple framework not only failed to measure risk accurately but also created incentives for banks to risk-shift. For example, a flat risk weight on all corporates meant that banks could increase their return by lending to riskier corporates; this would not have attracted higher capital requirements.

An additional concern was that of regulatory arbitrage – banks actively seeking loopholes in the rules rather than adjusting their desired portfolios because of the distortions created by regulation. This included the use of securitisation and innovative forms of credit enhancement (Jackson, 1999).

So by the turn of the millennium, the Basel I framework was thought to be insufficiently risk sensitive: both in an ex-ante and ex-post riskiness of banks by failing to measure risk adequately, providing imprudent incentives and allowing regulatory arbitrage.

[2] See also the papers and speeches presented at the 1998 conference hosted by the Federal Reserve Bank of New York and cosponsored by the Bank of England, the Bank of Japan and the Federal Reserve Board: 'Financial Services at the Crossroad: Capital Regulation in the Twenty-First Century', reproduced in the October 1998 issue of the Federal Reserve Bank of New York Economic Policy Review.

What followed was an attempt to replicate best practice in credit risk modelling (e.g. JP Morgan's CreditMetrics and Merton and Vasicek's KMV). This led to the Basel Committee publishing a revised set of rules for credit risk in 2004. The simple four-bucket approach was replaced with one that sought to tie capital requirements much more closely to risks. Banks were encouraged, subject to regulators' approval, to use the so-called Internal Ratings Based (IRB) approach, under which requirements were based on the outputs of their own internal rating systems. Banks lacking the capacity to model these risks were required to use the so-called Standardised Approach, under which capital requirements were flat, or based on observable benchmarks such external agency ratings.

The IRB approach requires some explanation. Capital requirements are calculated in three steps. First, banks use their internal models to produce an ordinal ranking of borrowers, grouped into a discrete set of rating grades. Second, banks estimate the average probability that borrowers within each grade will default, i.e. be unable to repay their debts over the course of the next year.[3] This is called the probability of default or PD – for some exposures this is floored such that it cannot be below 3 basis points. Third, the so-called 'IRB formula' sets the capital requirement such that stressed (unexpected) losses will not exceed the bank's capital up to a 99.9% confidence level. If the assumptions in the model are correct – in particular, the assumed default correlation and that the probability distribution function of credit losses is normal – the output is a capital requirement sufficiently large that a bank is only expected to become insolvent once every 1000 years.[4,5]

In addition to introducing internal model to credit risk modelling, the Basel Committee designed a modelling framework for operational risk called the Advanced Measurement Approach ('AMA'). The Basel Committee did not specify an approach or distributional assumptions in the AMA, in contrast to the IRB approach, where the Basel Committee specified the risk model in form of the IRB function (Basel Committee on Banking Supervision, 2006).

3 Has Basel II led to an improvement in the measurement of risk?

So on the eve of the financial crisis, the Basel Committee had introduced internal models in the global regulatory framework with the goal to increase risk sensitivity and reduce adverse incentives created by simpler regulation. This section assesses if it has succeeded in the first goal; we look at the second goal in Section 4.

[3] Banks on the so-called "Advanced IRB" approach must also estimate the expected loss given default, exposure at default and maturity.

[4] In practice the output of the formula has been scaled up by the Basel Committee on Banking Supervision so that capital requirements of the G10 banking system as a whole are, in aggregate, the same as under Basel I.

[5] An explanatory note of the Basel II IRB formula can be found at www.bis.org/bcbs/irbriskweight.pdf.

Note that it is difficult to assess the performance of Basel II in the crisis because most internationally active banks' capital ratios were calculated using Basel I risk weights at the time. So in this section we instead review the evidence on whether risk-weighted capital ratios per se have been a more effective predictor of banking distress than simple ratios that do not attempt to weight risk. And, since internal models were used in the market risk framework at the time, the performance of the Basel framework with respect to trading activities can help us gauge the effectiveness of the internal models used at the time.

It might seem naïve to ask if a more flexible approach can ever be less risk sensitive than a simpler one. But the answer is by no means clear cut. A great level of granularity – ex-ante risk sensitivity – does not necessarily lead to a better identification of risk. We first motivate why this may be the case from theory and then proceed to the empirical evidence.

3.1 When is ex-ante risk sensitivity not enough?

Models are a simplification of a complex reality. As a result their predictions will always be subject to some error. To be useful in an international capital standard an internal modelling framework needs to keep this error within a sufficiently narrow range. Ideally, the framework delivers predictions that turn out to be true on average; and, if predictions are wrong they should not be far from the truth. All else equal, this would result in prudent and comparable capital outcomes: the framework would correctly differentiate between different risks and assign similar capital requirements to similar risks – both within and across banks. Achieving this ideal would result in ex post risk sensitivity.

The mean-squared prediction error (MSPE) can be used to explore what causes internal models to perform well or badly. Mathematically, it is the sum of the squared bias and variance of a model.[6] Assume y is the fixed realisation of a data generating process, and $\hat{y}$ is a random variable that reflects the model's output.

$$\begin{aligned}\mathrm{MSPE} = \; & E[(y-\hat{y})^2] \\ & = E[y-E(\hat{y})]^2 + E[E(\hat{y})-\hat{y}]^2 \\ & = \mathrm{bias}^2 + \mathrm{variance}\end{aligned}$$

The first term is the model's squared bias: the extent to which the model is not correct on average. This happens if the model is misspecified. Zero bias is achieved if the predicted risk matches the real risk – on average.

The second term is the model's variance and reflects the nature of estimated values as random variables: depending on the sample used, the same model can make different predictions. The extent to which estimates differ depends on the quantity and quality of the available data.

[6] And an irreducible error if the data-generating process is subject to noise, which we can ignore for this purpose

Figure 1 illustrates the difference between bias and variance by using bull's-eyes (Neumann, 2015). The aim is for a modelling framework that 'hits' the centre often, resulting in a low MSPE. The left picture shows a framework that has a lot of bias but little variance, i.e. its estimates are close together but far from the centre. This can be interpreted as a standardised approach: it may be misspecified for particular portfolios because of its relative lack of granularity, but it will not exhibit variance. The right picture shows a framework with no bias but considerable variance, i.e. on average it is right but the hits are far apart. This might be reflective of some internal models.

The key point is that neither picture is particularly risk sensitive or otherwise desirable. This may be surprising since some would regard a non-biased model as the gold standard of risk assessment. But as Figure 1 illustrates, a lack of bias does not necessarily mean the model is correct most of the time. This insight is particularly crucial from a systemic perspective, where the crosses on the bull's-eyes can be interpreted as individual banks.

