BANK OF ENGLAND

Staff Working Paper No. 708

Mortgages: estimating default correlation and forecasting default risk

Tobias Neumann

February 2018

Staff Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This paper should therefore not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee, Financial Policy Committee or Prudential Regulation Committee.

BANK OF ENGLAND

Staff Working Paper No. 708

Mortgages: estimating default correlation and forecasting default risk

Tobias Neumann[1]

Abstract

Default correlation is a key driver of credit risk. In the Basel regulatory framework it is measured by the asset value correlation parameter. Though past studies suggest that the parameter is over-calibrated for mortgages — generally the largest asset class on banks' balance sheets — they do not take into account bias arising from small samples or non-Gaussian risk factors. Adjusting for these biases using a non-Gaussian, non-linear state space model I find that the Basel calibration is appropriate for UK and US mortgages. This model also forecasts mortgage default rates accurately and parsimoniously. The model generates value-at-risk estimates for future mortgage default rates, which can be used to inform stress-testing and macroprudential policy.

Key words: Mortgages, bank regulation, credit risk, default correlation, state space model, Basel Committee, stress testing, macroprudential policy.

JEL classification: G11, G17, G21, G28.

(1) Bank of England. Email: tobias.neumann@bankofengland.co.uk

The views expressed are those of the author and not necessarily those of the Bank of England. Many thanks to David Murphy from the Bank of England's Prudential Policy Directorate, Francesc R Tous from Cass Business School, Jon Hinder from the Prudential Regulation Authority's Regulatory Risk Specialists Directorate and participants at the Bank of England Policy Forum and Bank of England seminar participants for invaluable advice. I am also indebted to the advice from participants at the 2016 Paris Financial Management Conference.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/Working-papers

Publications and Design Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 email publications@bankofengland.co.uk

© Bank of England 2018
ISSN 1749-9135 (on-line)

1 Introduction

Bank capital requirements are a central pillar of banking regulation as they absorb losses, reduce the probability of bank failure and mitigate the associated negative externalities to society. A key element of the international regulatory capital framework, the Basel Accord, is the risk weight framework which matches capital requirements to the riskiness of banks' assets. Banks that fulfil certain requirements - in practice most larger banks - are allowed to assess this risk themselves using their internal models.

In the Basel internal ratings-based (IRB) framework for assessing credit risk, banks estimate certain parameters such as probability of default (PD) and loss given default (LGD) using their own data. They then feed these parameters into the regulatory-defined IRB function, which in turn transforms them into a 'risk weight'. This risk weight determines the minimum amount of capital a bank has to fund the asset with.

The asset value correlation (AVC) parameter within the IRB function is a key driver of how high the risk weight is for any given combination of PD and LGD. The intuition of this parameter is as follows: consider 100 loans with an annual probability of default of 1%. If correlation is zero then we would see about one default a year. If correlation is one then either all loans default together or no loan defaults. So we would see no defaults at all for 99 years and then 100 defaults in a single year. Clearly this all-or-nothing risk has different risk management and capital requirement implications than a steady procession of on average one default a year. The higher the AVC parameter, the closer the portfolio risk is to the all-or-nothing situation.

Because of the greater risk of extreme losses, regulators require more capital if default correlation is high. This effect is significant: in the IRB framework doubling correlation will more than double risk weights.[1] Since the AVC parameter cannot be directly observed it has to be estimated indirectly.

The focus of the paper is on the correlation parameter for mortgages, both because of their importance to the economy and banks' balance sheets. In the UK, for example, around a fifth of the six largest banks' balance sheets consists of exposures to mortgages: roughly £1 trillion. In the US mortgages are the second biggest asset class after Treasury and Agency securities (which themselves are largely mortgage-backed securities): $2 trillion representing about 18% of US commercial banks' balance sheets (Federal Reserve Board (2016)). In the Eurozone, banks hold mortgages worth roughly €4 trillion, representing approximately 15% of banks' balance sheets (European Central Bank (2016)). This makes mortgages a key asset class for regulators and risk managers alike.

[1] This is approximately true for a plausible range of PDs. For low PDs, increasing the AVC parameter will have a somewhat bigger impact than for higher PDs.

Given the importance of mortgages as an asset class and the AVC parameter to capital requirements, there is a surprisingly small number of studies that estimate the AVC parameter on mortgages.

The studies that have been done, used a variety of methods and data sources to estimate the AVC. The Basel Committee itself has estimated a 15% correlation parameter. The non-Basel studies tend to estimate AVC parameters below 15%, sometimes considerably so. This would imply that banks would be severely overcapitalised for mortgage defaults, in turn leading to inefficiently low or expensive mortgage lending.

But these studies relied on small sample periods leading to downward bias in the AVC estimate. They also assumed that the systematic shocks ('risk factors') causing defaults are normally distributed. The contribution of this paper is to explicitly adjust for small sample bias and to allow for non-Gaussian risk factors causing mortgage default.

Small sample bias might strike one as odd given that there are millions of mortgage borrowers in both the UK and the US. The problem is that the *time series* aspect is particularly important when estimating default correlation. One intuitive analogy is to consider estimating how retail sales increase because of Christmas: collecting data for one or a few days will not help much even if we interview millions of shoppers. What is needed is to ask many people over a long time to get a robust estimate of the change.

Typically, default rate time series are short relative to the extreme shocks regulators are interested in. To overcome this limitation, I conduct Monte Carlo simulations to adjust for small sample bias present in the methodologies I use to estimate the AVC parameter. I find that some of the estimators used in the existing literature perform poorly when time series are short. Though I include results using these methodologies to facilitate comparison with the literature, I rely on a maximum likelihood estimator and a non-Gaussian, non-linear state space model for preferred estimates. Both preferred models exhibit little small sample bias.

The Basel regulatory framework also assumes that the shocks that mortgage defaults are correlated with come from a normal distribution. I develop a generalised model similar to IRB that allows for fat tails; use a non-Gaussian, non-linear state space model to filter out the risk factor driving mortgage defaults; and adjust the AVC estimates for non-normality in the risk factors using Monte Carlo simulations.

Finally, I use the state space model to forecast default rates. I find that its forecasts are generally accurate in normal times. The model also generates prediction intervals for stressed times which can be used in value-at-risk style risk forecasting. This has a useful application in stress testing and macroprudential policy. For example, once a risk

Figure 1: Ranges of estimated AVC parameters[(1)]

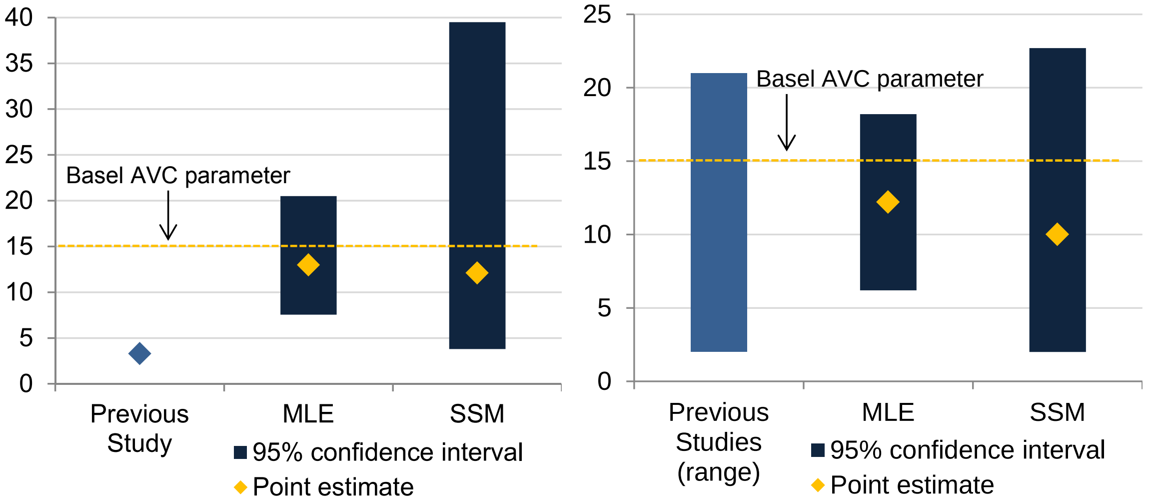

(a) UK (b) US

(1) MLE = maximum likelihood estimator; SSM = state space model

appetite is set (eg, a violation of the prediction interval once every ten years) then stress tests could be designed to deliver an appropriate capitalisation. Forecasting mortgage default risk can also be used as part of estimating the resilience of the financial system, for example in the context of macroprudential sectoral capital requirements or the countercylical capital buffer.

Figure 1 shows the estimated AVC parameters using the maximum likelihood estimator (MLE) and state space model (SSM) described in more detail below, and compares them to the range of estimates from previous studies (only one study has attempted to estimate the parameter for the UK, so there is no range). In conclusion, I cannot reject the hypothesis that the Basel parameter is appropriately calibrated for both the UK and US.

Section 2 of the paper describes how the IRB function is derived and gives an overview of the existing literature on the AVC parameter. Section 3 describes the four methods I use to estimate the AVC parameter; Section 4 describes the Monte Carlo technique I use to adjust for small sample bias and potential non-normality of the risk factors. Section 5 summarises the data and the results; and Section 6 concludes.

2 Review of the IRB function and the literature

The IRB model is an asymptotic single risk-factor model (ASRF) developed by Gordy (2003), building on Gordy (2000) and Vasicek (1997). It does two things: (i) it transforms a loan's unconditional PD into a PD conditional on a bad state of the world; and, (ii) it delivers marginal increases in capital that do not depend on the bank's current portfolio ('portfolio invariance').

Portfolio invariance is a convenient property because it means that the same loan is capitalised the same across banks regardless of what else they hold in their portfolios. This makes the framework more comparable across banks and jurisdictions, which is desirable in an international context. In order to achieve portfolio invariance, the framework assumes an infinitely well-diversified portfolio so that risk is not driven by idiosyncratic risk (the 'asymptotic' in ASRF). A second necessary assumption is that there is only one systemic risk factor affecting all loans (the 'single risk factor' in ASRF).

This section explains how the ASRF model works to the extent necessary to understand the empirical analysis; and gives a brief overview of the literature on estimating the AVC parameter.

2.1 The ASRF model

The AVC parameter (ρ) describes how the value (V) of a loan co-moves with the single systemic risk factor (Y). The remainder of its value is idiosyncratic risk (ε):

$$V = \sqrt{\rho} * Y + \sqrt{1-\rho} * \varepsilon \tag{1}$$

In the IRB framework systemic risk and idiosyncratic risk follow i.i.d. standard normal distributions, so the loan value as a whole also follows a standard normal distribution. The higher the correlation parameter, the more influence the systemic risk factor has on the value of the asset and the less idiosyncratic risk there is.

