Analysis of new Regulatory Proposals for Capitalisation of Securitisations*

Shalom Benaim†, Dherminder Kainth and Manuel Losi, QuaRC, RBS

August 7, 2013

Abstract

We explain and analyse the modelling and other assumptions behind the new regulatory proposals for capitalisation of securitisations, in particular the modified supervisory formula (MSFA) and the revised rating based approach (RRBA). We examine the impact of the different assumptions, and find that the underlying modelling approach is generally reasonable. However, we find that the additional overlays applied to the calculation, rather than the modelling assumptions per se, lead to excessively high estimates of the capital in many cases (in the sense that the requirement is significantly higher than any reasonable model would predict) and do not account for model risk in a risk sensitive way.

1 Introduction

The Basel Committee have (in [1]) proposed new rules for the calculation of capital requirements for securitisation exposures. In this paper, we discuss these rules and the assumptions behind them.

Ultimately, the source of the risk on an unsecuritised pool is the same as on a securitisation of a pool with the same characteristics, so the rules for securitisations should be extensions of the existing rules for the unsecuritised pool. The cashflows of the pool are the same irrespective of whether the assets are securitised or not, so the purpose of a securitisation model (and a regulatory framework) should be to allocate the cashflows, and the risk associated with those cashflows, generated by the pool to the different tranches of the securitisation.

Any regulatory framework for securitisations should therefore be consistent with what the regulatory rules would be for those same assets if they were held directly by the bank. As well as ensuring that the rules are self consistent and do not unduly[1] penalise or favour securitisations, this is necessary to avoid arbitrages between the two possibilities. This point is also made by the BCBS[2] and they state that the proposed revised framework attempts to be consistent with the Basel 2 IRB.

We therefore start by reviewing the capital requirements for portfolios of loans (and other collateral than can underlie securitisations) under Basel 2 (IRB), in section 2. Section 3 then discusses issues with the existing Basel 2 framework. We then examine in detail two components of the proposed new regulatory framework, the Modified Supervisory Formula Approach (section 4) and the Revised Rating Based Approach (section 5).

---

*The contents of this paper does not necessarily represent the views of RBS.

†E-mail address: shalom.benaim@rbs.com

[1]By unduly we mean for reasons that are not related to the relative risks of holding securitised or unsecuritised assets.

[2]See for example page 35 of [1]. The lack of consistency is cited as a shortcoming of the previous approach, see e.g. the first bullet point of page 6 of [1].

2 Basel 2

The Pillar 1 capital requirement[3] for a loan (often referred to as K_{IRB}) is taken to be the expected (negative) change in value of a loan conditional on a stress event over the capital horizon[4]. The following assumptions are made:

- The capital horizon is 1 year.
- Each individual exposure is assumed to be an infinitesimally small part of the banks portfolio (assumed to be well diversified), which ensures that capital is additive across exposures.
- The modelling is based on a one factor Gaussian model for firm value[5] (with the correlations between the different credits the bank is exposed to specified by regulators).
- The stress event is defined as the 99.9% quantile of the loss distribution of the bank's portfolio under this Gaussian model.
- It is assumed that expected margin income (both unconditional and conditional on a stress) is equal to the unconditional expected loss on the loan (i.e. that the loan originator would charge enough margin to cover expected losses and that, on average, the coupons on the loan are likely to be paid in a stress scenario). This income is assumed to offset losses (at least at the portfolio level).

In mathematical terms, asset value[6] for credit i is assumed to follow:

$$dA^i = \sqrt{\rho^i} dW^c + \sqrt{1-\rho^i} dW^i + \mu^i dt$$

where W^c is a Brownian motion representing systemic risk that affects all obligors and W^i is a Brownian motion representing risks idiosyncratic to this particular obligor[7]. ρ^i is the correlation and μ^i is the drift. Default occurs if this asset value goes below a particular barrier. Given the assumptions, the stress event can be expressed as the 99.9% quantile of the distribution of W^c.

For a 1 year exposure, it is straightforward to compute the capital requirement given the assumptions. The change in value for an exposure with a 1 year maturity in a particular stress scenario can be written as[8]

$$\mathbb{E}^{\mathbb{P}}[L|\text{Stress}] = V_0 - (1 - \mathbb{P}[\tau_d < 1|\text{Stress}]\text{LGD} + \mathbb{E}^{\mathbb{P}}[C|\text{Stress}])$$

where V_t is the value of the instrument at time t, τ_d is the default time, C is the margin income and $\mathbb{P}$ is the real world probability measure The current loan value value is assumed to be Par, and the given the assumptions about margin income we can write this as

$$\mathbb{E}^{\mathbb{P}}[L|\text{Stress}] = \mathbb{P}[\tau_d < 1|\text{Stress}]\text{LGD} - \mathbb{P}[\tau_d < 1]\text{LGD}$$

This is referred to as unexpected loss, i.e. total (principal) loss in a stress scenario minus loss that was expected.

---

[3]In this paper we will use "capital requirement" and 8% of risk weighted assets (RWA) interchangeably as that is the assumption that is used in the modelling we discuss. However, banks may hold a larger fraction of RWA than this in practice.

[4]Note that to the extent that any of these assumptions do not apply to an individual institution, the bank is required to hold additional pillar 2 capital.

[5]The model is often referred to as Asymptotic Single Risk Factor (ASRF).

[6]Equivalently, the log on the asset value can be assumed to have this distribution, or the right hand side can be multiplied by a scalar representing the volatility, neither change affects the outcome of the calculation.

[7]Strictly speaking, the risk modelled by W^i does not have to be idiosyncratic to a single obligor, but the proportion of obligors affected by individual risks modelled by W^i is assumed to be infinitessimal. This point is important for securitisations.

[8]We assume risk free rates are 0 for simplicity.

2.1 Maturity modelling

When the maturity of the asset exceeds the capital horizon, the modelling is more involved as we also need to calculate what the value of the loan will be in a stress scenario at the capital horizon when the obligor has not defaulted (but may have been downgraded).

For wholesale loans with a maturity greater than the capital horizon the change in value of the loan given a stress was computed by regulators (using a number of modelling approaches) and parameterised via a maturity adjustment (as a function of time and probability of default). Maturity was capped at 5 years, which presumably was based on an assumption that shocks in year 1 do not materially increase the probability of default after 5 years. For retail exposures (including retail mortgages), regulators incorporated maturity into the estimation of the correlation (which is then used as an input into the formula used for 1 year exposures) instead of an explicit maturity adjustment.

The representation of capital as (lifetime) unexpected loss still applies as the value of the loan at the capital horizon is a function of the probability of default at the time (if the value is interpreted as market value the probability needs to be interpreted as a risk neutral probability after the capital horizon). Mathematically,

$$\mathbb{E}^{\mathbb{P}}[L|\text{Stress}] = V_0 - \mathbb{E}[V_1|\text{Stress}].$$

Now, value after 1 year conditional on survival is:

$$1 - \mathbb{Q}[1 < \tau_d < M|\text{Stress}]\text{LGD} - \mathbb{P}[\tau_d < 1|\text{Stress}]\text{LGD} + \mathbb{E}^{\mathbb{Q}}[C|\text{Stress}]$$

where M is maturity and $\mathbb{Q}$ is the probability measure used to compute value (the risk neutral measure if "value" is interpreted as market value), and C refers to margin income after year 1.

If we make the same assumption about margin income being on average equal to expected loss and initial value being equal to Par, we get:

$$\mathbb{E}^{\mathbb{P}}[L|\text{Stress}] = \mathbb{Q}'[\tau_d < M|\text{Stress}]\text{LGD} - \mathbb{Q}'[\tau_d < M]\text{LGD}$$

where $\mathbb{Q}'$ is the probability measure where asset prices evolve as in $\mathbb{P}$ before time 1 and as in $\mathbb{Q}$ afterwards.

The question of whether the maturity adjustment overstates or understates maturity effects (and more broadly whether the regulatory approach overstates or understates capital on unsecuritised assets) is beyond the scope of this paper. As discussed in the introduction, the aim of the securitisation framework should be to measure the relative risk of securitised and unsecuritised exposures to the same underlying risk, and so we take the rules for unsecuritised exposures as given.

3 Existing rules for securitisations

As discussed by the BCBS, the rules for securitisation in Basel 2 were not fully consistent with the framework discussed in the previous section. There were two components to the approach, the "supervisory formula approach" (SFA) which modelled the risk on the securitisation as a function of the risk of the underlying assets (in particular the capital that would be held under IRB against those assets), and a rating based approach (RBA), that calculated capital as a function of credit rating.

