Procyclicality of CCP margin models: systemic problems need systemic approaches

Pedro Gurrola-Perez[1]

December 2020

Abstract: Margin requirements protect a central counterparty (CCP) and its users against potential losses generated by the default of any of its members. They have several components, one of them being the initial margin requirement, which is typically calculated using a market risk model to estimate the potential future exposure of each member's portfolio. By definition, market risk models - whether for centrally cleared or bilateral cleared trades - have to be sensitive to changes in market risk and, as a consequence, when market risk increases, initial margin requirements will tend to increase. After the 2008 crisis, regulators had concerns about this becoming "procyclical", in the non-technical sense of amplifying financial stress. As a result, CCPs have put in place different procyclicality mitigation tools. But when the markets are stressed and participants face larger margin calls, like in the recent events of March 2020, interest on the procyclicality of initial margin models seems renewed. In this paper we argue that the focus on initial margin models is misplaced. First, margin calls are largely driven by variation margin, not initial margin. Second, the inherent risk sensitivity of margin models, the stochastic nature of the problem, and the different trade-offs involved, constrain what can be achieved with model calibration. We illustrate why this is the case by empirically testing the performance of standard initial margin models during the recent March 2020 events and quantifying the different trade-offs. Therefore, if procyclicality of IM has been mitigated to the limit of what is practical and prudent, but fragilities in the system persist, how should these be addressed and who is responsible for addressing them? We argue that, since the ultimate objective is to minimize systemic propensities to adverse feedback loops, these questions demand a systemic perspective, focusing on the interactions between participants rather than on a single node.

Keywords: Central counterparties (CCPs), procyclicality, initial margin (IM), margin models, filtered historical simulation.
JEL classification: G17, C60, G23, G01.

---

[1] World Federation of Exchanges. Email: pgurrola@world-exchanges.org

Acknowledgements: I would like to thank Stefano Alderighi, Jeanne Balcom, Flavia Barsotti, Fernando Cerezetti, John Fennel, Richard Fenner, Gerardo Ferrara, Emily Hendrix, Michael McClain, Chia I Ming, Richard Metcalfe, David Murphy, Ketan B. Patel, Christer Rydberg, Nandini Sukumar and members of the WFE CCP Working Group for helpful comments and discussions. Any errors are solely mine.

WFE Working Papers: WFE Working Papers present research in progress by the author(s) and are published to share early insights, stimulate debate, and elicit feedback. Comments on Working Papers are welcomed, and may be sent to Head of Research, World Federation of Exchanges, research@world-exchanges.org. All rights reserved.

.

1 Introduction

During the first weeks of March 2020 the COVID-19 pandemic placed the global financial system under severe stress, with the volatility index (VIX) hitting levels higher than those seen during the Great Financial Crisis of 2008, asset prices plummeting, and a severe deterioration of market liquidity, including in assets - like US Treasuries - traditionally seen as the most liquid. Looking back at these events, there is no doubt that, as in 2008, financial market infrastructures (FMIs) proved to be resilient, demonstrating their contribution to systemic stability and to fully functioning markets.

The post-2008 reforms included, among many other things, considerations about the procyclicality of the margins that central counterparties (CCPs) require from their members. These margin requirements are a fundamental part of the CCPs' management of counterparty risk, ensuring that the exposure to a failing member will be sufficiently collateralized and, as consequence, the CCP will effectively act as a buffer against financial contagion.

Margin requirements include variation margin and initial margin. To avoid the building-up of risk and limit their current exposure, CCPs mark-to-market positions and call variation margin on a daily basis (and sometimes intraday). To protect against potential future exposure in the event of a member default, the CCPs also call for initial margin.

While it is recognized that adequate initial margins are fundamental for the soundness of the CCP and for the resilience of the system as a whole, there was a concern that, in a situation of financial stress, an unexpected increase in margin requirements would add liquidity pressure to the members at a time when they are already struggling, thus creating a feedback loop that could amplify the stress and potentially compromise financial stability. To reduce the effects of such "procyclical" behaviour with respect to the market, during the last 10 years or so, regulators and industry have put several measures in place; for example, an increased transparency of CCP risk management practices, including requiring CCPs to regularly disclose information about their initial margin models and their calibration, and the requirement for CCPs to monitor the procyclicality of their margin requirements.[1] Some jurisdictions have also prescribed specific requirements for the calibration of margin models, in an attempt to modulate their responsiveness.

Trying to reduce impacts on liquidity in the system through the calibration of CCPs' initial margin models has the obvious limitation that market-risk models - irrespective of whether they are used for central clearing, bilateral clearing or in the banking sector (notably in VaR models) - cannot simply *ignore* changes in market conditions. In the case of CCPs, the models' function is to alert risk managers to an increase in risk and to provide an estimate of the collateral needed to cover the CCP's potential future exposure that is consistent with the CCP's risk tolerance. If the model underestimates the level of risk and a member defaults, the CCP may find itself facing large partially uncollateralised losses. There is a three-way set of trade-offs between ensuring appropriate levels of risk sensitivity in the model, reducing its procyclicality, and keeping central clearing economically efficient. Regulation recognizes these trade-offs and requires that CCPs should reduce procyclicality but only to the extent possible, as it may be impractical and even imprudent for a CCP to establish initial margin requirements that are independent of significant or cyclical changes in price volatility (CPMI, 2012).