The loan is in default if its value falls below a certain threshold (K). This is more likely if either there is a bad realisation of the systematic risk factor or there is a bad idiosyncratic shock. A riskier asset will have a higher default threshold, so smaller shocks will make it default. The unconditional probability of this occurring is:

$$PD = P\left[V < K\right] \tag{2}$$

Put differently, it is the probability that a standard normal variable, V, is less than some fixed number, K. That is the definition of the standard normal cumulative distribution function (CDF), so Equation 2 becomes $PD = \Phi(K)$. We solve for K by taking the inverse CDF:

$$K = \Phi^{-1}(PD) \tag{3}$$

Substituting Equations 1 and 3 into Equation 2, and making the probability conditional on a realisation of $Y = y$ yields:

$$P\left[\sqrt{\rho} * Y + \sqrt{1-\rho} * \varepsilon < \Phi^{-1}(PD) | Y = y\right] \tag{4}$$

The Basel Committee decided to capitalise against a one-in-a-thousand year event, ie a realisation of Y that has a probability of 0.1%. The value of Y that corresponds to this probability is $\Phi^{-1}(0.001)$; and, since the normal distribution is symmetrical: $\Phi^{-1}(0.001) = -\Phi^{-1}(0.999)$. In practice, this is often referred to as the 99.9% confidence level (even though it has nothing to do with hypothesis testing). Substituting this into Equation 4 gives:

$$P\left[-\sqrt{\rho} * \Phi^{-1}(0.999) + \sqrt{1-\rho} * \varepsilon < \Phi^{-1}(PD)\right] \tag{5}$$

Re-arrange Equation 5 so that the only remaining random variable, ε, is less than a fixed number (the rest of the equation). Again, that is just a standard normal CDF, so:

$$P\left(\text{default} | Y = \Phi^{-1}(0.001)\right) = \Phi\left[\frac{\Phi^{-1}(PD) + \sqrt{\rho} * \Phi^{-1}(0.999)}{\sqrt{1-\rho}}\right] \tag{6}$$

Clearly, banks hold more than one loan. That is where the asymptotic element of the ASRF framework comes in. Gordy (2003) shows that assuming an 'infinitely fine-grained' portfolio - ie one where each individual exposure contributes vanishingly little to overall exposures - in addition to the single risk factor assumption means that the fraction of defaults should approach the conditional default probability by the law of large numbers.

Equation 6 is the 'core' of the IRB function. It calculates the probability of a loan defaulting conditional on a once-in-a-thousand year bad state of the world; or, in more precise terms it transforms the unconditional PD into a conditional PD. The IRB framework is concerned with losses rather than default rates, so it adds an estimate of downturn LGD but that is not relevant for this paper.

The probability of the proportion of loans, x, defaulting in the portfolio (ie, the portfolio loss distribution) is given by the following CDF, as shown in Gordy and Heitfield (2002):

$$P(P(\text{default}) < x) = \Phi\left[\frac{\sqrt{1-\rho}\Phi^{-1}(x) - \Phi^{-1}(PD)}{\sqrt{\rho}}\right] \tag{7}$$

The portfolio loss PDF is:

$$\frac{\sqrt{1-\rho}}{\sqrt{\rho}} * \exp\left[-\frac{1}{2\rho}\left(\sqrt{1-\rho}\Phi^{-1}(x) - \Phi^{-1}(PD)\right)^2 + \frac{1}{2}(\Phi^{-1}(x))^2\right] \tag{8}$$

In summary, note that the ASRF model makes three key assumptions: asymptoticity, a single risk factor, and normality.

Gordy and Lütkebohmert (2013) shows that violating *asymptoticity* can have a significant effect; for example, the authors estimate that the tail loss for portfolios of 500 - 999 obligors may be underestimated by 8 - 30% if granularity adjustments are not made. But this effect disappears quickly. The adjustment required for portfolios of 4000 to 8,999 obligors is only around 2%. Banks' mortgage portfolios are typically one, perhaps two, orders of magnitude greater than this. So, in practice, assuming asymptotic default behaviour is not a strong assumption in the context of mortgages.

Though the *single risk factor* assumption is a stronger assumption, it is sometimes misunderstood. It does *not* mean that there is only one (unobservable) risk driver for mortgages, which would be absurd. Both credit risk models used by practitioners and those proposed in the recent literature include several drivers of mortgage default probability: Kelly and O'Malley (2016) and Campbell and Cocco (2015) include, for example, the current loan-to-value ratio and the nature of the interest rate charged (fixed or variable).

Rather, the assumption is that the systematic risk factor is the only time-varying macroeconomic risk driver that affects all mortgages. The assumption does not rule out taking into account static drivers such as origination loan-to-value or loan-to-income ratios. Nor does it assume the default threshold for individual mortgages remains constant over time. In practice, these factors would precisely be taken into account when a bank segments its mortgage book and estimates unconditional PDs. So the single risk factor assumption really just means that once loans are sufficiently segregated, there is only one risk factor that is common to them over time.

Focusing on mortgages in single countries, the single risk factor is not an indefensible assumption. There are, of course, a number of macro-economic variables that may plausibly affect mortgage defaults - house prices, the level of unemployment and the interest rate to name but a few. But the ASRF model essentially makes the simplifying assumption that these can be condensed into one latent variable that affects the conditional probability of default for domestic mortgages: a 'frailty' factor as described in Duffie et al. (2009) or Koopman et al. (2012).

Given that I focus on individual countries and only one asset class, I interpret the systematic risk factor as a domestic mortgage frailty factor, which is weaker than the interpretation of it used in the Basel framework (which assumes a single risk factor across countries).

Finally, the assumption of *normally distributed* i.i.d.-distributed systematic shocks may underestimate the tail risk of systematic and idiosyncratic shocks. In Section 4.2 I develop a generalised version ASRF framework that can incorporate non-normality, and use it to adjust the empirical AVC estimates in Section 5.

2.2 Literature review

The Basel Committee calibrated the AVC parameter for mortgages at 15% (Basel Committee on Banking Supervision (2004)) using two approaches. The first treats banks' estimates of economic capital on their mortgage portfolios as if they had been arrived at by using the IRB function. The economic capital is used as estimates for the conditional default rate, which allows solving for the AVC parameter. The second approach relies on supervisory loss data on mortgages. The losses are split into PD and LGD, and an AVC parameter is solved for that would result in the same standard deviation as that observed in the empirical sample.

Superficially the first approach appears to be a considerable leap of faith because there is no reason why banks' economic capital should have been derived from anything looking like the IRB model. But it may have been a pragmatic approach. Assuming that banks' economic capital positions were an adequate reflection of the level of capital the Basel Committee thought prudent for mortgages, then the Committee just needed to find an AVC parameter that delivers that level of capital within the IRB framework. Given that these capital models did not drive regulatory capital at the time, the BCBS may have felt there is a low risk of these models exhibiting an imprudent bias. The second approach is the method of moments approach described in Section 3.1 of this paper and has been used elsewhere in the literature (see below).

There is a surprisingly small literature on the appropriate AVC parameter for mortgages, given the importance of mortgages as an asset class and the importance of the AVC parameter to capital requirements. Calem et al. (2003) find a range of 12.2 to 16.1% for the AVC parameter; so the Basel calibration falls within this range. The authors employ the Federal Reserve Board's credit risk model for residential mortgages and Monte Carlo simulation to estimate conditional PDs. Because these estimates are from 2003, they do not reflect the experience of the financial crisis and may therefore underestimate conditional PDs and the AVC parameter.

Comparing the results in Fitch Ratings (2008) to Fitch Ratings (2011) suggests that not reflecting the financial crisis experience biases the estimate of the AVC parameter downward. Fitch Ratings (2008) estimates the AVC parameter to be 2.07% for US (sample from 1991 to 2007) and 3.31% for UK (sample 1994 to 2007), suggesting the Basel calibration is much too high. In contrast, Fitch Ratings (2011) estimates an AVC parameter of 21% for the US using the same methodology as before but using a sample period from 1991 to 2011 Q1. Both studies fit a beta distribution to historical default rates to estimate conditional PD, and then solve for the AVC parameter.

There is a somewhat larger literature on the AVC parameter for corporate borrowers. The two Fitch reports mentioned above find lower AVC parameters for corporates than the Basel framework mandates. Gianfrancesco et al. (2011) use the same methodology and find that the Basel AVC parameter is overcalibrated for Italian corporate loans. Lopez (2004) uses a single-risk factor model owned by Moody's KMV to estimate conditional PD on credit portfolios comprised of US, Japanese and European firms. Looking at the whole sample suggests that the Basel AVC parameter is overcalibrated.

Düllmann and Scheule (2003) employ a method similar to the BCBS second approach mentioned above and apply it to a dataset of German corporates between 1987 and 2000 (similar to Gordy (2000)). The authors suggest that the Basel AVC parameter is undercalibrated for German corporates.

Düllmann and Scheule (2003) and Gordy and Heitfield (2002) recognise the problem of small sample and mis-specification bias, with the latter explicitly quantifying the bias for corporates. The authors find that the method of moments estimator suffers from considerable small sample bias, and suggest a less biased maximum likelihood estimator based on the binomial distribution.

An indirectly related literature to estimating the AVC parameter is that of estimating the systematic risk factor(s), also known as 'frailty factors'. Frailty factors are common latent risk factors that result in clustering of defaults over time. As discussed in Section 3 this is the same as the systematic risk factor of the ASRF models. Duffie et al. (2009) model the frailty factor using a default intensity model, in which the frailty factor follows the Ornstein-Uhlenberg process. The frailty factor is estimated using Markov Chain Monte Carlo estimation.

The literature building on Koopman and Lucas (2008) and Koopman et al. (2011), in contrast, uses states-space models to estimate the frailty factor in discrete time. The attraction of this approach is that the econometric framework can be mapped to the ASRF model fairly straightforwardly (if approximately), which makes it useful for estimating the AVC parameter directly.

3 Estimating the AVC parameter

At first glance calibrating the AVC parameter for an asset class looks easy. All that is needed is the PD in a stress event, ie the conditional PD, and the unconditional PD. One can then find the AVC parameter that solves Equation 6, similar to the first approach taken by the Basel Committee.

The practical problem is that we do not know the PD that would prevail in a once-in-a-thousand year crisis because we are unlikely to have observed it. At the same time, we cannot simply lower the confidence level to something that is more appropriate for the length of our time series (eg, a 1-in-10 years event). If a confidence level is chosen that is too low, there is no unique AVC parameter that matches unconditional and conditional PD in the IRB function.

In order for there to be a unique value, the conditional PD should increase monotonically as a function of the AVC parameter.[2] This can be shown by differentiating the term inside the CDF in Equation 6 with respect to ρ and simplifying, yielding the following condition:

$$\frac{\Phi^{-1}(\mathrm{P}D_{unconditional})\sqrt{\rho} + \Phi^{-1}(\mathrm{confidence})}{2\sqrt{\rho}(1-\rho)^{3/2}} > 0 \tag{9}$$

The condition in Equation 9 is true for all ρ and confidence levels above 50% if:

$$\mathrm{P}D_{unconditional} > 1 - \mathrm{confidence} \tag{10}$$

From this follows that the confidence level used has to be higher than $1 - \mathrm{P}D_{unconditional}$. Otherwise, the relationship between the AVC parameter and the conditional PD is hump-shaped as shown in Figure 2.

[2] The opposite is possible mathematically but would require a confidence level below 50% - ie looking at a stress that is more frequent than every two years. That is not the point of this type of risk model so we can disregard this situation.