Figure 1 – illustration of bias and variance

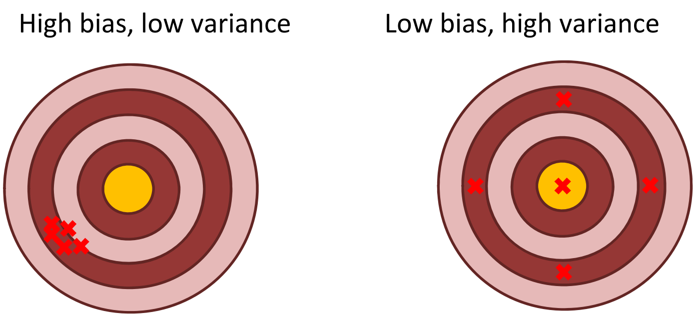

Source: Neumann (2015)

Decomposing the MSPE into bias and variance is a well-known insight in machine learning (Gigerenzer & Brighton, 2009). Its interest comes from the dilemma it presents. Often in order to reduce bias we need to make a model more complex. But the more parameters there are to estimate, which is one way of decreasing bias, the more likely it is that the model mistakes noise for signal in a given data set. This will tend to increase the variance of predictions. As such, optimising a model for bias often has to be traded off against greater variance (and vice versa).

Increasing ex-ante risk sensitivity is often shorthand for reducing bias. In the context of Basel and internal models, it was thought banks' superior data and modelling capabilities should result in outcomes that are closer to the truth than supervisory estimates. This is intuitive: many people would prefer a model with zero bias – i.e. a 'correct' model – to one with zero variance. At least such a model would be correct on average. A slight amount of bias would only be considered because of computational expediency or simplicity.

The problem with this narrow view is twofold: first, we might have too few realisations to even achieve a correct average. And, second, even if we are correct on average, every single prediction could be far off the truth. These concerns read across to banking regulation. A system with low bias but high variance might mean that some banks are undercapitalised while others are overcapitalised – but with few or no banks being adequately capitalised.

On average, the system might have enough capital, but it would contain a number of very weak links making it more fragile. The internal models framework aimed to improve bias, but it might have done so at the expense of variance. As variance increases, banks can have different capital requirements for the same risk, or the same capital requirements for different risks. This would not be a risk sensitive outcome, and may be one explanation of why internal models may fail to increase ex-post risk sensitivity, as discussed below.

Even though it is tempting to tackle a complex world with complex models, this strategy might not always result in the best outcomes if data are limited as they often are for certain asset classes. To illustrate this, consider a financial asset that is subject to market risk; for example an equity share in a company. In this example the share's riskiness is correlated with the last three days of risk: higher risk over the last three days means higher risk today.[7] We can model this risk using a range of complexity for models. The simplest model assumes that risk is only correlated with yesterday. The more complex models make the assumption that risk is correlated with up to the last five days.

We measure performance by estimating the mean squared prediction error and decomposing it into bias and variance. Intuitively, the models that assume fewer than 3 days' of correlation should have a greater bias because they are an incorrect description of the real world. But the performance of the more complex models can be more variable across different samples because they need relatively more data to keep variance small.

Charts 2 and 3 show the results for different amounts of data; the yellow diamond shows the overall error. For the smaller sample in Chart 2 (10 years of simulated data), the simpler but wrong model that only looks at yesterday has the lowest overall error – even lower than the error of the true model. This is because the true model does not have enough data to function reliably. Only for large amounts of data in Chart 3 (40 years) does the true model outperform all others. This shows – perhaps counterintuitively – that without sufficient data even models that reflect reality more accurately can fare worse than simpler approaches.

[7] To be specific, the asset follows a GARCH(3,3) process.

Chart 2: Modelling errors in market risk (10 years of data)

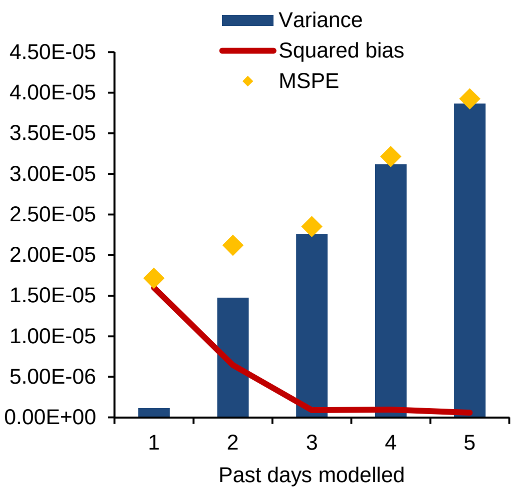

(a) True model is number 3. The more days are modelled, the more complex the model becomes.
(b) Error due to simplicity is squared bias, error due to complexity is variance. Error is defined as mean square prediction error.

Chart 3: Modelling errors in market risk (40 years of data)

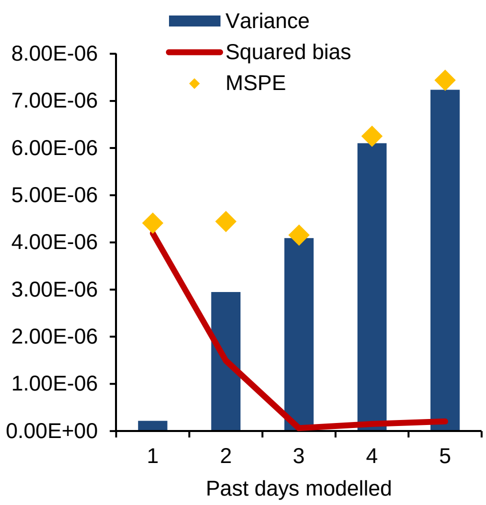

(a) True model is number 3. The more days are modelled, the more complex the model becomes.
(b) Error due to simplicity is squared bias, error due to complexity is variance. Error is defined as mean square prediction error.

The key point that this experiment illustrates is that accurate (and therefore risk-sensitive) modelling requires balancing both bias and variance. The claim that 'the world is complex, so simple regulatory models (such as simple standardised approaches determined by regulators) cannot be more accurate than more complex internal models' is not well-founded in the face of limited data, at least from a statistical point of view. The reality appears to be more nuanced. Simple approaches can outperform more complex modelling in the face of limited data, whereas more complex approaches can outperform simpler ones where data sets are sufficiently large.

Model variance does not only affect individual institutions but can show up system-wide. This leads to differences in modelled outcomes even if all banks were to use the same model, simply because of statistical error.

To see this, assume 100 banks use the same unbiased model but on different identically and independently distributed samples, each of which have a true average PD of 10bp. The task for each bank is to model the PD given the sample available.

Chart 4 shows the distribution of estimated PDs (using a simple average) for different random sample sizes. The dark blue area in each bar shows 95% of banks. As the sample size decreases, banks' estimates of PDs diverge ever more, even though they are using the same model. This variability would translate directly to risk weights and therefore capital. This is not an arcane statistical phenomenon. It is, instead, a basic concept of (frequentist) statistics: the very definition of a confidence interval. A confidence interval measures the distribution of estimators (in this case, average PDs) if they are applied to different identically and independently distributed samples. Generally the smaller the samples, the wider the confidence intervals.