There were a number of flaws with this approach

- In the SFA (supervisory formula approach), the possibility of further deteriorations in credit after the capital horizon was ignored, which lead to the risk being understated for tranches with a convex exposure to loss (such as senior tranches), and overstated for certain mezzanine tranches. In particular, it was assumed that, in a stress scenario, lifetime unexpected losses on the pool were deterministically equal[9] to K_{IRB}[10]. Thus, if you hold a thin mezzanine tranche just below K_{IRB}, it was assumed that the tranche would be wiped out, with 100% probability, in a stress, so the capital requirement would be 100% , whereas a slightly more senior thin tranche attaching above K_{IRB} was considered perfectly safe, with (almost) no capital requirement[11].

This led to cliff effects as a slight change in assumptions about the underlying loans could lead to a large change in the capital for a tranche (as it moved from being perceived as safe to risky)[12].

- In the ratings based approach, maturity was also ignored, and the capital per rating implicitly assumed that the underlying assets were lower risk (less correlated to systemic risk) than was assumed in the regulation of on balance sheet loans via IRB[13].

4 MSFA

To remedy the issues identified with the existing rules, the BCBS has proposed a number of new frameworks to capitalise securitisations.

The most important one is the introduction of a modified supervisory formula, which banks will use to calculate capital requirements as a function of the (IRB) PDs and LGDs of the underlying loans, when this data is available. Where this is not available, there are other frameworks, but these are also calibrated to the MSFA (we will examine one of these, the "Revised ratings based approach", in section 5).

It is based on a modelling framework that assumes the following[14]:

- It assumes a two period (1 year and asset maturity[15]) Merton model to model credit defaults.

- As in the calculation of capital for unsecuritised assets under Basel 2 (described in section 2), the asset values for each credit (whether in the securitisation pool or held directly by the bank) are driven by:

$$dA^i = \sqrt{\rho^i} dW^c + \sqrt{1-\rho^i} dW^i + \mu^i dt$$

  W^c is a common factor reflecting systemic (or bank-wide) risk, and W^i is an idiosyncratic factor for each credit. μ^i the drift (which is different in the risk neutral and real world measures).

- The correlations ρ^i are set as in IRB.

- It assumes that lifetime PD is a function of 1 year PD and maturity, with the function based on historical corporate default rates. (This is necessary as banks are not required to estimate lifetime PDs in the IRB framework so the data is not available).

[9]Before a small adjustment to incorporate "Uncertainty in Loss Prioritisation" and a random recovery assumption that had limited effect on large pools.

[10]For simplicity we ignore expected loss in this discussion.

[11]In reality, of course, the first tranche is only slightly less risky than the more senior tranche, as, even if lifetime losses are, on average, equal to K_{IRB} conditional on information available at the 1 year point, this is only an average, and credit conditions could change further between the 1 year point and the maturity of the pool. There are also reasons unrelated to maturity that have similar effects (i.e. idiosyncratic behaviour of the pool due to the fact that the pool is likely to be more concentrated in a specific region, sector etc. than two unrelated exposures).

[12]This issue is discussed on page 6 of [1].

[13]This issue is discussed in the first bullet point on page 6 of [1].

[14]The BCBS have produced a technical paper, [2], with the details of the calculation.

[15]All the assets underlying the securtsiation are assumed to have the same maturity.

- It makes an assumption about risk premium, in line with CAPM it assumes only the systemic factor attracts a risk premium (of 40% per annum),

- The stress is an expected shortfall (rather than VaR) stress, with the quantile chosen to be 99.7%.

- Recovery is assumed to be random. However, it is assumed to be idiosyncratic to each asset and hence the effect on large pools will be small due to the law of large numbers.

The model is then used to compute:

Par – Expected MTM of the tranche (excluding margin income) given a stress to W^c.

The reader will note that this is different to what was calculated for unsecuritised assets (i.e. "unexpected loss"), and the impact of this difference is material. This will be discussed in section 4.2.1.

Some numerical approximations need to be made to compute the loss without running a Monte Carlo simulation. In particular they use approximate moment matching to a Beta distribution. The second moment (equivalently the variance) is stressed as a model risk adjustment, and this is controlled by a parameter τ, set to 100.

Finally, two additional floors are added to the capital (and a couple of caps too). We will discuss these further below (the floors are discussed in sections 4.3.1 and 4.3.2).

4.1 Analysis of the MSFA model

We start by discussing the model regulators have assumed. We then discuss how the model has been used to calibrate the MSFA (in section 4.2). We then discuss the additional add-ons (floors and additional adjustments applied to retail exposures) that regulators have proposed in section 4.3.

Overall, the modelling approach seems relatively standard (and in line with Basel 2 IRB[16], which was calculated using a similar approach[17]). While there is nothing guaranteeing that the modelling of individual loans is consistent with Basel 2[18], it does seem relatively close (as we will see).

As has been discussed earlier, the model does not (directly) incorporate a couple of features that have important effect on the risk on securitisation.

1. It does not model the possibility that intra pool correlations could be higher than global correlations. In practice, assets within a given securitisation tend to share characteristics (they tend to be from the same region, sector, vintage, underwriter and so on), and hence are likely to be more correlated to each other than two unrelated credits. Mathematically, the simplest way to extend the model to incorporate this effect is to add an extra Brownian motion (W^p) representing shocks that apply to all the credits in the pool but not to credits outside the pool.

$$dA^i = \sqrt{\rho^i} dW^c + \sqrt{\rho_p^i - \rho^i} dW^p + \sqrt{1 - \rho_p^i} dW^i + \mu^i dt$$

   Thus, the correlation between two assets in the pool is ρ_p^i instead of ρ^i. The latter, as before, represents the correlation between assets in the securitisation and assets that the bank is exposed to outside the securitisation.

---

[16]As discussed earlier, the question of whether the calibration (of the systemic correlations and of the maturity treatment) of Basel 2 IRB is reasonable is beyond the scope of this paper.

[17]The documentation released with Basel 2 also mentions the use of structural models to calculate change in MTM.

[18]In the sense that the unexpected loss given a stress predicted by the model is relatively similar to that used for IRB wholesale exposures.

2. It does not model poolwide random recovery (random recovery is assumed to be purely idiosyncratic to each asset). This can have a material effect on senior tranches. Note that the risk of deterioration in recovery is incorporated to some extent as the IRB LGDs input to the model are supposed to "reflect economic downturn conditions where necessary to capture the relevant risks"[19].

Incorporating either of these extra effects would increase the variance of the loss distribution (conditional on a stress), transferring capital from mezzanine to more senior tranches. The current proposals recognise this, and they include a stress to the variance (controlled by τ), that is intended to mimic the effect of these two features and other model risks[20]. It should also be noted that the approximation for variance of the conditional loss distribution is intended to be "conservative", which means that it is intended to be higher than what the underlying model would imply.

4.1.1 Results

We have implemented the regulatory model underlying the MSFA[21]. Our implementation allows for more than one factor, and hence allows us to model the impact of intra-pool correlation directly. Our model uses Monte Carlo simulation instead of the approximations used in the MSFA, so it can be used to test the accuracy of these approximations.

Figures 2 and 1 show a comparison of the MSFA estimate of capital (before supervisory add-ons) and the capital calculated with the same model, in the same way, using Monte Carlo simulation. Tables 1 and 2 show the mean and standard deviation of the loss distribution (which determine the capital in the MSFA as the mean and variance are used to determine the parameters of the beta distribution used to model losses). Note that for the results in this subsection we use the models to compute the same proxies for capital as regulators have. The reasonableness of these proxies will be discussed in the following sections.

The results are for 100 names with a given rating and 5 year maturity. The results are shown for both VaR and Expected shortfall. We also show the results with a 2-factor model (with the second factor adding 10% to the intra-pool correlation). The MSFA results are shown with and without[22] the τ based adjustment. We use $LGD = 1$ in the tests in this section, however the results do not change materially with lower LGD (except for a scaling by LGD).

They show that:

- Numerical error is relatively limited for estimates of the mean, as shown in table 1.
- There is more numerical error in the variance of the distribution (table 2), and this increases the capital requirement for senior tranches (as is clear on figures 2 and 1). This is further increased by an adjustment to the variance to reflect model risk controlled by τ, although this adjustment is typically smaller than the numerical error.
- Increasing intra-pool correlation by 10% via the introduction of the second factor increases the risk on the more senior tranches in a similar way, and by a larger amount (in these examples) than the numerical error and model risk adjustment to the variance.

[19] Paragraph 468 of Basel 2.

[20] See paragraphs 39-40 of [2], where the calibration of the model risk adjustment to τ is related to the effect on the variance that increasing intra-pool correlation would have.

