![BANK OF ENGLAND]

An investigation into the procyclicality of risk-based initial margin models

David Murphy, Michalis Vasios and Nick Vause

An investigation into the procyclicality of risk-based initial margin models

David Murphy, Michalis Vasios and Nick Vause

david.murphy@bankofengland.co.uk
David Murphy, Bank of England, Threadneedle Street, London, EC2R8AH

michalis.vasios@bankofengland.co.uk
Michalis Vasios, Bank of England, Threadneedle Street, London, EC2R8AH

nick.vause@bankofengland.co.uk
Nick Vause, Bank of England, Threadneedle Street, London, EC2R8AH

Contents

An investigation into the procyclicality of risk-based initial margin models

David Murphy, Michalis Vasios and Nick Vause

The initial margin requirements for a portfolio of derivatives are typically calculated using a risk model. Common risk models are procyclical: margin requirements for the same portfolio are higher in times of market stress and lower in calm markets. This procyclicality can cause liquidity stress whereby parties posting margin have to find additional liquid assets, often at just the times when it is most difficult for them to do so. Hence regulation has recognised that, subject to being adequately risk sensitive, margin models should not be 'overly' procyclical. There is, however, no standard definition of procyclicality.

This paper proposes two types of quantitative measure of procyclicality: one that examines margin variation across the cycle and one that focuses on short-term margin increases. It then studies, using historical and simulated data, various margin models with regard to both their risk sensitivity and the proposed procyclicality measures. It finds that models which pass common risk sensitivity tests can have very different levels of procyclicality.

The paper recommends that CCPs and major dealers should disclose the procyclicality properties of their margin models, perhaps by reporting the proposed procyclicality measures. This would help derivatives users to anticipate potential margin calls and ensure they have adequate holdings of or access to liquid assets.

1   Introduction

Risk models should be accurate: we model risk in order to estimate potential losses, so it is important that those estimates are robust so that we are not misled about the real level of risk. Supervisors have recognised this since the early days of risk modelling. Thus, for instance, the Market Risk Amendment to Basel I (Basel Committee on Banking Supervision (BCBS) (1996)) requires that risk estimates from Value-at-Risk (VaR) models in the trading book meet certain standards: if these are not met, then additional capital penalties are imposed, or, in extremis, the model is de-recognised.

Risks models are used by both central counterparties (CCPs) and bilateral counterparties to estimate the margin requirements of portfolios of financial instruments. Here an additional concern arises: *procyclicality*. Broadly procyclicality refers to the tendency of any financial variable to move with the cycle. This is an undesirable property when the variable acts to intensify financial stress (Financial Stability Forum (2009)). For instance, if bank regulatory capital requirements are too procyclical, their increase in an economic downturn can depress lending activity and hence make economic recovery more difficult (Kashyap and Stein (2004)). As several authors have noted (Brunnermeier and Pedersen (2009), Heller and Vause (2012)), margins often behave in this way too. That is they tend to increase in periods of crisis, causing investors to face funding and market liquidity risk synchronously, which can be destabilising. The procyclicality of margin requirements refers to this tendency of margin requirements to rise in periods of market stress.

These two concerns come together when a risk model is used to estimate the margin required on a portfolio of financial instruments. We want risk sensitivity, so that margin estimates increase for a fixed portfolio as the market becomes riskier; but we do not want too much procyclicality. The motivation is the mitigation of funding liquidity risk. If initial margin increases substantially, the requirement to post margin on a timely basis may pose a substantial liquidity burden on the poster, often just at the time when they are least able to bear it. As requirements to post margin have been introduced into the post-crisis financial reforms, supervisors have recognised this. Thus for instance the relevant European Regulation (European Union (2012)) states that for CCPs:

> 'Margin calls and haircuts on collateral may have procyclical effects. CCPs, competent authorities and ESMA [the European Securities and Markets Authority] should therefore adopt measures to prevent and control possible procyclical effects in risk-management practices adopted by CCPs, to the extent that a CCP's soundness and financial security is not negatively affected.'

This text illustrates the key issue: there is potentially a trade-off between risk sensitivity and procyclicality. When markets become more volatile, risk is higher, and hence margin requirements *should* be higher. But it is undesirable for margin models to *overreact* to changing conditions.

The immediate difficulty which arises is in knowing what constitutes an overreaction, and hence when to reject a proposed model on the grounds of high procyclicality. First, though, we turn to the policy context.

Policy context 1: risk sensitivity

Clearly any model which is going to be relied upon to make estimates which are important for financial stability must be accurate. This holds for models which are used to estimate regulatory capital requirements and for models which are used to estimate margin requirements for important counterparties. There are various ways to achieve this goal, of which back testing model risk estimates against actual or hypothetical outcomes is an important one.

An early benchmark here was the original Basel market risk standard. This stated that a risk model which purported to calculate a 99% VaR estimate enters the 'red zone', where there is 'an automatic presumption that a problem exists', when there are ten or more back test exceptions (ie days when there was a loss of more than the risk estimate) over a one-year (*c.* 250-day) period (BCBS (1996)). Clearly we would expect on average that a model that reached a 99% confidence interval would display on average two or three exceptions in 250 days (a year of business days, roughly), so the idea that displaying ten or more indicates a problem is intuitive.

As knowledge of the limitations of 'simple' back tests has grown, so standards for an acceptable model testing regime have increased. Campbell (2005) discusses some of the tools which are available here. Some of these tools are more discriminating than simply comparing the claimed safety standard of a model (eg 99%) to the standard actually achieved. The simple test can however still be insightful, so we confine our estimate of risk sensitivity to its results.

It is important to note here that backstops to model-based risk estimates also have a role to play. For instance, a good stress testing regime provides model-independent estimates of possible losses in extreme but plausible conditions. This reduces the risk that an in-reality inaccurate model which happens to pass its back test leads to imprudence.[1] Still, within the gamut of 'acceptably risk sensitive' models, there are many choices, and these can exhibit quite different behaviour as markets move.

---

(1) Thus, for instance, the requirement to use stress tests to size default funds is an important backstop to CCP initial margin models.

Policy context 2: procyclicality

The broad concept of procyclicality was defined by the Committee on the Global Financial System (CGFS) (2010) as

> 'mutually reinforcing interactions between the financial and real sectors of the economy that tend to amplify business cycle fluctuations and cause or exacerbate financial instability.'

An important mechanism is the sense in which we use procyclicality: that is, the situation whereby margin requirements fall in 'good' low-volatility markets and rise in 'bad' high-volatility ones.[1] This is a potential issue for both bilateral markets and cleared ones, as both use risk models to calculate margin requirements. The latter are however particularly a concern, first since clearing of many products either is or soon will be mandatory by many parties (BCBS and International Organization of Securities Commissions (IOSCO) (2013)), and second since CCPs' margin calls are unilateral: the CCP makes the call and all its members must meet it.

The Commodity Futures Trading Commission (CFTC) pick up the story here, highlighting how a big margin call from a large US CCP (otherwise known as a 'systemically important derivatives clearing organisation' or 'SIDCO') could force its members to act in ways that are potentially destabilising (CTFC (2013)):

> '…in a stressed market where credit is tightening and margin calls are increased, a SIDCO's assessment of additional claims upon its clearing members may well exacerbate already weakened financial markets by potentially forcing clearing members and/or their customers to deleverage in falling asset markets, which will further drive down asset prices and stifle liquidity, or force clearing members to default on their obligations to the SIDCO. This in turn could start a downward spiral which, combined with restricted credit, might lead to additional defaults of clearing members and/or their customers, and would play a significant role in the destabilization of the financial markets.'