Chart 4: Ranges of estimated PDs depending on sample size[(a)]

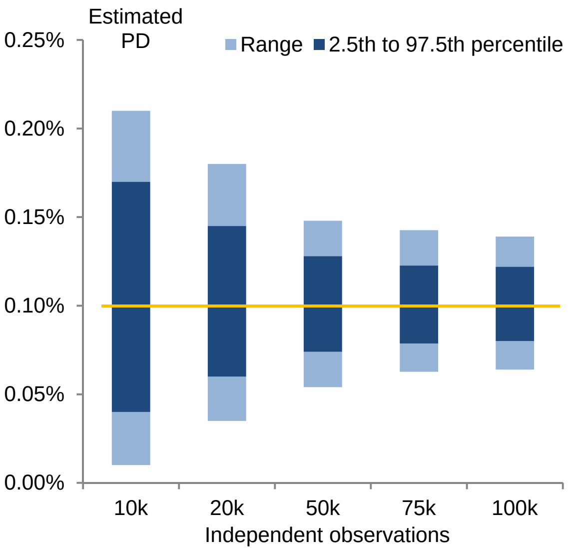

(a) True PD is 0.1%, signified by the yellow line.
Source: Neumann (2015)

Basic statistics, therefore, implies that capital requirements would be variable in a framework allowing internal models even if there were no gaming of models and no difference in the models banks use. This also means that variability is not necessarily an expression of differences in risk assessments (though it might), but that we would expect variability simply because of the random nature of banks' internal data samples.

These are not merely theoretical concerns. The Basel Committee's own analysis suggests that there exists considerable variability across capital estimates. It has conducted a series of hypothetical portfolio exercises, asking banks to estimate the risk weights on a portfolio designed by regulators using their regulatory capital models. This means that all participating banks estimate their capital requirements for identical risks.

The Basel Committee conducted two studies on the trading book, both of which found to vary considerably between banks. For the main diversified portfolio both studies found a ratio of most conservative to least conservative bank of 2.3 (adjusting for differences in supervisory multipliers) (Basel Committee on Banking Supervision, 2013b; Basel Committee on Banking Supervision, 2013c). The second study also found that variability increased with the complexity of the trade (Basel Committee on Banking Supervision, 2013c). Another study on the banking book showed that variability is greater where data are sparse because defaults are rare, for example for sovereigns and banks (Basel Committee on Banking Supervision, 2014).

Figure 2 illustrates how variability in risk weighted assets can impair system-wide risk sensitivity. The left-hand chart in Figure 2 shows the risk assessment of an individual bank, matching risk with a commensurate increase in capital. The right-hand chart shows the same picture from the system-wide perspective in the presence of risk-weighted assets variability. It illustrates that one implication of banks' holding different levels of capital for the same risks is that banks hold the same capital for different risks. So increasing ex-ante risk sensitivity for individual banks does not automatically lead to system-wide risk sensitivity if risk-weighted assets are variable, as they necessarily are when using internal models.

Figure 2: **Consequences of RWA variability on system-wide risk sensitivity**

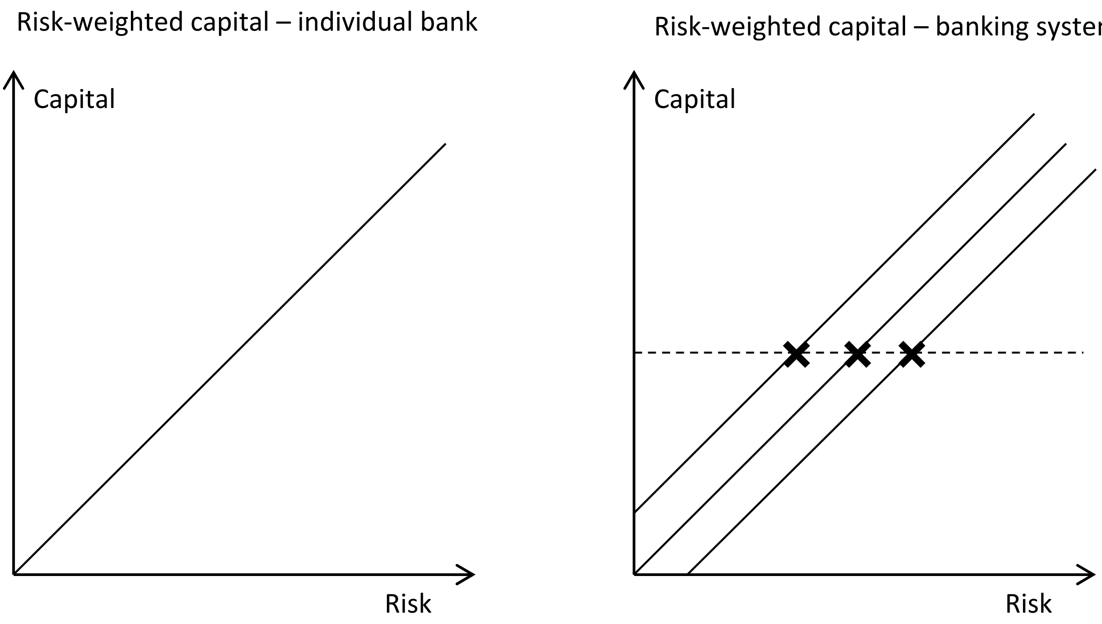

3.2 Empirical evidence

In this section, we document the empirical evidence on the relative accuracy of risk-weighted versus leverage standards in predicting failure during the financial crisis. As mentioned in the introduction to this section, IRB was not yet implemented at the time of the financial crisis; so this is a comparison of a risk-insensitive framework (the leverage ratio) to a mixed framework of a simple standardised approach for credit risk assets and internal models for trading assets (the Basel I risk-weighted standard). If ex-ante risk sensitivity were to necessarily lead to greater risk differentiation, we would expect the risk-weighted ratio to outperform the leverage ratio.

On balance, there is no overwhelming evidence that the risk-weighted ratio considerably outperformed the leverage ratio. This is surprising and confirms the theoretical consideration in the previous sub-section that we cannot assume that more sophisticated approaches ultimately deliver greater risk sensitivity.

Haldane & Madouros (2012) test whether a risk-insensitive leverage ratio outperforms a more sophisticated risk-weighted approach on a sample of about 100 large, global banks. They compare these banks' risk-weighted capital ratios (calculated on a Basel I basis) with their leverage ratios as of end-2006, sorted according to whether they survived or failed during the crisis (Laeven & Valencia, 2010). Chart 4, reproduced from their paper, summarises their results.

Chart 4: Average solvency ratios of major global banks, end-2006[(a)(b)]

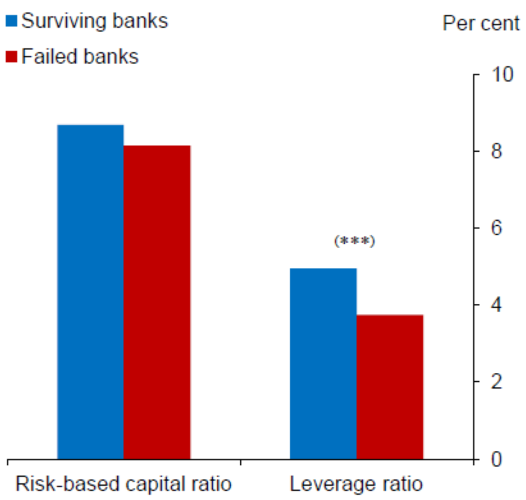

Source: Haldane and Madouros (2012).