Regardless of the fact that procyclicality mitigants are already in place, that the principle of preserving the safety of the CCP is recognized as prevalent, and that there has been an important amount of research on the topic, when the system is under stress and margin calls increase, as has happened during the COVID-19 crisis in March, the question of whether CCP's initial margin models are too procyclical seems

[1] The PFMI Quantitative Disclosures are reported quarterly by each CCP and are publicly available at the CCPs' websites.

to come back, including with renewed calls to revise the calibration of these models. These arguments often fail to recognize that margin calls include calls for initial margin (IM) and for variation margin (VM); that it is often VM (not IM) which tends to drive margin calls, usually by a much larger amount (Maruyama and Cerezetti, 2019; ESRB, 2020; FSB, 2020);[2] or that margin calls can also be driven by changes in the composition of the underlying portfolio or in the value of the collateral posted. In fact, when evaluating the role of CCPs during the COVID-9 crisis, policymakers note that CCPs remained resilient during the market turbulence and observed that while CCPs accept certain non-cash assets for IM, "anecdotal evidence indicates that firms chose to post cash instead, which raises questions as to whether IM calls added to cash demands and funding markets stress" (FSB, 2020).

In this note, we argue that the reason why a satisfactory answer to IM procyclicality seems to be eluding us is not because we haven't focused enough on margin models, or in their calibration, but rather the opposite; and that to move forward on this debate, it is essential to approach the problem instead from a systemic perspective. We can see at least three reasons for this change of approach.

First, IM models are inherently risk sensitive, which sets a limit to how much procyclicality can be reduced through model recalibration and without compromising the CCP's safety or the economic viability of central clearing.

Second, the stochastic nature of the problem implies that the models' output is not fully determined by the model parameters and by the initial conditions, and therefore one cannot expect simple deterministic relations to hold between inputs and output, let alone linear ones. It is not true, for example, that longer lookback periods *always* lead to more conservative (i.e. higher) margins, or that a lower number of breaches over a period is necessarily an indication of more conservative risk management. It is not that simple. Even if we were dealing with independent, identically distributed (*iid*) processes (and we are not), there is some degree of sampling error that needs to be accounted for. In addition, financial time series are known to have varying conditional volatilities, with periods of high and periods of low volatility, implying that properties that hold in the long run, converging to those of the unconditional distribution, may fail dramatically when considering market conditions at a given point in time. While longer lookback periods, for example, may contribute to the estimation accuracy for the unconditional distribution, they do not necessarily help with the accuracy of the conditional distributions. Crucially, the resilience of the CCP (and of the system) does not depend on long-run averages but on getting the numbers right locally, under the prevailing market conditions at a specific point in time.

To illustrate these two points, we will consider a simple empirical example to show how typical margin models respond to changes in the parameters or in the initial conditions. The example will also highlight why, because of the non-deterministic nature of the models, attempts to address procyclicality through the calibration of the margin model can only have limited success. The example will consider the S&P500 series as a sample (static) portfolio and will calculate the daily IM requirement of that portfolio using some of the standard models and risk measures used by CCPs: Value-at-Risk (VaR) estimated using historical (filtered and unfiltered) simulation.[3] While other studies have evaluated the performance of procyclicality mitigation tools using a simulation framework (Murphy, Vasios and Vause, 2016), our analysis will consider the COVID-19 events to perform the analysis. To avoid hindsight bias, we will calibrate the models using the period 2007-2019, and will keep the calibration fixed when testing the models' performance through the first months of 2020. We will also examine how the different models would have performed in a worst-case scenario (a member default), and what would have been the consequences for the CCP under different choices.

[2] According to ESRB (2020), for both bilateral and centrally cleared derivatives transactions, VMs account for approximately 80-90% of the net daily cash flow payments from margin calls. Similarly, a recent report by the FSB concluded that during this period of extraordinary volatility "daily demands due to additional initial margin requirements were notably smaller when compared to the same-day VM calls" (FSB, 2020).

[3] While individual CCPs may introduce variations of these models, the versions considered here are sufficiently representative to capture the nature of the problem.

Therefore, important questions remain: while we all agree that CCPs should use the least procyclical model, provided it is prudent, safe and practical, if fragilities in the system remain, as in the dash-for-cash experienced as part of COVID-19 crisis, how should these be addressed? Just as in the CCP resolution debate (Cerezetti et al., 2019), the question here seems to be ultimately the same: should authorities act as repo counterparty of last resort to provide liquidity, as it happened during March 2020; should CCPs sacrifice their safety to reduce participants' liquidity risk; or should clearing participants improve their liquidity management practices to deal with potential future margin calls?