In other words, margin calls which cannot be met from liquidity at hand either cause asset sales or defaults: both are destabilising. The example of AIG is relevant here, as inability to meet large margin calls caused this firm to require rescue (Murphy (2009)). This is an extreme example, admittedly, but it does provide a stark illustration of the issue.[2]

These considerations highlight the balance that market participants need to strike between risk sensitivity and procyclicality. As the Principles for Financial Market Infrastructures (Committee on Payment and Settlement Systems and Technical Committee of the International

Organization of Securities Commissions (2012)) put it, an entity

> 'should adopt initial margin models and parameters that are risk-based…[and these should] to the extent practicable and prudent, limit the need for destabilising, procyclical changes.'

Other context

Concerns about the procyclicality of margin appear in several other papers. For instance, Brunnermeier and Pedersen (2009) study the phenomenon in an economy with investors that face funding liquidity risk due to the risk of future margin fluctuations or losses on existing positions. They show that in this economy there is a clear link between funding and market liquidity which, under certain conditions, can destabilise markets through liquidity spirals. A recent paper by Abruzzo and Park (2013) provides supporting evidence. They show that futures margin requirements at the Chicago Mercantile Exchange (CME) rise quickly following volatility spikes, indicating the procyclical nature of margins. The link between margin requirements, market volatility and liquidity is also documented in empirical work by Mayhew, Sarin, and Shastri (1995) and Hedegaard (2011), who examine the impact of margin changes in equity options and commodity futures markets, respectively. They find significant increases in margin requirements when these markets become more volatile.

The policy considerations discussed in the previous two sections and the evidence from the literature above illustrate not just the ubiquity of concern about the procyclicality of margin, but also that several different techniques are available for addressing it. The toolkit includes:

- use of a model for calculating initial margin amounts that is not too procyclical in the first place;
- limiting the impact of procyclicality by a requirement that stressed periods must be used in calibrating the model, thereby reducing the tendency for margin estimates to fall too far;
- 'flooring' the output of the model at some level which explicitly limits how low margin can go; and
- limiting daily margin increases, so that a rapid change from low volatility markets to higher volatility ones only feeds through into margin increases more gradually.

In order to comment on these issues, we need some measures of various aspects of the procyclicality of margin, so we turn to

---

(1) This paper studies *initial* margin requirements: there is also the related problem of changes in the value of positions leading to variation margin calls. This is studied in Financial Stability Forum (2009), which points out that this effect can be substantial compared with the liquid asset buffers of large OTC derivatives dealers.

(2) It can be argued that limited available liquidity is just as much of a problem as large demands on it. Certainly more liquid firms will be better able to meet large collateral calls, *ceteris paribus*, than less liquid ones. We focus on liquidity demands in this paper, deferring an analysis of how they can be met to further work.

this question next. Section 3 discusses the proposed measures, then Section 4 sets out the framework in which we study them. Section 5 presents a number of different initial margin models. Each model's risk sensitivity and procyclicality is measured using both historical asset prices and randomly generated returns from a number of theoretical distributions: this throws some light onto the risk-sensitivity/procyclicality trade-off, as well as providing concrete measures of procyclicality in particular situations. It turns out that model behaviour is dependent both on the model's parameterisation and on the period studied: these effects are considered respectively in Sections 6 and 7. Finally Section 8 presents some tentative policy conclusions and suggestions for further work.

2   Measures of procyclicality

Risk estimates are usually a function of the market conditions used to calibrate the model which provides the estimate. Thus in particular, the initial margin required for a fixed portfolio depends on market conditions too. **Chart 1** illustrates the issue.

Chart 1 A financial asset return series and estimated margin for a position in that asset
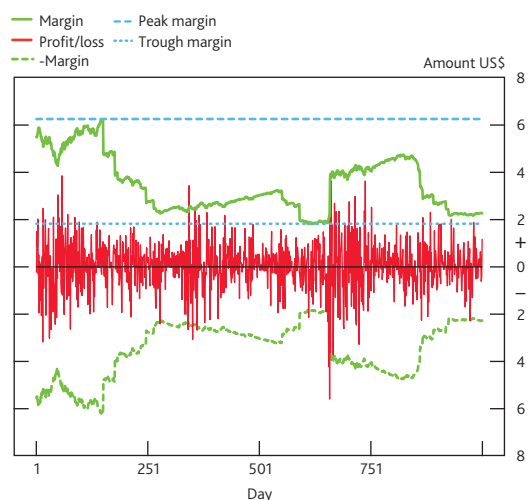
Here the red series is the daily log returns of a position in the S&P 500 total return index over a four-year period from 2009 to 2012 with an initial value of US$100, and the upper green line is the margin required for that position calculated by a popular initial margin model.[1] The reflection of this line is given to illustrate back test exceptions more clearly: recall that these occur when there is a loss bigger than margin, so they correspond to the occasions when the red line falls below the lower green line.

Peak-to-trough measure
The margin required over the illustrated period varies from a peak of about US$6.2 to a trough of about US$1.8 (as illustrated by the blue dashed lines). This range reflects the

variability of margin requirements for the same portfolio across the cycle, and thus is a measure of procyclicality:

Peak-to-trough measure. The peak-to-trough procyclicality of a margin model is the ratio of the maximum initial margin required for a constant portfolio to the minimum margin required over a fixed observation period.

If the period used is long enough to cover a boom and a recession then the peak-to-trough measure captures long-term procyclicality. However, if short-term liquidity drains caused by margin are the primary concern, then we need to examine changes in margin over a shorter horizon.

Measures of short-term margin rises
Chart 2 zooms in on the period around day 658 of our example to illustrate the issue.

Chart 2 Asset prices and margin requirements for a period generating a large one-day margin call
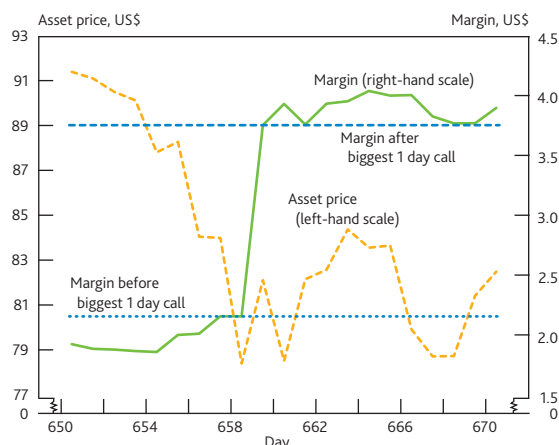
Here the green line (right-hand scale) is the margin required, and the dotted orange line (left-hand scale) is the price of the underlying asset. The asset has heightened volatility from day 655 or so, causing the model's estimate of volatility to increase. This behaviour causes the margin requirement to increase. For instance, the margin estimate goes from US$2.2 on day 658 to US$3.75 the following day. This is the largest one-day increase in the data window: the blue dashed lines in the chart illustrate 'before' and 'after' margin levels.

The biggest five-day increase in margin occurs over the same period and for the same reason: this is an increase from US$2.0 on day 655 to US$3.9 on day 660.

(1) That is, if the asset price at time t is denoted by S_t, we examine the log returns, $\log(S_t/S_{t-1})$, and the 99th percentile of profit/loss (P/L) estimated using a historical simulation model with a 200-day data window for a portfolio comprising 100 units of the asset. The margin is the 99% VaR, ie minus one times the 99th percentile worst P/L estimate.

This discussion suggests:

n-day procyclicality measure. The *n*-day procyclicality of a margin model is the largest increase in margin over an *n*-day period for a constant typical portfolio over a fixed observation period.

Thus for instance the 1, 5 and 30-day procyclicality measures with a ten-year observation period are indicative of the additional amounts margin posters might be required to fund in stressed conditions across the cycle (assuming, of course, that the observation period includes stress events).[1]

Increases from a high base

The two measures introduced so far are, strictly, measures of cyclicality rather than procyclicality:[2] they measure long and short-term margin variability, respectively, but they do not focus exclusively on margin increases in periods of stress. For this we need to examine the most difficult margin increases: those which occur when conditions are already stressed, and thus which might cause the forced sales or defaults that the CFTC quotation in Section 1 highlights as destabilising.

