BANK OF ENGLAND

Staff Working Paper No. 597

A comparative analysis of tools to limit the procyclicality of initial margin requirements

David Murphy, Michalis Vasios and Nicholas Vause

April 2016

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 Authority Board.

BANK OF ENGLAND

Staff Working Paper No. 597

A comparative analysis of tools to limit the procyclicality of initial margin requirements

David Murphy,[(1)] Michalis Vasios[(2)] and Nicholas Vause[(3)]

Abstract

The requirement to post initial margin on derivatives transactions is a key feature of the post-crisis reforms of the OTC derivatives markets. Initial margin requirements are usually determined by risk-based models. These models typically require increased margin in stressed conditions: they are *procyclical*. This procyclicality causes a liquidity burden on market participants which sometimes falls when they are least able to bear it. In this paper we study a variety of tools which have been proposed to mitigate the procyclicality of initial margin requirements. Three of these tools are proposed in European regulation; the other two are new proposals which offer attractive procyclicality mitigation features. The behaviour of all five tools is studied in a simulation framework. We examine the extent to which each tool mitigates procyclicality, and at what cost in demanding unnecessary margin compared to a benchmark unmitigated model. Our findings indicate that all of the tools are useful in mitigating procyclicality to some extent, but that the optimal calibration of each tool in a particular situation depends on the relative weights placed by the modeller on the objectives of minimizing procyclicality on the one hand and minimizing undesirable overmargining in periods of low volatility on the other. This suggests that it may be appropriate to consider moving from tools-based procyclicality regulation to one based on the desired outcomes.

Key words: Central counterparty, central clearing, initial margin, margin models, OTC derivatives, procyclicality.

JEL classification: G17.

(1) Bank of England. Email: david.murphy@bankofengland.co.uk
(2) Bank of England. Email: michalis.vasios@bankofengland.co.uk
(3) Bank of England. Email: nicholas.vause@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England.

Information on the Bank's working paper series can be found at
www.bankofengland.co.uk/research/Pages/workingpapers/default.aspx

Publications Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email publications@bankofengland.co.uk

© Bank of England 2016
ISSN 1749-9135 (on-line)

1 Introduction

Margin requirements are increasingly important features of the financial system. Both initial and variation margin are or soon will be required for both many cleared and many uncleared derivatives between larger counterparties. Thus one or both parties to a derivative often implicitly agree to post margin to each other throughout the life of the portfolio of trades between them.

1.1 The Impact of Margin Requirements on Liquidity

This agreement to post margin has implications for liquidity management. If a party has to post margin, it has to find and fund that margin, often either in the form of cash or in high quality securities. Margin has to be posted quickly after a margin call is made, sometimes within as little as two hours. Calls are sometimes made intra-day, and always at the end of the day for centrally cleared portfolios. Moreover, failure to meet a margin call is usually an event of default. Margin obligations can therefore be onerous, and the consequences of failing to meet them are usually severe.

Initial margin requirements are often risk-based. That is, they depend on the risk of a portfolio of instruments as estimated by a *margin model*. These models are typically based to some degree on historical data. For instance, a Value-at-Risk-based (VaR-based) margin model might determine an initial margin requirement based on the 99th percentile of the estimated loss distribution of the portfolio over some assumed liquidation horizon.

Risk estimates for a constant portfolio change as the data used to calibrate the model which estimates them change. Thus for instance if we use a fixed length lookback window to calculate a 99% historical simulation VaR, and update this window every day as new data arrive, the model estimate will change as a new day's return comes in and an old day's return leaves. Figure 1 illustrates the variability of the resulting margin requirement for a particular fixed portfolio using a 500 day window. It can be clearly seen that margin varies substantially over the period: indeed, the maximum required is about 2.3 times the minimum.

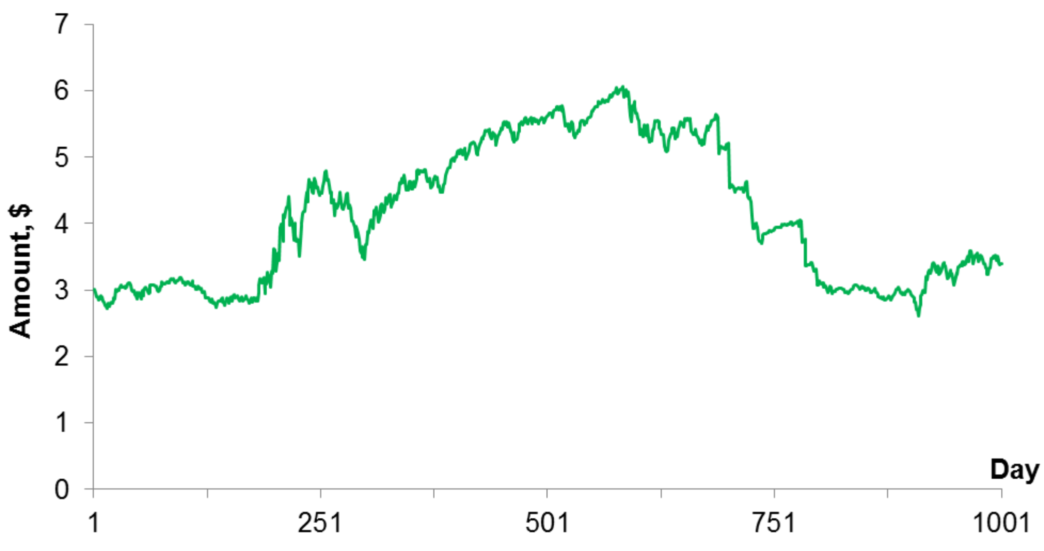

Figure 1: Margin requirements for a fixed portfolio over a four year period calculated using a historical simulation VaR model estimating risk at the 99th percentile using a 500 day lookback window.

This illustration evidences the need for parties to understand the liquidity requirement created by the obligation to post margin. A party with the portfolio used to produce Figure 1 might have to post 2.3 times the current level of initial margin, and thus would have to include this contingency in its liquidity planning.

1.2 Initial Margin Procyclicality

The variability of initial margin requirements is a property of both the portfolio and the initial margin model. Typically as markets become stressed, risk factor volatilities and correlations change, and this usually increases margin requirements. Thus margin requirements are procyclical: they tend to increase in times of stress, and thus put a greater burden on margin posters when they are less able to bear it.

In extremis, this phenomenon could cause a feedback loop whereby one or more market participants are unable to fund the margin requirements for their portfolios, and are thus forced either to default or to deleverage themselves [3, 11]. This deleveraging would typically take place in a stressed market, and thus has a greater propensity to move prices further. Moreover a default would mean that the non-defaulting party (for a bilateral portfolio) or the central counterparty (for a cleared portfolio) has to both close-out the defaulter's portfolio and to liquidate their collateral, actions which could further increase stress. Therefore if margin requirements are excessively procyclical, they have the potential to threaten financial stability.

1.3 Regulatory Context

This destabilising mechanism is recognised in OTC derivatives policy. Thus for instance the relevant European Regulation, EMIR, [5] states that for central counterparties ('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.[1] When markets become more volatile, risk is higher, and hence margin requirements *should* be higher. But it is undesirable for margin models to *over-react* to changing conditions. In order to constrain potential over-reactions, the EMIR implementation therefore prescribes three tools for *procyclicality mitigation*. All CCPs in the EU must use at least one of these three tools in their initial margin models.

The three anti-procyclicality ('APC') tools and the context around their use are set out in Article 28 of the EMIR Regulatory Technical Standards ('RTS') [6] as follows:

1. A CCP shall ensure that its policy for selecting and revising the confidence interval, the liquidation period and the lookback period deliver forward looking, stable and prudent margin requirements that limit procyclicality to the extent that the soundness and financial security of the CCP is not negatively affected. This shall include avoiding when possible disruptive or big step changes in margin requirements and establishing transparent and predictable procedures

[1] See also the earlier work from the CGFS [4] on the broader problem of procyclicality across both margin requirements and collateral haircuts.

for adjusting margin requirements in response to changing market conditions. In doing so, the CCP shall employ at least one of the following options:

(a) applying a margin buffer at least equal to 25% of the calculated margins which it allows to be temporarily exhausted in periods where calculated margin requirements are rising significantly;

(b) assigning at least 25% weight to stressed observations in the lookback period calculated in accordance with Article 26;

(c) ensuring that its margin requirements are not lower than those that would be calculated using volatility estimated over a 10 year historical lookback period.

2. When a CCP revises the parameters of the margin model in order to better reflect current market conditions, it shall take into account any potential procyclical effects of such revision.

We will use the shorthand 'margin buffer', '25% stressed observations' and 'ten year floor' respectively for these three approaches to procyclicality mitigation.

1.4 Under- and Over-margining

The potentially trade-off between risk sensitivity and procyclicality can be understood in two dimensions. First, an initial margin model should not *under-margin* too much, because that leaves the posted party over-exposed to the failure of its counterparty. Second, it may however be acceptable sometimes to *over-margin* to some degree, if this prevents margins from falling too far, and thus rising too fast when conditions change.