(***) Denotes null hypothesis of mean equality rejected at the 1% significance level.

(a) The classification of bank distress is based on Laeven and Valencia (2010), updated to reflect failure or government intervention since August 2009.

(b) For the purposes of the leverage ratio calculation, total assets have been adjusted on a best-efforts basis to achieve comparability between institutions reporting under US GAAP and IFRS.

The striking finding of this exercise is that pre-crisis levels of risk-weighted capital of failed and surviving banks are statistically indistinguishable, while pre-crisis leverage ratios of failing banks were statistically significantly lower than survivors at the 1% level, by on average 1.2 percentage points.

The authors also consider the performance of these metrics in predicting the failures of FDIC-insured banks during the crisis. This covers 8,500 institutions. In contrast to the approximately 100 global banks used in the first part of their study, the majority of the FDIC sample are small, regional banks.

Intriguingly, here, their findings are reversed: risk-weighted capital ratios are statistically significantly lower for failed banks than for survivors, whereas leverage ratios of the two groups are indistinguishable. The authors offer two explanations for this result. One is that, during the sample period, US banks were already subject to a leverage ratio. This may have encouraged them to seek higher-risk assets, which would tend to be better reflected in risk-weighted capital ratios. Another is that simple rules might perform better in an environment of complex risks faced by larger banks.

These findings have been corroborated by other studies of bank failure during this crisis. Demirguc-Kunt, Detragiache, & Merrouche (2013) examine the relationship between banks' capital ratios and stock returns as an indicator of performance. They find that ex ante better capitalised banks experienced a smaller decline in their equity value during the crisis. But this effect was stronger for large banks and when capital ratios were measured on an un-weighted basis rather than for risk-adjusted Basel ratios.

Mariathasan & Merrouche (2012) analyse the explanatory power of the 2005 and 2006 vintages of risk-weighted and un-weighted capital ratios in predicting distress over the crisis using a sample of international banks. Interestingly, the authors find that levels of risk-weighted capital in 2005, but not 2006, were informative about bank failure; whereas un-weighted leverage ratios have predictive power in 2006 but not 2005. They put this down to evidence of risk-weight manipulation – systemic problems were visible by 2006, they argue, but not in 2005. This theme is picked up again in the section on regulatory arbitrage below.

Blundell-Wignall, Atkinson, & Roulet (2014) model the distance-to-default of a sample of 90 international banks over the period between 2005 and 2012. They also examine the relative performance of the risk-weighted ratio and the leverage ratio in explaining bank default risk – and find the latter to perform more strongly than the former.

How robust are these findings to looking at other crisis periods? One prominent earlier study is Avery & Berger (1991), who compared the relative performance of the pre-Basel leverage ratio regime in the US with the then soon-to-be-introduced risk-weighted capital standards in predicting bank failure. Their sample includes all US commercial banks with assets in excess of $10 million between 1982 and 1989. More than 40% of the banks that failed both standards in 1987 were bankrupt by the end of 1989. But, in stark contrast to the results in Haldane & Madouros (2012), failing the risk-weighted standards was a significantly better predictor of future poor performance than failing the pre-Basel I leverage ratio.

Other studies of the pre-2007 period also cast risk-weighted ratios in a more favourable light. Estrella, Park, & Peristiani (2000) study failures of FDIC-insured commercial banks between 1989 and 1993, a period that captures the initial implementation of Basel I risk-weighted capital standards. They found risk-weighted ratios to be more effective predictors of failure than leverage ratios over long time horizons (more than two years), but little difference between the two metrics at shorter horizons.

In a wide-ranging analysis of bank failures between 1984 and 2010, Berger & Bouwman (2013) find that unweighted leverage ratios and risk-adjusted capital ratios have comparable predictive power. Mayes & Stremmel (2012) report similar results in their study of US bank failures between 1992 and 2012. But in an interesting twist, they find leverage ratios to be more informative for FRB-supervised banks, whereas risk-weighted ratios more informative for FDIC and OCC-supervised banks. The rationale they offer is that FRB-supervised banks (bank holdings companies) tend to be more complex – consistent with Haldane and Madouros (2012).

In summary, there is no conclusive evidence that the risk-weighted ratio outperforms the leverage ratio in predicting bank failure. This is a stark result. The very reason the risk-weighted ratio of introduced was because regulators were hoping that its greater ex-ante risk capture can distinguish between banks likely and banks unlikely to fail. Some of the evidence suggests that the risk-weighted ratio fails precisely where it should be needed most: for the most complex banks.

The findings are consistent with the theoretical arguments above that there is no automatic link between ex-ante risk sensitivity and ex-post risk sensitivity. Two studies that try to assess such a direct link between ex-ante and ex-post risk sensitivity are Vallascas & Hagendorff (2013) as well as Barakova & Palvia (2014). The papers agree that banks are more risk sensitive under Basel II than under Basel I, though they disagree on the extent to which the Basel II requirements appropriately capture banks' underlying portfolio risk.

These findings hint at another, perhaps complementary explanation: banks are adept at blunting the effect of any individual regulation – Goodhart's law. A striking feature is that the studies finding the leverage ratio outperforms the risk-weighted ratio used international banks. This could be because they are more complex, as mentioned above; but it is also true that the leverage ratio was not a regulatory capital measure in most countries at that time. In contrast, the studies based on US data find ambiguous evidence, which may reflect that both the risk-weighted and the leverage ratio were in use there. This suggests that the ex-post risk sensitivity of any measure is impaired as soon as it is used as a regulatory target.

4 Has Basel II mitigated banks' incentives for risk shifting and regulatory arbitrage?

In the previous section, we have established that a more granular capital framework does not necessarily lead to greater risk sensitivity. In this section, we examine the potential outcomes of a risk-sensitive capital framework, in particular whether it reduces regulatory arbitrage and risk shifting.

4.1 Risk shifting vs regulatory arbitrage

The conceptual difference between risk shifting and regulatory arbitrage can be illustrated by drawing on the literature on taxation. It is common in this literature to distinguish the concepts of tax avoidance from the incentive effects on behavioural decisions, e.g. labour supply (see Allingham & Sandmo, 1972, the seminal paper introducing the distinction, and Hanlon & Heitzmann, 2010, and Slemrod & Gillitzer, 2014, for recent reviews).

Tax avoidance refers to a situation where people with similar incomes end up paying very different tax rates, mainly because of loopholes.[8] In this case, behaviour does not change in a real sense but is just presented differently to save tax. This is conceptually very distinct to incentives to change underlying behaviour, for example to work less because labour is taxed or to take on debt because it provides corporate tax relief.