Figure 2: Conditional PDs by confidence level (unconditional PD = 0.1%)

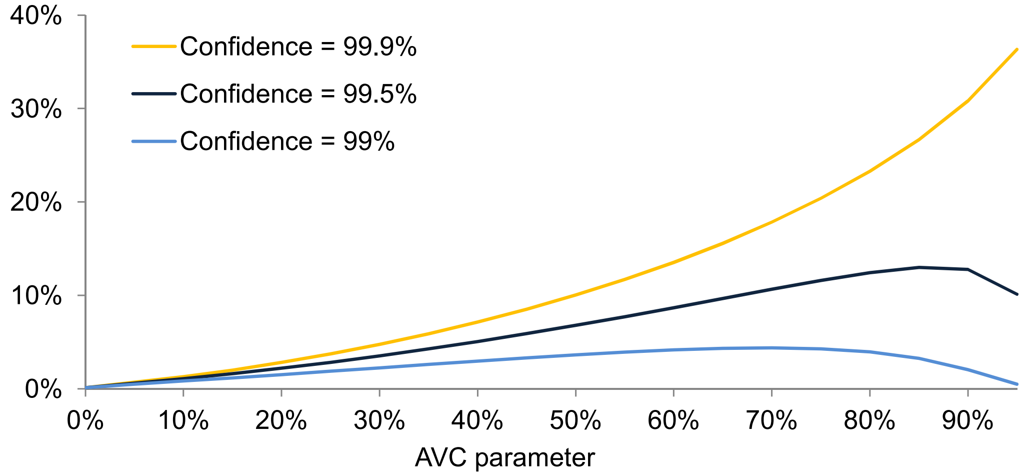

Not being guaranteed a unique solution makes it particularly challenging to estimate correlations for low risk assets unless we assume a very high confidence level. This matters in practice. Based on banks' own estimates of unconditional PD, 31% of UK mortgages have a PD below 0.1%; that means a correlation parameter cannot even be estimated for these mortgages given the Basel confidence of 99.9%.

This means the AVC parameter has to be estimated by different means. The first method I use is similar to that used in Gordy and Heitfield (2002) using a maximum-likelihood estimator (MLE) of the AVC parameter based on the IRB function. Though Gordy and Heitfield (2002) also use MLE estimators I base the likelihood function on the portfolio default rate probability distribution function; this somewhat more direct approach is appropriate given the large portfolios I analyse. Their approach is more relevant for smaller portfolios such as corporates, which is indeed what they use it for.

A novel approach is to use a non-linear, non-Gaussian state space model as introduced by Koopman and Lucas (2008) to estimate both the level of the latent systematic risk factor and the correlation parameter; this is described more in Section 3.1. The main advantage of this model is not that it gives more accurate results or is less computationally intensive (in fact, it is as accurate as the MLE method but more computationally intensive), but that it also allows me to estimate the systemic risk factor itself. This is useful for risk forecasting for policymakers and practitioners alike to forecast default rates as well as to correct mis-specification bias in the AVC parameter estimates that may arise from non-normality in the risk factors. This paper is a contribution to the literature on this topic.

The two methods I use to verify my results and compare them to the literature are fitting a beta distribution, following Fitch Ratings (2008, 2011); and a method of moments (MM) that relies on the sample variance of default correlation, as in Gordy and Heitfield (2002) and Düllmann and Scheule (2003). The relative performance of these models is analysed in Section 4 using Monte Carlo simulation.

The remainder of this section first sets out the three non-state space methods: MLE, MM and fitting a beta distribution; Section 3.2 sets out the state space model.

3.1 Non-state-space models

The main non-state space method I use is *maximum likelihood estimation*. I derive the likelihood function from the portfolio loss distribution in equation (8), and then maximise it. The log-likelihood function is:

$$LL(\rho|D) = \frac{1}{2}\log(\frac{1-\rho}{\rho}) + \sum_{t=1}^{T}\left[-\frac{1}{2\rho}\left(\Phi^{-1}(d_t)(1-\rho) - \Phi^{-1}(PD)\right)^2 + \frac{1}{2}\left(\Phi^{-1}(d_t)\right)^2\right]$$

Where D is a vector of observed default rates d_t at time t.

Gordy and Heitfield (2002) use maximum likelihood estimation to derive the AVC parameter for rated corporates. The portfolios they analyse are small, which violates the asympoticity assumption in Section 2.1. Defaults in this case follow a binomial distribution, which accordingly forms the basis of their method. My samples are large cross-sectionally so I can base the MLE estimator directly on the ASRF PDF.

A second approach, the *method of moments*, derives the AVC parameter from the sample variance of default rates. Gordy (2000) shows that the variance of the conditional default rate is:

$$Var(PD_{conditional}) = \Phi_2(\Phi^{-1}(PD_{unconditional}), \Phi^{-1}(PD_{unconditional}), \Sigma) - PD^2_{unconditional} \tag{11}$$

where Φ_2 signifies the bivariate standard normal cumulative distribution function with covariance matrix $\Sigma = \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}$. The variance of the conditional default rate can be estimated by using the sample variance of unconditional default rates.[3] This estimator is unbiased if the number of obligors (*not* years) is sufficiently large. My samples cover mortgages in the whole of the UK and US, so this limitation is of no practical concern (though as shown in Section 4.1, the method is also biased for small time series samples).

[3] Assuming that realisations of the systemic risk factor and obligor defaults conditional on those realisations are independent, as they are in the IRB model.

The AVC parameter is estimated by using the sample variance of default rates as an estimator for $Var(\mathrm{PD}_{conditional})$ and the average sample default rate as estimator for $\mathrm{PD}_{unconditional}$. Though the Basel Committee is not explicit in having used this method, its description of using observed standard deviations to estimate the AVC parameter suggests that it used this, or a similar, method.

The final non-state space approach is to estimate the 99.9% default rate directly from the data. This is the approach taken by Fitch Ratings (2008) and Fitch Ratings (2011), who fit a beta distribution to a default rate time series. This is a natural distribution to use since it is bounded between 0 and 1. The estimate of the conditional default rate can then be used as described in the introduction to Section 3. In Section 4 I show that this systematically underestimates the AVC parameter.

Using other distributions, for example the generalised Pareto distribution used in extreme value theory, is problematic in practice. The first drawback is the lack of data in the short default time series available.

The 'peak-over-threshold' method used to estimate tail events, for example, requires one to determine a threshold value above which a generalised Pareto distribution can be fitted. The inherent trade-off of the threshold being high enough to appropriately describe the tail but still leaving enough data to fit a distribution is a particularly difficult one to strike in samples of only 25 - 50 years. Even looking at only one half of the empirical distribution - which one would hardly call the tail in most applications - would correspondingly halve the sample size. This makes the resulting confidence intervals very wide.

At the same time, the results are very sensitive to the choice of threshold, for which there is no standard methodology (the most common being visual inspection of a chart). See Gomes and Guillou (2015) and Rocco (2014) for recent surveys. Clearly, any methodology relying on visual inspection or other sample-dependent judgements would make a Monte Carlo evaluation unreliable.

An additional challenge with extreme value theory is that the predicted conditional default rates are not bounded by 1. This is important in practice. I simulate 10,000 samples of 50 defaults using the ASRF model and fitted a generalised extreme value distribution; this distribution does not need a threshold value so is easier to fit than the generalised Pareto distribution. I then use the fitted distribution to estimate the 99.9% default rate: 14% of results were above 1 even though the underlying simulated data was based on normally distributed shocks. Simulating fat-tailed shocks would only increase this number. So extreme value theory cannot be used reliably for the purpose of this paper.

Figure 3: Partial Autocorrelations and 95% confidence intervals for default time series

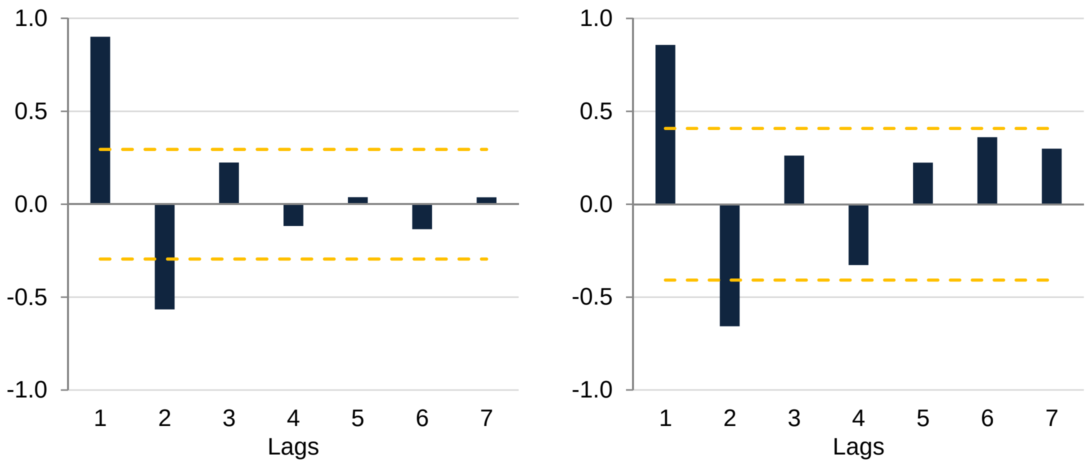

(a) UK (b) US

3.2 State-space model

State-space models (SSM) are used to estimate latent - ie unobservable - variables. Since the AVC parameter is a measure of the correlation of defaults with the unobserved systemic risk factor, state space modelling is a natural statistical framework to use. Being able to estimate a model for the mortgage frailty factor is the key advantage of using a state space model compared to the three methods used in the literature so far. I use the estimated frailty factor to adjust the AVC estimates for potential fat tails in the data generating process. The estimated model can also be used to forecast mortgage default rates.

State space model are very flexible. I exploit this by modelling the mortgage frailty factor as an AR(2) process with an unconditional mean of zero, and an unconditional variance of 1. This is consistent with the assumptions in the ASRF model. In contrast to the ASRF model, where the systematic risk factor is presented as white noise, an autocorrelated process is more realistic. Figure 3 plot the partial autocorrelation function for the UK and US default time series, which suggests an AR(2) process. Given that in the ASRF model defaults over time are driven by variation in the systematic risk factor, I also model the mortgage frailty factor as an AR(2) process.