Both of these issues are clearly a matter of degree: a small degree of under-margining versus a benchmark may be acceptable if the model has sufficient risk coverage, while a large degree of over-margining is economically inefficient.[2]

1.5 This Paper and Related Work

This paper will review the behaviour of these three tools to limit procyclicality in conjunction with their impact on the two dimensions of risk sensitivity. Briefly, our findings indicate that all of the tools are useful in mitigating procyclicality to some extent, but that the optimal calibration of each tool in a particular situation depends on the relative weights placed by the modeller on the objectives of minimizing procyclicality on the one hand and minimizing undesirable over-margining in periods of low volatility on the other. This suggests that it may be appropriate to consider moving from tools-based procyclicality regulation to one based on the desired outcomes.

Concerns about the procyclicality of margin appear elsewhere in the literature. For instance, in earlier work [12] we propose quantitative metrics of the procyclicality of the initial margin requirements, and show that common risk-based margin models can be very procyclical depending on the assumptions about the model's parameters, the calibration strategy, and the

[2] A key point here is that we do not know what the 'right' process to model the comovement of many risk factors is, so any benchmark model for real returns must itself have model risk. This in turn implies that any model can undermargin some of the time, as the 'right' level of margin is not known. Hence regulation demands that margin models meet certain standards of risk coverage, such as (but not limited to) the requirement to pass backtests at a fixed level of confidence.

data used in the calibration among others. In a theoretical setting, [3] studies the funding liquidity risk created by the requirement to post margin in a stylized economy and demonstrates the possibility of a destabilising liquidity/solvency spirals. A related paper [1] shows that the actual levels of margin at the Chicago Mercantile Exchange have risen quickly following volatility spikes, indicating procyclicality. Similarly, [9] shows further evidence of the procyclical behaviour of margins. Thus there is evidence from the literature both that margins are procyclical and that this procyclicality can have destabilising effects.

The rest of the paper proceeds as follows. Section 2 reviews the two classes of procyclicality measure proposed in [12], describes the implementation of the three EMIR procyclicality mitigation tools, and introduces two additional tools. Section 3 then sets out the simulation framework used to analyse them. The impact of each of these five tools on under and over-margining is then illustrated in Section 4. The paper ends with some tentative policy conclusions.

2 Defining Procyclicality and Procycliality Mitigation Tools

Both the term 'procyclicality' and the three EMIR procyclicality mitigation tools lack precise definitions. Therefore we turn next to procyclicality measures, selecting an appropriate one for our context. The procyclicality mitigation tools studied are then defined by setting out exactly how much margin a model which uses each form of mitigation would require at each point in time. These definitions will be the key ingredients for our analysis in the following sections.

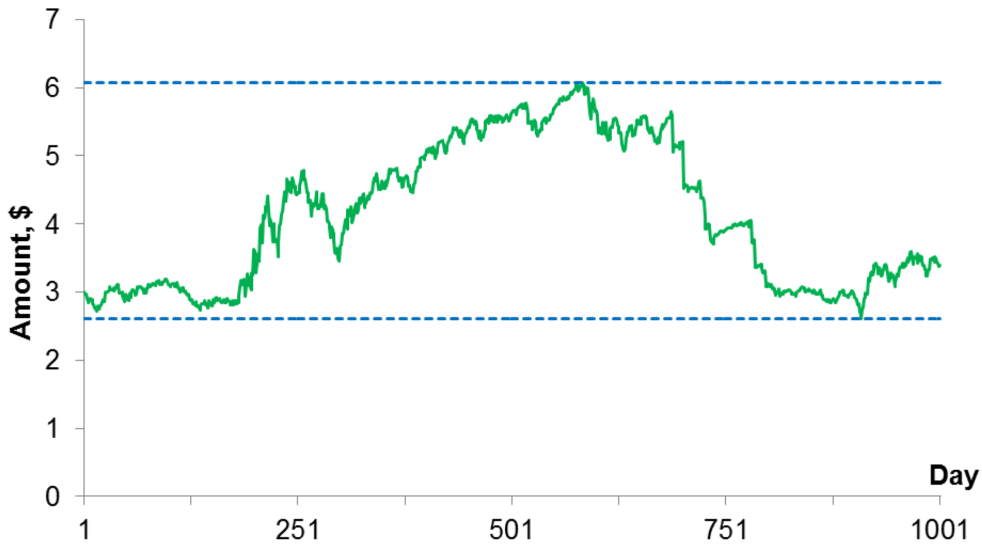

Figure 2: The peak-to-trough measure of procyclicality is the ratio of maximum to minimum margin across the cycle. The chart illustrates the maximum and minimum levels of margin for a fixed portfolio. The peak-to-trough ratio here is approximately 2.3.

2.1 Measures of Procyclicality

There are two key aspects of initial margin variation:

1. Over the economic cycle, how much do margin requirements vary?; and
2. In the short term, if markets suddenly become stressed, how large could the increase in margin be?

These were called 'peak-to-trough' and 'n-day' procyclicality in our previous work [12]. The two measures are illustrated in Figures 2 and 3.

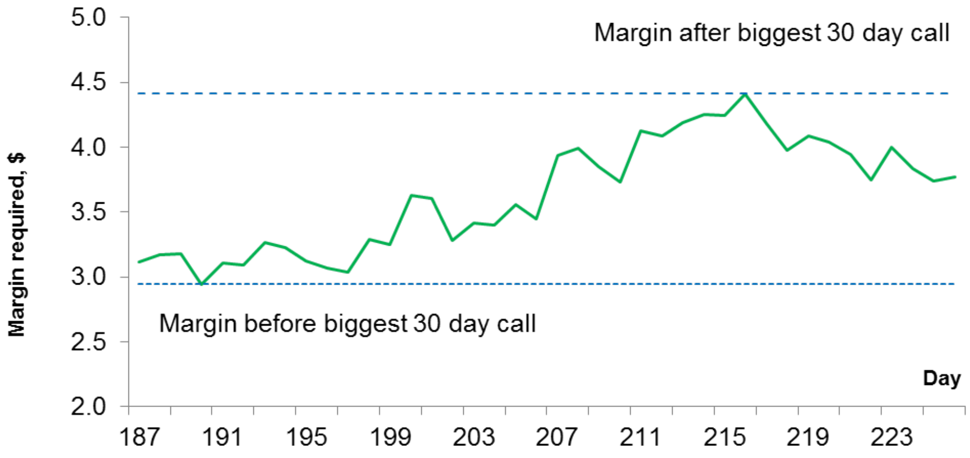

Figure 3: The n-day of procyclicality is the largest margin call over an n-day period. The chart illustrates the 30-day measure for the previous portfolio: the 30 day period with the largest net call ends on day 216. This figure therefore 'zooms in' on this area from Figure 2.

In this paper we will use the 30-day measure for procyclicality. There are several reasons for this choice. First, discussions with prominent clearing members suggest that 30 days is a common liquidity planning horizon; and second, 30 days is enshrined in the regulation of liquidity risk as the period over which banks must resource themselves for a stressed outflow. For example, the Basel 3 Liquidity Coverage Ratio [2] requires that Banks maintain "an adequate level of unencumbered, high-quality liquid assets that can be converted into cash to meet its liquidity needs for a 30 calendar day time horizon under a significantly severe liquidity stress scenario".

2.2 Overview of the EMIR Procyclicality Mitigation Tools

The three procyclicality mitigation tools discussed above were:

1. Applying a *margin buffer* at least equal to 25% of the calculated margins which it allows to be temporarily exhausted in periods where calculated margin requirements are rising significantly;
2. Assigning at least 25% weight to *stressed observations* in the lookback period; and
3. Ensuring that margin requirements are not lower than those determined by the *10 year VaR floor*.

We examine each of these options, together with an additional mitigation tool, in the following subsections. First, though, it is helpful to set out how the tools are used in the calculation of margin.

2.3 The Use of Procyclicality Mitigation Tools

The situation we are concerned with is where we have a risk-based margin model. This model is calibrated to market data reflecting different risk factors, so the model takes these data, as well as information on the sensitivity of a given portfolio to changes in these risk factors, and produces a margin requirement. The margin call this model creates is just the difference between one day's margin requirement and the next. Figure 4 illustrates this.

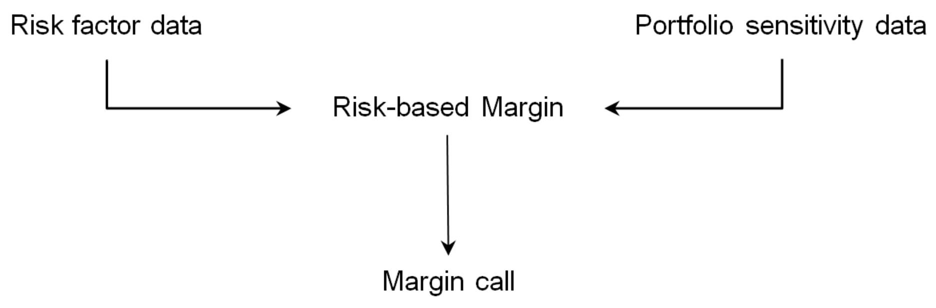

Figure 4: Risk-based models take market data on risk factor returns and information on the sensitivity of a given portfolio to changes in these risk factors, producing a margin requirement.

Procyclicality mitigation can occur in two places:

- The buffer and floor requirements modify the call made; while
- The stressed observation technique modifies the risk factor data used.

Figure 5 illustrates both approaches.