- We can think of regulatory arbitrage as the analogue of tax avoidance. We define it here as re-structuring a bank's assets or liabilities, or their legal presentation, to minimise capital or liquidity requirements holding risk (entirely or very nearly) constant.

- Risk shifting is the analogue of the effect of tax on underlying behaviour. We define it here as the impact of the regulatory regime on the riskiness of banks' desired portfolios.

The degree of tax avoidance vis-à-vis behavioural changes can be affected by the complexity of the tax code. Even if, as recently put by Slemrod & Gillitzer (2014) taxes 'magically collected themselves', they would still create behavioural distortions. In fact, this is the implicit assumption of classical optimal taxation analysis in the vein of Diamond & Mirrlees (1971). Tax arbitrage, in contrast, is a result of the imperfect enforceability of the letter (in the case of tax evasion) or the spirit (tax avoidance) of the law.

Seen in this light, it is not surprising that simple linear tax schedules tend to be more robust to the problems of tax avoidance than complex rules (see, for example Hindriks, Keen, & Muthoo, 1999, and Richardson, 2006). The more complex a tax framework, the more loopholes there are to avoid it.

For example, a flat value added tax on all products cannot be avoided. But in a differentiated system of tax curious classification conundrums can arise. In the United Kingdom, for example, the consumer goods firm Procter & Gamble won a court case in 2008 exempting Pringles from VAT because they could not, as argued by the firm, be considered crisps. Though the ruling was later overturned, it shows the extent to which agents will go to change their tax burden without changing the substance (literally, in this case) of their behaviour (Procter & Gamble UK v HM Revenue and Customs , 2009).

While complex rules may create more scope for avoidance, simpler and coarser rules might have

[8] Tax evasion – the criminal act of not paying taxes due – is another distortion created by the tax system.

greater behavioural effects. Consider, for example, a simple flat tax where everybody has to pay the same fixed percentage of income. This flat rate might be difficult to avoid. But because it is likely to be considerably higher than the rate low-earners would pay in a progressive system, low earners might respond by decreasing labour supply (high earners might increase it). Similar logic would suggest that simple constraints such as leverage ratios are likely to be more robust to the problem of regulatory arbitrage than risk-weighted capital requirements. We explore the evidence supporting this proposition before turning to the question of whether the increase in risk sensitivity has reduced incentives for risk shifting.

4.2 Has the quest for risk sensitivity reduced regulatory arbitrage?

The CRD4 / CRR regulation that implements Basel III in the European Union has a word count of over 250,000 – longer than the Old Testament. And this does not include the technical standards issued by the European Banking Authority, or additional rules domestically issued by EU member states. As with complex tax codes this has probably increased banks' ability to arbitrage the regime, compared to the 10,000 words or so of the 1988 Basel Accord.

For example, in a capital management survey of European banks in 2012, Babel et al. (2012) found that more than 65% of banks surveyed had conducted programs to "optimise" RWAs and had seen capital savings of about 5-15%. Respondents also said they saw "large outstanding RWA optimisation opportunities".

This is supported by the JP Morgan 'whale' case study from 2012, which suggests that the complexity of the framework allowed the firm to arbitrage it. The US Senate's hearing into a loss of $6.2bn at JPMorgan Chase that followed a change of the bank's VaR model found that 'a key motivation for developing the new VaR model was to produce lower VaR and Risk Weighted Assets (RWA) results' (p169) and that 'efforts to manipulate RWA results to artificially lower the bank's capital requirements were both discussed and pursued by the bank's quantitative experts' (p196) (Levin & McCain, 2013).

Academic studies point in a similar direction. Behn, Haselmann, & Vig (2016) compare risk weight estimates of two types of banks: one set uses internal models for capital purposes; another set of banks which have developed internal models and submitted a model application but do not yet use them for capital purposes. The banks using internal models for capital purposes systematically underestimate risk weights compared to the control group. But what makes the result stand out is that the interest rates charged on the loans were in-keeping with risk for both groups. So banks that were able to price risk correctly seem to have underestimated it for capital purposes.

In a similar vein, Plosser & Santos (2014) look at probability of default estimates from banks that own the same syndicated loan – i.e. the same credit risk at the same time. As explained above, probability of default is only one of several parameters that advanced banks have to estimate in their credit risk portfolios. But it is likely the one that should vary the least across banks, because though banks can influence some other risk parameters themselves they should have less influence on when a counterparty chooses to default. In contrast, some risk parameters such as loss given default depends on how successful a bank pursues recovery after bankruptcy.

The authors find that some banks systematically report probabilities of default below the median estimate. These banks tend to be the least well capitalised ones and when the obligor represents a larger part of the bank's portfolio. Those are exactly the banks with most to gain from underreporting.

Firestone & Rezende (2016) conduct a similar study and find that banks assign lower probability of default estimates to syndicated loans in which they hold a larger share. Berg & Koziol (2017) find similar evidence and add that banks' PDs estimates on new clients increase significantly after capital increases, corroborating a causal link between low capitalisation and more aggressive modelling. Finally, Begley, Purnanandam, & Zheng (2016) find similar incentive effects for trading book assets.

All in all, though the increased sophistication of Basel II may have put stop to some specific arbitrage concerns of the time, it has not put a stop to regulatory arbitrage. But this should not be surprising. Regulatory experience – and that of tax authorities the world over – suggests it would be naïve to think that any metric could be made so sophisticated that it would be at the same time implementable, enforceable and un-arbitrageable.

4.3 Has the quest for risk sensitivity reduced risk shifting?

In addition to tackling problems of regulatory arbitrage at the time, Basel II attempted to combat risk-shifting by increasing the granularity of risk weights. This meant more buckets in the Standardised Approach and a continuous function mapping risk to capital in the form of the IRB and AMA approaches: if capital increases sufficiently with risk, banks should not be incentivised to shift into risky portfolios.

What empirical evidence is available on the importance of risk shifting? While there exists an extensive literature documenting this phenomenon at the level of financially-distressed firms (Eisdorfer, 2008), to date there has been very little empirical analysis addressing the question of whether increasingly granular risk-weighted approaches generally curbs risk shifting incentives.

One exception is the study by Furlong (1988) who examined how the behaviour of 98 US bank holding companies was affected by the introduction of the leverage ratio in 1981. He measured bank riskiness using the volatility of the return on assets, as implied by the volatility of the return on equity using the Black-Scholes option formula. While Furlong found that the riskiness of US banks increased after the leverage ratio was introduced, he found no difference between the banks constrained by the regulation and unconstrained banks. However, he also found that constrained banks reduced their holdings of low-risk, liquid assets by more than unconstrained banks – consistent with some degree of risk shifting.

Sheldon (1996) employed a similar approach to study the impact of Basel I on bank risk. Using a sample of 219 banks across 11 countries, he found that the volatility of asset returns fell following its introduction (indicating reduced risk taking), but without any discernible difference between the banks constrained by the regulation and those that were not.[9] Overall, the Basel study commissioned to look into the arbitrage and risk shifting incentives of Basel I concluded: 'Owing to the great difficulties in measuring bank risk-taking with available data, the very limited academic literature in this area is inconclusive' (Jackson, 1999).