This brings us to the third reason why there needs to be a change of perspective. It is well-known that financial systems are complex systems (Johnson, 2003). Complexity is characterised, among other features, by feedback loops, hierarchical structures, strong interdependencies, and non-linearity. The propensity to feedback loops is a property of the interactions in the system and not of its individual elements. Since concerns around procyclicality are fundamentally about the generation of adverse feedback loops that could amplify a stress, then the focus should be directed to understand and control the system interactions and not solely to change the behaviour of one single node or agent. Complex system failures are not prevented by ensuring that each component works properly but by defensive strategies that can mitigate negative interactions.[4] Drawing parallels with how system safety is addressed in complex systems will help us to understand why, focusing on one component of the system (the IM model) may have limited impact in addressing the safety of the system (financial stability). There is limited gain in having a correctly calibrated and minimally procyclical IM model if other parts of the system are too sensitive to or inadequately prepared for a sharp increase in risk (and, consequently, in margin calls), let alone to one of the magnitude seen during the March 2020 events. While the need for a systemic view is far from new, its importance is being increasingly recognized by regulators.[5]

The rest of the paper is organised as follows. Section 2 is a brief summary of the main regulatory requirements concerning procyclicality. Section 3 discusses some of the related literature. In Section 4 we discuss the different components of margin requirements, including on an intraday basis. Section 5 discusses the methodology and the metrics that will be applied. In Section 6 we present and discuss the results of the analysis. Section 7 presents a discussion of the results from a complex systems theory perspective. In Section 8 we summarize the conclusions.

2 Regulatory context

According to the Principles for Financial Market Infrastructures, or PFMIs, (CPMI, 2012), the global principles set by CPMI and IOSCO, procyclicality typically refers to:

> (3.6.10) changes in risk-management practices that are positively correlated with market, business, or credit cycle fluctuations and that may cause or exacerbate financial instability. For example, in a period of rising price volatility or credit risk of participants, a CCP may require additional initial margin for a given portfolio beyond the amount required by the current margin model.

While this initial definition places the emphasis on changes in risk management required when the margin model has under-performed, the PFMI adds that margin requirements themselves can constitute another

[4] The fact that a system can be unsafe even when all its components are reliable, and the other way around, is a fundamental issue when approaching safety in complex systems (Leveson, 2011).

[5] For example, Jon Cunliffe, the Bank of England Deputy Governor for Financial Stability, recently recognized that "given that clearing and margining are important risk mitigants in the system, the answers may lie more in ensuring that financial market participants, be they hedge funds or pension funds, understand how margin call can evolve in a stress and have the resilience to manage the consequent liquidity pressures" (Cunliffe, 2020). Similarly, in their analysis of the March 2020 events, the FSB notes that while increases in margin are to be expected in volatile markets, "some market participants may not have anticipated the size of the timing of the increase in margin requirements, and so needed to utilise cash buffers or obtain further funding to obtain those margin requirements" (FSB, 2020).

source of procyclicality:

> These adverse effects may occur without any arbitrary change in risk-management practices. To the extent practicable and prudent, a CCP should adopt forward-looking and relatively stable and conservative margin requirements that are specifically designed to limit the need for destabilising, procyclical changes. To support this objective, a CCP could consider increasing the size of its prefunded default arrangements to limit the need and likelihood of large or unexpected margin calls in times of market stress.

The trade-off between risk-sensitivity and procyclicality is recognized and it should be resolved in favour of the first, as it may "be impractical and even imprudent for a CCP to establish margin requirements that are independent of significant or cyclical changes in price volatility" (CPMI, 2012). And there is a recommendation for the CCP to consider increasing its prefunded arrangement "to limit the need and likelihood of large or unexpected margin calls in times of market stress".

It is clear that there are two types of situations that could lead to excessive procyclicality. One is derived from an IM model overreacting to a change in market conditions and generating larger margin calls in times of stress and too low requirements in tranquil times. And the other arises from models under-reacting to an increase in risk, forcing the CCP to recalibrate or override the model to respond to the new circumstances which the model failed to anticipate. Decreasing the risk sensitivity of the model to address the first type leads to an increase in the likelihood of facing the second one. Of course, one can argue that procyclicality can be reduced without affecting risk sensitivity by increasing the prefunded requirements in tranquil times,[6] but this option can have negative systemic implications, increasing the cost of central clearing (which could potentially conflict with the incentives underpinning the clearing mandate), and reducing the overall liquidity.

Some jurisdictions took a more prescriptive approach to procyclicality. In EMIR (ESMA, 2013), for example, it is assumed that procyclicality can be addressed through an adequate choice of model parameters and inputs: the confidence interval, the liquidation period, and the lookback period used for the calculation of historical volatility (EMIR, par. 5, Art 41). Moreover, CCPs are required to alter their model applying at least one of three specific anti-procyclicality (APC) measures (EMIR TS, Article 28):

(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.