One way to do this is to look at the volatility of volatility ('vol of vol'). For our data series, 90-day volatility averages 19%, but with a standard deviation of 7%. If we restrict analysis to periods when 90-day volatility was 26% or more (ie one standard deviation or more over the average), then the *n*-day measures over these periods reflect margin increases which are more likely to cause stress. Thus we suggest:

n-day stressed procyclicality measure. The *n*-day procyclicality of a margin model is the largest increase in margin over an *n*-day period for a constant portfolio over a fixed observation period, restricted to those sub-periods where volatility is elevated.

Obviously a clear definition of 'elevated' is required: our suggestion of 90-day sample volatility in excess of the average long-term level plus one standard deviation is one possible choice.

Chart 3 illustrates the phenomenon: the largest 30-day call during a stressed period occurs closer to the start of the dataset. Here US$1.2 of additional margin is required from day 49 to day 77.[3] The 30-day margin call occurs from the relatively high base of US$4.3 of margin already posted reflecting the volatile market in which these calls occur. It is margin calls like this that potentially cause the most liquidity stress on market participants and hence may be a significant risk to financial stability.

This chart also illustrates the phenomenon of margin calls caused by price changes rather than volatility estimates. The margin required for our portfolio is typically roughly

Chart 3 Margin requirements generating a large 30-day margin call during a stressed period
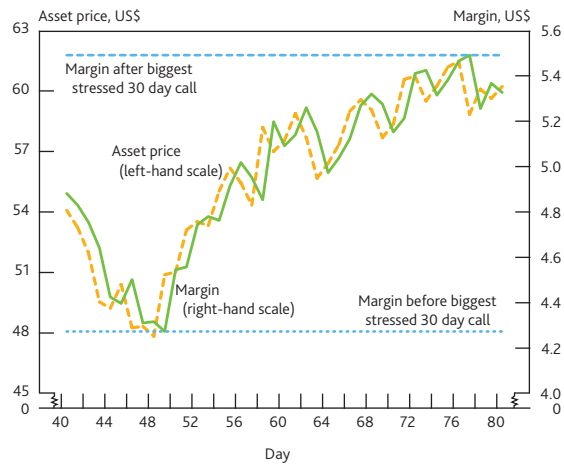
proportional to some measure of volatility times the asset price, so rising asset prices cause rising margin requirements if volatility does not change. The increase in **Chart 3** is almost entirely derived from this price effect, rather than from changes in volatility. There is therefore a case for studying margin calls as a percentage of the asset value to isolate the impact of non-price factors on margin changes.

3   Discussion of the measures

The next section applies the suggested measures of procyclicality to particular margin models. First, though, we say a little more about the measures themselves and the context in which they are used.

The measures

Our three classes of measure are motivated by different aspects of procyclicality.

• The peak-to-trough measure captures the variation of margin over the observation period. Thus for instance if one were interested in the total margin that a large CCP were to require from its clearing members — perhaps because one was interested in the sufficiency of collateral available to meet those requirements (Committee on the Global Financial System (2013)) — then this measure would serve to estimate how much higher requirements might go.

---

(1) There is a delicate issue as to whether to report margin increases as a percentage of asset value or absolutely. We use asset value processes which do not mean revert, reflecting our study of the S&P500, so we prefer to report margin changes as percentages; for mean-reverting asset classes, such as interest rates, absolute changes may be more informative.

(2) The peak-to-trough metric correctly measures long term changes in margin requirements, but it cannot distinguish whether these changes come from a procyclical or an anticyclical model. This is not a practical concern because all the initial margin models of interest are risk-based. That is, because margin is a function of the underlying risk or volatility, the models in use by leading institutions are by construction procyclical.

(3) This illustrates an important definitional issue: the largest *n*-day call is the largest increase in margin in any *n*-day period, whether it spans the whole period or not. Here the increase occurs over 28 of the 30 days in the period.

• The *n*-day measures capture the amount of extra margin that market participants would need to fund on a short-term basis, and hence they measure an important aspect of the liquidity risk of collateral measures. Thus for instance market participants often estimate their liquidity outflows on a 30-day basis, so the likely worst 30-day increase in margin is a useful measure in this context.[1]

• It might be argued that short-term increases in margin are of most concern from a financial stability perspective when conditions are already stressed. The stressed *n*-day measures capture this.

Models and calibration

It is important that any measurement of the procyclicality of initial margin includes all sources of increase of margin estimates. Thus in particular increases caused by model recalibration must be included. This is important because some commonly used margin models are not significantly procyclical until they are recalibrated, but the overall level of margin demanded can nevertheless be procyclical. The well-known SPAN model (Chicago Mercantile Exchange (2010)) is a good example here: the margin calculated by SPAN for a single unit of an asset is a fixed percentage of its value known as the 'scanning range'. However these scanning ranges are regularly recalibrated, often by reference to another, meta-model.[2] This recalibration changes the margin requirements and hence can introduce procyclicality. We therefore distinguish:

• the intrinsic procyclicality of the margin model itself without recalibration; from
• the meta-model procyclicality, which includes the model recalibration strategy.[3]

Both of these aspects must be included in a useful measure of procyclicality.

A similar consideration applies to risk sensitivity tests: the object of study here too should be the model together with its recalibration strategy.[4] Thus in what follows we use 'model' as shorthand for an initial margin model, its starting parameters, and the recalibration strategy for those parameters.

4　An illustrative framework

Our illustrations thus far have used a simple margin model and historical asset returns. While the use of historical returns gives an insight into the behaviour of the different margin models in real market conditions, it does not allow the assessment of the procyclicality measures across a large number of different scenarios. For the latter, one has to rely on simulation techniques to generate new return series by assuming an underlying data generating process. Hopefully,

the new return series will capture the different market conditions. As the true underlying process is not known, we use various processes that have been extensively used in prior literature.[5]

Return dynamics

Therefore in addition to the historical returns we pick three different theoretical processes, each of which we use repeatedly to generate four years (1,000 days) worth of data. These give us:

i. a historical return series from the S&P 500 total return index, beginning in January 2009;
ii. return series generated by a normally distributed process with a constant volatility;
iii. return series generated by a mixture of two normally distributed processes;[6] and
iv. return series generated by a two-state normally distributed regime switching model.[7]

The three theoretical processes represent different features of the historical data. The constant volatility process captures just the historical data's long-run variance; the mixture process in addition captures the fatness of its tails; while the regime process also captures some information on the clustering of returns. Further details of the processes and their calibration are given in the Appendix.

Theoretical processes are used in addition to historical data partly because some of the procyclicality measures are quite sensitive to the data window. We run each theoretical process 1,000 times and average the procyclicality measures to get a stable estimate.

Margin models

Four different types of margin model representing a range of practice are studied:

• a constant volatility model;
• an exponentially weighted moving average (EWMA) model;
• a historical simulation VaR model; and

---

(1) This is in part motivated by the Basel III definition of 30 days as the stressed period for liquidity coverage ratio purposes.
(2) For instance, the scanning range could be determined by a 99% VaR estimate from a historical simulation (HS) model. Thus here HS VaR is the meta-model.
(3) A SPAN model whose scanning range was set based on a very reactive volatility estimation model would have little intrinsic procyclicality but significant meta-model procyclicality.
(4) This consideration suggests that back testing should similarly examine the results of the model with a specified recalibration strategy. Note though that prescription does not extend to measures to address procyclicality such as floors or the use of stressed calibration. These should be *excluded* from risk sensitivity testing unless they reduce risk sensitivity.
(5) For instance, an exponentially weighted moving average model will be admirably risk sensitive if the underlying asset dynamics are auto-regressive in the same sense: it may however perform markedly less well in other situations.
(6) Log normal mixture processes (McLachlan and Peel (2000)) have been extensively studied as simple extensions to the log normal models which incorporate fatter tails: see the Appendix for more details.
(7) Regime processes (Gray (1996)) represent another generalisation: again, please see the Appendix for more details.