I use the non-Gaussian, non-linear state space model introduced by Koopman and Lucas (2008) to model the systematic risk factor. Default counts, y_t, are modelled as a binomial distribution determined by the size of the portfolio, k_t, and the conditional probability of default, π_t. The conditional default probability is determined by a 'signal', θ_t, which in turn is a linear function of the mortgage frailty factor α_t. The logistic function transforms the signal into the conditional probability of default. This step has no economic interpretation, but it ensures that the linear transformation α_t is bounded between 0 and 1, as the unconditional probability of default has to be. As discussed above, α_t is an AR(2) process; the scalar S is chosen such that the process has an unconditional variance of one. Formally:

$$y_t \sim \text{Binomial}(\pi_t, k_t) \tag{12}$$

$$\pi_t = \frac{1}{1+\exp(-\theta_t)} \tag{13}$$

$$\theta_t = \lambda + \beta\alpha_t \tag{14}$$

$$\alpha_t = \gamma_1\alpha_{t-1} + \gamma_2\alpha_{t-2} + S\eta \tag{15}$$

$$S = \sqrt{\frac{(1+\gamma_2)(1-\gamma_1-\gamma_2)(1+\gamma_1-\gamma_2)}{(1-\gamma_2)}} \tag{16}$$

$$\eta \sim \text{N}(0,1) \tag{17}$$

This econometric model can be linked to the structural ASRF framework. Note that in the ASRF framework: $\pi_t = \Phi\left[\frac{\Phi^{-1}(PD)-\sqrt{\rho}\alpha_t}{\sqrt{1-\rho}}\right]$. Combining this with Equation 13 and Equation 14 yields:

$$\frac{1}{1+\exp(-\lambda-\beta\alpha_t)} = \Phi\left[\frac{\Phi^{-1}(PD)-\sqrt{\rho}\alpha_t}{\sqrt{1-\rho}}\right] \tag{18}$$

First, notice that the left-hand side of Equation 18 is the CDF of the logistic distribution with mean zero and scale parameter, s, of one evaluated at $-\lambda - \beta\alpha_t$; this means it has a variance of $\frac{\pi^2}{s^2 3} = \frac{\pi^2}{3}$ (the number π). Since the logistic distribution is very similar to the normal distribution (the right hand side of the equation) it can be used to approximate the standard normal distribution, using a scale parameter of $\frac{\sqrt{3}}{\pi}$ to achieve a variance of one.[4]

This adjustment is approximately correct for the centre of the CDFs, ie 0.5. But since most of the default rates will be far lower than 0.5 and the logistic distribution approximates the normal distribution less well in the tails, the approximation is not precise enough.

To gain precision in the tail I fit a function, f, such that it approximately transforms the normal distribution into a logistic distribution for a given vector x:

$$[1+\exp(f(x))]^{-1} \approx \Phi[x] \tag{19}$$

[4]Recall that the normal CDF is a scaled error function, which cannot be evaluated in closed form. This explains the interest in approximating the normal CDF (and hence the need to rely on tables in statistics text books); using the logistic distribution is one way of doing so.

The domain over which the approximation is chosen affects how useful it is for the purpose of this paper. The domain must include default rates we can expect to observe if generated from the ASRF model; and the approximation should be most accurate for the most frequent default rates we expect to observe. I therefore generate random numbers by sampling from $x = \frac{\Phi^{-1}(PD) - \sqrt{\rho}\alpha_t}{\sqrt{1-\rho}}$ to ensure the range is relevant.

Next, I calculate the values $f(x) = y$ must take for each element of x such that $[1 + \exp(y)]^{-1} = \Phi[x]$. The solution to this equation is $y = \log \frac{1-\Phi[x]}{\Phi[x]}$. I approximate this function linearly by fitting the regression $f(x) = y = b_0 + b_1 x$.

Using this approximation, we can relate $-\lambda - \beta\alpha_t = b_0 + b_1 \frac{\Phi^{-1}(PD) - \sqrt{\rho}\alpha_t}{\sqrt{1-\rho}}$; and thus:

$$\lambda = -b_0 - b_1 \frac{\Phi^{-1}(PD)}{\sqrt{1-\rho}}$$

$$\beta = b_1 \frac{\sqrt{\rho}}{\sqrt{1-\rho}} \tag{20}$$

Equation 20 relates the parameter β, estimated in the state space model, to the AVC parameter in the ASRF model. This is analogous to similar relations in Koopman and Lucas (2008).

Rearranging this equation yields:

$$\rho = \frac{\beta^2}{b_1^2 + \beta^2}$$

Though the fit of the linear approximation is usually excellent over the domain we are likely to observe (with an R^2 generally greater than 98%), the estimated regression coefficients vary somewhat depending on the choice of ρ. The reason is that the AVC parameter affects the tail-thickness of the default rate distribution compared to the logistic distribution.

To test to what extent this imprecision may be problematic, I fit the approximation to a sample generated from an ASRF model with $PD = 1\%$ and $\rho = 15\%$ .This yields $b_0 = -2.04$ and $b_1 = -2.88$. The range of $b_1$corresponding to AVC values between 1% and 30% (the highest AVC parameter in the Basel Accord) is between -2.7 and -3.2. This translates to a variability of the estimate of ρ of approximately $\pm 10\%$ around the parametrisation of $b_1 = -2.88$. This is considerably lower than the typical confidence intervals of the estimate for β so the added imprecision is acceptable.

To estimate the model parameters I use the importance sampling methodology introduced by Durbin and Koopman (1997). The aim is to simulate the signal vector, θ, conditional on the vector of observations, y. The simulated data are then used to calculate a maximum likelihood estimate $\hat{\theta}_t$ for each period t. A Kalman filter and smoother is applied to Equation 14 and Equation 15 to extract a smoothed estimate of the state vector, α.

There is no closed-form expression for the probability density $p(\theta|y)$; so it has to be evaluated using numerical methods in order to draw samples of the signal vector. First, the mode of the distribution is estimated by using a linear Gaussian model to approximate the relation between the observations and the signal:

$$\begin{aligned} \widetilde{y}_t &= \theta_t + \mu_t \\ \mu_t &\sim N(0, H_t) \end{aligned} \tag{21}$$

This is matched to the binomial distribution in Equation 12 by using the Taylor expansion suggested in Koopman and Lucas (2008) and Durbin and Koopman (2012), which yields the following set of conditions:

$$H_t = k_t^{-1}(1 + \exp(\theta_t))^2 \exp(-\theta_t) \tag{22}$$

$$\widetilde{y}_t = \theta_t + H_t y_t - 1 - \exp(\theta_t) \tag{23}$$

The estimation starts with using the vector of default observations, y, and an initial guess for the vector θ^0. This is then used to estimate a vector θ^1 applying the Kalman filter and smoother, which in turn is used to estimate a new set of H_t and $\widetilde{y}_t$. This process is continued until the sequence converges.

The converged vectors $\widetilde{y}$ and $\widetilde{\theta}$ are then used in Equation 21 to draw a random sample of signals using a simulation smoothing algorithm. This vector of simulated θ_t^i (a vector of M simulated signals per period t) is then used to arrive at the importance sampling estimator $\hat{\theta}_t$:

$$\hat{\theta}_t = \frac{\sum_{i=1}^{M} \theta_t^i p(y|\theta_t^i)/p_G(\widetilde{y}_t|\theta_t^i)}{\sum_{i=1}^{M} p(y|\theta_t^i)/p_G(\widetilde{y}_t|\theta_t^i)}$$

Where $p(y|\theta_t^i)$ is the binomial distribution and $p_G(y|\theta_t^i)$ is the Gaussian approximating model in Equation 21. The final step is to use $\hat{\theta}_t$ to estimate Equations 14 and 15 using the Kalman filter and smoother.

4 Monte Carlo analysis of model performance

This section assesses the performance of the four models used in estimating the AVC parameter. In Section 4.1 I assume the default time series is generated by an ASRF process and test for small sample bias arising from short default rates time series.

Section 4.2 develops a framework for non-Gaussian risk factors. I develop a generalised ASRF model - G-ASRF - based on the stable distribution for this analysis. Finally, in Section 4.3 I test the performance of the four models when asset value correlation varies across mortgages in the portfolio.

4.1 Small sample bias

As mentioned in Section 2.2, Gordy and Heitfield (2002) show that the method of moments estimator exhibits small sample bias. This section estimates the degree of sample bias in the four methods I use to estimate the AVC parameter. I use the results of the analysis to calculate correction factors that I apply to the raw estimates for US and UK AVC parameters in Section 5.

The intuition behind small sample bias is as follows: high levels of correlation result in default time series with very low default rates for most of the time and rare events of very high default rates. In short time series relatively more samples will suggest a very calm default process, whereas a smaller number will suggest a very violent process.

Recall the example used in the introduction, where loans in a portfolio with unconditional PD of 1% have a correlation of nearly one. So over 100 years there would be 99 years with near-zero defaults and 1 year with a near-100% default rate. Consider the case where we have five non-overlapping samples of 20 years. There will be one sample that includes the near-100% default rate, the other four will not. The sample with the spike in the default rate may estimate the AVC parameters accurately or even over-estimate it, but the average and median results are brought down considerably by the four samples where the AVC parameter would be estimated to be very low (because there are no large spikes in default rates). It may be that for some AVC / sample length combinations there may be no bias or a bias to overestimate the AVC parameter, but intuition suggests that the effect is to underestimate asset value correlation.

To assess the extent of this bias, I run two Monte Carlo simulations of portfolio default rates with 1,000 draws each. The length of the simulated series as well as their unconditional default rates correspond to that in the UK (48 years, PD = 0.6%) and US (26 years, PD = 1.21%) samples, respectively. In each run, the AVC parameter varies randomly between 1% and 50%, since we do not know the right AVC parameter. To put this into context: the highest AVC parameter used elsewhere in the Basel framework is 30% for large financial institutions. Any correlation parameters much larger would imply that we observe default rates that are unrealistically severe.

The choice of upper bound could in principle bias the results: if it is chosen too low then the adjusted AVC parameters may also be too low. This is not the case here. Given the AVC parameters estimated from the data (and their confidence levels), not a single simulated run suggests that an underlying true AVC parameter of 50% is consistent with the AVC estimates in the data. So in practice the choice of 50% as an upper bound does not affect the results, but it ensures that sufficient relevant data are available in the Monte Carlo simulations.

After simulating the samples I estimate the AVC parameter using the four methods described in Section 3. This results in two sets of 1,000 simulated small-sample realisations of the difference between the estimated AVC parameter and the underlying true AVC parameter. To gauge the degree of underestimation, I divide the estimated AVC parameter by the underlying AVC parameter to arrive at 'underestimation factors' for each estimation method. For example, a median underestimation factor of 0.8 suggests that the median estimator is only 80% of the true AVC value.

Table 1 summarises the results of the simulation study. The maximum likelihood estimator and the state-space model exhibit little bias. In the case of the state-space model, the bias may also lead to overestimation for high AVC values. The variability of the two methods is also similar suggesting that both estimators are appropriate for small sample estimation.

In contrast, fitting the beta distribution and the method of moments exhibit greater bias; in the case of the beta distribution considerably so. I investigate whether this is caused by the small samples available or by a more general bias by running a Monte Carlo simulation for the two estimators using simulated samples of 1,000 years and a probability of default of 1%. In this case, the MM approach exhibits no bias and little variability. The beta distribution method, in contrast, exhibits the same bias as in the small sample simulations. This suggests that the beta distribution is a generally biased estimator for the ASRF model. Though the MM approach does not exhibit as extreme a bias as the Beta approach, it has the widest ranges of estimates of all approaches.

Table 1: AVC underestimation factors based on ASRF process

	US				UK			
	MLE	**SSM**	**Beta**	**MM**	**MLE**	**SSM**	**Beta**	**MM**
Mean	0.94	0.96	0.56	0.84	0.93	1.07	0.52	0.87
Median	0.92	0.95	0.54	0.78	0.92	1.05	0.51	0.82
Interquartile range	0.28	0.30	0.22	0.42	0.23	0.24	0.17	0.35
95% range	0.88	0.88	0.66	1.24	0.68	0.69	0.56	1.13
St. Dev.	0.22	0.32	0.17	0.32	0.17	0.47	0.14	0.3

Notes: Ratio of true AVC to estimated AVC. Based on 1,000 simulations each for US-type sample (length: 26 years; PD = 1.21%) and UK-type sample (length: 48 years; PD = 0.6%). The underlying model is the ASRF model. The AVC parameter is capped at 50% to avoid unrealistically large default rates.
MLE = Maximum likelihood estimator based on ASRF model.
SSM = State space model estimator.
Beta = Fitted beta distribution estimator.
MM = Method of moments estimator.