Different authors have pointed out the limitations of these approaches (which we will refer to as the "buffer", the "stress", and the "floor").

On the buffer approach, Glasserman and Wu (2016) observed that the 25% is arbitrary and it should instead be linked to the "persistence" and "burstiness" of the underlying process and, consequently, to the asset class considered.[7] On the other hand, (Murphy, Vasios and Vause, 2016) pointed out that, since no one has foreknowledge of how big a given shock is going to be, it seems impossible in advance to know the right time to release the buffer, so there is the risk of constructing an approach which releases the buffer too early failing to absorb the most intense stress, or constructing one which releases it too late,

[6] Clearly, requiring a uniform increase of IM *at all times* will not have any impact on procyclical behaviour.

[7] "Burstiness" is a concept usually associated to the dynamics of a wide range of real systems, from email patterns to earthquakes, which display a bursty, intermittent nature, characterized by short timeframes of intensive activity followed by long quiet periods. While somehow related, this seems different from the meaning intended in Glasserman and Wu (2016).

and so provides too little procyclicality mitigation. There is also the question of when the buffer should be replenished.

The stress approach does not identify which observations should be considered "stressed", and it does not define what it means to give these observations a 25% weight. It could mean artificially adding stressed observations to the lookback period (which would distort the backtesting process), or weighting the current volatility with some stressed volatility estimate, although this could lead to disproportionate margin increases in periods of low volatility (Maruyama and Cerezetti, 2019). The requirement has also been interpreted as mixing VaR with a "stressed VaR". In this case, the risk is that, if conditions become even more stressed than those that were used to calibrate 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 (Murphy, Vasios and Vause, 2016). In fact, the example we present illustrates precisely this situation.

On the floor approach, Glasserman and Wu (2016) observed that the rule would be more effective if it referred to estimating a quantile over a long lookback period, rather than a volatility. On the other hand, Murphy, Vasios and Vause (2016) have suggested using a percentile of the distribution of shorter term VaRs, rather than the unconditional 10-year VaR, to allow the floor to be tuned to give the desired degree of procyclicality mitigation.

In the US, as it currently stands, no prescriptive legal text exists regarding how CCPs should address procyclicality of their margin models. However, the CFTC, as a principles-based regulator, has highlighted how a big margin call from a large US CCP could force its members to act in ways that are potentially destabilising (CFTC, 2013). Mitigating such behaviour would be considered necessary pursuant to the obligations for initial margin requirements. Moreover, the US and EU equivalence agreement for CCP requirements lists, as one of the conditions for mutual recognition, the inclusion of measures to mitigate the risk of procyclicality.[8]

3 Related work

Broadly speaking, we can distinguish three strands in the analysis of the procyclicality of CCP margins. The first attempts to theoretically model the interactions of the financial network in which CCPs operate and the dynamics derived from participants' behaviours and incentives. The role of margins in amplifying a spiral between funding liquidity and market liquidity was modelled by Brunnermeier and Pedersen (2009). In contrast with that view, Biais et al. (2016) emphasize the positive effect of margins on risk-prevention incentives. Using also a stylized CCP model, Raykov (2018) finds that less procyclical margins do not unambiguously improve financial stability, and that the net effect depends on the market characteristics and on the CCP members risk aversion: a combination of high volatility plus insufficient risk aversion can lead to additional risk taking, creating exposures that cannot be predictably controlled or managed. Domanski et al. (2015) provide preliminary econometric evidence suggesting that banks operating in systems where a larger portion of transactions were cleared by CCPs were less likely to suffer a significant deterioration of solvency during the Great Financial Crisis of 2008. Faruqui et al. (2018) re-examine the liquidity link between CCP margin and banks' balance sheets and describe the potential interactions between banks and CCPs in periods of stress, concluding that, because these interactions could potentially lead to destabilising feedback loops, the risks of banks and CCPs should be considered jointly, rather than in isolation. King et al. (2020) argue that, because of the complexity of the interconnections and the fact that CCP-related liquidity needs are inherently procyclical, liquidity-focused macroprudential stress tests are needed to help manage systemic liquidity risk.

A second strand of work attempts to empirically model the CCPs financial network, which then is stressed to examine how the shock propagates and what are the consequences, including detecting

[8] https://www.cftc.gov/PressRoom/PressReleases/cftc_euapproach021016

situations that could generate adverse feedback loops. It requires access to very granular data covering a sufficiently large set of participants, plus the computational capacity to process big data sets. It is therefore not surprising that most of the research is done by regulators. Bardoscia et al. (2019), for example, investigate whether margin calls on derivative counterparties could exceed their available liquid assets and spread liquidity shortfalls through the market. They simulate variation margin calls in a stress scenario and compare these with the liquid-asset buffers of the institutions facing the calls. They find an aggregate liquidity shortfall equivalent to only a small fraction of the average daily cash borrowing in international repo markets. With a different approach, O'Neill and Vause (2018) look at the impact of a macroprudential buffer on top of microprudential initial margin requirements as a way of avoiding fire sales externalities. Interestingly, they find that, depending on how the buffer is calibrated, it may have radically different and even harmful effects.