Box 1
Margin models studied

The *constant volatility* model assumes that changes in portfolio value are normally distributed, and that the variance of returns on the underlying asset is constant. The parameters of the normal distribution are obtained from historical data. In our analysis, we only need to estimate the long-term portfolio variance (and no covariances) as our portfolio contains a single asset and the average return has been set to zero. In this model, the margin is given by the 99th percentile of the normal distribution, ie it is given by the asset price times the daily volatility times a constant. The constant is found by applying the inverse normal distribution to the chosen confidence interval.

The *exponentially weighted* moving average model is similar. It also assumes that changes in portfolio value are normally distributed, but it computes the daily variance of returns using an exponentially weighted moving average with a decay factor λ.[1] The decay factor controls how rapidly recent information is incorporated into the model's estimate of volatility: at $\lambda = 1$, the model is never updated, while for smaller values of λ, recent information increasingly quickly updates the volatility estimate (at the cost of giving less weight to past behaviour).

We also examine a variation of the EWMA model: the *floored* version. Here volatility is set at the maximum of the EWMA estimate and the long-run (ten-year) historical volatility.

The *historical simulation* model is based on a non-parametric approach. Here we assume that the distribution of historical returns is a good indicator of the distribution of future returns we will face over the next period. In our analysis we implement this approach using a two-year window of historical returns. Specifically, every day we compute the 99% VaR of a portfolio using the distribution of the previous 500 returns and by selecting the fifth worst loss.

Finally, we examine an extension to the historical simulation technique suggested by Hull and White (1998) and Barone-Adesi, Giannopoulos and Vosper (1999). These *Hull and White approaches* scale historical returns based on short-term volatility: they are also known as filtered historical simulation techniques. In models like this, if current volatility is higher than usual, returns are scaled up; while in a quieter period, they are scaled down. They therefore combine a non-parametric technique with reactivity to current conditions, although they differ in the detail of precisely how the scaling is defined. Our version uses the ratio of 60-day to long-term volatility.

(1) CCPs tend to use larger decay factors than the 0.94 proposed in the original Risk Metrics standard (JP Morgan/Reuters (1996)). Reflecting that, we use a decay (or 'λ') factor of 0.99.

- a historical simulation model with returns scaled based on short-term volatility, in the fashion of Hull and White.

Finally, to give a flavour of the effect of a procyclicality mitigation tool, we also examine the EWMA model with a floor set by the long-run volatility of the historical series. This gives five initial margin calculation methodologies in total which we refer to, respectively, as CV, EWMA, HS, HW and fEWMA. Box 1 gives further details of the models studied.
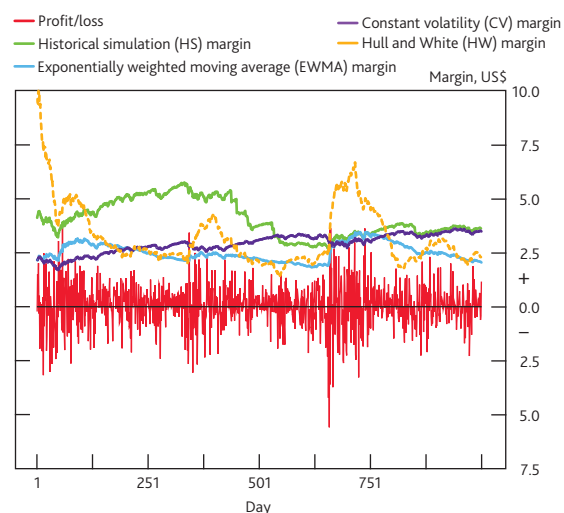
We will proceed as follows.

- Each model will be used to calculate margin on a linear position of starting size 100 in an asset whose price is determined by a given return.

- The models' risk sensitivity will be measured using back test exemptions: we will study how often the loss on the portfolio is bigger than the margin required.

- The procyclicality measures proposed above will also be calculated to give insight both into behaviour of the models themselves and the risk sensitivity/procyclicality trade-off.

5   The performance of margin models

We begin our analysis by illustrating the differing levels of margin required by the different models. **Chart 4** shows the initial margin estimates for the historical returns as calculated by the CV, EWMA, HS and HW models. As we might expect from such contrasting behaviour, the five models will have different properties.

Chart 4 Sample initial margin estimates from four of the models based on historical returns

Legend:
- Profit/loss
- Historical simulation (HS) margin
- Exponentially weighted moving average (EWMA) margin
- Constant volatility (CV) margin
- Hull and White (HW) margin

Measures of risk sensitivity

There is a statistic which measures whether the number of back test exceptions a margin model demonstrates is plausible, given the confidence level it purports to respect. This is the Kupiec proportion of failures (POF) statistic (Kupiec (1995)). The 'more than ten exceptions for a 99% model in a year is bad' standard discussed in Section 1 corresponds to a POF statistic of more than 6.6: the POF statistic increases when a model has fewer or more exceptions than we would expect, with larger POF statistics therefore being worse.[1]

Each of the models we study has some measure of risk sensitivity based on this statistic. **Table A** shows the POF statistics for each of the models based on the historical returns: it can be seen that each of them pass the standard test easily (Campbell (2005)).[2]

Table A Kupiec statistics for the models using historical returns

Data process	Margin model				
	CV	EWMA	fEWMA	HS	HW
Historical	0.4	2.2	0.1	4.7	1.0

Turning from back testing against historical data to back testing against simulated returns, we find that the three theoretical processes pose different challenges to the margin models. The regime-switching process in particular is demanding in that a model might be lulled into low margin estimates by the low-volatility regime, then presented with much higher volatility conditions which it has to adapt to. If the model does not do this quickly enough, there will be many back test exemptions and hence the model will fail to be sufficiently risk sensitive. As illustration of these phenomena, the average POF statistics for the 99% risk estimate over 1,000 runs of each theoretical process are presented in **Table B**.

Table B Average Kupiec POF statistic for the models based on returns from three different processes

Data process	Margin model				
	CV	EWMA	fEWMA	HS	HW
Normal	5.6	0.9	5.6	0.5	0.5
Mixture	2.7	12.6	2.7	0.5	0.7
Regime	5.4	3.6	2.2	2.7	1.1

The poor performance of the EWMA model on the mixture process is evident here: essentially these models can react well to sustained changes from one volatility regime to another, but they cannot react to the mixture process which can switch every day from one volatility to the next. The model smoothes out the two processes in the mixture, and hence has numerous exceptions when the higher volatility process in the mixture is realised.

Peak-to-trough procyclicality

Table C presents the peak-to-trough measure for the five models studied using the historical return series. This illustrates some differences between the models. The simplest model, constant volatility, only demands a little over twice as much margin at the peak as at the trough (which is all a result of price level changes, since the volatility of returns is assumed constant), whereas the more reactive Hull and White model has roughly three times this range, ie six times as much margin at the peak as at the trough. Furthermore, the floor in the fEWMA model somewhat reduces its peak-to-trough range relative to the EWMA model. This pattern is (mostly) repeated for returns generated from the three theoretical processes, as **Table D** demonstrates.