[21] The model is identical to the one described in [2] except that it does not incorporate random recovery. As discussed in section 4.1, the latter has a limited effect for large pools. Note also that the results in this section are for LGD=1 where the regulatory model also assumes deterministic recovery.

[22] Setting $\tau = \infty$ removes this adjustment.

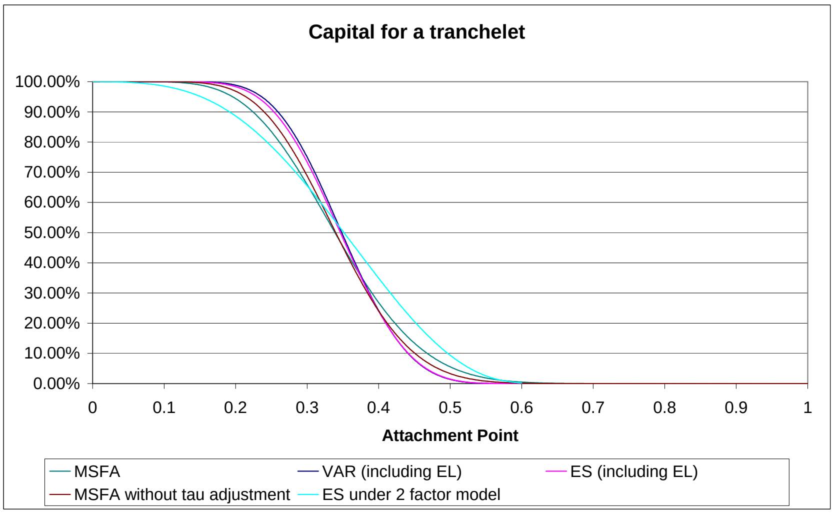

Figure 1: Capital requirements under various assumptions. 1 year PD 3.9%, M=5, LGD=1.

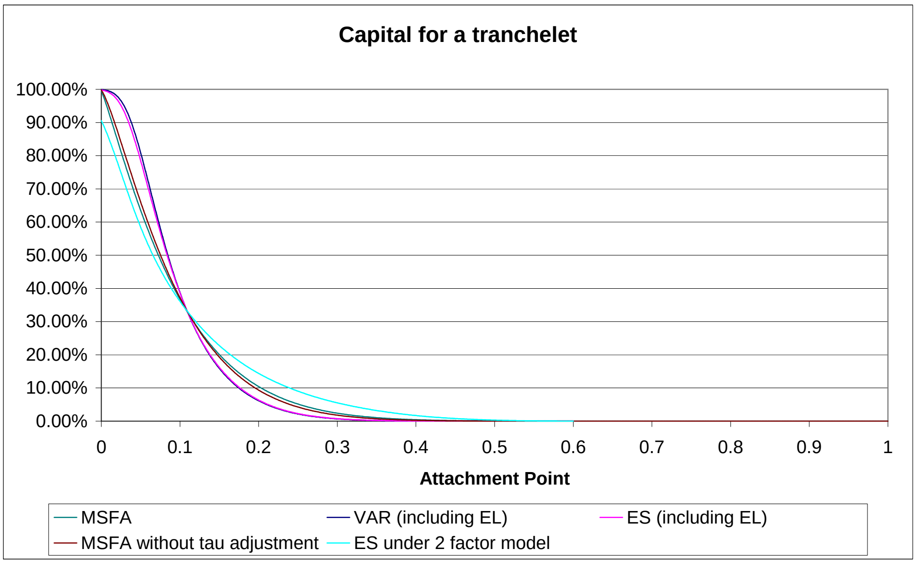

Figure 2: Capital requirements under various assumptions. 1 year PD 0.2%, M=5, LGD=1.

Maturity	Rating	VaR (99.9%)	ES (99.7%)	MSFA formula
2Y	A	3.9%	3.9%	3.9%
2Y	BBB	8.8%	8.7%	8.5%
2Y	BB	21.9%	21.6%	22.0%
2Y	B	36.4%	35.9%	36.5%
5Y	A	7.7%	7.7%	7.1%
5Y	BBB	16.2%	16.1%	15.8%
5Y	BB	37.6%	37.2%	37.5%
5Y	B	58.2%	57.8%	57.4%

Table 1: Mean conditional on a stress calculated under the regulatory model using VaR and ES stresses. The approximate value assumed by the formula is also shown. Note that the differences are small. PDs are based on historic default rates for that rating and LGD is 1

Maturity	Rating	VaR (99.9%)	ES (99.7%)	MSFA formula
2Y	A	2.7%	3.0%	4.3%
2Y	BBB	4.3%	4.9%	6.3%
2Y	BB	6.6%	7.6%	8.7%
2Y	B	7.5%	8.3%	9.0%
5Y	A	6.1%	6.3%	8.7%
5Y	BBB	9.4%	9.7%	12.2%
5Y	BB	12.6%	13.0%	14.9%
5Y	B	11.4%	11.8%	13.2%

Table 2: Standard deviation of losses calculated under the regulatory model using VaR and ES stresses. The approximate value assumed by the formula is also shown (calculated before the τ based adjustment for model risk). Note that the numerical error in the approximation (i.e. the difference between the second and third columns) has a larger effect than the choice of risk measure. PDs are based on historic default rates for that rating and LGD is 1.

4.2 Representation of capital requirement

As well as the details of the model, it is important to understand the financial quantity the MSFA is calculating in order to determine the capital requirement. There are a number of elements to this.

4.2.1 Treatment of expected loss

As discussed earlier, the model seeks to compute:

$$\text{Par} - \text{Expected MTM of the tranche (excluding margin income) given a stress}$$

Now, capital should aim to compute the value today minus the expected value of the position at the capital horizon[23] conditional on a stress.

The MSFA differs from this in an important way. In reality, the value of the tranche includes margin income, but these are specifically excluded for the MSFA. This is a departure from the treatment for loans where regulators assume that the spread margin income has a value (on average) equal to the expected loss on the tranche, even in a stress scenario, and hence only "unexpected loss" is included. However, regulators have chosen to ignore these payments for modelling securitisations, giving a number of reasons:

- They are concerned that the margin income may have little value is a stress scenario (as the tranche is wiped out before the payments can be made). This does seem to

[23] We include cashflows that occur before the capital horizon, such as interest and principal payments, in the "value" at the capital horizon.

happen in some circumstances, such as for certain equity tranches of for which coupons are paid proportionally to outstanding tranche notional, although it is dependent on how the deal is structured. Note however that if we take the sum over all tranches this effect should cancel out as the margin income not paid to the equity tranche will be paid to another tranche instead (the expected margin income over all tranches in a stress is the same as for the portfolio of underlying loans), the issue seems to be that it is not clear which tranche benefits from the margin income on the underlying collateral. Note also that the risk that margin income may not be paid should only account for a fraction of the risk, as the risk of principal losses should be the main driver of risk.

- They may also be concerned that the expected loss may not be included in the valuation. This can happen if there has been deterioration of credit conditions that has not been reflected in the accounts, or simply because it is a retained tranche that has never traded but is held at Par even though it is more likely to experience losses than the margin it attracts suggests. They therefore want to ensure that enough capital is held to ensure that these issues do not lead to overstatement of regulatory capital.

- They also say they that calculating expected loss would lead to extra complexity. This does not seem a very compelling reason to substantially change the capital requirement, and in practice expected loss could be calculated with the same numerical techniques as the existing calculation.

Tables 3, 4 and 5 show the effect of including expected loss on the total capital requirement across all tranches. The effect is shown as a proportion of K_{IRB}, i.e. what the capital requirement for the pool would be if it were not securitised (which in theory should be equal to the total capital requirement across all tranches after securitisation, as the underlying risk is the same, see the discussion in section 4.3). As is clear a large proportion of the capitalisation, particularly for retail collateral and for low rated wholesale collateral, comes from the expected loss. The results are particularly extreme for retail securitisations, where this assumption increases the capital requirement by a factor of 4.

Figures 3 and 4 also illustrate the effect. The top two lines show the capital requirement under the MSFA (with and without the supervisory add ons discussed later). The area under the bottom line of the graphs is the extra capital that comes from the inclusion of expected loss in the calculation.

PD	0.20%	0.50%	1%	5%
Incease due to “real world” EL	58%	70%	80%	111%
Increase due to risk premium	31%	32%	34%	29%
Total	89%	102%	114%	140%

Table 3: Increase in total capital requirement across all tranches of a securitisation due to the inclusion of expected loss (split into real world expected loss and the risk premium) for 5 year RMBS securitisations, for various probabiltiies of default (LGD does not affect the results as the ratio is independent of LGD). Note that we have computed risk premium using a lower correlation than required by the MSFA as proposed (i.e. with implicit maturity effects stripped out), otherwise the effect would be higher (see the discussion in section 4.3.3).