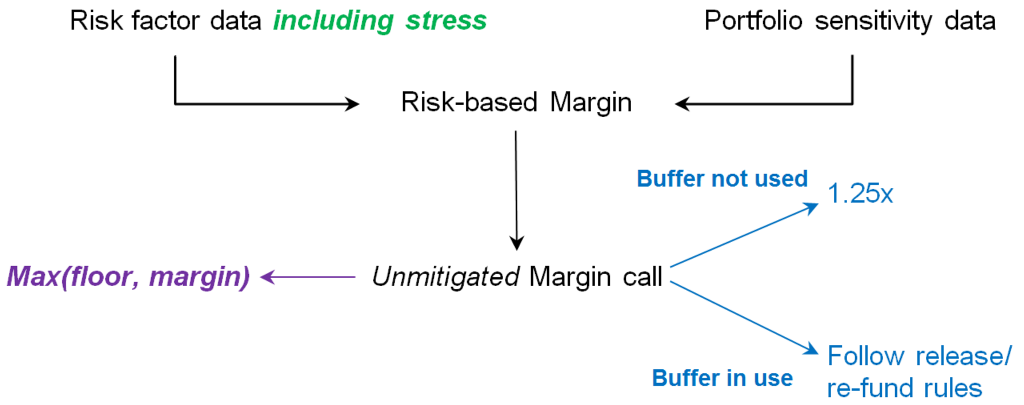

Figure 5: Risk-based margin models with various forms of procyclicality mitigation shown in colour.

2.4 Margin Buffers

The first tool envisaged by EMIR is a margin buffer. Here the margin taker calculates margin requirements using some model. In ordinary market conditions this margin requirement is multiplied by a fixed amount: 1.25 is suggested. This extra margin is then available to reduce the size of the call that needs to be made in stressed conditions.

There are therefore three key features of a buffer:

1. The extra or buffer amount that is taken in ordinary conditions;
2. The criteria used for *releasing* the buffer, i.e. how it is decided not to require 1.25 times the amount determined by the margin model but rather some smaller amount; and
3. The criteria used for determining when the buffer is built back up or *re-funded* after it has been used.

It should be noted a margin policy which simply calls 1.25 times a model output without any possibility of release is by definition 25% more procyclical in terms of the size of the margin call than the model without a buffer. Moreover, if the buffer is released too soon, it may be exhausted before market conditions become highly stressed. It can therefore be seen that release policy is key: an optimal approach would release the buffer just in time to smooth out a shock. However, no one has foreknowledge of how big a given shock is going to be, so model designers run the twin risks of constructing an approach which releases the buffer too early and so does not absorb the most intense stress, and constructing one which releases too late and so provides too little procyclicality mitigation.

2.5 The Distribution of Conditional Volatilities

For the reasons discussed above, policy makers want to mitigate overly procyclical margin calls. By definition these occur in stressed periods. Therefore we need some definition of a stressed period. The obvious choice here is a period in which (conditional) volatility is elevated. Thus one obvious way to define a stressed period is:

- Define a data window for the risk factor return x, e.g. 500 days going back from today, $x(t-499), \ldots x(t)$;
- Define some notion of the conditional volatility of x, such as the EWMA volatility with some decay constant λ, ${}^{\lambda}\sigma(t)$;
- Define some threshold S such that we say that conditions at time t for risk factor x are stressed if ${}^{\lambda}\sigma(t) > S$.

The question then becomes how to define the threshold S for each risk factor. One way to do this is to review the level of EWMA volatility across the cycle. Suppose we take an extended period of data, $1, \ldots,$ T and calculate the EWMA volatility ${}^{\lambda}\sigma(t)$ for each day t in $1, \ldots,$ T where as usual

$${}^{\lambda}\sigma(t)^2 = \lambda \cdot {}^{\lambda}\sigma(t-1)^2 + (1-\lambda) \cdot x(t-1)^2,$$

and we 'seed' the induction by setting ${}^{\lambda}\sigma(0)$ to some long term average volatility.

The time series of EWMA volatilities can be gathered into a distribution, such as Figure 6. This shows the distribution of volatilities of the USD/JPY FX rate from the start of 1987 to the end of 1996, calculated using an EWMA model with decay factor λ = 0.97. The percentiles of this distribution offer a convenient definition of 'stress'. For instance we could say that conditions at time t in risk factor x are stressed if the EWMA volatility ${}^{\lambda}\sigma(t)$ exceeds the 90$^{\text{th}}$ percentile of historical observations. For the example above this corresponds to volatility in excess of 0.906% .

2.6 Release and Re-fund Rules for the Margin Buffer

A definition of stressed conditions allows us to specify the margin buffer in detail. Let $Mar(t)$ be the amount of margin we hold at time t, and suppose the risk-based margin model determines that the amount of margin required is $Mod(t)$. The buffer we have at any given time is $Buf(t) = Mar(t) - Mod(t)$.

Suppose that initially ($t = 0$) conditions are not stressed. Then we set $Mar(0) = 1.25 \cdot Mod(0)$. We calculate a typical 'large' margin level, *StressedMargin*, for instance using some percentile of the distribution of daily margin for the portfolio. The buffer (if any) will be used

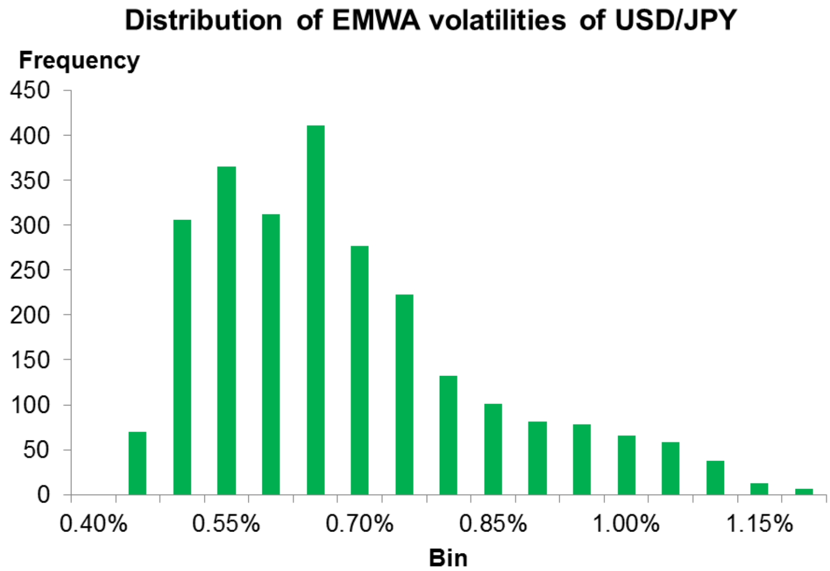

Figure 6: The distribution of daily EWMA volatilities of USD/JPY, 1987-1996.

when the unmitigated model would cause margin to go above *StressedMargin*. For each day t, we calculate $Mod(t)$. Then,

- If the margin is not 'too big', by which we mean $1.25 \cdot Mod(t) \leq StressedMargin$, then make the call, and set $Mar(t) = 1.25 \cdot Mod(t)$.
- If on the other hand $1.25 \cdot Mod(t) > StressedMargin$ then use the buffer to the extent that it is available, setting $Mar(t) = \max(StressedMargin, Mod(t))$.

This rule has the effect of releasing the buffer when 1.25 times the model margin goes above the ceiling level *StressedMargin*, and re-funding (i.e. refilling) the buffer when model margin goes below the ceiling level. If the unmitigated margin requirement $Mod(t)$ rises above the ceiling, then the required amount, not the ceiling, is called for.

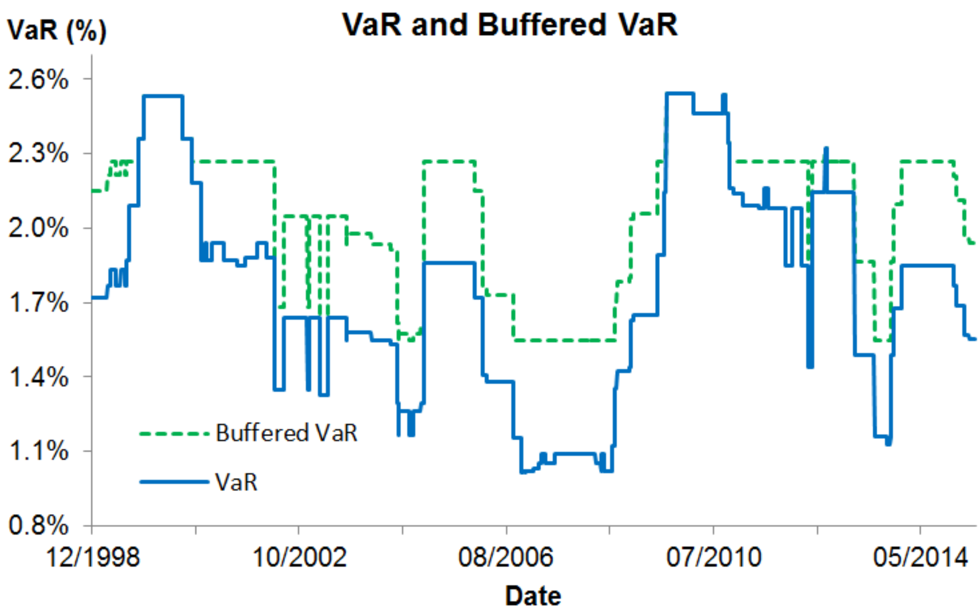

Figure 7: An illustration of the 99% VaR of a USD/JPY FX position with and without the use of a 25% margin buffer which is used if margin goes above the 90^{th} percentile of historical margin levels.