Table C The peak-to-trough measure for each model based on historical data

Data process	Margin model				
	CV	EWMA	fEWMA	HS	HW
Historical	2.1	1.9	1.7	2.2	6.2

Table D The average peak-to-trough measure for each model based on returns from three theoretical processes

Data process	Margin model				
	CV	EWMA	fEWMA	HS	HW
Normal	2.7	3.0	2.8	2.7	3.0
Mixture	5.0	4.6	5.0[a]	6.1	7.0
Regime	5.1	5.4	5.4	5.9	8.5

(a) This may seem anomalous: how can the floored EWMA model have a bigger range than an unfloored one? The answer is that the floor is a floor on volatility. For the mixture process, the floor is not biting in troughs as these are periods in which returns are volatile (while price levels are low), so both models have the same minimum margin. In contrast, peak margin occurs when return volatility is lower (but price levels are higher). The result is that the fEWMA model has a higher peak margin than the unfloored model and thus a bigger peak-to-trough range.

All the models perform similarly for the simplest process. The fat tails of the mixture process are needed to differentiate the models: here we see that the risk sensitivity of the Hull and White model comes at the cost of increasing the already significant peak-to-trough procyclicality of the historical simulation approach. The results for the regime switching process are similar, indicating that an extended period of higher volatility does not make the peak-to-trough measure much worse.

n-day procyclicality

Table E presents 1, 5 and 30-day measures for the historical return series. Again we see that the risk sensitivity of the more reactive model comes at the price of significantly bigger liquidity risk for margin posters. Moreover the floor in the

(1) Note that the regulatory test does not penalise over-conservatism (too few exceptions). For a 99% risk estimate on a one year basis, that is not a serious constraint, as even an accurate model could easily have zero exceptions here, but for eg the 90% risk estimate, too few exceptions can indicate that there is a model problem too.
(2) It can reasonably be argued that the standard is quite weak, and that one should examine Kupiec statistics for multiple confidence intervals, and conduct more sophisticated tests such as the Christoffersen test as well (Campbell (2005)).

fEWMA model does not limit short-term swings in margin to a significant degree: this model performs rather similarly to the unfloored version.

Table E The largest 1, 5 and 30-day margin calls for each model based on the historical data as a percentage of asset value (starting asset value = 100)

	Margin model				
n-day	CV	EWMA	fEWMA	HS	HW
1-day	0.24%	0.56%	0.56%	0.45%	0.94%
5-day	0.39%	0.95%	0.77%	0.75%	2.18%
30-day	1.01%	1.81%	1.81%	1.91%	3.99%

Table F presents the analogous results for the three theoretical processes: normal, mixture and regime.

Table F The largest 1, 5 and 30-day margin calls for each model based on the historical data as a percentage of portfolio value at the start of the period, averaged over 1,000 runs of each process

	CV	EWMA	fEWMA	HS	HW
1-day normal	0.15%	0.27%	0.16%	0.18%	0.33%
1-day mixture	0.19%	0.40%	0.24%	0.29%	0.69%
1-day regime	0.19%	0.42%	0.37%	0.40%	0.74%
5-day normal	0.30%	0.42%	0.31%	0.28%	0.55%
5-day mixture	0.33%	0.55%	0.37%	0.40%	0.98%
5-day regime	0.39%	0.68%	0.61%	0.51%	1.18%
30-day normal	0.77%	0.84%	0.77%	0.69%	1.14%
30-day mixture	0.82%	1.00%	0.84%	0.92%	1.87%
30-day regime	0.93%	1.49%	1.28%	1.11%	2.84%

Here we can see the clustering of large returns almost always makes procyclicality worse: the margin calls for the regime process are higher than for the mixture. The constant volatility model is the least, and the more reactive Hull and White model the most procyclical, while the floor is again not having a major impact on the size of large margin calls.

n-day stressed procyclicality

Table G presents 1-day, 5-day and 30-day stressed measures for the historical data. As these are the same as the analogous results in **Table E** for all the models except Hull and White, we can conclude that at least for this data set, restricting the measure to high volatility periods makes little difference: all large margin calls occur in these periods anyway. This suggests that all the margin models under investigation are, as expected, pro (rather than anti) cyclical. This finding gives comfort that our metrics are indeed capturing the right phenomena.

### The whole of the margin call distribution

The measures we have suggested rely on the extremes of the relevant distributions: the largest and smallest margin amounts; the largest 1, 5 or 30-day changes in margin. These

Table G The largest stressed margin calls as a percentage of portfolio value at the start of the period for each model based on the historical data

	Margin model				
n-day	CV	EWMA	fEWMA	HS	HW
1-day	0.24%	0.56%	0.56%	0.45%	0.50%
5-day	0.39%	0.95%	0.77%	0.75%	2.18%
30-day	1.01%	1.81%	1.81%	1.91%	3.99%

measures can vary significantly depending on the observation window chosen, so it is sensible to examine more of the distribution. **Chart 5** does this, presenting the right-hand tail of the 1-day margin call distribution (ie those days where large margin calls were made). This shows the number of days that each model made a call in a given size range, and validates the intuitions suggested by **Table F**.

Chart 5 The right hand tail of the 1-day margin call distribution for the five models studied
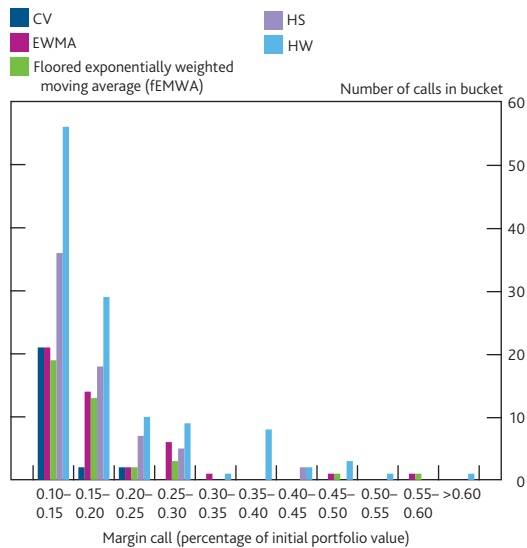
### 6 Parameter sensitivity analysis: the example of the decay factor

We now turn to the role of model parameterisation on the behaviour of margin models. The parameters used in margin models can vary substantially over time and across bilateral and cleared relationships as models are recalibrated. Thus, for instance, one solution to the problem of a model whose risk sensitivity dips is to recalibrate its parameters.

As a simple illustration of the impact of model parameterisation, we examine the features of one model, the EWMA, as its key parameter changes. Specifically we let the decay factor λ vary from 0.90 to 0.99, and examine the effect on risk sensitivity and procyclicality. The decay factor determines the weight of short-term volatility when computing the next period's margin, as discussed in Box 1; the larger λ is, the less sensitive margin is to short-term market

fluctuations. Therefore, one would expect that an EWMA model with a large decay factor would be less procyclical and less risk sensitive.

To test if this assertion holds, we use the same four years of S&P 500 historical returns starting in 2009 and compute initial margin requirements from three variations of the EWMA model: one with λ equal to 0.99; one with λ equal to 0.96; and one with λ equal to 0.90. In **Chart 6** we plot the margin requirements of the three models together with the daily profit/loss (P/L) function of a fixed portfolio that consists of an initial investment of US$100 in the S&P 500 total return index.