This assumption thus has a material effect on the capital requirement, especially for retail exposures and low quality wholesale exposures. It seems implausible that the risk that the coupons of a securitisation exposure may not be paid in full in a stress could lead to a doubling or quadrupling of the risk across the capital structure, and so ignoring margin income completely seems excessive. We would therefore argue that expected loss/margin

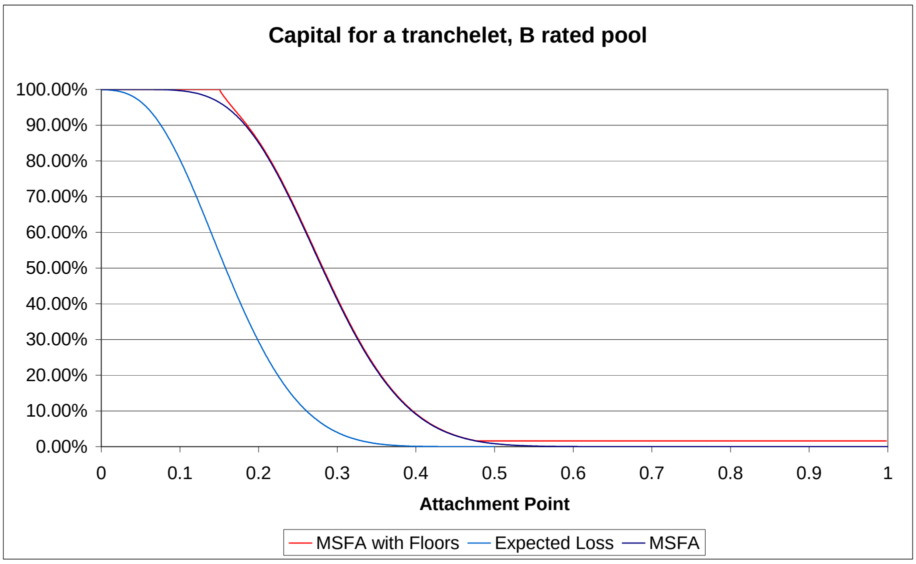

Figure 3: Effect of inclusion of EL and floors on capital. 1 year PD 3.9%, M=5, LGD=50%. Note that more than half of the capital comes from expected loss and not "unexpected loss".

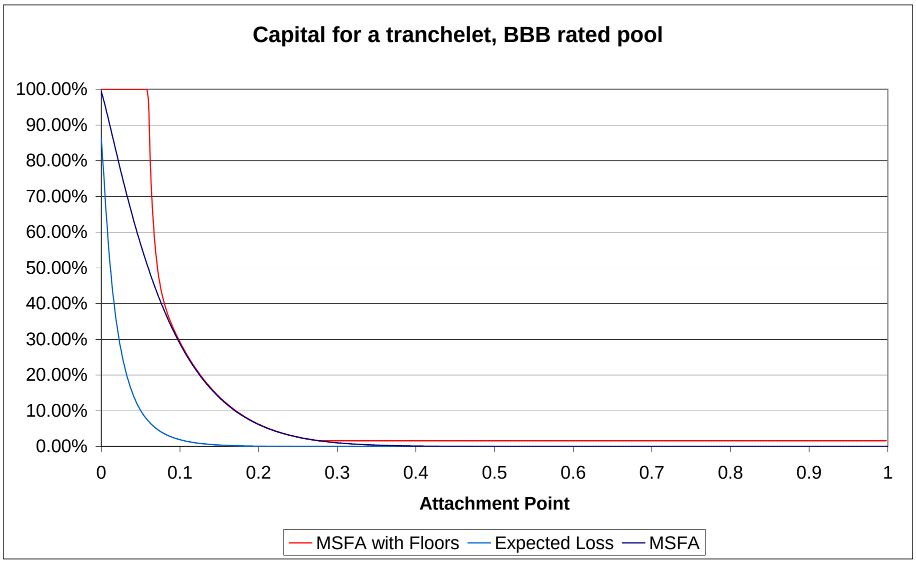

Figure 4: Effect of inclusion of EL and floors on capital. 1 year PD 0.2%, M=5, LGD=50%.

PD	A	BBB	BB	B
Incease due to "real world" EL	9%	16%	39%	82%
Increase due to risk premium	13%	17%	25%	27%
Total	22%	33%	64%	109%

Table 4: Increase in total capital requirement across all tranches of a securitisation due to the inclusion of expected loss (split into real world expected loss and the risk premium) for 5 year wholesale securitisations, for various ratings of the collateral (LGD does not affect the results as the ratio is independent of LGD).

PD	0.20%	0.50%	1%	5%
Increase due to "real world" EL	219%	243%	263%	302%
Increase due to risk premium	46%	48%	48%	37%
Total	265%	292%	311%	338%

Table 5: Increase in total capital requirement across all tranches of a securitisation due to the inclusion of expected loss (split into real world expected loss and the risk premium) for 5 year non-mortgage retail securitisations (e.g. auto loans, credit cards), for various probabilities of default of the collateral (LGD does not affect the results as the ratio is independent of LGD). Note that this assumption quadruples capital requirements.

income should be included in the calculation, at least where banks could demonstrate that regulator's concerns have been addressed or are not pertinent.

4.2.2 Use of Mark to Market (MTM)

We also note that the measure of value used is MTM, which may be considered to be inappropriate for a banking book/credit risk position. It is certainly true that a number of issues not directly related to credit risk affect the market value of securitisations, such as changes in risk premia and liquidity. This arguably led to material changes in the market value of securitisations during 2007-2009, which were subsequently reversed. It should be noted however that the model underlying the MSFA does not model changes to risk premia or liquidity, so it is, in reality, only measuring the change in MTM due to changes in credit fundamentals (downgrades, credit cycle etc.), and not other factors that may affect MTM.

Note however that, as well as the change in MTM, the exclusion of margin income means that the risk premium associated with the credit risk (which would normally be offset by margin income) is capitalised in full. In other words, the risk premium charged by the market (after year 1) for holding the risk is included in the capital requirement, in addition to both expected loss[24] and unexpected loss. This seems difficult to justify being included in a credit risk measure (see tables 3, 4 and 5 for a quantification of the impact).

4.2.3 Switch to Expected shortfall

The definition of capital also departs from existing rules by the use of expected shortfall [ES] instead of Value at risk [VaR] as the risk measure to be calculated. The reason given for this is that it eliminates cliff effects present in the old rules. However, it should be noted that these cliff effects were, as discussed earlier, a consequence of the shortcomings of the model. These shortcomings have partly been remedied in the new approach, due to the more rigorous modelling of maturity effects, so the cliff effects are no longer present for maturities greater than one year. The other important reason why there should not be any cliff risk in a reasonable model is the presence of idiosyncratic behaviour within the

[24] In the "real world measure".

pool, which the model risk adjustment controlled by τ is intended to adjust for[25]. This should further reduce cliff effects for all maturities. Ironically, as will be discussed later, the supervisory add-ons super-imposed on the MSFA in the current proposal have the effect of re-introducing cliff effects (in a way which is not affected by the switch to ES) as one of the add-ons[26] is implicitly based on a modelling framework similar to the old SFA (supervisory formula approach).

Expected shortfall also has some attractive properties that make it more attractive (in theory), although the reasons for this are not necessarily relevant here:

- Unlike VaR, it is a "coherent risk measure" so the risk of a portfolio is never more than the sum of the risk of the individual components of the portfolio.
- It does not ignore the tail of the distribution beyond the confidence level used[27]. This should be relevant for securitisations (particularly large holdings of super-senior tranches) which can be exposed to low probability events.

However, neither of these issues are relevant when the loss distribution being calculated is approximately Gaussian (in which case 99.9% VaR and 99.7% ES behave similarly). Conversely, the more discrete or fat tailed the loss distribution is the more the two measures should diverge. For example, the loss distribution for a portfolio consisting only of a thin tranche at its maturity (or a thin tranche plus other exposures perfectly correlated to it[28]) is discrete as it takes a value of either 0 or 1 (hence the "cliff effects" in VaR but not in ES). Conversely, if we are modelling the distribution of the value a portfolio that consists of a 5 year maturity transaction (and other exposures) in 1 years time, the distribution is more continuous (the probability of default before year 5 conditional on information available at year 1[29] can be anywhere between 0 and 1). In the MSFA model, this MTM at 1 year would be a relatively smooth function of the Gaussian common factor driving the model, and so we are likely to be in the region where 99.9% VaR behaves coherently and similarly to 99.7% ES.