[9] As discussed in section 2, some countries had risk-weighted regimes prior to the introduction of Basel I whereas others (such as the US) did not. This in turn influences the expected sign of the impact of its introduction on risk-taking incentives.

Laeven & Levine (2009) examine empirically the interactions between banks' ownership structure, regulation, and banks' risk-taking behaviour. They determine that banks with more powerful owners tend to take greater risks. The stringency of capital regulation reduces banks' risk-taking if their ownership is widely dispensed. In contrast, a large owner can increase banks' risk-taking behaviour if capital regulations tighten.

Becker & Ivashina (2015) provide more recent evidence of risk shifting from the insurance sector. The authors find that insurer's corporate bond portfolios appear to be systemically biased towards higher yield, higher risk bonds within each regulatory risk weight bucket. This result is more pronounced for insurance firms for which capital requirements are more binding. The authors also study the portfolios of pension and mutual funds – neither of which are subject to capital regulation – and find no evidence of risk shifting for these firms.

Benetton, Eckley, Garbarino, Kirwin, & Latsi (2017) provide evidence for risk shifting in the mortgage market. With the introduction of Basel II, lenders who adopted IRB models benefitted from a relative advantage in risk-weighted capital requirements for low-risk mortgages. This comparative advantage encouraged large IRB banks to specialise in low-risk mortgages, and helped them win market share against smaller ones. A greater proportion of higher-risk mortgages ended up in smaller banks that did not adopt IRB models.

We complement these studies with an exercise of our own (Box 1) exploiting regulatory changes in capital requirements in the wake of the transition from Basel I to Basel II. Our analysis is one of the first we are aware of to show that moving from Basel I to Basel II may indeed have mitigated the previous framework's risk shifting incentives. This suggests that overly simple approaches on their own may indeed lead to risk shifting.

Box 1

How did Basel II affect the riskiness of UK banks' mortgage portfolios?

We use a novel approach and dataset to overcome two shortcomings that empirical studies of banks' risk shifting in response to changes in capital requirements have often been confronted with. First, a lack of granular data has meant that risk shifting has often been identified by changes between relatively broad asset classes on banks' balance sheets, such as government securities and mortgage loans. But this potentially misses the possibility that banks adjust the composition of risky and relatively safe sub-categories of assets within each class. By merging three different sources, our new dataset (described in more detail below) allows us to examine banks' risk shifting behaviour within the UK mortgage market.

Second, the relationship between banks' individual capital requirements and the riskiness of their assets goes both ways: capital requirements depend on the riskiness of the respective assets, but at the same time, capital requirements also influence which assets banks choose to hold on their balance sheets. This endogenous relationship necessitates a carefully designed econometric strategy in order to isolate the effect of capital requirements on banks' risk-taking. We examine banks' risk-shifting behaviour in the UK mortgage market by exploiting regulatory changes in capital requirements in the wake of the transition from Basel I to Basel II. We view this transition as a quasi-natural experiment that allows us to observe shifts in banks' behaviour that are the consequence of an event outside of their control. We hypothesise that, under the Basel I framework with its flat risk weights, banks had an incentive to grant relatively more higher-risk mortgages than under a more risk-sensitive regime. This is because risky mortgages are potentially more profitable, while banks were not required to hold more capital against them to reflect their higher riskiness. Therefore, we expect 'reverse risk shifting' to take place when the Basel II framework was introduced: risk weights on mortgages became a function of risk as proxied by the corresponding loan-to-value (LTV) ratios. This was the case especially for banks on internal models (banks under the IRB approach described in Section 2), and less so for banks on the standardised approach. Hence, relative to banks on the standardised approach, the incentives for IRB banks to grant higher-risk mortgages were reduced, so we expect the frequency of lower-risk mortgages by IRB banks to increase at the expense of higher-risk mortgages.

This will be true as long as IRB banks' capital requirements are more risk-sensitive than the ones of standardised approach banks: Up to an LTV of 80%, risk weights on mortgages for banks following the standardised approach are flat, and above an LTV of 80%, they are less risk sensitive for the majority of such banks.

Data

We create a unique dataset by combining several sources:

- The Financial Conduct Authority's (FCA's) confidential mortgage Product Sales Database (PSD) of individual mortgages from April 2005 to December 2013. PSD captures details of loans for house purchases and re-mortgages, such as the property value, loan amount, and location of the property. It excludes products such as second-charge lending, commercial, and buy-to-let mortgages.

- A confidential bank panel dataset that merges regulatory returns and bank balance sheet data.[10] It includes information on banks' assets, liabilities, and capital adequacy at the UK-consolidated level.

- Information on whether banks follow the standardised or IRB approach. A bank is assumed to follow the IRB approach once it has received permission by the regulator to do so.

As described below, this new dataset allows us to examine the effect of the introduction of Basel II on banks' risk shifting behaviour

Methodology

We consolidate the dataset at the group level to be able to properly account for capital requirements set at that level. This leaves us with a relatively small number of banking groups - too few to make any valid statistical inference using them as the units of observation. Hence, we exploit variation in banks' lending practices across postcodes. Postcodes have different shares of mortgages extended by IRB banks or standardised approach banks both in the cross-section and over time. This fact allows us to test whether postcodes in which there are on average relatively more loans granted by banks on IRB will have a higher share of lower-risk mortgages, all else equal. To the extent that banks' lending decisions are not perfectly correlated across postcodes, we can utilise this source of variation. This approach has the additional benefit of enabling us to control for time-invariant differences in loan demand between postcodes by using fixed effects, recognising that loan demand is closely tied to the local economic conditions in a certain geographic area.

We construct quarterly averages of all independent variables at the 3-digit postcode level. We create several dependent variables. They represent the number of loans falling into different 5-percentage-points wide LTV buckets for a given postcode and quarter, expressed as the share of all loans in that postcode and quarter (e.g. the share of all loans with an LTV ratio between 80% and less than 85% in a given postcode and quarter).

We use a fixed-effects regression model with standard errors clustered at the postcode level. We include a full set of quarterly time dummies to account for unobserved factors affecting all

[10] We are grateful to Jon Bridges, Courtney Escudier and Amar Radia for their help in compiling the bank panel data set used in this paper.

postcodes uniformly. Since some postcode-quarters only have a small number of loans in a specific LTV bucket, this could bias the results. We therefore only use postcode-quarters for which we have at least a certain number of observations in a specific LTV bucket (more than 4 in the baseline scenario).

For each separate LTV band, the regression equation is therefore

$$LTV_{band\,i,t} = \alpha_i + \beta . IRB_{i,t} + \gamma . controls_{i,t} + \eta . time_i + \varepsilon_{i,t}$$

with 'IRB' the share of LTV mortgages in a specific band granted by IRB banks, 'controls' a set of bank control variables, 'time' a set of quarterly time dummies, and 'ε' the residuals.