Chart 6 Initial margin estimates from the EWMA model using historical returns and different decay factors (λ)
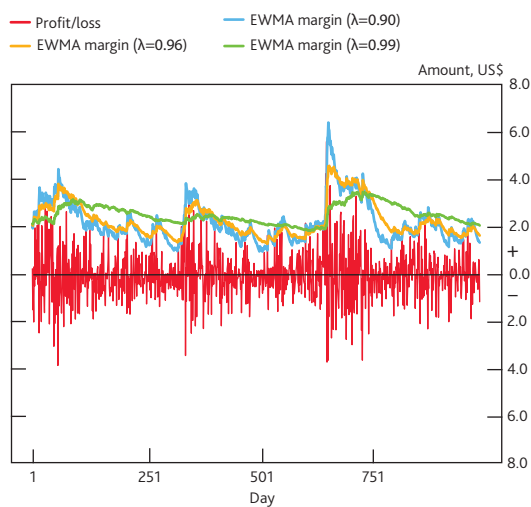

This chart shows that the three EWMA models behave qualitatively similarly, as margin requirements increase or decrease around the same points of time. For example, in the second half of the chart, we observe sharp increases in margin requirements for all models around day 662. This period is in summer 2011, when negative economic news caused price falls and volatility increases in many asset classes. Such comovement of EWMA models is to be expected as they all share a common underlying mechanism.

The chart also shows that the magnitude of the different models' margin fluctuations can vary substantially, especially during periods of market stress. The EWMA model with the decay factor of 0.90 requires a peak margin of US$6.4, while the one with the decay factor of 0.99 demands only US$2.8. The margin requirement of the model with the intermediate decay factor of 0.96 is US$4.5. Hence, we see that even a modest change of the tuning parameter can have a large impact on the model's behaviour. Moreover, as we can also see from the chart, margins which fluctuate less are less sensitive to changes in risk: the higher the decay parameter, the more we observe lengthy periods both with too many and too few P/L breaches.

To elaborate upon this issue, we compute procyclicality metrics for different values of the λ parameter. **Chart 7** presents the values of the peak-to-trough and n-day metrics as a function of λ. When we increase the decay factor, the EWMA model becomes less procyclical. Intuitively, a large λ means that the weight on short-term volatility is small: this makes the margin model less responsive to market fluctuations, and thus, less procyclical. Box 2 discusses this issue further, showing how the choice of λ impacts the model's ability to react to changes in volatility.

Chart 7 Procyclicality measures as the decay factor, λ, in the EWMA margin model varies

Peak-to-trough measure

***n*-day measure**
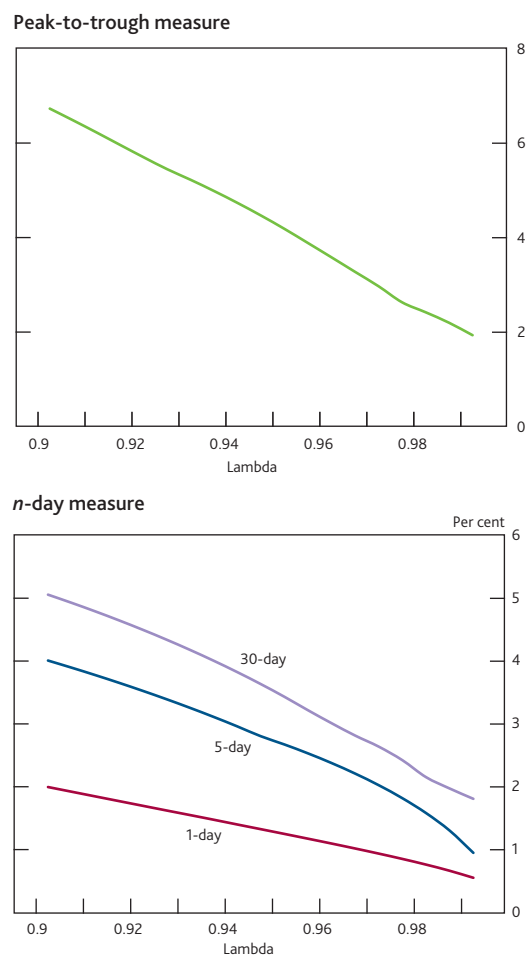
This example demonstrates that changes to model parameters can affect the procyclicality of margin estimates significantly. This helps to emphasise that the assessment of margin models should always be parameter-specific: it is the model, its parameters, and the strategy for recalibrating them that should be studied.

7   Performance in extreme conditions

In order to give some insight into the performance of the models in extreme conditions, we now study the data for 2008. The market moves in the period after the default of

Box 2
Reacting to change

There is no perfect measure of the instantaneous volatility of real financial returns. All estimates require some data to make their estimate, and gathering that data takes time. The more data we have, the more reliable the estimate (unless of course the volatility we wish to estimate has shifted during the time taken to collect the extra data). This gives rise to a fundamental challenge for margin models: we want them to ignore 'temporary' changes in volatility caused by noise or estimation error, and we want them to react quickly to permanent or structural changes in volatility. These two goals are in opposition to each other.

We illustrate this phenomenon by generating two time series of log-normal returns that start off at 20% volatility then jump to 40%. This type of jump in volatility is sufficiently common place that many models capture some form of volatility variation explicitly, as surveyed in Christoffersen, Heston and Jacobs (2009). The main difference of the two time series is the length of the volatility jump. We generate a shorter term jump to proxy for a temporary increase in volatility (ie noise), and a longer term jump to proxy for a more permanent fundamental change in volatility. In particular, for the first time series, the 40% period lasts 20 days, after which volatility falls back to 20%. **Chart A** shows the reaction of three EMWA models with λ = 0.99, 0.97 and 0.93. The largest lambda model shows little reaction to the 'bump'; the smallest one does, but it reacts to a lot of noise too.

Chart A The reaction of three EWMA models to a short-term volatility bump

- – λ = 0.99   ···· λ = 0.93
- – λ = 0.97   —— Actual
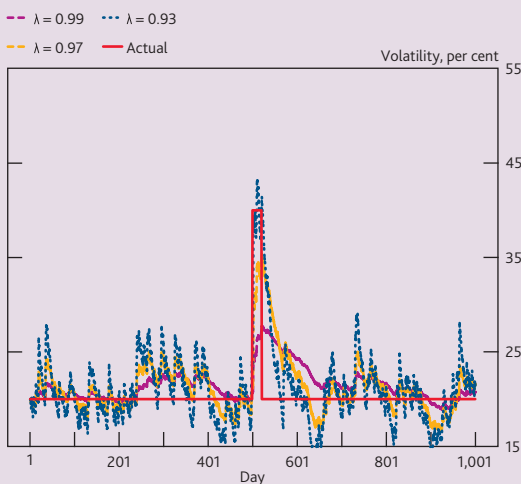
In **Chart B** we lengthen the 40% period out to 100 days, a period more reflective of 'minor' market stress events. Here the period of elevated volatility lasts long enough for all of the models to react, but the largest lambda model takes the whole period to get close to the new volatility. Meanwhile the other two EMWA models quickly adjust to the volatility change, but they tend to overshoot: they react to fleeting volatility increases. This illustrates that for EWMA models without returns filtering at least, there is no 'perfect' balance between reactivity to a real change and oversensitivity to noise. If we want to detect 'bumps' of a given duration, lambda can be chosen optimally: but if we want both to react fast to changes in volatility and to have stable margin when conditions are stable, the problem cannot be solved by the right choice of decay parameter alone. We *have* to react fast in order to respond quickly enough to what might turn out to be an important change in the volatility environment, but doing that necessarily makes the model fairly procyclical. Hence some procyclicality mitigation measures — such as a margin buffer — are required.

Chart B The reaction of three EWMA models to a medium-term volatility bump

- – λ = 0.99   ···· λ = 0.93
- – λ = 0.97   —— Actual
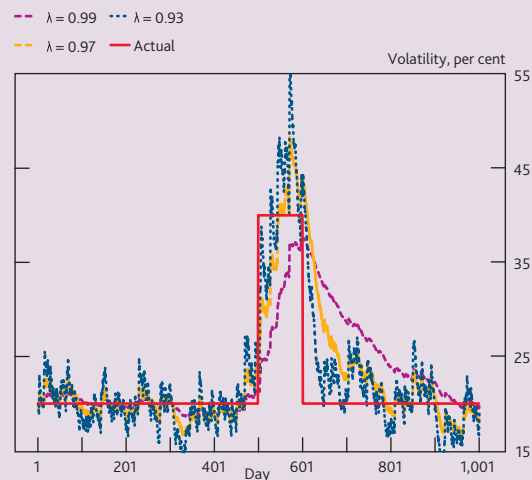
Lehman Brothers were very large, and many risk models failed to perform adequately in this period. How will our models do when faced with this historical data?