Our results (tables 1 and 2) confirm this intuition. They show that the switch to ES has little effect on the capital for securitisations, at least for large homogenous pools and long maturities. For example, on graphs 2 and 1 the capital requirements based on VaR and ES are almost indistinguishable. There is a larger effect on shorter maturities, where the switch to ES increases the capital for more senior tranches. Even in this case, the expected loss conditional on a stress does not seem to change materially between VaR and ES, so this does not directly create arbitrages with the rules for non-securitised exposures.

It should be noted as well that, within the modelling framework, the change from VaR to ES affects the final result via an increase in the variance of the loss distribution. The changes that are observed come from the fact that the variance of the conditional loss distribution is increased with ES, although even here the change is smaller than the increase arising from numerical error in the approximation of the variance.

4.2.4 Other risk mitigants

It should be noted securitisations often have additional structural components that mitigate the risk that holders of securitisation tranches face. For example, they often have "excess spread" that can be used to reduce principal losses. They also often have features that divert

[25]See paragraph 40 of [2].

[26]The capital deduction below K_{IRB}, see section 4.3.2.

[27]If, for example, there is a 0.05% probability of a very large loss then this will not affect VaR, but it will affect Expected shortfall.

[28]This is effectively how the model underlying the SFA behaves for a large portfolio, and, for exposures of 1 year or less the MSFA model behaves in the same way.

[29]This is the key determinant of the value of the portfolio at that time, and in fact in the MSFA approach it is the only determinant.

income from junior tranches to senior tranches when pool performance deteriorates. This would have an effect equivalent to reducing pool expected loss (while keeping "unexpected loss" constant), hence redistributing risk and capital requirements from more senior to more junior tranches.

These are not incorporated in the calculation (presumably because it would have been impossible to include features like this in a single simple formula that applies to all securitisations), but it is important to note that these reduce the risk investors in tranches face (often significantly).

4.3 MSFA and capital neutrality

As discussed in section 2, for the purpose of calculating capital for credit risk under Pillar 1, it is assumed that each exposure is an infinitesimally small part of the banks exposure. Consequently, the capital requirement for a portfolio is equal to the sum of capital requirements of each component of the portfolio. If we apply this to the tranches of a securitisation, then the sum of the capital requirement over all the tranches should be equal to the capital requirement for the 0-100 tranche, i.e. to the capital requirement for the underlying pool (K_{IRB}). We refer to this property as capital neutrality.

This property can break down in the presence of model risk. If we think of the capital requirement for a given exposure is the highest estimate obtained from the set of plausible models then the capital is no longer additive as the most conservative model for one exposure is not necessarily the most conservative model for a different exposure. However it seems unlikely that model uncertainty around features specific to securitisation[30] is very large.

As proposed, the MSFA capital requirement is not capital neutral (indeed it is a long way from being so). The main reasons for this are:

1. Expected loss (on a tranche) is not deducted, so capitalisation is based on total (partly risk neutral) expected loss conditional on a stress, rather than unexpected loss . This largely affects more junior tranches and is more important for lower quality collateral. (This was discussed in section 4.2.1).

2. Capital is floored at 100% below K_{IRB}. Clearly this only affects more junior tranches.

3. Capital is floored at 1.6% globally. This floor only affects more senior tranches and is particularly important for higher quality collateral.

There are a couple of other effects that apply in some cases:

- Inconsistent use of correlation for retail transactions (the fact that the correlation already contains a maturity adjustment is ignored). This has a large effect when it is applicable.

- Use of tranche maturity as maturity of the underlying collateral. We will not discuss this in detail but this will have a material impact in practice as tranche maturity is typically longer than the maturity of the collateral due to the way securitisations are structured.

Tables 6, 7 and 8 shows the effect of various issues on total capital for different collateral types. The differences are shown as a proportion of K_{IRB}. Note that the effect of EL deduction has been calculated after removal of floors, otherwise the effect would be minimal. The effect of the floor has been calculated assuming the securitisation is split into thin tranches.

Graphs 3 and 4 also illustrate the effect of these increases in capital.

- The line in the centre shows the MSFA before the floors are added.

[30] Any uncertainty that affects capital for unsecuritised assets as well would already be reflected in K_{IRB}.

	A	BBB	BB	B
100% Floor (with no EL deduction)	57%	32%	9%	2%
1.6% floor	37%	18%	8%	5%
No EL deduction	22%	33%	64%	109%
Numerical error	-8%	-3%	1%	-2%
Modelling assumptions	-6%	2%	0%	-16%
Total	102%	83%	81%	98%

Table 6: Effect of different factors on capital neutrality for a 5 year wholesale transaction. The impact is shown as a proportion of K_{IRB}, i.e. the capital that would be held on the pool if it were not securitised. The tests were done on a pool of 100 names with a given rating (i.e. a PD based on historic default rates for that rating) and 50% LGD.

1 year PD	0.2%	0.5%	1%	5%
Deduction below K_{IRB}	23%	8%	4%	0%
1.6% floor	86%	84%	65%	22%
No EL deduction	89%	102%	114%	140%
Inconsistent Correlation	131%	101%	76%	28%
Numerical error	7%	1%	-1%	-4%
Modelling assumptions	26%	27%	21%	-18%
Total	362%	324%	278%	170%

Table 7: Effect of different factors on capital neutrality for RMBS exposures (5+ year maturity). The impact is shown as a proportion of K_{IRB}, i.e. the capital that would be held on the pool if it were not securitised. The tests were done on a pool of 100 names with the PDs in the top row and 20% LGD. These results, particularly the results for low PD, should be seen as a proxy for high quality RMBS. The effect of the inconsistent correlation was calculated by using an equivalent correlation that returned the same IRB capital for a given exposure when used with the maturity adjustment. Note that the results do not depend on assumptions made about intra-pool correlation, as this only redistributes capital between tranches without changing the overall capital. Note also that we have assumed that the attachment point of the senior tranche is chosen so the capital is floored at the minimum of K_{IRB} and 1.6%. Without this cap, the effect of the 1.6% floor would be larger for the low PD pools (219% for 0.2% PD and 108% for 0.5% PD.)

1 year PD	0.2%	0.5%	1%	5%
Deduction below K_{IRB}	12%	1%	0%	0%
1.6% floor	89%	85%	81%	21%
No EL deduction	265%	292%	311%	338%
Inconsistent Correlation	152%	134%	104%	40%
Numerical error	91%	15%	-2%	-18%
Modelling assumptions	47%	51%	47%	4%
Total	655%	577%	541%	385%

Table 8: Effect of different factors on capital neutrality for non-mortgage retail exposures (5 year maturity). The impact is shown as a proportion of K_{IRB}, i.e. the capital that would be held on the pool if it were not securitised. The tests were done on a pool of 100 names with the PDs in the top row and 50% LGD. Note also that we have assumed that the attachment point of the senior tranche is chosen so the capital is floored at the minimum of K_{IRB} and 1.6%.

- The top line shows how much extra capital the floors add.

- The bottom line shows how much of the capital is from expected loss (as assumed by the MSFA model) and not from unexpected loss.

Once we adjust for these effects, the rules are (approximately) capital neutral. This implies that the following are not important drivers of the lack of capital neutrality:

- Underlying modelling assumptions and definition of capital. These are largely consistent with Basel 2 IRB when applied to loans (note that the Basel 2 doc describes maturity adjustment as having been calculated using KMV portfolio manager, which is a similar modelling approach, a structural model calculating MTM losses). This includes the risk premium which may have have been modelled similarly for Basel 2.
- Numerical error. E_G, which determines the total capital, is reasonably well calculated.
- Switch from VaR to ES.

4.3.1 Global Floor

Capital is floored at 1.6% globally. This floor only affects more senior tranches and is particularly important for higher quality collateral. This is intended as a safeguard against model risk. In other words, even though the MSFA model (or other modelling approach) predicts a capital requirement close to 0 for a particular exposure, regulators are concerned that this could be due to a flaw in the model and another more realistic model could imply a higher amount of risk.

We agree that there is some justification for a model risk adjustment in some cases. The model risk comes from a number of sources (high intra pool correlation, random recovery, uncertainty in PD and LGD estimates)[31]. These all have the effect of increasing the variance of the stressed loss distribution, and this would only affect the tranches that the MSFA predicts are borderline low risk. However, it seems hard to justify for exposures where this is not the case (such as more senior tranches of high quality pools).