Results

The table below shows our baseline regressions for six different LTV buckets (e.g. Reg_LTV80 for the 80-85% bucket). IRB banks have an incentive to reduce their exposure to high-LTV lending and to increase their low-LTV lending due to the more favourable capital treatment of the latter under the IRB approach. Also, the coefficients tend to decrease as the LTV buckets increase – this is consistent with the fact that the difference in capital requirements between the standardised and the IRB approach tends to shrink as LTV ratios become higher.

Table A Regression results by loan-to-value band

Variable	Reg_LTV70	Reg_LTV75	Reg_LTV80	Reg_LTV85	Reg_LTV90	Reg_LTV95
IRB	3.5055***	3.5103***	1.9910***	-2.8066***	-3.8450***	-3.7064***
Loan-to-income ratio	0.4572***	0.4392***	0.3907***	0.3513***	-0.2784***	0.5794***
Interest rate	-0.1877*	0.8481***	2.2611***	3.8942***	3.3833***	0.8067**
Capital trigger	-6.9820***	1.5841*	-0.8777	-4.0889***	-8.2978***	4.2042**
Assets	-0.0067***	-0.0105***	-0.0066***	-0.0085***	-0.0042***	0.0025**
Investments	0.0034	0.0486***	-0.0571***	-0.0245**	0.0229	-0.1660***
Deposits	0.0217***	0.0283***	0.0233***	0.0221***	0.0201***	0.0003
Time dummies	yes	yes	yes	yes	yes	yes
Constant	14.7452***	-0.8480	-2.9295*	6.0035***	8.7867***	-1.3815
Number of observations	53855	54479	47168	49955	42014	18374
Number of groups	2.4e+03	2.4e+03	2.3e+03	2.3e+03	2.3e+03	2.1e+03
Average group size	22.8878	23.1137	20.5078	21.4491	18.4191	8.9021
R-squared (within)	0.3290	0.4929	0.3985	0.2144	0.1534	0.1090
R-squared (between)	0.0001	0.0861	0.0562	0.0621	0.1466	0.0049
R-squared (overall)	0.2633	0.4055	0.3369	0.2233	0.1774	0.0441

legend: * p<.1; ** p<.05; *** p<.01

The IRB coefficients in the different LTV regressions have the expected sign: For LTV buckets below and including 80-85%, the coefficient is positive, while it is negative for LTV buckets above 80-85%. In other words, postcodes with a higher share of IRB loans have a higher share of safer mortgages; and they also have a lower share of riskier mortgages. This evidence is in line with 'reverse risk shifting' behaviour. As for the control variables, the loan-to-income (LTI) ratio serves as a proxy for risk that is based on a borrower characteristic. Higher LTI ratios tend to be associated positively with the dependent variables (the 90-95% LTV regression being the exception). The coefficients of the mortgage interest rate tend to increase along with the LTV ratios. This shows that, as expected, riskier lending is correlated with a higher average interest rate, with the "risk-neutral" interest rate located between an LTV of 70 and 75%. Capital constraints ("capital trigger") – a measure of how close a bank comes to breaching its minimum capital requirements – don't show a clear pattern. Bank size ("assets") is significantly negative for all but the highest bucket, indicating that only the most risky mortgages are positively associated with larger banks. While investments are not clearly linked to the dependent variables in a specific way, deposits are usually associated with higher lending in the respective buckets.

Discussion

Our results are in line with the hypothesis of 'reverse risk shifting' in the wake of Basel II. This is one of the first direct pieces of evidence we are aware of that the introduction of more risk-sensitive IRB approaches in Basel II may have reduced risk shifting incentives. The new dataset has allowed us to look in more detail into banks' risk shifting behaviour within one broad asset class. To our knowledge, both our approach and dataset are novel in this literature.

One caveat with respect to our analysis is that, despite the included control variables, the regressions might not account for all important factors, for example some of those related to the financial crisis. If there are missing variables that influence banks on the IRB approach vs the standardised approach in a systematically different way, this could be picked up by the IRB variable even if differences in capital regulation do not play any role. But the risk of this being the case is mitigated by including the control variables, time dummies, and postcode fixed effects in our regressions.

5 Conclusion

The analysis in this paper shows that the twin forces of changing economic activity and regulatory arbitrage chip away, and always have chipped away, on the effectiveness of even a well-designed rule. The German banking crisis of the 1930s showed a simple leverage ratio was not enough to capture the risk German banks were taking, giving rise to the risk-weighted capital ratio. The simple Basel Accord of 1988 was not enough to capture increasing sophistication in banks' trading activities, giving rise to the first internal models approach. And while the Basel II reforms seem indeed to have achieved the result of decreasing risk shifting, the financial crisis and its aftermath showed that ever-complex models alone cannot ensure banks have enough capital; neither could the models banish regulatory arbitrage.

The rapid financial innovation we are currently experiencing is only likely to strengthen the forces that demand greater complexity. In accommodating this change regulators have learned, and must remember, that matching this complexity by making rules more complex may not result in a safer system. This may be because the optimal response to a complex problem is not always a complex solution; or because added complexity increases the scope for arbitrage. So, overall, our analysis suggests the need for a robust regulatory framework with several complementary standards interacting and reinforcing each other, even if, prima facie, subjecting banks to a number of regulatory constraints adds to complexity (Rule, 2015; Aikman et al., 2018).

At the same time, history suggests that regulators cannot assume that their regime is infallible. This speaks to the need for much greater emphasis on active monitoring of the behavioural responses to regulation and regular 'running repairs' of the regime to deal with unintended consequences (Woods, 2016).

References

Aikman, D., Haldane, A., Hinterschweiger, M., & Kapadia, S. (2018). Rethinking financial stability. *Bank of England Staff Working Paper* No. 712.

Alfriend, M. C. (1988). International risk-based capital standard: history and explanation. *Economic Review*. Federal Reserve Bank of Richmond.

Allingham, M., & Sandmo, A. (1972). Income tax evasion: A theoretical analysis. *Journal of Public Economics*, 323-338.

Avery, R., & Berger, A. (1991). Risk-based capital and deposit insurance reform. *Journal of Banking & Finance,* 847-874.

Babel B., Gius D., Gräwert A., Lüders E., Natale A., Nilsson B. & Schneider S. (2012). Capital Management: Banking's New Imperative. *McKinsey Working Papers on Risk*, Number 38.

Bank of England (1980). The measurement of capital. Bank of England Quarterly Bulletin: September 1980, 324-330.

Barakova, I., & Palvia, A. (2014). Do banks' internal Basel risk estimates reflect risk? *Journal of Financial Stability, 13*, 167-179.

Bardos, J. (1987). The risk-based capital agreement: a further step towards policy convergence. *FRBNY Quarterly Review*.

Basel Committee on Banking Supervision (1988). International Convergence of Capital Measurement and Capital Standards. Basel.

Basel Committee on Banking Supervision (1996). Minimum capital requirements for market risk. Basel.

Basel Committee on Banking Supervision (2006). Basel II: International Convergence of Capital Measurement and Capital Standards: A Revised Framework - Comprehensive Version. Basel.