The third approach focuses on the IM model itself and its input/output relation. Abruzzo and Park (2014) analysed reported futures margin requirements in the US and observed that they increased quickly as volatility increased beyond a threshold but they did not immediately decrease when volatility dropped down. Based on this observation they argued that, although risk-based margins help protect the CCP, they could worsen the funding conditions of clearing members during economic downturns, and suggested that regulators should consider introducing some tools, such as through-the-cycle margin or countercyclical buffers, to dampen procyclical behaviour of margins. Murphy, Vasios and Vause (2014) propose to address the conflict between procyclicality and risk-sensitivity by restating the question as a choice between different models with the same risk sensitivity but different procyclical behaviour. In such a case, one should select the less procyclical model, or avoid models that are "overly" procyclical. They propose three measures of procyclicality and, using these measures, they find that models which pass common risk sensitivity tests can have very different levels of procyclicality. Also using these measures, Gurrola-Perez and Murphy (2015) analyse the impact on procyclicality of filtered historical simulation (FHS) models. Glasserman and Wu (2016) investigate procyclicality assuming that the margin requirement is proportional to volatility, and that volatility follows a GARCH(1,1) process. Under these assumptions, they show how the coefficients of the GARCH specification characterise the persistence and burstiness of the process, implying that different asset categories should be treated differently.

Murphy, Vasios and Vause (2016) study a variety of tools which have been proposed to mitigate the procyclicality of initial margin requirements. The paper examines, within a simulation framework, the extent to which each tool mitigates procyclicality and at what cost (where the cost is obtained by comparing with a benchmark unmitigated model). Their results 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 trade-off between minimizing procyclicality and minimizing the costs of overmargining in periods of low volatility. The authors conclude that it may be appropriate to consider moving from tools-based procyclicality regulation to one based on the desired outcomes. Maruyama and Cerezetti (2019) look closely at how some of the antiprocyclicality measures work in practice and conclude that the best mitigation for procyclicality may be achieved through the establishment of an outcome-based approach and enhanced disclosure of margining practices.

4 Margin calls: IM, VM and add-ons

When discussing procyclicality of margin requirements, it is important to note, first, that margin requirements (and margin calls) consist of different components (mainly variation margin, initial margin and add-ons); and, second, that these requirements can also change due to changes in the composition of the portfolio.

Variation margin (VM) is the result of daily (and sometimes intra-day) mark-to-market of positions to cover any *current exposures* (CE).

Initial margin (IM) is the collateral required to cover *potential future exposure* (PFE); that is, the exposure that the CCP would face in the event of a member default. It is estimated at least daily and often also intraday.

IM Add-ons correspond to additional margin that is required, on top of IM, to cover other risks that could also crystallize in the event of a member default. These add-ons can include concentration/liquidity risk add-ons (to cover the risk that a clearing member's portfolio is too large or concentrated, and therefore more difficult and costly to liquidate in a default scenario), delivery risk add-ons (for securities with physical delivery), wrong-way risk add-ons, or credit risk add-ons (for when a particular member's inherent risk exposure is not captured within the IM model).

While VM is paid/received in cash and its exchanged between the counterparties of a trade, IM and the add-ons can be met with non-cash collateral and are lodged at the CCP. Differentiating between these components is important for various reasons:

- Contrary to what is sometimes assumed, VM (and not IM) is usually the one that reacts first and accounts for the largest bulk of the margin calls (Maruyama and Cerezetti, 2019; ESRB, 2020). This is a consequence of VM being the immediate result of portfolio losses or gains, no matter how small or large, while IM is mostly a function of events at the tail of the distribution. In an HS VaR IM model, for example, IM will only significantly change when changes are large enough to fall within the tail of the distribution.

- Some of these add-ons may be procyclical, depending on their function and their calibration. An example would be a liquidity add-on calibrated to respond to changes in volumes or in bid-ask spreads. In some cases, the percentage of total margin requirements attributable to the add-ons could be significant, up to 20%.

- Portfolios change over time, and as a consequence, their exposures also change. Without reference to the portfolio, it is difficult to say whether a margin call is attributable to a change in volatility or a change in portfolio composition.[9]

- Margin calls can also result from changes in the value of collateral posted.

Margin erosion and intraday calls

In this section we present a schematic description of how daily and intraday margin frameworks generally operate. While specific implementations vary across CCPs, reflecting the diversity of markets and of risk management approaches, a description of the general mechanisms can help explain the importance of having the capacity to respond to intraday risks promptly (a capacity enshrined in the PFMIs) and the operational complexities involved; and will also justify some of the modelling assumptions we make in Section 6.4.