For example, the green line on figure 5 shows the requirement for a high quality RMBS pool (0.5% PD and 20% downturn LGD), calculated with the model underlying the MSFA[32]. The orange line shows the capital requirement calculated with the MSFA (which is more conservative as it uses a conservative estimate of the variance of the loss distribution and contains an adjustment, controlled by τ, for model risk, as discussed in section 4.1). In order to illustrate the effect of model risk, we stress two of the assumptions in the MSFA model. In particular, we stress the intra-pool correlation by 20%, and we assume that there is a 50% chance that, in a downturn, the LGD doubles to 40%[33], keeping expected LGD constant. We calculate the capital requirement in this model, and this is the blue line on figure 5. As is clear, if we believe that this is a plausible alternative model, then the capital requirement is indeed under-estimated for the 10-20 tranche (and to some extent for the 20-30 tranche), which can be considered "borderline low risk". The yellow line shows the MSFA capital requirement after the floor has been applied. The floor does indeed partially mitigate the under-capitalisation of the 10-20 tranche, although the capital is still less than half what the stressed model predicts. However, the vast majority of the effect of the floor (around 80%), is on the 30-100 tranche, that even the stressed model predicts has negligible risk. Thus, in this example, the floor is a very crude instrument for preventing under-capitalisation due to model risk.

[31] Some of these are discussed in section 4.1.

[32] For this exercise we use a systemic correlation of 8% that has been adjusted to remove double counting of maturity effects, see section 4.3.3. As proposed, the capital requirement is more conservative than we use here, so the need for extra adjustments is less.

[33] If we assume LTV is 60% and foreclosure costs are 20% then this corresponds to a decline in house prices of just over 50%.

As discussed in the previous paragraph and in section 4.1, the existing proposal already contains adjustments to increase the variance to incorporate model risk. This is done explicitly via the adjustment controlled by τ, and implicitly by the use of a conservative approximation for the variance in the MSFA (this implicit adjustment is typically several times larger than the explicit adjustment). If supervisors believe that these two existing adjustments are insufficient they could further increase the variance. This would achieve the same result where it is necessary (i.e. borderline low risk tranches) without a global distortion on all tranches, the majority of which are unlikely to be modelled incorrectly by a material amount.

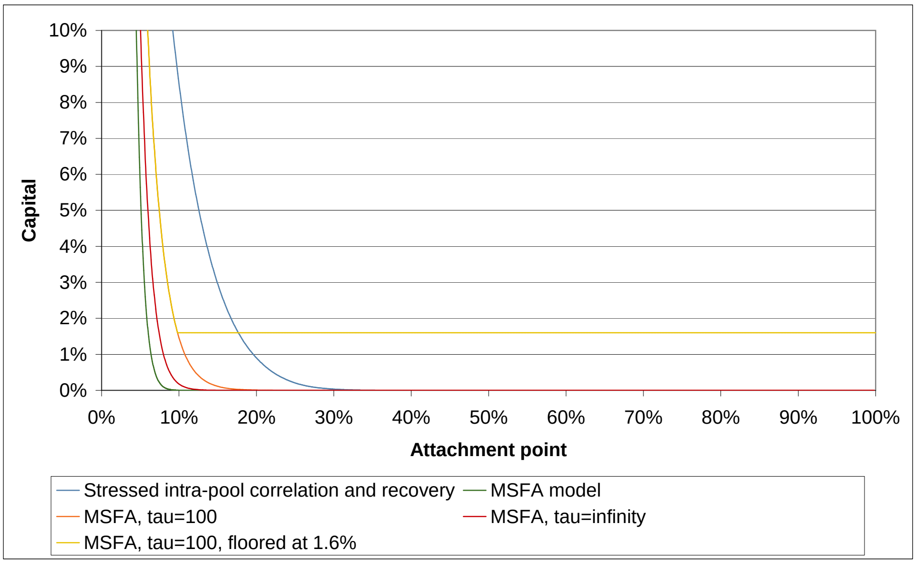

Figure 5: Illustration of effect of floor and of model risk on a prime RMBS securitisations. Note that the effect discussed in section 4.3.3 has been stripped out of these results , so capital under current proposals would be more conservative than shown here.

Relationship to crisis losses

It should further be noted that this extra capital, which has the largest effect on the most senior tranches of securitisations of high quality collateral, would have done little to mitigate the losses in the recent crisis. The crisis losses were driven largely by securitisations that were known to have relatively low quality (e.g. subprime) or highly correlated (ABS) collateral, or by the more junior tranches of securitisations of better quality assets. A floor whose main effect is to increase capital for the highest quality securitisations would therefore would have done little to mitigate the losses that were experienced in the crisis.

As an illustration, consider the fact that even the worst US prime RMBS securitisations from 2006/2007 suffered losses of less than 20% (losses on European RMBS have been much smaller than this). However, more than 80% of the extra capital imposed by this floor on a prime securitisation would have been imposed on the 20-100 tranche which did not suffer any losses despite a severe housing downturn.

Relationship to covered bonds

For another relevant comparison, we can also note that covered bonds can achieve a risk weight half as high as the highest quality securitised tranches with far less credit enhancement.

Recall that a covered bond is a bond issued by a bank secured on certain assets, with some over-collateralisation. Covered bonds are thus similar to the senior tranche of a securitisation but with an additional guarantee from the bank that issued the bond to pay the principal and interest irrespective of the performance of the assets. The highest rated covered bonds can achieve a capital requirement of 0.8% (10% risk weight) under current regulations.

If we take, say, the 50-100 tranche[34] of a high quality RMBS deal, in the extreme scenarios that would required to realise losses on this exposure (i.e. a scenario far more extreme that the recent US mortgage crisis where losses on the worst performing prime pools were less than half this), it is unlikely that the bank guaranteeing the covered bonds (which most likely is heavily exposed to similar mortgages to those underlying the RMBS), will be able to honour the guarantee. We would argue therefore that a covered bond with, say, 10% credit enhancement is riskier than the senior tranche of a securitisation of those same assets with 50% credit enhancement.

Potential unintended consequences

We also note that the floor may have the effect of increasing systemic risk. For example, it is likely that banks seeking to generate even a moderate return on equity will find holding the safest securitisations unattractive as the margin on these is likely to be small. They are more likely to prefer to invest in the "border-line low risk" tranches. If on a particular securitisation (such as the example on figure 5), the market is pricing a risk that the MSFA does not include, there may be a tranche that is modelled as "borderline low risk tranche" within the MSFA but is in fact risky (e.g. the 10-20 tranche on figure 5) and which will pay a moderately high coupon but which would attract the same capital requirement as the safest tranches. Regulated institutions (at least those with risk management failings) could therefore end up disproportionately invested in these borderline tranches, with the genuinely and unequivocally safe tranches being transferred to unregulated institutions[35].

4.3.2 Capital deduction below K_{IRB}

As mentioned above, there is a supervisory add-on that assigns capital requirement of 100% to all exposures below K_{IRB}. There is also an "omega adjustment" that smoothes the transition between the region where full deduction applies and where the underlying model based estimate applies.

As discussed in section 3, in the existing regulations (SFA), this behaviour arises naturally because of the underlying modelling assumptions, i.e. that lifetime pool losses in a 1 year stress scenario are assumed to be exactly equal to K_{IRB}. This however is not a realistic assumption, as discussed in section 3 and in the BCBS papers [1] and [2]. The MSFA now makes more realistic assumptions and, as a consequence, it produces a more realistic estimate of capital that does not have this feature. However, supervisors have superimposed this requirement on the MSFA, effectively requiring that, where the old (flawed) modelling assumptions give a higher capital requirement than the MSFA, the former should be used. There is little discussion in the regulatory papers about the rationale for this particular

[34]We discuss the 50-100 tranche since this will account for more than 50% of the notional of senior tranches with less than 50% credit enhancement, and so the risk of this tranche is relevant even for securitisations that do not have as much credit enhancement for the senior tranche. This is particularly relevant for this issue as this floor is proportional to tranche notional. Note also that swap counterparties whose exposure is senior to all the tranches effectively have significantly more than 50% credit enhancement, and they would also be affected by this floor.

[35]Of course, it is impossible to forecast with any certainty what the effects will be, but it seems plausible to us that this could happen.

add on, and it should be noted that the assumption that future margin income should be ignored already imposes a significant amount of conservatism on the modelling of more junior tranches[36], so an extra add-on seems unnecessary.

4.3.3 Double counting of maturity effects for retail exposures

As acknowledged in the BCBS paper (pages 36-37 of [1]), the proposed rules for retail securitisations double count maturity effects because the IRB correlation parameters contain an implicit maturity adjustment (i.e. they have been adjusted upwards to incorporate maturity effects), but the way they are used in the MSFA assumes that these correlations are unadjusted and hence predicts far higher losses on retail exposures in a stress scenario than the IRB rules predict.