Basel Committee on Banking Supervision (2013a). The regulatory framework: balancing risk sensitivity, simplicity and comparability. Basel.

Basel Committee on Banking Supervision (2013b). Regulatory Consistency Assessment Programme (RCAP) – Analysis of risk-weighted assets for market risk. Basel.

Basel Committee on Banking Supervision (2013c). Regulatory Consistency Assessment Programme (RCAP) - Second report on risk-weighted assets for market risk in the trading book. Basel.

Basel Committee on Banking Supervision (2014). Regulatory consistency assessment programme (RCAP) – Analysis of risk-weighted assets for the banking book. Basel.

Becker, B., & Ivashina, V. (2015). Reaching for yield in the bond market. *The Journal of Finance, 70(5),* 1863-1902.

Begley, T. A., Purnanandam, A. K., & Zheng, K. K. (2016). The strategic under-reporting of bank risk. *Ross School of Business Working Paper* No. 1260.

Behn, M., Haselmann, R., & Vig, V. (2016). The limits of model-based regulation. *ECB Working Paper* No. 1928.

Benetton, M., Eckley, P., Garbarino, N., Kirwin, L., & Latsi, G. (2017). Specialisation in mortgage risk under Basel II. *Bank of England Staff Working Paper* No. 639.

Berg, T., & Koziol, P. (2017). An analysis of the consistency of banks' internal ratings. *Journal of Banking & Finance*, 27-41.

Berger, A. N., & Bouwman, C. H. (2013). How does capital affect bank performance during financial crises? *Journal of Financial Economics, 109(1),* 146-176.

Blundell-Wignall, A., Atkinson, P., & Roulet, C. (2014). Bank business models and the Basel system: Complexity and interconnectedness. OECD Journal: Financial Market Trends 2014.

Demirguc-Kunt, A., Detragiache, E., & Merrouche, O. (2013, September). Bank Capital: Lessons from the Financial Crisis. *Journal of Money, Credit and Banking, 45(6),* 1147-1164.

Deutsche Bundesbank. (1973). Geschäftsbericht der Deutschen Bundesbank für das Jahr 1972. Deutsche Bundesbank.

Diamond, P. A., & Mirrlees, J. A. (1971). Optimal taxation and public production I: Production efficiency. *The American Economic Review, 61(1),* 8-27.

Eisdorfer, A. (2008). Empirical evidence of risk shifting in financially distressed firms. *The Journal of Finance*, 609 - 637.

Estrella, A., Park, S., & Peristiani, S. (2000). Capital Ratios as Predictors of Bank Failure. *Federal Reserve Bank of New York Economic Policy Review.*

Firestone, S., & Rezende, M. (2016). Are banks' internal risk parameters consistent? Evidence from syndicated loans. *Journal of Financial Services Research*, 211 - 242.

Furlong, F. (1988). Changes in Bank Risk-Taking. *Federal Reserve Bank of San Francisco Economic Review*.

Gigerenzer, G., & Brighton, H. (2009). Homo Heuristicus: Why Biased Minds Make Better Inferences. *Topics in Cognitive Science,* 107 - 143.

Goodhart, C. (2011). The Basel Committee on Banking Supervision - a history of the early years 1974 - 1977. Cambridge: Cambridge University Press.

Haldane, A., & Madouros, V. (2012). The Dog and the Frisbee. London: Bank of England.

Hanlon, M., & Heitzmann, S. (2010). A review of tax research. *Journal of Accounting and Economics*, 127-178.

Hindriks, J., Keen, M., & Muthoo, A. (1999). Corruption, extortion and evasion*. Journal of Public Economics*, 395 - 430.

Inter-Bank Research Oranisation. (1978). The Regulation of Banks in the Member States of the ECC. London: Sijthoff & Noordhoff.

Jackson, P. (1999). Capital requirements and bank behaviour: the impact of the Basel Accord. BCBS Working Papers.

Laeven, L., & Levine, R. (2009). Bank governance, regulation and risk taking. *Journal of Financial Economics*, *93(2),* 259-275.

Laeven, L., & Valencia, F. (2010). Resolution of banking crises: the good, the bad, and the ugly. International Monetary Fund Working Paper.

Levin, C., & McCain, J. (2013). JPMorgan Chase Whale Trades: A case history of derivatives risk and abuse. Washington: Permanent Subcommittee on Investigations, United States Senate.

Mariathasan, M., & Merrouche, O. (2012,September). The manipulation of Basel risk-weights. Evidence from 2007 - 2010. *University of Oxford, Department of Economics Discussion Paper* Number 621.

Mayes, D. G., & Stremmel, H. (2012). The effectiveness of capital adequacy measures in predicting bank distress.

McDonough, W. J. (1998). Conference Overview: Major Themes and Directions for the Future. Financial Services at the Crossroads: Capital Regulation in the Twenty-First Century, (pp. 3-6). New York.

Neumann, T. (2015, August 12). It's a model - but is it looking good? When banks' internal models may be more style than substance. Retrieved March 2016, from http://bankunderground.co.uk/2015/08/12/its-a-model-but-is-it-looking-good-when-banks-internal-models-may-be-more-style-than-substance/

Neumann, T. (2016, February). From Berlin to Basel: what can 1930s Germany teach us about banking regulation? Bank Underground Blog. Retrieved from Bank Underground Blog.

Neumann, T., & Haldane, A. G. (2016). Complexity in Regulation. In Palgrave Handbook of European Banking (pp. 323-347). Palgrave Macmillan UK.

Norton, J. J. (1995). Devising International Bank Supervisory Standards. Kluwer Academic Publishers.

Plosser, M., & Santos, J. (2014). Banks' Incentives and the Quality of Internal Risk Models. *Federal Reserve Bank of New York Staff Reports*.

Procter & Gamble UK v HM Revenue and Customs , 407 (EWCA Civ 2009).

Richardson, G. (2006). Determinants of tax evasion: A cross-country investigation. *Journal of International Accounting, Auditing and Taxation*, 1069–1075.

Rule, D. (2015). What is left to do on the post-crisis bank capital framework? Bank of England.

Sheldon, G. (1996). Capital Adequacy Rules and the Risk-Seeking Behavior of Banks: A Firm-Level Analysis. *Swiss Journal of Economics and Statistics*, 709 - 734.

Sheppard, D. K. (1971). The growth and role of U.K. financial institutions, 1880-1962. Methuen.

Slemrod, J., & Gillitzer, C. (2014). Multiple Tax-System Instruments. In Tax Systems (pp. 93–112). MIT Press.

Vallascas, F., & Hagendorff, J. (2013). The risk sensitivity of capital requirements: Evidence from an international sample of large banks. *Review of Finance, 17(6),* 1947-1988.

von Bissing, D. W. (1933). Die Schrumpfung des Kapitals und seine Surrogate. *Untersuchung des Bankwesens, 1 (1)* (pp. 57 - 112). Berlin: Heymann.

Woods, S. (2016). The revolution is over. Long live the revolution! Bank of England.