Both VM and IM are calculated at least daily, at the end of the day. The new margin requirements are then communicated to the members, who usually have to meet the requirements by some fixed time next day in the morning. In addition to this end-of-day cycle, a CCP also monitors its exposures through the day, because large price movements or changes in the composition of the portfolio may result in the level of coverage being eroded by increases in current exposure, or in potential future exposure.

If new trades arrive during the day, for example, and they significantly increase the risk of the portfolio, the CCP should have the option to collateralize the additional exposures without the need of waiting till the end-of-day margin calculation and the margin collection next day.[10] For this reason, CCPs often run

[9] For this same reason, using reported margin breaches to derive conclusions about the margin model or about the risk management as a whole, as in Huang and Takats (2020), may be misleading when the underlying portfolios are not static.

[10] This may be particularly important to adequately capture the risk brought by day traders; that is, traders that mainly operate during the day and close a large part of their positions at the end of the day.

their margin model several times a day, on a routine basis, and have the capacity to call margin intraday if they perceive that their exposure to risk has increased beyond some tolerance level.[11] They can also run ad-hoc margin calculations in case of extreme volatility or large changes in positions (Wendt, 2006). Developing the capacity to respond to intraday risks promptly has also been enshrined in regulatory standards, as CCPs "should have the authority and operational capacity to make intraday [initial and variation] margin calls, both scheduled and unscheduled, from participants" (CPMI, 2012, 3.6.6).

On the other hand, CCPs also consider the potential impact of their intraday margin collections (and payments, in the case of variation margin) on the liquidity positions of its participants (CPMI, 2012, 3.6.11). For example, intraday changes in the composition of a portfolio do not necessarily trigger a margin call: CCPs generally allow members to put on positions that increase risk, without requiring them to prefund the associated increase in initial margin, provided that the increase is within certain tolerances defined by the CCP. Some CCPs may prefer to avoid hard thresholds for calling for additional resources because of the potential for increasing procyclicality. The thresholds are generally firm-specific (Wendt, 2006). In general, regulatory standards are not explicit on the level of coverage that initial margin should provide intraday.

As mentioned at the beginning of this section, not all CCPs follow the same approach, due to the potential negative liquidity implications ad hoc intraday calls can have on market participants. Consequently, some CCPs run standard intraday cycles for certain product types (e.g., exchange-traded derivatives), where they collect initial margin and collect and pay-out variation margin. This mitigates the likelihood of needing to make ad hoc calls, unless it is determined absolutely necessary to protect the stability of their markets. Where a CCP makes such ad hoc intraday calls, it typically will consider allowing the calls to be met with non-cash collateral to further mitigate potential liquidity strains on the market.

Intraday margin calls imply opportunity costs, collateral borrowing costs, and additional back-office processes. While for an end-of-day margin call clearing members have to deposit their collateral early the next day and have time to run back-office reconciliations overnight, in the case of intraday margin calls the funds and collateral have to be deposited within a much shorter timeframe.

5 A case study

We will now consider a simple empirical example to show how the limitations of a model-calibration approach to procyclicality played out during the March 2020 market stress.

5.1 Data

We will use the S&P500 daily prices P_t from from January 3rd, 2005 to May 26th, 2020 (15 years, 3,875 observations) and their daily log-returns, $r_t = \log(P_{t+1}/P_t)$. The period under consideration captures both the Great Financial Crisis (2008-2009), the Eurozone crisis (2010-2011), and the COVID-19 events of March 2020.

To prevent hindsight bias, we will consider two separate periods: the first, from January 3rd, 2005 to December 31st, 2019 (3,775 observations), will be used for calibration purposes. In the second period, from January 2020 to May 2020 (100 observations), we will test how the models respond to the March 2020 shock, assuming no recalibration takes place.

It is worth noting that, since we are working with a constant portfolio, the volatility observed is only caused by changes in market prices.

[11] While CCPs retain the right to run ad hoc settlement cycles, in a period of stress such a prerogative is often used in a manner that considers any procyclical implications.