The BCBS say that they have done this to offset other risks that they believe are present in retail securitisations, in particular high (intra pool) correlation. However, this seems a very crude way of incorporating this effect. In particular:

- Correlation increases the variance, not the mean, of the conditional loss distribution. As discussed in section 4.3.1, a more risk sensitive way of incorporating this into the model is to increase the variance (as has already been done to some extent).
- This extra conservatism is only applied to the MSFA, the other rules, i.e. the SSFA (Simplified supervisory formula approach), and the Revised rating based approach (RRBA) do not double count maturity effects in the same way[37]. This is also clear from the graphs on page 26 of the BCBS paper [1] which show that the SSFA generates much lower estimates than the MSFA for retail exposures, particularly for lower quality pools.

Graph 6 shows how incorporating higher intra-pool correlations affects capitalisation in a qualitatively different way to what has been proposed, i.e. increasing intra-pool correlation increases capital for more senior tranches while keeping the total capital constant, whereas the proposed change increases capital for all tranches (including mezanine and junior tranches whose risk should stay the same or decrease as intra-pool correlation increases). It also shows that, even if we adjust the correlation to be consistent with Basel 2 (which gives us a systemic correlation of around 8%), the numerical error and τ based model risk adjustment in the MSFA are equivalent to increasing the intra-pool correlation significantly (to around 28%).

It should be noted that this does have a material effect on the total capital requirement across all tranches[38]; for high quality RMBS, the total capital more than doubles due to this effect (before other add-ons are taken into account).

5 Revised Rating based Approach (RRBA)

As mentioned earlier, the MSFA can only be used where information[39] on the underlying pool is available. Where this is not available, regulators have proposed a revised rating based approach (RRBA) that would require banks to hold capital based on the credit rating (produced by a rating agency) of the securitisation (this framework would only apply in certain jurisdictions as in some countries, e.g. the USA due to Dodd-Frank, the use

[36] Junior tranches have higher expected loss so a disproportionate amount of margin income is paid to the holders of those tranches.

[37] The BCBS paper states that maturity is also double counted in the same way in the RRBA, but this is not true as the calibration is done purely based on corporate exposures. Note also that, as is discussed in 5.2, increasing intra-pool correlation should in fact materially reduce, not increase, the capital in a rating based approach (assuming of course that the ratings are correct and incorporate the high correlation).

[38] This would of course be unaffected by a higher intra-pool correlation.

[39] PD and LGD for each exposure calculated using an internal regulator approved model.

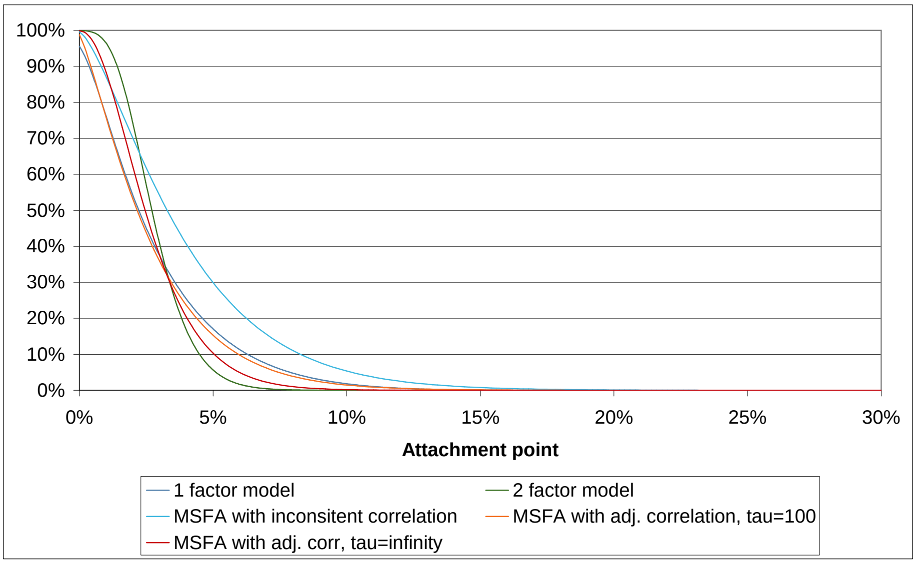

Figure 6: Capital requirements under different assumptions for a high quality RMBS securitisation. We assume 0.5% 1 year PD and 5 year maturity. Other than the top line, which uses 15% correlation as proposed by regulators, the other results use a systemic correlation chosen so that it gives the same capital, when used with a maturity adjustment, under IRB rules (around 8%). For the two factor model, we have added an extra factor so that the intra-pool correlation is 20% higher (i.e. 28%), and this brings it close to the MSFA with 8% correlation.

of credit ratings in regulation is illegal). As well as rating, the capital requirement also depends on tranche thickness, maturity and whether it is a senior tranche.

The rating based approach is based on the MSFA, and was calculated using the following algorithm.

1. Identify a tranche that the model underlying the MSFA predicts would have a given rating (and the other required characteristics). The rating is assumed to be determined by the (unconditional) expected loss on the tranche, and is chosen so that the tranche has the same expected loss as corporate bonds with the same rating have had historically. The details of how this is done are important and will be discussed in the following section.

2. The capital requirement for this tranche is then calculated using the MSFA. This is taken to be the capital requirement for tranches with the same characteristics (the final capital is floored at 1.6% as with the MSFA[40]).

[40] However, this does not have much effect as the requirements calculated before the floor is applied are almost always above 1.6%, except for 1 year AAA senior exposures where the requirement is slightly below the floor.

The general approach seems reasonable to us. However, we believe there are a number of issues with the details of what has been done which lead to the capital estimates being excessively conservative in some scenarios.

5.1 Choice of tranche

One issue we see is that the calibration takes the worst possible tranche with a given rating (in the Merton model on which the MSFA is based) and calculates the capital for that tranche. In particular

1. It assumes the attachment point is the lowest attachment point that would be attract a given rating.

2. It also assumes the pool is a low quality pool (of B rated corporates).

This capital requirement is then used for all tranches with this rating, even though many of these will not be as bad as the worst possible tranche.

The first point is a particular issue for AAA tranches as there is a large amount of variation within this category, and therefore the worst possible tranche is not representative of the risk on most tranches. Note for example that securitisations are often split into a number of different AAA tranches[41], and so the more senior of these AAA tranches has considerably more credit enhancement than the most junior AAA tranche, and hence will have less risk than the "riskiest possible" AAA tranche. Indeed, if we repeat the calibration exercise for a senior AAA tranche with another AAA tranche below it, then we will get a much lower capital estimate[42].

The second point is relevant because the capital for the riskiest senior tranche with a given rating increases with the PD of the pool. For example, on high quality RMBS securitisations the IRB capital requirement per unit notional for the entire pool is often well below the RRBA capita for the senior AAA tranches alone, and hence the latter is much higher than it would be if calculated otherwise.

Note also that with low rated pools (as used for this exercise), more than half of the capital comes from "expected loss", not unexpected loss, so this significantly increases the capital requirement.

5.1.1 Alternative Possibilities

It is clear that the reason why regulators have chosen to calibrate the RRBA in this way (i.e. based on the riskiest possible tranche) is that this is the best that can be done if the inputs used are the only information available. However, the requirements could be made more risk sensitive if further information was available. For example, rating agencies could release ratings of the underlying pool, which could be used either by making the rating of the pool an input to the RRBA rules or to be used in a formula like the MSFA.

5.2 Choice of model

It is important to note that, because of the way the model is used in calibrating the RRBA, the results depend on the relative predictions it makes on (real-world) expected loss (which determines rating) and loss given a stress (which determines capital), and is in fact largely

[41] Note also that the counterparties to swaps facing the securitisation vehicle that are often super-senior so face much less risk than the most senior tranche.

[42] Even if we assume that there is an infinitessimally thin AAA tranche below the senior tranche, the methodology used to calibrate the RRBA leads to an estimate of 5.7% instead of 58%. This may seem counterintuitive, however it is because a thin tranche needs far more credit enhancement to have an expected loss commensurate with a AAA rating than a senior tranche so the fact that there is a more junior AAA tranche means that the senior AAA tranche is much safer than the riskiest possible senior AAA tranche.

unrelated to whether the capital requirement predicted by the model for a given tranche is conservative or aggressive in absolute terms.