The simplest way to adjust for small sample bias is to take the median (or mean) adjustment factor and apply it to all estimates. But this means that observations from the simulation study may be used in the adjustment process which have no relation to the estimates based on the empirical default time series. For example, if the highest AVC estimate across all models including their confidence intervals is 20%, then it makes no sense to include simulated AVC estimates greater than 20% to derive the adjustment factor.

So, I fit a linear regression model to map the AVC estimates from the simulations to their corresponding underlying AVC values. The relevant domain extends from the minimum to the maximum AVC estimate (including their confidence intervals), discarding the irrelevant observations from the simulation study. This regression model is then applied to the AVC estimates from the real-world data. Note that this can change the adjustment somewhat compared to Table 1. For example, over the relevant domain of low estimated AVC parameters, the SSM approach shows a small underestimation bias rather than overestimation when looking at the entire simulation.

Table 2: AVC small-sample adjustment based on ASRF process

	US				UK			
	MLE	**SSM**	**Beta**	**MM**	**MLE**	**SSM**	**Beta**	**MM**
Intercept	nil	0.01	nil	nil	nil	nil	-0.02	nil
Slope	1.15	1.06	1.97	1.48	1.12	1.06	2.21	1.30
RMSE	0.040	0.035	0.067	0.079	0.011	0.010	0.025	0.020
R2	0.78	0.81	0.77	0.60	0.85	0.86	0.81	0.73

Notes: $AVC = \beta_0 + \beta_1 A\hat{V}C$. Based on 1,000 simulations each for US-type sample (length: 26 years; PD = 1.21%) and UK-type sample (length: 48 years; PD = 0.6%). The underlying model is the ASRF model. The range of estimated AVC parameter is the same as the range of parameters (and confidence intervals) estimated from the empirical data.
MLE = Maximum likelihood estimator based on ASRF model.
SSM = State space model estimator.
Beta = Fitted beta distribution estimator.
MM = Method of moments estimator.

Table 2 shows the results of the regressions. Over the relevant domains (which differs for the US and UK), all estimators suffer from some small sample bias. The MLE and SSM estimators, though, have the least bias. They also exhibit the least variance - roughly half that of the Beta and MM estimators - and provide the best fit overall. Given their smaller bias and greater accuracy compared to the other two approaches, the MLE and SSM approach are the preferred methods for estimating the AVC parameter.

4.2 Bias from non-normality

If the shocks to the systematic and idiosyncratic risk factors do not follow a normal distribution, the normal ASRF model is mis-specified. In principle, the best solution would be to adapt the ASRF model by assuming a different distribution and then to estimate an appropriate AVC parameter.

Not any distribution can be used to do this. Most non-normal distributions commonly used in economics and finance are not additive. For example, adding two Student-t distributions will not generally result in another Student-t distribution, but rather in an unknown distribution. This means that Equation 3, the default threshold, cannot generally be expressed analytically and would have to be derived from simulations on a case-by-case basis which means a general model cannot be derived.

An exception to the generally non-additive non-normal distributions is a Student-t distribution with one degree of freedom - the Cauchy distribution. The Cauchy distribution, in turn, is one member of the family of 'stable distributions'. Stable distributions can be defined so that the sum of two stable distributions (with appropriate parameters) is also a stable distribution.

Stable distributions are described by four parameters; I follow Nolan (2015) in using the notation $S(\alpha, \beta, \gamma, \delta)$. The index parameter $\alpha \in (0, 2]$ determines the tail fatness. The smaller the index parameter the more heavy-tailed the distribution. The index parameter must be the same across the systematic and the idiosyncratic risk factors for the sum of two stable distributions to also be a stable distribution.

The parameter $\beta \in [-1, 1]$ indicates skewness. Note that the skewness parameter only impacts the tail behaviour for $\alpha < 2$. The parameter $\gamma \in \mathbb{R}^+$ is a scale parameter and the parameter $\delta \in \mathbb{R}$ indicates location. The standard normal distribution can be represented in this way as $S(2, 0, \sqrt{0.5}, 0)$.

In this generalised ASRF model - G-ASRF - the variables Y and ε in Equation 1 are distributed $S_Y(\alpha, \beta_Y, \gamma_Y, \delta_Y)$ and $S_\epsilon(\alpha, \beta_\epsilon, \gamma_\epsilon, \delta_\epsilon)$. Note that the α parameter must be equal across both. The distribution of the asset value, and therefore of the default threshold, is $S_V(\alpha, \beta_V, \gamma_V, \delta_V)$ where:

$$\beta_V = \frac{\beta_Y(\sqrt{\rho}\gamma_Y)^\alpha + \beta_\epsilon(\sqrt{1-\rho}\gamma_\epsilon)^\alpha}{(\sqrt{\rho}\gamma_Y)^\alpha + (\sqrt{1-\rho}\gamma_\epsilon)^\alpha}$$

$$\gamma_V = \left((\sqrt{\rho}\gamma_Y)^\alpha + (\sqrt{1-\rho}\gamma_\epsilon)^\alpha\right)^{\frac{1}{\alpha}}$$

$$\delta_V = \delta_Y + \delta_\epsilon$$

In addition, I assume $\alpha > 1$, as otherwise the mean of the distribution is undefined. If the mean is undefined the law of large numbers does not apply (as there is no mean to converge to). This would violate the asymptoticity assumption as we would not be able to assume that the average default rate of the infinitely fine grained portfolio converges to the real conditional default rate.

Following analogous steps to Section 2.1, the G-ASRF model is:

$$P(\text{default}|Y = S_V^{-1}(0.001) = S_\epsilon \left[\frac{S_V^{-1}(PD) - \sqrt{\rho}S_Y^{-1}(0.001)}{\sqrt{1-\rho}}\right] \tag{24}$$

$$P(P(\text{default}) < x) = S_{-Y}\left[\frac{\sqrt{1-\rho}S_{\epsilon}^{-1}(x) - S_{V}^{-1}(PD)}{\sqrt{\rho}}\right] \tag{25}$$

Where Equation 24 is the conditional loss function and Equation 25 is the CDF of portfolio losses. Note that the term surrounding the systematic risk factor is subtracted and that the inverse CDF is a function of 0.1% rather than 99.9% as in the normal ASRF model. This is because the stable distribution is not generally symmetrical so we cannot generally write $S_Y^{-1}(0.001) = -S_Y^{-1}(0.999)$ as we can with the normal ASRF model.

Note also that in the portfolio loss CDF, the distribution is S_{-Y} not S_Y, which is distributed $S_{-Y} \sim S(\alpha, -\beta_Y, \gamma_Y, -\delta_Y)$.

For $\alpha = 2$, $\beta = 0$, $\gamma = 1$, $\delta = 0$, Equation 24 is the same as the ASRF model. For $\alpha < 2$, the G-ASRF model results in a fatter-tailed default distribution than the normal ASRF model.

To address any suspected non-normality, one either needs to develop an estimator that is robust to potential non-normality in the G-ASRF model or to adjust the existing estimators using Monte Carlo simulations. Unfortunately, stable distributions generally have no closed-form solution for their CDF or PDF so a maximum likelihood estimator analogous to the one used here cannot be derived. The variance of stable distributions other than the normal is infinite, rendering the method of moments also invalid.

The method I use in this paper is to instead derive underestimation factors using Monte Carlo simulation, similar to the analysis done to adjust for small sample bias. To do this the distribution of the systematic and idiosyncratic risk factors must be estimated. This is where state space modelling is particularly useful since it is designed to estimate latent variables.

Having estimated the underlying state process, I fit a stable distribution to the innovations associated with the systematic risk factor (ie Equation 17). These determine the unconditional distribution of the systematic risk factor.

I use the program STBLFIT (see Veillette (2012)) to fit a stable distribution to the state innovations derived from estimating the AR(2) state space model in Section 3.2 to the US and UK samples. The state innovations in the US sample follow approximately a normal distribution, whereas those of the UK sample exhibit fatter tails. The fitted stable distribution is $S_{UK}(1.77, -0.8, 0.64, -0.11)$. So for the US, the ASRF model appears to be an appropriate framework, whereas for the UK the non-normality of the mortgage frailty factor has to be taken into account.

Table 3: AVC underestimation factors based on G-ASRF process for UK sample

	MLE	SSM	Beta	MM
Mean	0.81	0.96	0.64	0.81
Median	0.66	0.76	0.38	0.31
Interquartile range	0.44	0.57	0.18	0.13
95% range	2.40	2.48	3.07	6.06
St. Dev.	1.1	1.47	2.03	3.01

Notes: Ratio of true AVC to estimated AVC. Based on 1,000 simulations for UK-type sample (length: 48 years; PD = 0.6%). The underlying model is the G-ASRF model. The distribution is S(1.77,-0.8,0.64,-0.11). The simulated AVC parameter is capped at 16% to avoid unrealistically large default rates (ie, more than half the population of mortgages defaulting).
MLE = Maximum likelihood estimator based on ASRF model.
SSM = State space model estimator.
Beta = Fitted beta distribution estimator.
MM = Method of moments estimator.

I use these estimates for the UK for the distribution of the systematic risk factor in the G-ASRF model. For the idiosyncratic risk factor I must assume the same α parameter as for the systematic risk factor as described above. For the remaining parameters I choose the same value as that for the systematic risk factor, keeping the symmetry between the two risk factors of the normal ASRF.

Assuming a range of possible AVC parameter of 0 - 50% for the Monte Carlo simulation, as above, is unrealistic: given the severe shocks in the G-ASRF an AVC value of 50% would imply virtually all mortgages would default together at the 99.9% confidence level. In addition to being unrealistic, it is also not consistent with the approach taken in the small sample simulation in Section 4.1. An upper limit of 16% results in a stressed default scenario of roughly half of all mortgages defaulting, which is a severe shock.

Table 3 summarises the results. Generally, the MLE and SSM methods perform best as their median estimate is the closest to the true AVC while exhibiting the least variability. Note that these underestimation factors will also capture any bias caused by small samples in addition to the mis-specification of the shock distributions.

I perform similar regression analysis to Section 4.1 to adjust for small sample bias. Table 4 shows that the MLE and SSM methods perform best in terms of bias. In contrast to the normal ASRF, all methods exhibit similar levels of variability as shown by the RMSE and R^2. Again, given the performance characteristics, the MLE and SSM estimators are preferred compared to the other two.