Figure 7 illustrates the effect of this form of procyclicality mitigation using the same data as before, and setting *StressedMargin* equal to the 90th percentile of the historical distribution of VaRs. The mitigated model keeps more margin over most of the cycle, and this buffer is used to mitigate the impact of margin rises for instance in July and August 2009.

2.7 Stressed Data in the Observation Period

Suppose we can identify a historical period of elevated volatility in most or all risk factors, perhaps by using the method discussed above. The second procyclicality mitigation tool requires the use of a period (or periods) of stressed volatility in the model input of no less than 25% of the total. Thus if a historical simulation VaR model is being used with a 1,000 day data window, at least 250 of those days must be from a period of stress.

There are various methods to calculate a VaR from this point. One approach would be to construct a composite period composed of, say, 250 days stressed data and 750 days current data, and use this to perform a historical simulation VaR calculation.[3] Another would be calculate two VaRs, a pure stressed VaR based on the stressed data, and a pure current VaR based on the current data, and calculate margin as a blend of the two VaRs using

$$25\% \textit{ stressed VaR} + 75\% \textit{ current VaR}$$

For illustrative purposes, we take the maximum VaR for our USD/JPY position in the 1987-1996 sample period as the stressed VaR. Figure 8 then illustrates the effect of blended stressed VaR for the same position in the subsequent period.

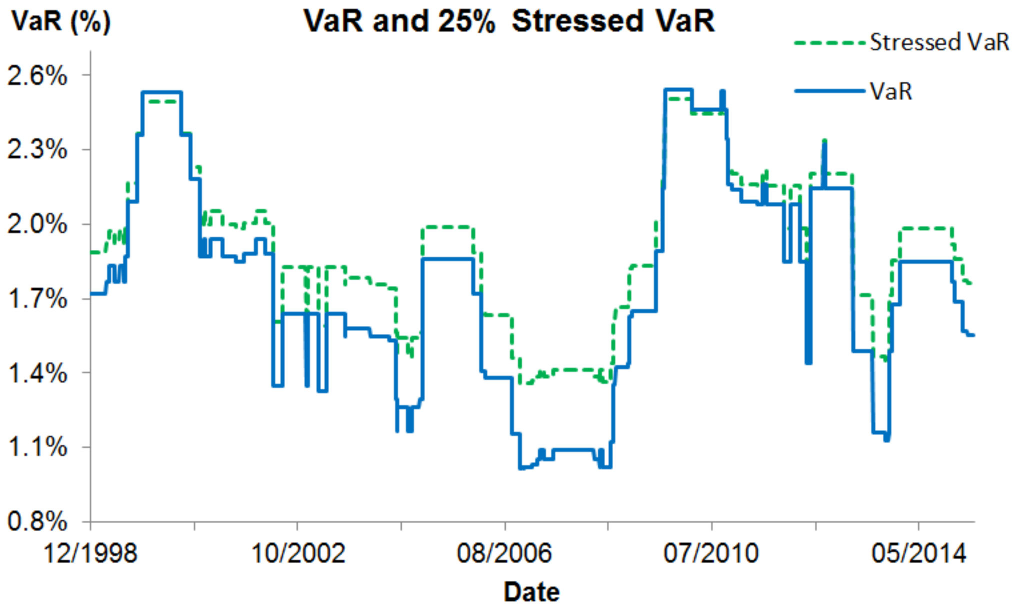

Figure 8: An illustration of the 99% VaR of a USD/JPY FX position with and without the blending of 25% stressed VaR.

This example illustrates that if conditions become even more stressed than those that were used to calibrated the stressed VaR, then the use of a 25% blended stressed VaR can result in a total margin estimate which is *lower* than the margin calculated without procyclicality mitigation. This can for instance be seen in the period November 1999 to August 2000 in Figure 8.

[3] Neither of one nor the other of these two approaches dominates the other: depending on conditions and portfolios, the 'composite data' method produces a VaR estimate which is sometimes above and sometimes below the 'blend of two VaRs' method. The 'blend' method was selected simply because it makes the implementation of adaptive stressed VaR below somewhat less complex.

2.8 Floors on Margin

Broadly speaking, margin follows volatility: if volatility is lower, margin tends to be lower too. This suggests that we can use approach discussed in section 1.3 to set a floor on margin. The 10^{th} percentile of the distribution in Figure 6, for instance, corresponds to a volatility of 0.69%. We might decide that the margin model should not use a volatility lower than this. Therefore in a simple EWMA margin model which set margin as a constant times the EWMA volatility, we would set volatility as $\max(\mathit{floor}, {}^{\lambda}\sigma(t))$.

The same idea can be used for a historical simulation VaR model, albeit with a slightly more complicated implementation. We calculate the VaR for each day in the calibration period, gather these into a distribution, and set the floor based on some percentile of that distribution. Thus for instance Figure 9 shows the floored and unfloored historical simulation VaRs for a long position in USD/JPY. A floor could for instance be set at the 20^{th} percentile of the distribution of historical VaRs using the same data as used for Figure 6; the out-of-sample performance in the period 1998-2014 is illustrated.

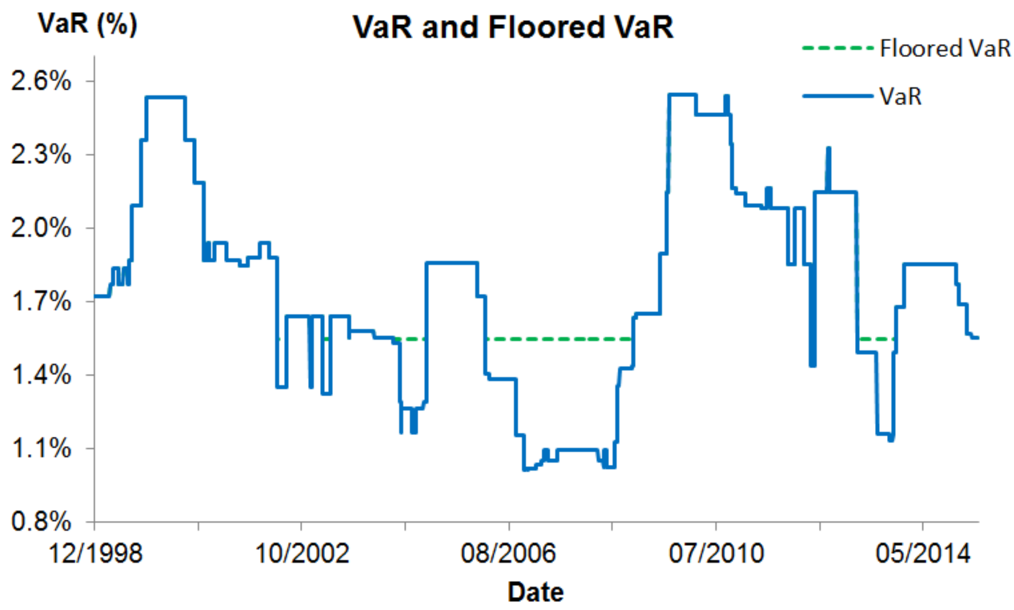

Figure 9: An illustration of the 99% VaR of a USD/JPY FX position with and without a floor set at the 20^{th} percentile of historical VaRs.

This is a slightly more flexible idea of a floor than EMIR anticipates: using a percentile of the distribution of shorter term VaRs rather than the unconditional ten year VaR allows the floor to be tuned to give the desired degree of procyclicality mitigation. Thus for instance if we had used the 20^{th} percentile of the distribution, the floor in Figure 9 would have been 1.55%: in contrast the ten year VaR is 1.68%.

2.9 Adaptive Stressed Volatility

Figure 6 suggests that there are some periods where volatility is materially lower than average, and Figure 8 reminds us that adding in a stressed VaR is only conservative if the stressed period is more stressed than the current one. One approach to address both of these issues is instead of holding constant the amount of stressed VaR, to vary the contribution depending on current conditions.

If α is this variable blending factor, the margin after procyclicality mitigation would be determined by

$$\alpha \cdot stressed\ VaR + (1-\alpha) \cdot current\ VaR$$

The idea would then be:

- If conditions are not stressed, as measured by the current EWMA volatility, then use more stressed VaR, i.e. $\alpha > 25\%$; but
- If conditions are more stressed than the ones used to determine the stressed VaR, as measured by the current EWMA volatility, then use less stressed VaR, i.e. $\alpha < 25\%$.

A reasonable approach is to use a negative exponential, so that $\alpha = \beta \exp(-\kappa\sigma)$, where σ is the EWMA volatility. In the limit where conditions are completely placid, $\sigma = 0$ so here $\alpha = \beta$. We therefore set $\beta = 0.5$ so that margin is 50% of the stressed VaR in the most placid markets. To ensure $\alpha = 25\%$ when $\sigma = \sigma^{stress}$ we set

$$\kappa = \left(\frac{\ln(0.25) - \ln(\beta)}{\sigma^{stress}} \right)$$

Figure 10 illustrates the blending factor this method produces as a function of the current level of volatility. Figure 11 illustrates the performance of the adaptive approach: relative to 25% stressed VaR, it suffers less from under-margining when current conditions are very stressed, and it keeps margin higher in placid periods.