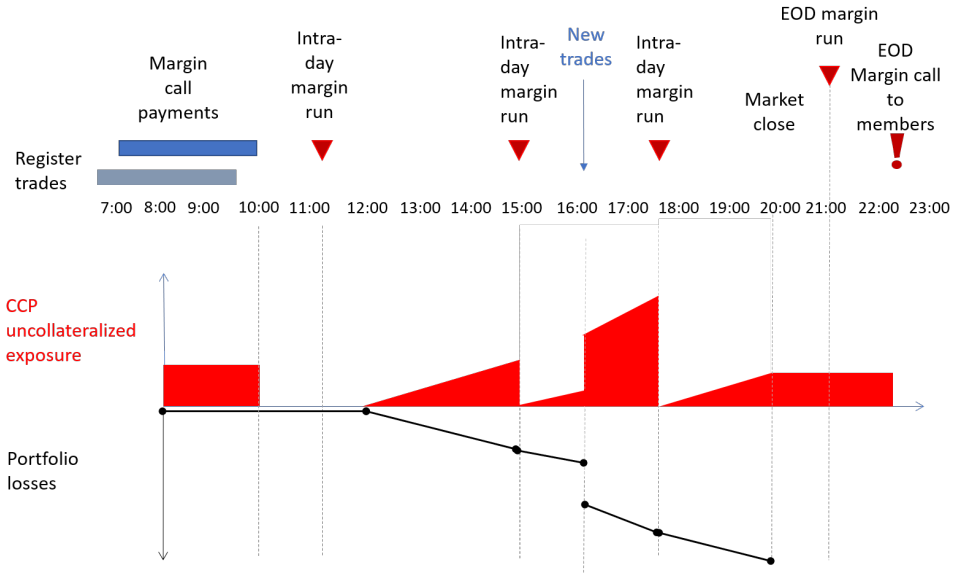

Figure 1: The figure represents a stylized intraday margining process, showing the margin erosion derived from the changes in the risk of the portfolio (either by changes in market prices or because of new trades), the corresponding increase in exposures, and the reduction of exposures after collateral is posted. We assume intraday calls are met immediately and that there is no tolerance for increments in exposure. In reality, there is a gap before payment arrives (e.g. an hour) and CCPs may prefer not to have hard thresholds. Individual CCPs may have different schedules for intraday margin runs and different rules for processing the corresponding margin calls.

5.2 IM models

CCPs use a variety of models to estimate potential future exposure. It is frequent to estimate the distribution of returns and use quantile-based risk measures like Value-at-Risk (VaR) or Expected Shortfall (ES).[12] These estimates are used directly to determine the IM, or are sometimes used as an input to a scenario-based model. To keep things simple, we will estimate VaR using empirical distributions derived from historical simulation (HS) and from Filtered Historical Simulation (FHS) processes. These models (or some variants of them) are common across CCPs.[13] Unless otherwise stated, we will estimate percentage VaR and we will follow the convention of quoting VaR (and portfolio losses) as positive numbers.

The calibration of these models involve a small set of parameters. For HS VaR, we need to specify the confidence level, the lookback period, and the margin period of risk (MPOR). In addition to these, FHS models require the specification of a decay factor λ. For the analysis, we will choose values which are commonly used across the industry (or are prescribed by regulation).

Confidence level: We will calibrate the models at a (single-tailed) 99% confidence level, which is the minimum required by the PFMIs (CPMI, 2012). In other words, the models' VaR estimate is expected to be exceeded by portfolio losses only 1% of the time. Cases in which the observed loss is larger than the VaR estimate are called breaches.

It is worth noting that we are modelling events far away in the tail, which already imposes limitations to the accuracy of the estimations. At a 99.5% confidence level, for example, even with 10 years of data (2500 observations) one would only have 12 points in the tail of the distribution.

Lookback period: An important choice to make when calibrating a model is the selection of the sample

[12] VaR is the maximum expected loss of a portfolio over a given time horizon and with a given probability and ES is the expected loss beyond a given VaR. See Dowd (2005) for a general reference.

[13] CCPs disclose information about their initial margin models as part of their quarterly Quantitative Disclosures.

used to construct the empirical distribution at any point on time (the lookback period). First, as we will discuss later, there are different trade-offs to consider when choosing between shorter and longer periods. There may also be a case for capturing periods of stress. For these reasons, CCPs sometimes prefer to use a combination of different periods or to apply some filtering approach, like in the FHS model discussed below. According to the CCPs quantitative disclosures data, lookback periods range from one up to 12 years of data, depending on the asset class.[14] For the analysis, we will consider lookback periods of 1-, 10- and 12-year rolling windows. The 10-year window is often chosen by regulators as a way of ensuring that a full set of economic conditions is captured. In our particular example, the 10-year rolling window will only capture the 2008-crisis until the end of 2018, while the 12-year period will still capture those events in full.

Margin period of risk (MPOR): The estimation of potential future exposure requires making an assumption about the number of days, n, it will take to liquidate a portfolio, a period usually known as the margin period of risk (MPOR). The MPOR varies from 1 to 7 days across jurisdictions and asset classes, reflecting the different characteristics of products and markets, as well as the differences between members' and clients' accounts. As the MPOR increases, there is a loss of accuracy in the estimation because the sample is reduced by a factor of n (the reason being that for a n-day MPOR we would have to sample at n-day intervals to avoid artificially inducing autocorrelation in the return series). To deal with this problem, volatility is sometimes simply rescaled by assuming that it increases following the square root of time, but this implies assuming that observations are independent. To avoid these issues we will restrict the analysis to a 1-day MPOR.

FHS and decay factor: For each day T, let $\mathcal{D}_T$ be the distribution of returns at time T. In the case of historical simulation models, $\mathcal{D}_T$ is the empirical distribution of the observations $y_T = \{r_{T-N}, ..., r_{T-1}\}$, where N is the length of the look-back period. In other words, we are assuming that the returns in the lookback period have the same distributional properties.