Table 4: AVC non-normality adjustment based on G-ASRF process

	MLE	SSM	Beta	MM
Intercept	0.03	0.03	0.03	0.04
Slope	1.04	0.98	1.83	1.89
RMSE	0.025	0.023	0.025	0.022
R2	0.49	0.50	0.53	0.58

Notes: $AVC = \beta_0 + \beta_1 A\hat{V}C$. Ratio of true AVC to estimated AVC. Based on 1,000 simulations for UK-type sample (length: 48 years; PD = 0.6%). The underlying model is the G-ASRF model. The distribution is S(1.77,-0.8,0.64,-0.11). The simulated AVC parameter is capped at 16% to avoid unrealistically large default rates (ie, more than half the population of mortgages defaulting).
MLE = Maximum likelihood estimator based on ASRF model.
SSM = State space model estimator.
Beta = Fitted beta distribution estimator.
MM = Method of moments estimator.

Ideally, the IRB framework would be changed to reflect the potential non-normality identified in this paper. But this may be impractical because any single specification of the stable distribution is unlikely to be appropriate for all countries it would be applied to. Alternatively, it would also require the identification and estimation of the 'global' systematic risk factor, which may be challenging. For practitioner risk assessment this constraint is, of course, not present. So the G-ASRF framework can be used to estimate economic capital for specific asset classes in specific countries.

A practical way for policy makers to take account of this insight, though, is to assess how far away the capital required by the IRB framework is from that implied by the calibrated G-ASRF model. The IRB AVC parameter can then be increased or decreased to reflect the additional capital, in effect absorbing the mis-specification. I calculate the 99.9% tail loss associated with the US and UK G-ASRF calibrations for a range of AVC values in the G-ASRF model. I then calculate the AVC needed in the normal ASRF model to achieve the same coverage against tail losses. Figure 4 shows the mapping from G-ASRF to ASRF AVC parameters. For example, the loss absorbency provided in a G-ASRF model with an AVC parameter of 8% is equivalent to roughly that in an ASRF model with a 16% AVC parameter.

Figure 4: Mapping from G-ASRF AVC to equivalent ASRF (UK calibration)

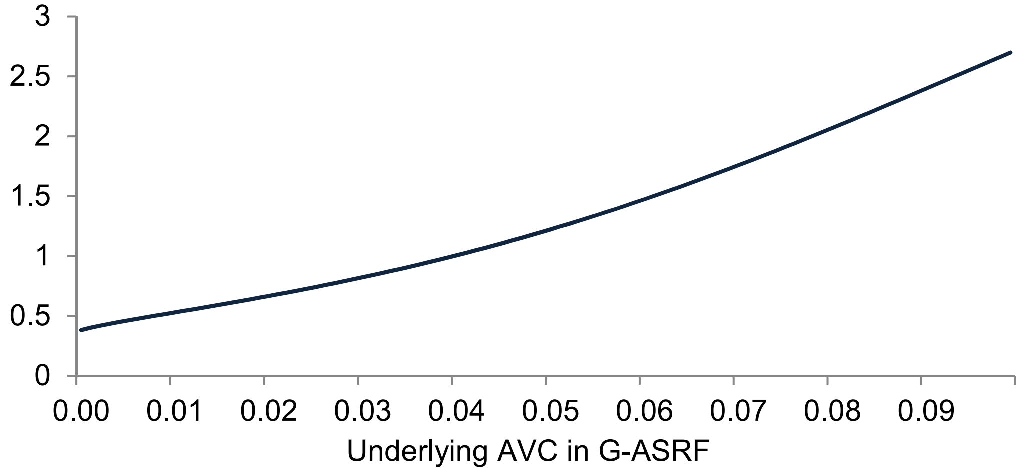

Notes: Based on G-ASRF with $s_{UK}(1.77, -0.8, 0.64, -0.11)$ and PD = 0.6%.

4.3 Bias from mixed-correlation portfolios

The Basel Committee assumes that the AVC parameter is constant across mortgages. This is different to other asset classes, such as corporates, where the AVC parameter is assumed to vary in the probability of default. This section analyses how biased AVC estimates might be when a single AVC parameter is fitted to a portfolio consisting of mortgages with varying asset value correlations.

The economic argument the Basel Committee uses to justify varying the AVC parameter is intuitively as follows: safer assets tend to default only in a very bad state of the world, and not default at all in all other states. This implies a high correlation with the systemic risk factor. Riskier assets, on the other hand, always default to some degree. Though there may be more defaults in a bad state of the world, the relative increase in defaults in risky assets due to a systematic stress event would be expected to be lower than for safer assets. In short: risky asset default is expected to be more influenced by idiosyncratic risk than safer assets and therefore riskier assets have lower asset value correlation.

The evidence on this assumption regarding corporate defaults is ambiguous. Whereas Lopez (2004) and Das (2007) find evidence supporting the hypothesis, Vozzella and Gabbi (2010), Dietsch and Petey (2002) and Düllmann and Scheule (2003) find the opposite: high-risk assets have higher asset value correlations than low risk assets.

Table 5: AVC underestimation factors from mixed-correlation portfolios

	AVC decreases in PD				AVC increases in PD			
	MLE	**SSM**	**Beta**	**MM**	**MLE**	**SSM**	**Beta**	**MM**
Mean	0.70	0.71	0.49	0.76	0.98	1.05	0.54	0.92
Median	0.70	0.68	0.49	0.7	0.97	1.04	0.54	0.87
Coeff. of Variation	0.28	0.80	0.2	0.39	0.2	0.22	0.18	0.32

Notes: Based on 1,000 simulations each for the increasing and the decreasing AVC simulations. The portfolio consists of equal amounts of a safer mortgage with PD = 0.2% and a riskier mortgage with PD = 2%. The length of all simulations is 50 years.
MLE = Maximum likelihood estimator based on ASRF model.
SSM = State space model estimator.
Beta = Fitted beta distribution estimator.
MM = Method of moments estimator.

The granularity of the mortgage data available does not allow me to test the hypothesis directly. Instead, I run Monte Carlo analyses where portfolios are composed of two types of mortgages: low and high risk. In the first analysis I set the AVC parameter to be higher for the safer asset, following the Basel Committee's argument. In the second simulation I assume that higher risk assets have higher AVC parameters as found in parts of the literature. The results are summarised in Table 5.

The Monte Carlo analysis suggests that the composition of the mortgage portfolio can affect the performance of the AVC estimators. If the AVC parameter increases with risk then the estimators perform adequately: the single AVC estimate results in a tail loss estimate close to the real tail loss of the mixed portfolio. But if asset value correlation is higher for safer assets than riskier assets the stressed default rate may be underestimated significantly.

The explanation for this is related to that given for small sample bias: a high level of asset value correlation for safe mortgages means that the default rate of those assets is very low most of the time, with very rare spikes. So in the aggregate portfolio, the default rates in normal times are somewhat high because of the low asset value correlation of the riskier loan. As such, the default process - for short time series - often resembles that of a low-AVC process. Indeed, increasing the sample size to 100 years eliminates the performance issue. Further research is needed to ascertain to what extent this bias is of practical concern given the lack of granularity in the available data.

5 Data and results

I use data from the Council of Mortgage Lenders (CML), which covers UK mortgage defaults from 1969 to 2016 (48 years). For the US, I use data from the Federal Reserve Board on annual charge-off rates for residential mortgages from 1991 to 2016 (26 years).

Default in the UK data set is defined as being 6 - 12 months in arrears, which is a less stringent definition than the 90 days past-due definition used in Basel. The choice of time series is motivated mainly by the fact that it is by far the longest time series provided by CML, which is the most crucial aspect in estimating default correlation. It is also motivated by being more comparable to the US time series, which is based on charge-offs, ie when mortgage go off banks' books. This usually occurs later than 90 days past-due (federal regulation only requires a charge-off for loans that are 180 days past-due, see eg Federal Deposit Insurance Corporation (2000)). The data are annualised and published quarterly; I choose the default rate as of Q1 each year.

Charge-off rates are defined as the loss amount, net of recoveries, of defaulted mortgages divided by the total outstanding mortgage balance. In other words, they are the product of the default rate and the loss-given-default (LGD) rate. In order to obtain default rates for the US sample I divide the charge-off rate by LGD rates. There is no corresponding time series for LGD rates so they need to be estimated.

To do so, I distinguish between stressed LGD and non-stressed LGD. As discussed in Qi and Yang (2009), LGD rates tend to be higher during stress as house prices are falling. This means the value of the collateral is lower which, all else equal, should imply greater losses on defaulted loans. Assuming that LGD is higher during a downturn is a realistic assumption, and in fact required by the Basel framework.

I define the crisis period to be between 2008 and 2012, which spans the period of falling house prices according to the Federal Housing Finance Agency's House Price Index. For this period I use the average LGD values five large US mortgage lenders use in their IRB models. According to Basel this LGD is meant to represent a downturn LGD so should reflect the banks' experience during the recent crisis. The data are taken from Pillar 3 reports for Wells Fargo, JPMorgan Chase, Citi, Bank of America and US Bank as of end-2015. The downturn LGD derived this way is 46.6%. I assume a non-downturn LGD of 30.5% as estimated in Qi and Yang (2009) for the period from 1990 to 2003. I apply this LGD to charge-off rates between 1991 and 2007, as well as 2013 and 2016. The annual default rates for the UK and US are shown in Figure 5. The average default rate for the UK is 0.6%; it is 1.21% for the US.

Figure 5: Annual default rates on residential mortgages

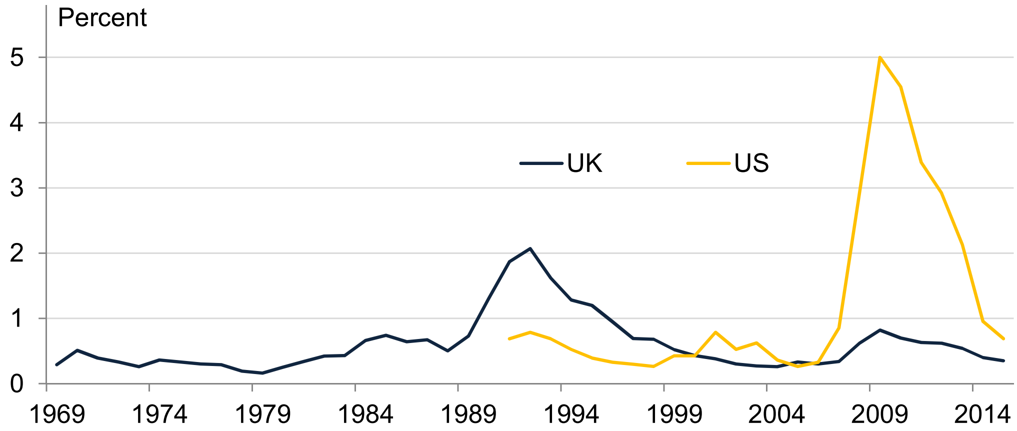

Sources: Council of Mortgage Lenders, Federal Reserve, Federal Housing Finance Agency, Banks' Pillar 3 reports

5.1 Estimated AVC parameters

The results of the empirical analysis are summarised in Table 6. I adjust the estimated AVC parameters and their associated confidence intervals using the regression model in Table 2. Given the non-normality of the UK model, I separately adjust the UK results using the model in Table 4. Finally, I adjust the estimated UK AVC parameters assuming that the normal ASRF model continues to be used (making it essentially absorb the mis-specification of assuming normally distributed shocks when they are non-normal).

Given the poor performance of the beta and relatively poor performance of the method of moments estimators in terms of bias and variability I put most weight on the MLE and SSM estimators. I present the results for the less preferred estimators for completeness and to facilitate comparison with previous estimates.