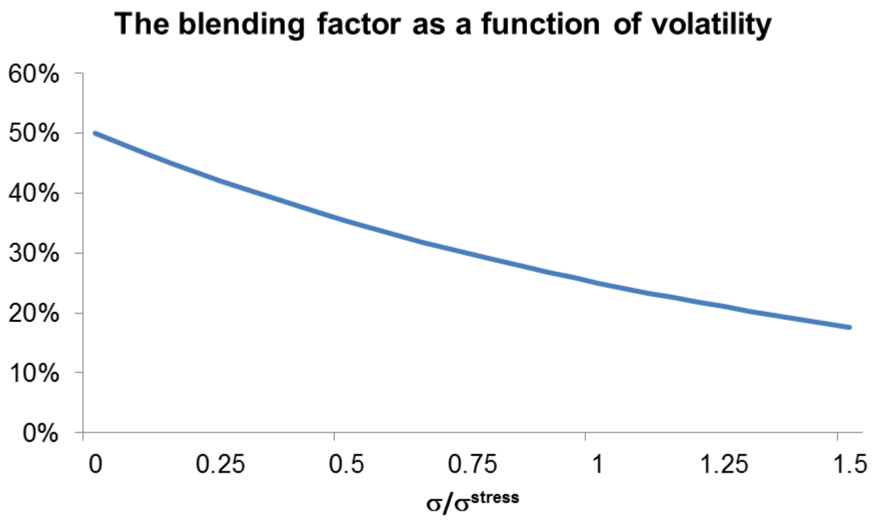

Figure 10: The adaptive stressed volatility blending factor as a function of the ratio between current volatility and the volatility of the stressed period.

2.10 Speed Limits

The final procyclicality mitigation tool considered is a limit on the speed of margin increases. This seeks to mitigate over-reaction to market conditions by limiting the size of one day margin calls. Thus for instance we could limit each day's call to, say, the 90th percentile of the historical distribution of one day changes in modeled margin requirements for the portfolio.

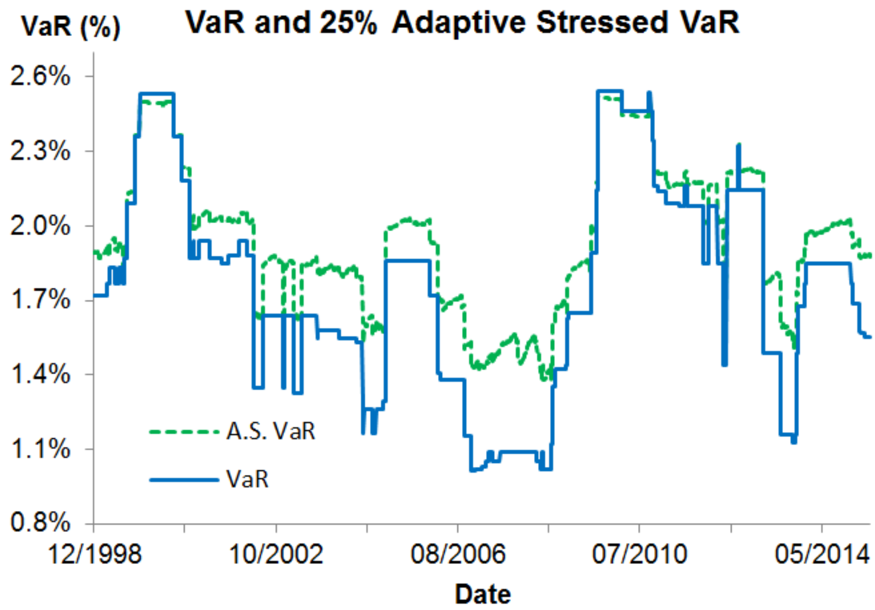

Figure 11: An illustration of the 99% VaR of a USD/JPY FX position with and without blending in an adaptive stressed VaR.

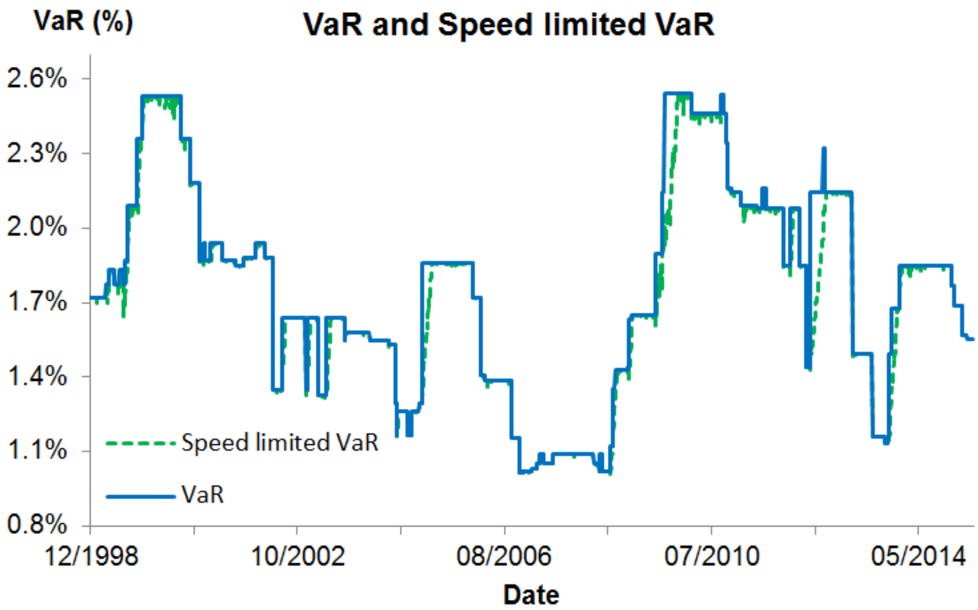

Figure 12: An illustration of the 99% VaR of a USD/JPY FX position with and without a speed limit on margin increases set at the 90th percentile of the historical distribution.

Figure 12 illustrates the effect of this tool: large margin increases, such as those beginning in June 2009, are smoothed out. Margin also takes a little longer to reach the peak level.[4]

The prudence of the speed limit is clearly a function of calibration: a high speed limit may help to prevent over-reactions while only delaying appropriate reactions by a day or two; a lower one, though, may result in too little margin being called just when it is most needed

3 Simulating a Representative Portfolio

Our illustrations thus far have used a single-asset portfolio and historical data. In what follows we rely instead on a portfolio of OTC derivatives that represents the exposures between two major dealers. The behaviour of this portfolio is analyzed in a simulation framework.

3.1 The Portfolio and Its Risk Factors

The OTC derivatives portfolio we analyse is based on public data from the Bank of International Settlements and the Depository Trust & Clearing Corporation. Together this data allows us to determine the aggregate exposures of a typical large dealer in three major classes of OTC derivative: interest rates, credit, and FX. These asset classes collectively account for around 90% of the OTC market, thus they ensure the representativeness of the portfolio. We then apply relative entropy maximization techniques (as in [14]) to reconstruct the network of exposures of this major dealer against 50 counterparties, including 3 CCPs, 15 dealers, and smaller banks. We estimate the matrix of exposures that minimize the cross entropy of each bilateral position of this dealer with each counterparty subject to a number of constraints given by the data. Our representative portfolio consists of the net exposures between the major dealer and the largest other dealer in this network of bilateral exposures.

Although derivatives portfolios can be highly complex, often consisting of many thousands of bespoke individual products, for modelling purposes it is sometimes sufficient to use a few key risk factors per asset class to capture the price and risk dynamics of the portfolio. The simulation is therefore limited to a small number of risk drivers: the 5Y USD, EUR, JPY and GBP interest rate swap par rates; 2Y USD/EUR, USD/JPY and USD/GBP currency forwards; and 3Y CDX.NA.IG and iTraxx Europe credit default swap spreads. These risk factors are all benchmarks within their asset class, and give us a reasonable range of key profit and loss ('P/L') drivers without the need to model the joint dynamics of thousands of risk factors.

3.2 The Simulation

Our modeling strategy is based on a multivariate simulation framework. Broadly, the advantage of using simulation techniques instead of historical data is that they allow the testing of a model's performance across a large number of market scenarios, and over a long time horizon. Both of these features are crucial in our analysis, because:

- We want to assess the trade-off between a margin model's procyclicality and risk sensitivity in various market conditions; and

[4] The 'noise' in the speed limited margin, for instance in April 2000, is due to the fact that the speed limit has been defined as an absolute ($) value, and margin calls come about as a result of changes in the asset price as well as changes in the relative risk estimate. The speed limit as currently defined can affect the call whatever the cause. If this is viewed as undesirable, the speed limit should instead be expressed as a percentile of the distribution of relative VaRs, i.e. VaR expressed as a percentage of the portfolio value.

- The third tool proposed in EMIR requires the use of a ten year floor. The simulation approach allows us to calibrate to the (varying lengths of) data available on each risk factor, then simulate the necessary path length many times based on that dynamics.
- In addition, one can use the data generating process of the simulation as a benchmark to evaluate the properties of different margin models.