Risk sensitivity

The Kupiec POF statistics for the estimates of one-day 99% initial margin in the extended period 2008-2012 are presented in **Table H**.

Table H Kupiec statistics for the 99% margin estimates using historical returns

Data process	Margin model				
	CV	EWMA	fEWMA	HS	HW
Historical	16.0	3.0	0.0	7.2	3.3

The historical returns challenge the risk sensitivity of most of the models: only EWMA and HW react fast enough to avoid an excess of margin breaches, with fEWMA doing even better as it starts its reaction to the events of late 2008 from a higher base. We would unequivocally reject both the constant volatility model and the historical simulation approach on risk sensitivity grounds based on the Kupiec statistics for the historical returns.

Additional insight into the risk coverage issues here can be gained from examining the size of margin shortfalls. These indicate the size of a typical breach, ie the degree of under-coverage of margin estimates. **Table I** presents this data.

Table I Expected shortfalls for each model using historical returns 2008–12

Data process	Margin model				
	CV	EWMA	fEWMA	HS	HW
Historical	1.3	1.3	1.5	1.2	0.8

This shows that while the fEWMA model gives margin levels which have the 'right' coverage, at least based on the POF statistics for 99% margin, when there are exceedances, they are bigger than those from all the other models. It could be that in practice one might prefer a model that 'just' fails its Kupiec test such as HS, but has small expected shortfalls, against one that passes but gives much larger albeit infrequent exceedances. This simply highlights that one test of risk sensitivity is not enough: multiple tests are needed, and the reasons for the behaviour studied must be understood.

The peak-to-trough measure

Table J presents the peak-to-trough measure for the extended period. These are rather similar to those given in **Table C** indicating the 'echo' of 2008 in the data from 2009.[1]

Table J Peak-to-trough measures for each model and process including 2008

Data process	Margin model				
	CV	EWMA	fEWMA	HS	HW
Historical	2.2	2.8	2.0	2.2	7.7

The *n*-day measure

The *n*-day measures in **Table K** show that while the extremely stressed period of 2008 produced much higher 1-day margin calls for some of the models, the intensity of the liquidity demands moderate on a 30-day basis for all of the models except HW.

Table K The largest *n*-day margin calls for historical returns including 2008

Period	Margin model				
	CV	EWMA	fEWMA	HS	HW
1-day	0.4%	1.1%	1.1%	0.7%	1.5%
5-day	0.6%	1.4%	1.4%	1.0%	3.0%
30-day	1.0%	2.3%	2.3%	2.2%	6.7%

8   Conclusions

The analysis in this paper has shown that initial margin models which pass a standard risk-sensitivity test can vary quite widely in their degree of procyclicality. This variation may come from different models (eg HS versus EWMA), different parameter settings in a given model (eg the value of the decay factor in the EWMA model), or different strategies for updating parameter values in the light of new data. More specifically, the analysis suggests that model calibrations which give higher weight to recent data are more procyclical. This could be because calibration is based on a short look-back period or because the weight given to observations in the look-back period declines rapidly as they age. In the case of the HW model, it could also be due to the scale factor that resizes observations in the look-back period being based only or mainly on the recent past.

Although the specification of initial-margin models can make a substantial difference to their procyclicality and hence to the potential size of margin calls, guidance to CCPs, market participants and supervisory authorities on how to mitigate procyclicality is limited. For example, international guidance on calibrating initial-margin models for non-centrally cleared trades (BCBS and IOSCO (2013)) says that an episode of financial stress should be included in the look-back period, but does not provide details such as the proportion of observations

---

(1) The HS and HW models for 2009 and 2010 still retain the 2008 data in their 500-day window, and when this turmoil starts to fall out of the model in 2011, coincidentally a new period of stress begins. The effect of 2008 on HW in particular is significant, as the first 60 days or so of **Chart 4** illustrates.

that this should represent or whether it is acceptable for the intensity of this episode to be reduced by scaling.[1] Equivalent guidance for models for centrally cleared trades states only that procyclicality should be limited to the extent practical and prudent after satisfying risk-sensitivity criteria. That said, the technical standards for European CCPs make this more precise.[2] However, this does not guarantee that CCPs will pick the least procyclical of these options.

How should authorities and market participants respond to this issue? First, they could compute procyclicality measures to better understand, assess and compare the procyclicality of existing and potential future initial-margin models. In this paper, we have found three such measures to be useful. One is the *n*-day measure, which might be calculated over a few weeks, and shows how much extra margin could be called over this period. This measure indicates the potential liquidity stress market participants could encounter in meeting short-term increases in collateral requirements. A variant of this is the *n*-day stressed procyclicality measure. Since this measures potential increases in collateral requirements from already high levels, it might be computed over a slightly shorter period, as this could still be enough to generate liquidity stresses given the relatively high levels of margins already posted. The third measure is the peak-to-trough ratio, which shows the range of margin requirements for a portfolio over the business cycle. This could help market participants to evaluate whether the variation of collateral requirements was compatible with their business models. It could also help authorities to evaluate the scope for margin requirements across the financial system to rise relative to the available supply of collateral-eligible assets.[3]

CCPs and large dealers could disclose both peak-to-trough and *n*-day procyclicality measures for their margin models. This would help clearing members and their clients to anticipate potential increases in margin requirements, and to prepare accordingly.

Third, armed with measures of procyclicality, international authorities would be better able to investigate the extent to which it can be reduced, without breaching risk-sensitivity requirements, through policies such as stressed calibration, margin floors and speed limits on margin calls. More detailed guidance on limiting procyclicality could potentially then be issued. There are advantages to doing this at international level: system-wide margin requirements over the cycle would be less volatile, and the risk of trading activity moving across jurisdictions from time to time to take advantage of different regulatory regimes would be avoided.

Further work

The analysis presented above was based on a very simple single-asset 'portfolio'. One obvious extension is to examine the effects of using more realistic portfolios composed of assets from multiple classes. This would allow investigation of how procyclicality is affected by the degree of correlation between multiple assets. A second extension is to investigate the effectiveness of other policies — in addition to the margin floor (appended to the EWMA model above) — in reducing procyclicality and to consider whether this comes at the cost of reduced risk sensitivity. Other such policies include investigating speed limits on margin increases, calibration of margin models to include periods of stress and calibration based on long look-back periods.

---

(1) An 'Industry Standard Margin Model' is being developed to provide market participants with a standard for bilateral initial margining. If such a model is widely used, then its procyclicality is potentially a systemic issue. It is therefore encouraging that procyclicality mitigation is explicitly discussed in the White Paper proposing this model (ISDA (2013)), and it is to be hoped that this concern will suitably inform the design of the final model.

(2) The options are (i) to include an episode of stress in the look-back period that accounts for at least 25% of observations; (ii) to impose a minimum margin requirement based on a ten-year look-back period; or (iii) to raise margin requirements by at least 25% in non-stressed times and allow this buffer to fall during periods of financial stress, thus reducing overall increases in margin requirements.