The unadjusted estimates across all methods are generally comparable to those in the previous literature. All unadjusted estimates are below the Basel AVC calibration of 15%, with the estimated AVC parameter being consistently higher for the US than for the UK. These results hold once small sample adjustments are made. For the US, this suggests that we cannot reject the hypothesis that the Basel AVC parameter is appropriate.

For the UK, the adjustments for non-normality suggest a point estimate for the AVC that is slightly higher than in the US (focusing on the preferred methods) but still below the Basel levels. The results from this analysis considerably exceed those in previous studies, and the Basel parameter is within the 95% confidence interval. As a result, I cannot reject the hypothesis that the Basel calibration is appropriate for the UK.

Table 6: Estimated AVC parameters (Basel calibration: AVC = 15%)

In percent	MLE	SSM	Beta	MM
US				
Unadjusted	10.6 (5.4 – 15.7)	8.5 (1 – 20.5)	7.1 (4.1 – 20.6)	13.5 (8.5 – 24.5)
Small sample	12.2 (6.2 – 18.2)	10.0 (2 – 22.7)	13.5 (8.1 – 38.4)	20.0 (12.5 – 36)
UK				
Unadjusted	4.1 (2.5 – 5.6)	4.1 (0.9 – 9.1)	3.4 (2.5 – 5.2)	5.3 (4.1 – 7.3)
Small sample	4.6 (2.8 – 6.3)	4.3 (1 – 9.7)	5.5 (3.5 – 9.6)	6.9 (5.3 – 9.5)
Non-normality (G-ASRF)	7.2 (5.6 – 8.9)	7.0 (3.9 – 11.9)	9.2 (7.6 – 12.7)	14.1 (11.7 – 17.8)
(ASRF)	13.0 (7.55 – 20.5)	12.1 (3.8 – 39.5)	22.4 (14.4 – 44.4)	54.1 (37.8 – 74.8)

Notes: Central estimate and 95% confidence interval shown. The ‘non-normality (G-ASRF)’ adjustment reflects the appropriate AVC parameter when using the G-ASRF model. The ‘non-normality (ASRF) row shows the AVC parameter needed in the ASRF model to deliver the same amount of capital as that produced in the ‘Non-normality (G-ASRF)’ row.
MLE = Maximum likelihood estimator based on ASRF model.
SSM = State space model estimator.
Beta = Fitted beta distribution estimator.
MM = Method of moments estimator.

Note that the somewhat higher AVC parameter for the UK compared to the US is not driven by higher asset value correlation as such. Instead, it is a result of using a normal ASRF model when a G-ASRF model should be used. In contrast, comparing the asset value correlations adjusted for small sample bias suggests that it is considerably higher for the US.

What might drive the very different asset value correlations for mortgages between the UK and the US? A main feature of the US mortgage market is that banks in a number of states do not have recourse to a borrower’s assets other than the mortgaged property if the borrower defaults. Ghent and Kudlyak (2011) classify 11 states as ’non-recourse’, including California, the most populous state in the US. In addition, the US has a more generous bankruptcy law than other countries. Bankruptcy offers another way out of paying the mortgage: even where the bank would have recourse to other assets bankruptcy would simply extinguish the debt in the first place. In practice, the effect of the more

Figure 6: Estimated AVC parameters across US states

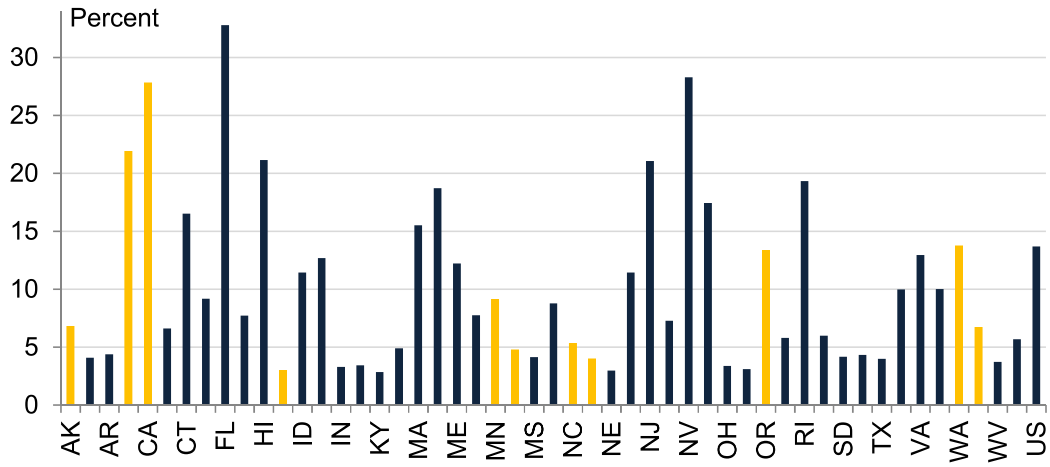

Non-recourse states shown in orange.

lenient bankruptcy code in the US makes the entire US mortgage market closer to a non-recourse market than most other mortgage markets, see Harris (2010). This structural difference in mortgage markets suggest that US mortgages are more correlated with the systemic risk factor than UK mortgages. If the cost of defaulting on your mortgage is low, then any headwinds from the systemic risk factor that affects mortgage borrowers (eg, unemployment, negative equity) translates into default more easily, implying a higher AVC parameter.

To investigate this proposition further I use default series from all 50 US states from Center for Microeconomic Data at the Federal Reserve Bank of New York. The data cover mortgage delinquencies (more than 90 days due) between between 2003 and 2015 (13 years). I estimate the AVC parameters using the MLE method, and adjust for small sample bias. As would be expected, the small sample bias is more severe than in the US and UK samples because of the short time span: the regression intercept is 0.01 and the slope coefficient is 1.22. I use the MLE method because it is more accurate at very small samples, but the results remain qualitatively the same when using the SSM method.

Next, I use the classification in Ghent and Kudlyak (2011) to distinguish between recourse and non-recourse states. The results are shown in Figure 6. The AVC estimate for the entirety of the US is 13.7%, close to that from the longer time series used for the main analysis.

Though the mean AVC estimate of non-recourse states is higher than that of recourse states, this is not statistically significant. So further research is needed to investigate the link between incentives to default (eg, because of non-recourse laws or looser bankruptcy codes) and asset value correlation.

The state-level data allows some additional insights, however: first, there is considerable heterogeneity of asset value correlation across the various states, which is relevant for the US where many banks have local rather than genuinely national lending business. A bank with concentrated exposures in Florida (where the AVC exceeds 30%), say, may be undercapitalised in the Basel framework. Note, however, that in practice the US agencies floor the Basel IRB framework using a US standardised approach. The floored capital requirements tend to be higher than those derived from IRB.

The heterogeneity in asset value correlation across states may also not be a problem for nationwide diversified banks. In Section 4.3 I show that portfolios in which the AVC parameter increases with risk show little estimation bias. The correlation between the AVC parameter and the average PD in the sample is indeed strongly positive (79%), suggesting that there we should expect little bias from this heterogeneity.

In addition, the data allow me to explain a perhaps startling result in this paper: why are the risk factors in the US normally distributed whereas those in the UK are strongly non-normal? Given the heterogeneity observed in estimated AVC parameters across the United States, one explanation may be that the US mortgage frailty factor is the weighted average of several distinct state frailty factors. As such, the sum of realisations of these distributions would approach a normal distribution by the central limit theorem. In the UK, in contrast, the theorem may fail to apply because there is no, or less, aggregation of sub-risk factors driving the UK-wide frailty factor.

5.2 Forecasting default rates

The section sets out the performance of the model in predicting default rates by dividing the sample into a training sample and then assessing the forecast default rates out-of-sample against the realised default rates. As described in Section 2.1, the Basel Committee assumes a global systematic risk factor that affects all asset classes. Though this assumption is needed for the framework to be the same across the world, it is not consistent with the observed heterogeneity across countries. Especially for quintessentially local lending such as mortgages.

In this paper, I interpret the single risk factor in the ASRF framework as country and asset class specific: a US or UK mortgage frailty factor. Though there may still be regional systematic risk factors my interpretation of the systematic risk factor appears more plausible than the Basel Committee's, at least for making national assessments of default rates. Forecasting this factor, and modelling its relationship with default rates, is useful for policy makers for macro-stress testing purposes and macro-economic forecasts.

Table 7: Fitted state-space model

State process	**US**	**UK**
γ_1 (First lag)	1.35 (1.00 – 1.71)	1.25 (0.99 – 1.51)
γ_2 (Second lag)	-0.56 (-0.87 - -0.26)	-0.39 (-0.66 – -0.11)
λ	-4.95 (-5.72 - -4.18)	-5.40 (-5.97 – -4.84)
β	-0.88 (-1.46 - -0.29)	-0.59 (-0.91 – -0.28)
Tests for state innovations (p-values)		
H0: No autocorrelation	0.81	0.95
H0: Homoskedasticity	0.22	0.33
H0: Normal distribution	0.47	0.035

Notes: Central estimate and 95% confidence interval shown.
Test for autocorrelation is Ljung-Box Q test Ljung and Box (1978); test for homoskedasticity is Engel's ARCH test Engle (1982); test for normality is Lilliefors test Lilliefors (1967).

Table 7 shows the fitted models for the UK and the US (Equations 12 to 17), as well as certain test statistics. The coefficient on both α_{t-1} and α_{t-2} are statistically and economically significantly different from zero. I cannot reject the null hypotheses that the state innovations (the residuals from the model) show no autocorrelation and are homoskedastic for both the UK and US, indicating a well-specified model. As discussed above, I can reject the hypothesis of normally distributed state innovations for the UK but not for the US.

Generally, the models perform very well in out-of-sample prediction. To show this I train the model on a part of the empirical default rate and then let it forecast the next period default rate, as well as 90% prediction intervals based on the estimated distribution of the state innovations. These predictions intervals have a similar interpretation as value-at-risk: a 90% interval suggests that we can expect ten instances where the realised default rate lies outside the prediction interval in one hundred forecasts. Figure 7 shows the results.