We model the dynamics of the OTC derivatives portfolio with a copula-based multivariate model that allows for a great degree of flexibility in the model specification.[5] In particular, we allow each risk factors to follow a separate ARMA(1,1)-GARCH (1,1) process:

$$Y_t = \mu + \phi Y_{t-1} + \epsilon_{t-1} + \epsilon_t$$
$$\epsilon_t = \sigma_t Z_t$$
$$\sigma_t^2 = \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2$$

where Y_t represents the log changes in the risk factor, the innovations Z_t are white noise, and σ_t^2 is the conditional variance of Y_t. To model the conditional distribution of the standardised residuals we rely on the Extreme Value Theory (EVT). We assume a Generalised Pareto Distribution in the upper and lower tails and a non-parametric Gaussian kernel in the interior, as in [10]. Finally, a Student's t copula is used to introduce the dependence structure that links the marginal distributions of the standardised residuals for each risk factor.

The key advantage of this approach is that we can separately model the marginal distributions and the dependence structure. It thereby ensures that the simulated data reflect the real market dynamics observed in the historical data, including the comovements between different assets.

The calibration of the model uses historical data from January 2009 to June 2014, which we obtain from standard sources. Having calibrated the model, we then repeatedly (1,000 times) simulate a 12 year path of daily returns. The change in value of the portfolio over a ten day horizon is calculated each day on each path: to map the changes in the risk factors into changes in the value of the portfolio, we follow [7], and all valuations use OIS discounting as proposed in [8]. Overall, the simulation produces 1,000 12-year series of 10-day portfolio P&Ls. Collectively these offer a rich variety of market conditions to test the procyclicality mitigation tools.

3.3 Benchmark and Estimated Models

An important element in our analysis is the evaluation of the different margin models. This requires a benchmark model in order to determine for instance if the margin level is too high or too low. In reality, volatility – the main driver of margin – is unobservable: therefore, there is no obvious benchmark model for margin. The simulation framework, however, is based on a process with known specification, and a Monte Carlo simulation of this can be used as the benchmark. More specifically, at each point in the simulated data we obtain the benchmark 10-day VaR by generating 1,000 returns from the true model, and take the benchmark VaR as the 99% percentile of the distribution of these returns.

The margin model which we will use as a starting point for the tools in contrast is not based on the true returns process. Rather it estimates volatility using the well-known exponentially weighted moving average ('EWMA') approach, in our case with a decay constant of 0.97. An EWMA model is a realistic choice: EWMA volatility estimation is commonly used in margin models by participants in the OTC derivatives markets.

[5] See [13] for a review of copula-based multivariate models.

4 Results

In this section the performance of the procyclicality mitigation tools is evaluated with respect to:

- Their effectiveness in making margin less procyclical, and
- Their impact on margin levels.

Procyclicality is measured by the average 30-day margin increase (in percentage points) in the top 10% of the distribution of all margin changes. Essentially, it is the average large margin increase observed in the simulation. Over- and under-margining is measured against the benchmark model. That is, over-margining (under-margining) is the average positive (negative) difference in percentage points from the benchmark margin. A higher value of these metrics indicates a less accurate model.

We use the first 10 years of every simulation to calibrate the procyclicality mitigation tools, and then the last 2 years (500 days) to evaluate their out-of-sample performance.

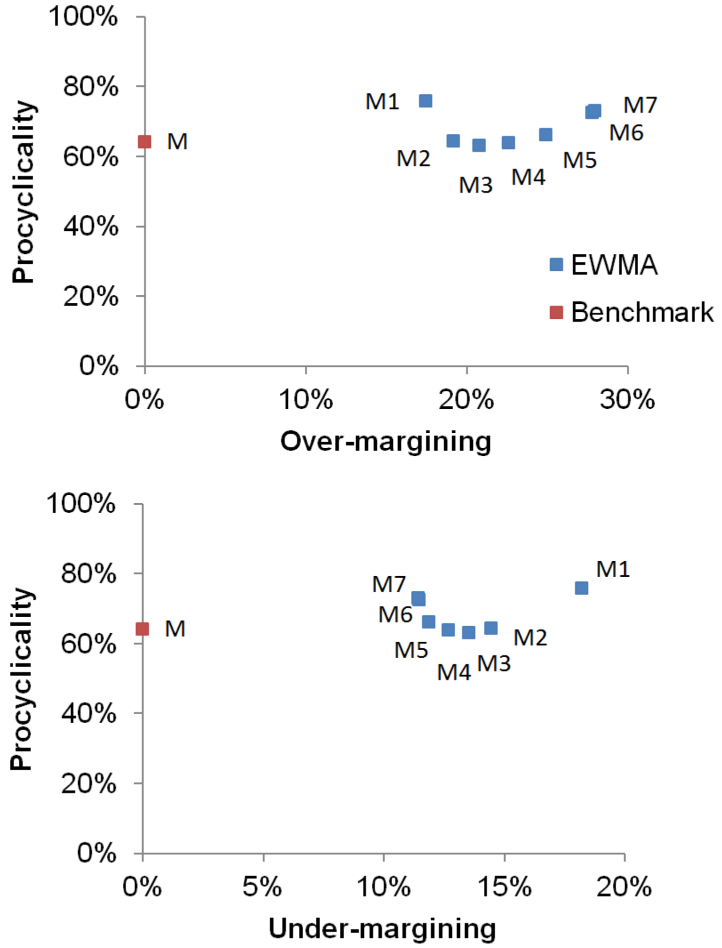

Figure 13: The procyclicality, under- and over-margining properties of the true margin model (in red) and seven models with various buffers (in blue).

4.1 Margin Buffers

We start the evaluation with the margin buffers. Here the basic idea is to ask for more margin most of the time, but to release the buffer in periods of stress.

This tool has two tuning parameters to calibrate: the buffer size and the size of a typical large margin, *StressedMargin*. The former controls how much more margin to ask in calm conditions. In our analysis we follow the proposal in EMIR RTS 28 and set the buffer equal to 25%. That is, in periods of low or average volatility we ask for 25% extra margin which can be used to smooth out any big margin increases in stress.

'Stress' is determined using some percentile of the empirical distribution of margin in the first 10 years of historical data. Intuitively, *StressedMargin* plays the role of a trigger: we release the buffer when the unmitigated margin requirement is higher than *StressedMargin*. A low *StressedMargin* will release buffer too fast, and a high one risks using the buffer too seldom.

Figure 13 shows the performance of the models in procyclicality, over- and under-margining dimensions. The models M1 to M7 (in blue), are the original unmitigated EWMA model (M1), then six mitigated models with a *StressedMargin* at the 30^{th}, 50^{th}, 70^{th}, 90^{th}, $1.1 \times 100^{\text{th}}$, and $1.2 \times 100^{\text{th}}$ percentile of the margin over the 10 year lookback period, respectively. In all cases the buffer is set at 25%. The benchmark model M (shown in red) is based on simulations of the true returns process as discussed above. All results presented are averages over the 1,000 simulation scenarios.

The differences in margin required between the benchmark, M, and the unmitigated model, M1, reflects the inherent estimation error involved in using M1.

The figure shows that for plausible levels of *StressedMargin*, the buffer approach can mitigate the large margin calls that are typically observed in periods of stress to some degree. In particular, any *StressedMargin* between the 30^{th} and 70^{th} percentile of historical margin seems to strike a reasonable balance between the mitigation of procyclicality on one hand, and the loss in risk sensitivity, on the other hand. For instance, model M4, reduces procyclicality by 24 percentage points and does not, at least for this portfolio and these risk factor dynamics, dramatically increase over-margining.

In addition, the results demonstrate the importance of the release and re-fund rule for margin buffers. Adapting a less responsive rule, such as one with a higher trigger point, does not necessarily result in a mitigated model which is less procyclical. Indeed, models M6 and M7 are as procyclical as the original unmitigated model, while at the same time asking for more margin than is necessary.

4.2 Stressed VaR

The second tool is stressed VaR, which assigns a 25% weight to data from a stressed period. This approach aims to limit procyclicality by preventing margin from plummeting during periods of low volatility. The key calibration parameter is the definition of the stressed period. The obvious choice is to use some percentile of historical volatilities, and this is what we do.

We review 6 mitigated models, M2-M7: these are the original EWMA model with 25% of a stressed volatility equal to the 30^{th}, 50^{th}, 70^{th}, 80^{th}, 95^{th}, and $1.2 \times 100^{\text{th}}$ percentile of the distribution of EWMA volatilities in the 10 year lookback period, respectively. As before, M1 is the unmitigated EWMA model, while M is the benchmark model.

Figure 14 presents the performance of these models. It can be seen that the stressed VaR is effective at mitigating procyclicality at least for the range of stressed volatilities examined here. However, this tool can come at the expense of asking for more or less margin than that of the benchmark model depending on calibration. M6 and M7 illustrate the former: M2 shows the latter phenomenon.