Thus, a model that slightly under-estimates capital requirements for a given tranche, but significantly under-estimates expected loss, will significantly over-state the capital requirement for a given rating. This is because the model will predict that even a tranche with a small amount of credit enhancement will attain a high rating, and the model will then predict that, because of the lack of credit enhancement, the capital requirement for this tranche is very high. Conversely, if a model overstates expected loss significantly, but only moderately overstates capital, it will predict that a very large amount of credit enhancement is required to attain a given rating, and this large amount of credit enhancement will ensure the capital requirement is small.

Given this, we note that the model used is, in a sense, the most conservative possible model that could have been used for the calibration of the RRBA (irrespective of whether it is a conservative model when used in isolation). To see this, note that we can broadly divide the risks in securitisations into two categories:

1. Risks that are perfectly correlated to systemic (or bank-wide) risk.
2. Risks that are independent of systemic risk, but nevertheless affect the risk on an individual securitisation.

Note that expected loss (which determines rating) depends equally on risks in either of the two categories, as it is based on an unconditional expectation (i.e. all that matters is the probability of loss, it does not matter whether this loss occurs in a crisis or in normal market conditions). However, for the purpose of calculating capital requirements, the first category is significantly more important as capital is based on losses that would be realised in a systemic crisis (and indeed if there were no risks in the first category the capital requirement would be 0 in the Pillar 1 framework[43]). The risks in the second category do have some effect on capital (which can be positive or negative), and it is important not to ignore these, but they only have a second order effect on capital requirements[44] Thus, a model which assumes there is an equal amount of risk in both categories will lead to significantly lower capital requirement than a model that assumes all risks are in the first category. At the other extreme, a model that places all the risk in the second category will predict a capital requirement of 0 for all ratings.

Now, the MSFA, as used in the calibration of the RRBA, only incorporates risks in the first category[45]. In particular, the only component that affects the calculations within the model is the systemic common factor. There is no modelling of factors idiosyncratic to the pool, such as pool specific risk. As has been mentioned earlier in this paper and in the BCBS papers this is unrealistic as in reality the assets in a given pool will be concentrated in a given sector, region, vintage, originator etc. so the behaviour of the securitisation will be dependent on pool specific risk, which falls in the second category.

To quantify this effect, we have replicated the calculation regulators have carried out[46].

[43]This is because the framework assumes that each exposure is an infinitesimal part of a diversified portfolio.

[44]They only affect capital due to what would be called cross-gamma effects in a derivatives pricing context. So, for example, a tranche with a large amount of subordination will have a fairly low exposure to intra-pool correlation as it is unlikely to experience a loss irrespective of correlation. However, if expected losses increase substantially then there is some possibility of the tranche taking a loss, but the probability of this happening will be sensitive to intra-pool correlation. The size of the loss in a stress thus depends on the intra-pool correlation.

[45]The full MSFA model does incorporate some effects that are in the second category. In particular, for finite sized pools the idiosyncratic risk to each asset does have some effect on the loss distribution (including via the random recovery assumption), and the model also indirectly incorporates intra-pool correlation via the τ based adjustment. However, these features are not included in the calibration of the RRBA as this is done on an infinitely granular pool and without the τ based adjustment.

[46]We have used 100 equally weighted names names instead of an infinitely granular pool as that was easier to calculate with our models.

We have used both the MSFA model and the two factor model described in section 4.1, setting the intra pool correlation to twice the systemic correlation (26.2% and 13.1%[47] respectively). We use both the MSFA formula and a Monte Carlo implementation of the model (in all cases we use the Monte Carlo implementation to calculate expected loss and hence rating). The results are in tables 9 and 10. The results with the two factor model are as low as a quarter of the results in the one factor model and an even smaller fraction of the results from the MSFA formula[48].

Rating	MSFA formula	MSFA model	2-factor model
AAA	113%	103%	26%
AA+	159%	146%	42%
AA	231%	216%	71%
AA-	266%	251%	89%
A+	316%	302%	111%
A	362%	351%	136%
A-	425%	417%	171%
BBB+	498%	495%	216%
BBB	572%	576%	271%
BBB-	704%	722%	373%
BB+	835%	868%	501%
BB	956%	1000%	640%

Table 9: Risk weights (i.e. capital times 12.5) for thin tranches calculated using the MSFA formula, the MSFA model (without the approximations in the formula), and a 2 factor model (with intra-pool correlation double systemic correlation). The methodology is otherwise as used to compute proposed risk weights in the RRBA, as described in [3].

Rating	MSFA formula	MSFA model	2-factor model
AAA	43%	41%	15%
AA+	59%	57%	24%
AA	84%	83%	39%
AA-	95%	95%	48%
A+	110%	111%	59%
A	125%	126%	71%
A-	145%	148%	88%
BBB+	167%	170%	110%
BBB	191%	196%	135%
BBB-	233%	238%	183%
BB+	278%	284%	239%
BB	325%	331%	303%

Table 10: Risk weights for senior tranches calculated using the MSFA formula, the MSFA model (without the approximations in the formula), and a 2 factor model (with intra-pool correlation double systemic correlation). The methodology is otherwise as used to compute proposed risk weights in the RRBA, as described in [3].

[47] This is the IRB correlation for large wholesale exposures with the PD used in this exercise (4.729%).

[48] For reasons that are not clear to us our results are different (of the order of a third lower) than in the BCBS technical paper [3]. It is not clear in the paper what computational technique regulators have used to calculate expected loss so we cannot be sure what the source of the discrepancy is.

[h!]

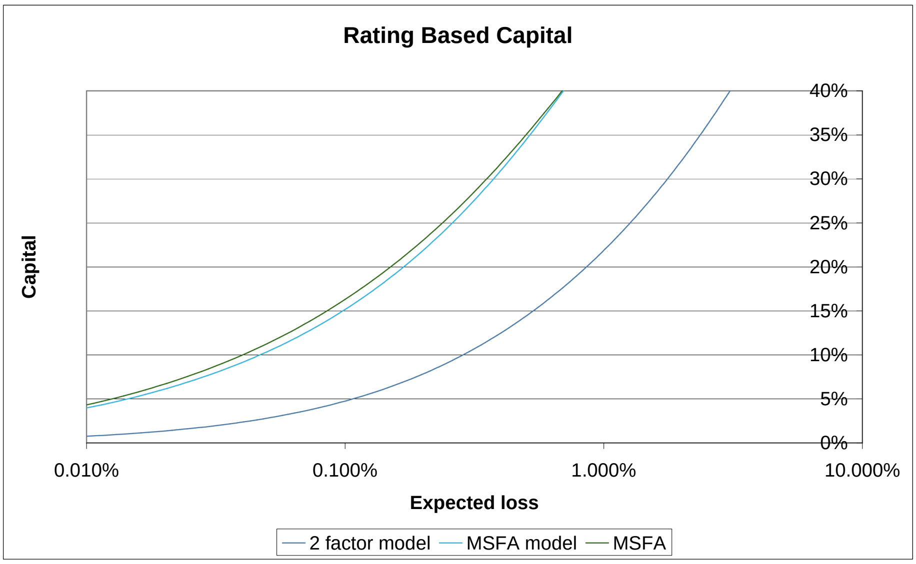

Figure 7: Capital of a thin tranche for a given expected loss (and hence rating) with the three approaches described in the text.

6 Conclusion

We have analysed the modelling and as well as examined the other assumptions behind the MSFA and RRBA. We find that the underlying modelling approach for the MSFA is a reasonable extension of the rules for unsecuritised assets. However, we find that the additional add-ons, floors and other conservative assumptions applied to the calculation lead to excessively high estimates of the capital in many cases (in the sense that the requirement is significantly higher than any reasonable model would predict) and do not account for model risk in a risk sensitive way.

In particular, the following factors lead to an unreasonably high capital requirement for some tranches:

- It is assumed that all margin income is not paid in stress scenarios, and so expected loss is not deducted as it is for unsecuritised assets.
- The capital requirement is floored at 1.6% (20% RWA) for all tranches.
- The capital requirement is floored at 100% below K_{IRB}, i.e. the pre-securitisation capital requirement.
- For retail securitisations (including mortgages), the correlation parameter is used in a way that is not consistent with how it is used and was estimated for the capitalisation of unsecuritised assets. This leads to maturity effects being double counted.

We also explain why the RRBA (revised ratings based approach), leads to the capital requirement being over-estimated in many cases.

References

[1] Revisions to the Basel Securitisation Framework, BCBS, Dec 2012, http://www.bis.org/publ/bcbs236.pdf.

[2] Foundations of the Proposed Modified Supervisory Formula Approach, BCBS, Jan 2013, http://www.bis.org/publ/bcbs_wp22.pdf.

[3] The Proposed Revised Ratings-Based Approach, BCBS, Jan 2013, http://www.bis.org/publ/bcbs_wp23.pdf.