Figure 7: US and UK default forecast performance

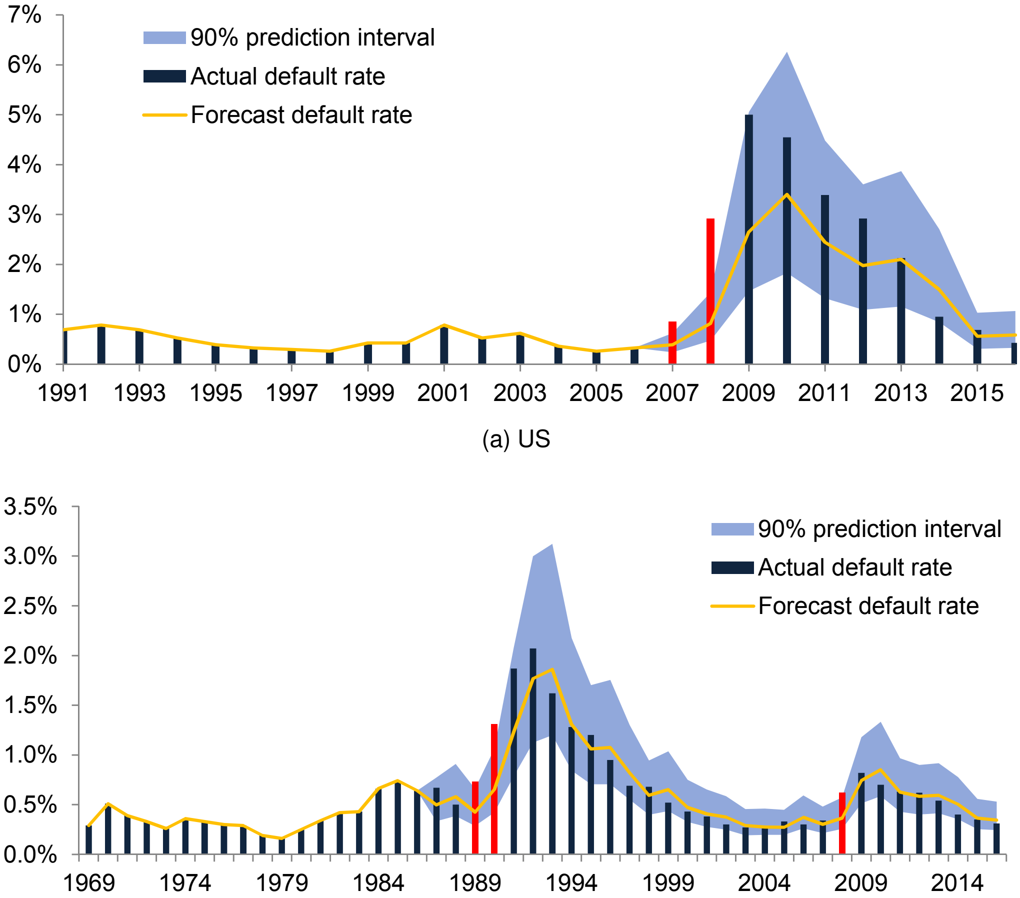

(b) UK

For the US I start the out-of-sample assessment in 2007, which means that there are ten out-of-sample predictions. The model is initially trained on the observations from 1991 to 2006 for the 2007 forecast; the 2008 forecast is then be based on a model trained on observations from 1991 to 2007 and so on. For 2007 and 2008 the realised default rate is outside the confidence interval, in the case of 2008 substantially so. The model adapted reasonably quickly and had no further violations.

Overall, there are two violations of the prediction interval when one violation is expected - given the very small sample this is a good performance. That said, increasing the confidence level further to 99% does not prevent these violations, suggesting that the model cannot overcome the limitations of the short and benign default time series it is trained on just before 2007.

The model performs very well for the UK. Here, I start the out-of-sample test in 1987, giving 30 years of out-of-sample observations. Three realised default rates fall outside the prediction interval, which is exactly the expected number. Moreover, in contrast to the US simulation, increasing the confidence level reduces the violations appreciably: increasing it to 95% results in two violations (1.5 would be expected) and increasing it to 99% results in zero violations (0.3 would be expected). In addition to the good value-at-risk type performance of the model, it also predicts next period's default rate very well in 'normal' times.

Given the good performance in forecasting the default rate and estimating associated prediction intervals the model may be useful for regulators and risk managers for stress testing purposes, or macroprudential assessments of the economy. The performance of the model is perhaps surprising because it is very parsimonious: it does not include macroeconomic or other variables, as for example Koopman et al. (2011). Only the default time series. Further research would be needed to assess whether the performance could be further improved by adding more variables and whether this would be justified with the associated increase in complexity and risk of overfitting.

6 Conclusion

The paper estimates the Basel default correlation parameter for UK and US mortgages, which is a key driver in the Basel credit risk framework. Its main contribution is to adjust for biases that have received little attention in past estimates of the AVC parameter. The paper also provides a model that can be used to forecast default rates and associated prediction intervals.

Estimation biases surrounding default correlation can be significant: for example, small sample bias may underestimate the AVC parameter by 10 - 15% in my analysis. More significant is the adjustment for non-Gaussian risk factors, which may increase initial estimates by a factor of three. This highlights the importance of taking into account small sample and (where relevant) mis-specification bias when assessing credit risk.

The results also suggest that the Basel AVC parameter of 15% is appropriate for the UK and US mortgage market, in contrast to previous studies which tend to suggest the parameter has been overestimated. That said, given the paucity of data available there are considerable error bands around these estimates.

For example, the analysis suggests that the risk factors in the US mortgage market are Gaussian whereas those in the UK market are not. There may be reasons for this related to the very different size and make-up of the two countries. But given that the US sample length is fairly short it might be that very severe shocks have not been observed yet. If the US risk factors follow a distribution similar to that of the UK, then the Basel AVC parameter would likely be too low for the US.

The opposite is true for the UK: if the risk factors are Gaussian then the Basel parameter would be considerably too high. But given that the UK time series is fairly long (48 years) and that both normality tests and fitting a stable distribution suggests that the risk factors are non-normal, this may be less likely than the previous scenario.

A limitation of the analysis generally is that it does not distinguish between the asset value correlation of mortgages within a mortgage market. Taking the characteristics of a market, eg non-recourse versus recourse, as given, there is an economic argument that suggests low-risk mortgages should be more correlated with the systemic risk factor than high-risk mortgages. Safe assets tend to default only in a bad state of the world, and not just randomly regardless of the state of the world.

This implies a high correlation with the systemic risk factor. Risky assets, on the other hand, always default to some degree; they may also default more in a bad state of the world, but default behaviour appears to be driven more by idiosyncratic risk than for safe assets. Indeed, the IRB function assumes this inverse relationship between the correlation parameter and PD for wholesale exposure (though not for mortgages).

If this is the case, the average AVC would underestimate tail losses. There is very tentative evidence that suggests that this is not a problem for the US sample, where there is a positive - not negative - relationship between average default rates and the AVC parameter. That said, more granular data and further research are needed to establish this conclusively.

The state space model used to estimate default correlations can also be used to forecast default rates. I show in out-of-sample analysis that the model performs very well. In particular, the model generates prediction intervals that can be useful to estimate the risk of sharply rising default rates. This may be useful in stress testing and macroprudential policy.

References

Basel Committee on Banking Supervision (2004). An Explanatory Note on the Basel II IRB Risk Weight Functions.

Calem, P., J. Follain, et al. (2003). The asset-correlation parameter in Basel II for mortgages on single-family residences. *Board of Governors of the Federal Reserve System*, 27.

Campbell, J. Y. and J. F. Cocco (2015). A model of mortgage default. *The Journal of Finance 70*(4), 1495–1554.

Das, S. R. (2007). Basel II: correlation related issues. *Journal of Financial Services Research 32*(1-2), 17–38.

Dietsch, M. and J. Petey (2002). The credit risk in SME loans portfolios: Modeling issues, pricing, and capital requirements. *Journal of Banking & Finance 26*(2), 303–322.

Duffie, D., A. Eckner, G. Horel, and L. Saita (2009). Frailty correlated default. *The Journal of Finance 64*(5), 2089–2123.

Düllmann, K. and H. Scheule (2003). Determinants of the asset correlations of German corporations and implications for regulatory capital. *Deutsches Bundesbank*.

Durbin, J. and S. Koopman (2012). *Time Series Analysis by State Space Methods: Second Edition*. Oxford Statistical Science Series. OUP Oxford.

Durbin, J. and S. J. Koopman (1997). Monte Carlo maximum likelihood estimation for non-Gaussian state space models. *Biometrika 84*(3), 669–684.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica: Journal of the Econometric Society*, 987–1007.

European Central Bank (2016, October). Statistical bulletin.

Federal Deposit Insurance Corporation (2000). Uniform retail credit classification and account management policy.

Federal Reserve Board (2016, October). Assets and liabilities of commercial banks in the United States (weekly) - h.8.

Fitch Ratings (2008). Basel III correlation values: An empirical analysis of EL, UL and the IRB model.

Fitch Ratings (2011). Basel III ccorrelation: An empirical analysis reflecting the financial crisis.

Ghent, A. C. and M. Kudlyak (2011). Recourse and residential mortgage default: evidence from US states. *Review of Financial Studies*.

Gianfrancesco, I., D. Curcio, and A. Malinconico (2011). Investigating implied asset correlation and capital requirements: Empirical evidence from the Italian banking system. *Banks and Bank Systems 6*(2), 116–125.

Gomes, M. I. and A. Guillou (2015). Extreme value theory and statistics of univariate extremes: a review. *International Statistical Review 83*(2), 263–292.

Gordy, M. and E. Heitfield (2002). Estimating default correlations from short panels of credit rating performance data. *Unpublished Working Paper*.

Gordy, M. B. (2000). A comparative anatomy of credit risk models. *Journal of Banking & Finance 24*(1), 119–149.

Gordy, M. B. (2003). A risk-factor model foundation for ratings-based bank capital rules. *Journal of Financial Intermediation 12*(3), 199–232.

Gordy, M. B. and E. Lütkebohmert (2013). Granularity adjustment for regulatory capital assessment. *International Journal of Central Banking 9*(3), 33–70.

Harris, R. (2010). Recourse and non-recourse mortgages: Foreclosure, bankruptcy, policy. *Policy (April 10, 2010)*.

Kelly, R. and T. O'Malley (2016). The good, the bad and the impaired: A credit risk model of the Irish mortgage market. *Journal of Financial Stability 22*, 1 – 9.

Koopman, S. J. and A. Lucas (2008). A non-gaussian panel time series model for estimating and decomposing default risk. *Journal of Business & Economic Statistics 26*(4), 510–525.

Koopman, S. J., A. Lucas, and B. Schwaab (2011). Modeling frailty-correlated defaults using many macroeconomic covariates. *Journal of Econometrics 162*(2), 312–325.

Koopman, S. J., A. Lucas, and B. Schwaab (2012). Dynamic factor models with macro, frailty, and industry effects for US default counts: the credit crisis of 2008. *Journal of Business & Economic Statistics 30*(4), 521–532.

Lilliefors, H. W. (1967). On the Kolmogorov-Smirnov test for normality with mean and variance unknown. *Journal of the American statistical Association 62*(318), 399–402.

Ljung, G. M. and G. E. Box (1978). On a measure of lack of fit in time series models. *Biometrika*, 297–303.

Lopez, J. A. (2004). The empirical relationship between average asset correlation, firm probability of default, and asset size. *Journal of Financial Intermediation 13*(2), 265–283.

Nolan, J. P. (2015). *Stable Distributions - Models for Heavy Tailed Data*. Boston: Birkhauser. In progress, Chapter 1 online at academic2.american.edu/~jpnolan.

Qi, M. and X. Yang (2009). Loss given default of high loan-to-value residential mortgages. *Journal of Banking & Finance 33*(5), 788–799.

Rocco, M. (2014). Extreme value theory in finance: A survey. *Journal of Economic Surveys 28*(1), 82–108.

Vasicek, O. (1997). The loan loss distribution. *KMV Technical Document*.

Veillette, M. (2012). STBLFIT. http://math.bu.edu/people/mveillet/html/alphastablepub.html.

Vozzella, P. and G. Gabbi (2010). Default and asset correlation: An empirical study for Italian SMEs.