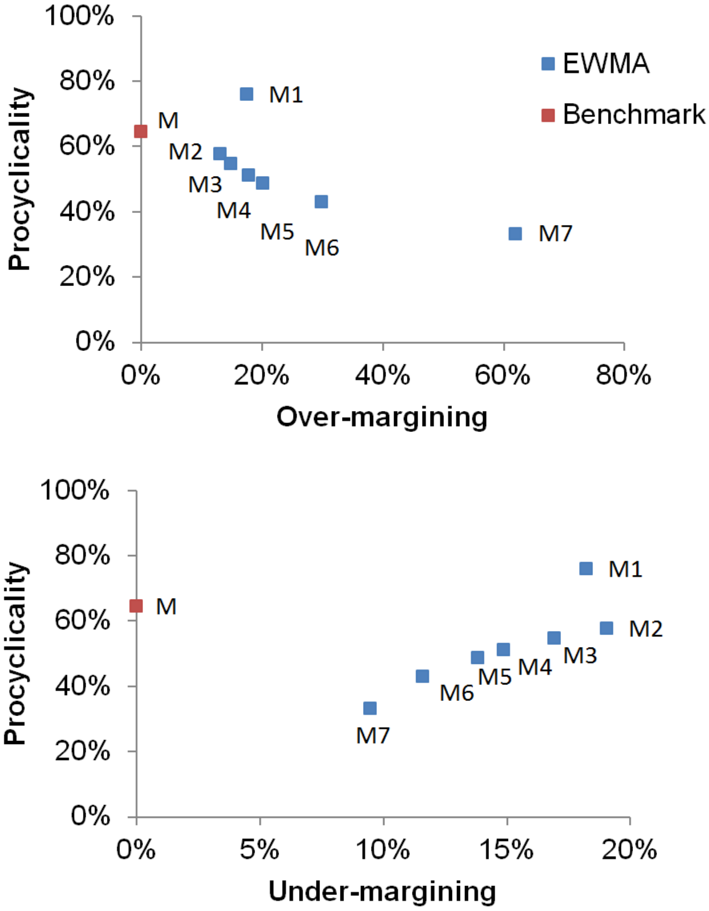

Figure 14: The procyclicality, under- and over-margining properties of the true margin model (in red) and seven models with 25% of different stressed periods (in blue).

4.3 Adaptive Stressed VaR

The results in the previous section demonstrate that a model with stressed VaR procyclicality mitigation can exhibit some degree of under-margining for low levels of stressed volatility. The reason is that when the stressed volatility is not stressed enough, the mitigated margin model can underestimate the underlying volatility. As we saw, this is the case for model M2 in Figure 14, for example. The inclusion of a mis-specified stressed period makes the margin after mitigation less prudent during market turmoil. This is obviously an undesirable property.

As discussed in Section 2.9, one approach to address this issue is instead of applying a constant 25% weight to stressed volatility, to apply a variable blending factor which adapts to the prevailing market conditions. In Figure 15, we present 6 models which apply a variable weight to stressed volatility. Similarly to the stressed VaR case discussed above, model M1 is the original EWMA model, and M2 to M7 are this model mitigated with a stressed volatility equal to the 30th, 50th, 70th, 80th, 95th, and 1.2×100^{th} percentile of the distribution of EWMA volatilities in the 10 year lookback period, respectively. The difference is that this time we allow the blending factor to vary with current volatility as described in Section 2.9.

The results presented in Figure 15 show that adaptive stressed VaR can both mitigate procyclicality and reduce under-margining. It is an improvement over the standard stressed VaR. In addition, the decrease in procyclicality using adaptive stressed VaR is greater for 'more stressed' volatilities: however in these cases, over-margining is greater too.

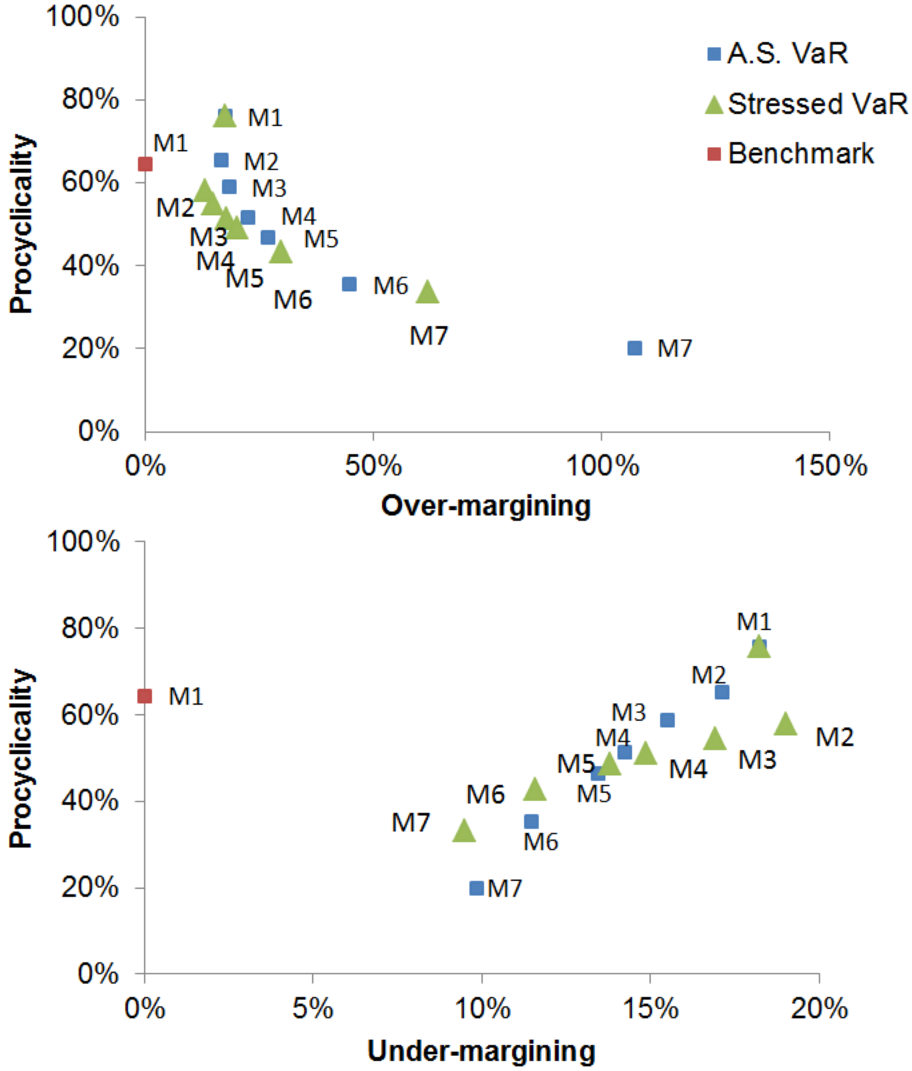

Figure 15: The procyclicality, under- and over-margining properties of the true margin model (in red), seven models with 25% of different stressed periods (in green), and seven models with an adaptive amount of the same stressed periods (in blue).

4.4 Floor on Margin

The next tool is a floor on margin. The basic intuition of this tool is that margin should not fall too far in periods of low volatility. As a result, the transition to a stressed period in a floored model yields smaller margin increases.

The empirical distribution of margin from the first 10 years of the stimulation was used to calibrate the floor. We estimated 6 different margin models, M2-M6: these are the original EWMA model with a floor at the 10^{th}, 30^{th}, 50^{th}, 70^{th}, 90^{th}, and 100^{th} percentile of the margin over the 10 year historical lookback period, respectively. Figure 16 shows that even a low floor can mitigate procyclicality. For example, by setting the floor equal to the 10^{th} percentile of the margin in the historical 10-year period, we can reduce procyclicality by almost 10 percentage points on average. We also find that the higher the floor, the less procyclical the model is. The drop in procyclicality comes at the cost of asking for more margin than needed, as the top panel illustrates. This is especially the case during calm times. Finally we note from the bottom panel that the risk of asking too little margin decreases with the size of the floor.[6]

[6] It should be noted that the under-margining of the EWMA models is observed because we have assumed that it is

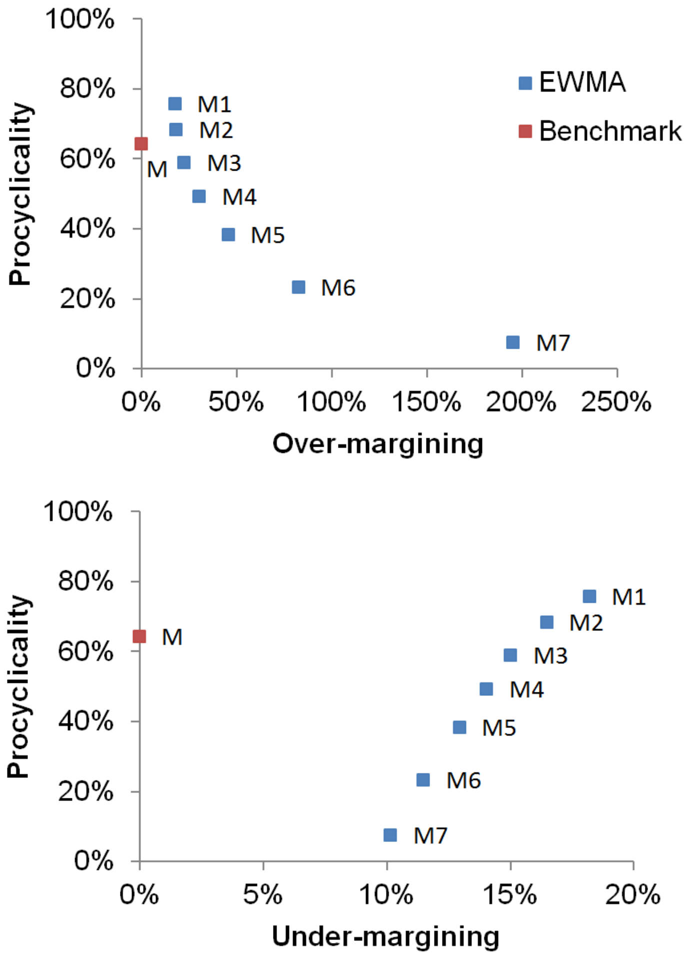

Figure 16: The procyclicality, under- and over-margining properties of the true margin model (in red) and seven models with different floors (in blue).