(3) The Macroeconomic Assessment Group on Derivatives' recent impact assessment of OTC derivatives reforms has a central-case initial-margin requirement across the financial system for both bilateral and cleared exposures of €886 billion (Macroeconomic Assessment Group on Derivatives (2013)). If we take this as the middle of the band, and assume a peak-to-trough measure of 2, low margins would be €590 billion, high margins would be €1,180 billion and the trough-to-peak funding requirement across the system would be €590 billion. If instead we assume a peak-to-trough measure of 4, low margins would be €354 billion, high margins would be €1.4 trillion and the trough-to-peak funding requirement across the system would be €1.1 trillion. These are rough calculations: it is unlikely that all counterparties' peaks and troughs would synchronise for instance. Nevertheless, they do indicate that the amounts concerned could be large.

Appendix
Theoretical processes

The 'normal' data generating process (series (ii) in Section 4) draws returns from a single normal distribution with a zero mean and variance that is equal to the observed historical variance of the log returns of the S&P500 total return index as defined in footnote 1 on page 6. This is $\sigma_1^2 = 0.02\%$.

The 'mixture' process (series (iii) in Section 4) is generated by a mixture of two (log) normal distributions. Specifically, returns are drawn from one of two normal distributions, which have different means and variances. An identically and independently distributed random variable determines each day the distribution from which returns are drawn. We estimate the distribution parameters using S&P500 historical total returns and maximum likelihood estimation.[1] The values of the parameters (daily mean and variance) of the two normal distributions are: $\mu_2 = -0.05\%$, $\sigma_2^2 = 0.3\%$, $\mu_3 = 0.13\%$ and $\sigma_3^2 = 0.003\%$ respectively. The probability that returns are drawn from the first normal distribution is 46.2%.

One nice feature of the normal-mixture distribution is that it allows for skew and kurtosis that can better fit historical returns, including their fat tails, compared to a single normally distributed process (Brigo and Mercurio (2002), McLachlan and Peel (2000)).

Finally, we use a Markov regime switching model to generate the last return series. In this case, we assume that there are two possible regimes and that the regimes switch according to a Markov chain. Returns are drawn from a distinct normal distribution for both regimes. We estimate the distribution parameters and the transition probabilities using S&P500 historical total returns, maximum likelihood estimation and the algorithm described in Hamilton (1990). The values of the parameters (daily mean and variance) of the normal distribution in the two regimes are: $\mu_4 = -0.03\%$, $\sigma_4^2 = 0.04\%$, $\mu_5 = 0.08\%$, $\sigma_5^2 = 0.01\%$, respectively. The probability transition matrix is determined by the two 'state persists' probabilities: $p_{44} = 98.7\%$ and $p_{55} = 99.4\%$.

The mixture normal process is a special case of the Markov switching model. The key difference is that the latter allows for longer lasting regimes and is therefore more suitable for describing correlated data that exhibit distinct behaviours during different time periods (Gray (1996)). This is usually the case for portfolio returns as they typically comove with business cycles, which might last for several years.

---

[1] We obtain the maximum likelihood estimates using two techniques. The first one is a numerical optimisation and the second one is the expectation-maximisation (EM) algorithm. The techniques produced very similar results.

References

Abruzzo, N and Park, Y (2013), 'An empirical analysis of future margin changes: determinants and policy implications', Working Paper, Federal Reserve Board.

Barone-Adesi, G, Giannopoulos, K and Vosper, L (1999), 'VaR without Correlations for Portfolios for Portfolios of Derivative Securities', *Journal of Futures Markets*, Volume 19, pages 583–602.

Basel Committee on Banking Supervision (1996), 'Supervisory framework for the use of 'backtesting' in conjunction with the internal models approach to market risk capital requirements', January.

Basel Committee on Banking Supervision and International Organization of Securities Commissions (2013), 'Margin requirements for non-centrally cleared derivatives', September, available at www.bis.org/publ/bcbs261.pdf.

Brigo, D and Mercurio, F (2002), 'Lognormal-mixture dynamics and calibration to market volatility smiles', *International Journal of Theoretical and Applied Finance*, Vol. 5, No. 4, pages 427–46.

Brunnermeier, M and Pedersen, L (2009), 'Market liquidity and funding liquidity', *Review of Financial Studies*, Vol. 22. No. 6, pages 2,201–38.

Campbell, S (2005), 'A Review of Backtesting and Backtesting Procedures', *Federal Reserve Board Finance and Economics Discussion Series 2005-21*, April, available at www.federalreserve.gov/pubs/feds/2005/200521/200521pap.pdf.

Chicago Mercantile Exchange (2010), 'CME SPAN: standard portfolio analysis of risk', available at www.cmegroup.com/clearing/span-methodology.html.

Christoffersen, P, Heston, S and Jacobs, K (2009), The shape and term structure of the index option smirk: why multifactor stochastic volatility models work so well', *Management Science*, Vol. 55, No. 12 pages 1,914–32.

Commodity Futures Trading Commission (2013), 'Enhanced risk management standards for systemically important derivatives clearing organizations', Final Rule 17 CFR Part 39, August, available at www.cftc.gov/ucm/groups/public/@lrfederalregister/documents/file/2013-19791a.pdf.

Committee on the Global Financial System (2010), 'The role of margin requirements and haircuts in procyclicality', *CGFS Paper No. 36*, March, available at www.bis.org/publ/cgfs36.htm.

Committee on the Global Financial System (2013), 'Asset encumbrance, financial reform and the demand for collateral assets', *CGFS Paper No.49*, May, available at www.bis.org/publ/cgfs49.htm.

Committee on Payment and Settlement Systems and Technical Committee of the International Organization of Securities Commissions (2012), 'Principles for financial markets infrastructures', April, available at www.bis.org/publ/cpss101a.pdf.

European Union (2012), 'Regulation (EU) No. 648/2012 on OTC derivatives, central counterparties and trade repositories', July, available via www.esma.europa.eu/page/European-Market-Infrastructure-Regulation-EMIR.

Financial Stability Forum (2009), 'Report of the financial stability forum on addressing procyclicality in the financial system', April.

Gray, S (1996), 'Modeling the conditional distribution of interest rates as a regime-switching process', *Journal of Financial Economics*, Vol. 42, No. 1, pages 27–62.

Hamilton, J (1990), 'Analysis of time series subject to changes in regime', *Journal of Econometrics*, Vol. 45, No. 1, pages 39–70.

Hedegaard, E (2011), 'How margins are set and affect asset prices', Economics Department, NYU Stern School of Business, available at http://econ.as.nyu.edu/docs/IO/22510/Hedegaard_12122011.pdf.

Heller, D and Vause, N (2012), 'Collateral requirements for mandatory central clearing of over-the-counter derivatives', *BIS Working Paper No. 373*.

Hull, J and White, A (1998), 'Incorporating volatility updating into the historical simulation method for value at risk', *Journal of Risk*, Vol. 1, No. 1, pages 5–19.

ISDA (2013), 'Standard Initial Margin Model for Non-Cleared Derivatives', White Paper, December, available at www.isda.org.

JP Morgan/Reuters (1996), RiskMetrics technical document.

Kashyap, A and Stein, J (2004), 'Cyclical implications of the Basel II capital standards', *Federal Reserve Bank of Chicago Economic Perspectives*.

Kupiec, P, (1995), 'Techniques for verifying the accuracy of risk management models', *Journal of Derivatives*, Vol. 3, No. 2, pages 73–84.

Macroeconomic Assessment Group on Derivatives (2013), 'Macroeconomic impact assessment of OTC derivatives regulatory reforms', August, available at www.bis.org/publ/othp20.pdf.

Mayhew, S, Sarin, A, Shastri, K (1995), 'The allocation of informed trading across related markets: an analysis of the impact of changes in equity-option margin requirements', *The Journal of Finance*, Vol. 50, No. 5, pages 1,635–53.

McLachlan, D and Peel, D (2000), *Finite mixture distributions*, Wiley.

Murphy, D (2009), *Unravelling the credit crunch*, Chapman Hall.