4.5 Speed Limits

The last procyclicality mitigation tool is the speed limit. The main characteristic of this tool is that it seeks to smooth out margin peaks by limiting the acceptable size of one day margin increases.[7] As a result, a speed limit may adversely affect a model's risk coverage, particularly in periods of elevated volatility.

We consider 6 models with a speed limit. Models M2-7 are the original EWMA model with a speed limit equal to the 90^{th}, 70^{th}, 60^{th}, 50^{th}, 40^{th}, and 30^{th} percentile of historical margin increases in the 10 year lookback period, respectively.

Figure 17 shows that a high speed limit may help to limit procyclicality without overly increasing under-margining. As we decrease the speed limit, however, the adverse consequences increase too. For example, under-margining increases by approximately 10% when we move from the unmitigated EWMA model (M1) to the model with a speed limit equal to the 30^{th} percentile of historical margin increases (M7).

---

a less accurate model than the benchmark model. Hence, there will be instances when the margin held is too low relative to the "true" margin levels predicted by the benchmark model.

[7] A speed limited reactive margin model is similar in some respects to a less reactive one. Therefore one way of thinking about the trade-off between more and less reactive margin models – EWMA models with low λ compared to those with higher λ for instance – is to think of the latter as the former plus a speed limit.

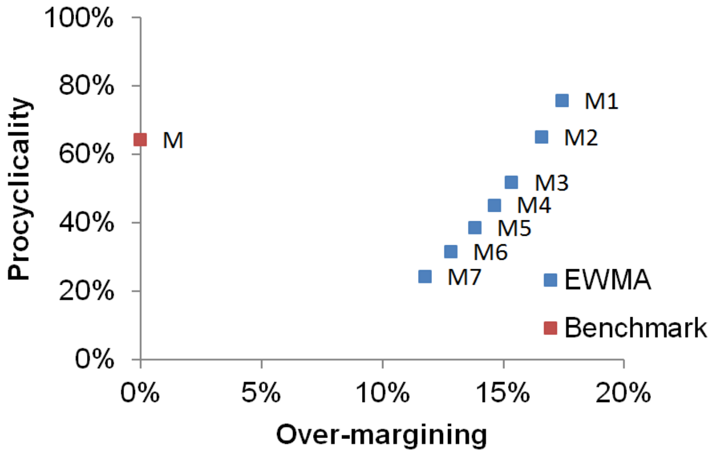

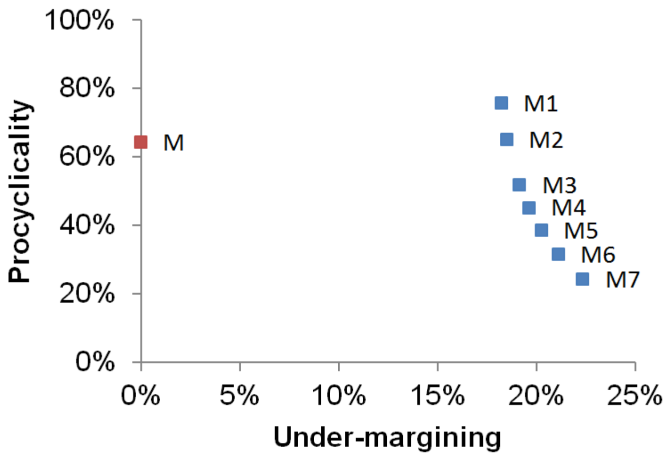

Figure 17: The procyclicality, under- and over-margining properties of the true margin model (in red) and seven models with different speed limits (in blue).

5 Conclusions and Further Work

We have distinguished between two issues:

- An initial margin model should not *under-margin* too much, as evidenced for instance by passing a suitable suite of backtests. This ensures that the level of margin called does not leave the posted party over-exposed to the failure of its counterparty;
- It may however be acceptable sometimes to *over-margin* to some degree, if this prevents margins from falling too far, and thus rising too fast when conditions change.

The under- and over-margining behaviour of an initial margin model with various forms of procyclicality mitigation has been presented. Our results indicate that EMIR's focus on procyclicality is important, and that broadly all of the EMIR tools do some good in mitigating one or both forms of procyclicality.

That said, the EMIR Regulatory Technical Standards leave unspecified certain key details of the tools, such as the release and re-fund rules for the buffer, or the precise definition of how the stressed market conditions used in the stressed VaR should be determined, or what should be done with them once determined. Several points are relevant here:

- The impact of a tool depends on the particularities of the situation (i.e. the model being mitigated, the portfolio under consideration, and the risk factor history used). For our situation at least, even a relatively small amount of any of the tools studied can provide significant procyclicality mitigation often without substantially worse under- or over-margining.

- The optimal calibration of each tool in a particular situation depends on the relative weights placed by the modeller on the objectives of minimizing procyclicality on the one hand and minimizing undesirable over-margining in periods of low volatility on the other.
- Different tools have different behaviour in different dimensions. The ten year VaR floor, for instance, is good at mitigating across the cycle procyclicality, but it will do nothing to reduce the impact of large margin calls when margin is already over the floor level.
- There are tools such as adaptive stressed VaR which are as good or better than the one of the prescribed tools at mitigating procyclicality for a given cost in both under- and over-margining.

These considerations lead to several tentative policy conclusions. First, there is an unequivocal case for CCPs calculating and appropriately disclosing some measure of the across-the-cycle procyclicality and the short term procyclicality of their models. Our earlier work [12] presents several measures which might be appropriate here.

Second, procyclicality mitigation may be required either in ensuring margin does not fall too far; or in ensuring that it does not go up too far; or in both aspects. This suggests that CCPs should continue to be free to determine the best tool given their modelling strategy. It does however suggest that a stated policy on their tolerance for procyclicality may be appropriate.

Third, innovation in procyclicality mitigation should be encouraged. We have examined two tools which are not contemplated in the RTS but which may be of interest: there are doubtless others. At a minimum, there is a case for a statement that any tool which performs at least as well as a permitted one on suitable metrics such as the ones presented above should be permitted too.

There may also be an argument as tools are developed regulation should eventually move to an outcome-based approach rather than a tools-based one. That is, the requirement could simply be, for instance, that a CCP must demonstrate that the peak-to-trough procyclicality of all of its initial margin models is less than 3 for the five riskiest portfolios it clears, or that the 30-day call for the same portfolios was not more than 50% of the starting margin amount for the same portfolios. More research in the appropriate standard is needed before it would be appropriate to move to this approach, but we see it as a promising policy direction.

References

[1] N. Abruzzo, Y. Park, *An empirical analysis of future margin changes: determinants and policy implications*, Federal Reserve Board Discussion Paper 2014-86 (2014)

[2] Basel Committee on Banking Supervision, *Basel III: International framework for liquidity risk measurement, standards and monitoring*, December 2010

[3] M. Brunnermeier, L. Pedersen, *Market liquidity and funding liquidity*, Review of Financial Studies Volume 22 No. 6 (2009)

[4] Committee on the Global Financial System, *The role of margin requirements and haircuts in procyclicality*, CGFS Papers No. 36 (March 2010)

[5] European Union, *Regulation (EU) No. 648/2012 on OTC derivatives, central counterparties and trade repositories* (July 2012)

[6] European Union, *Regulation (EU) No. 153/2013 supplementing Regulation (EU) No. 648/2012 with regard to regulatory technical standards on requirements for central counterparties* (December 2012)

[7] J. Hull, *Options, Futures and Other Derivatives*, Prentice Hall (2014)

[8] J. Hull, A. White, *Libor vs. OIS: The derivatives discounting dilemma*. The Journal of Investment Management, Forthcoming

[9] S. Mayhew, A. Sarin, K. Shastri, *The allocation of informed trading across related markets: an analysis of the impact of changes in equity-option margin requirements*. The Journal of Finance, Volume 50, No. 5 (1995)

[10] A. McNeil, R. Frey, *Estimation of tail-related risk measures for heteroskedastic financial time series: an extreme value approach*. Journal of Empirical Finance, Volume 7, No. 3 (2000)

[11] D. Murphy, *OTC Derivatives: Bilateral Trading and Central Clearing*, Palgrave (2013)

[12] D. Murphy, M. Vasios, N. Vause, *An investigation into the procyclicality of risk-based initial margin models*, Bank of England Financial Stability Paper No. 28 (2014)

[13] A. Patton, *Copula methods for forecasting multivariate time series*. Handbook of Economic Forecasting Volume 2 (2012)

[14] C. Upper, A. Worms, *Estimating bilateral exposures in the German interbank market: Is there a danger of contagion?* European Economic Review Volume 48, No. 4 (2004)