In the filtering approach proposed by Hull and White (1998) the empirical distribution $\mathcal{D}_T$ is rescaled so that its volatility matches the most recent volatility estimate. Returns are first standardized to obtain residuals which are assumed to be approximately stationary. This residual sample is then rescaled using a conditional volatility estimate that can be derived, for example, from an EWMA volatility updating scheme or from a GARCH process. In other words, on each day T, instead of directly calculating a percentile from y_T, each one of the historical returns r_t in y_T is divided by the volatility estimate σ_t for day t. The resulting standardized returns are then multiplied by the conditional volatility σ_T to obtain the rescaled returns

$$R_t = r_t \times \frac{\sigma_T}{\sigma_t} \tag{1}$$

In the case of using EWMA volatility estimates (as is often the case with CCPs), the volatility forecast at time t for time $t+1$ can be calculated using the recursive formula

$$\sigma^2_{t+1} = \lambda\sigma^2_t + (1-\lambda)r^2_t \tag{2}$$

The decay factor λ determines how the volatility estimates incorporate past and new information. The sensitivity to new information increases as λ gets close to 0, while the persistence of the past will increase as λ is closer to 1. We would therefore expect the FHS model to be more reactive as the parameter λ decreases. While reacting quickly to new information can be desirable feature of a risk model, it may also lead to undesirable outcomes if the model overreacts or if it quickly forgets about past periods of stress. In particular, with lower values of λ, periods of low volatility will tend to produce too benign risk estimates, potentially underestimating risk. To avoid this, when using

[14] See also (Huang and Takats, 2020).

these models, CCPs tend to apply decay factors in the higher range (> 0.97). Our example will use $\lambda = 0.985$.

Some observations may be relevant here. First, CCPs using a FHS model usually keep λ fixed. In contrast with what happens in a GARCH estimation, λ is not recalibrated as the lookback period is updated. Second, neither HS or FHS produce estimates which are necessarily directly proportional to volatility; while they do react to volatility, they do it in a more subtle way.

5.3 Performance measures

We now turn to the standard coverage and procyclicality metrics that we will use to compare model performance.

Coverage

To evaluate model performance we will report backtesting breaches. Readers interested in backtesting can refer, for example, to Campbell (2005) or, for the case of FHS models, to Gurrola-Perez and Murphy (2015), Gurrola-Perez (2018), or Laurent and Firouzi (2017).

Tail behaviour

To gather more information about the behaviour at the tails, we will include the expected shortfall score function introduced by Dowd (2005), which penalizes deviations from the expected value at the tails. To estimate the score, we first consider the loss function

$$\mathcal{L}_t = \begin{cases} r_t & \text{if } r_t < -\text{VaR}_t \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

and we apply a quadratic error function that penalizes deviations from the expected shortfall ES_t at time t,

$$S = \frac{1}{N} \sum_{t=1}^{T} (\mathcal{L}_t - \text{ES}_t)^2 \tag{4}$$

The lower value of S, the more accuracy in estimating the tail.

Procyclicality

While there is a general idea of how procyclical behaviour looks like *ex-post* (that is, over a particular realisation of the random process), finding a metric that captures it *ex-ante* (independent of specific realisations of the random process) has proved to be an elusive task. Peak-to-trough ratios and n-day procyclicality measures, introduced in Murphy, Vasios and Vause (2014), are useful to rank models when comparing them over the same period and underlying process. In particular, we will consider:

Peak-to-trough (PT) ratio is the ratio of the maximum initial margin required for a static portfolio to the minimum margin required during a fixed observation period I, for a given underlying process.

n-day measure is the largest increase in margin over an n-day period assuming a constant portfolio over a fixed observation period. It is a relative measure of the speed of change: a higher n-day value means that the model produced a larger n-day reaction to the same sequence of events. By restricting to those subperiods $\tilde{I}$ where volatility is elevated, we will also measure the speed of change under stress. We will consider $n = 1$ and 5. When applied to IM, we estimate these measures translating VaR into monetary terms.

These are relative metrics, in the sense that they depend on the period and underlying process chosen. Having these, or similar, measures in place is part of the toolkit recommended in ESMA (2018).

Table 1: The table reports the PT ratios, the n-day maximum increases, and the volatility estimates for the conditional volatility of S&P500 returns, for the stressed periods of 2007-2012 and January-May 2020.

	2007-2012	2020 (Jan-May)
PT ratio	8.52	9.49
5-day max	149%	177%
1-day max	102%	80%
Vol of vol	0.072	0.104

Cost

To measure the trade-off between sensitivity and cost, we need to have some measure of the cost of overestimating margins. Assuming that the pre-defined target coverage has been met, one could argue that from a cost/benefit perspective, the perfect model would be that which, except for when there are breaches, accurately predicts the return observed and calls precisely for that amount of collateral. In other words, a model that does not over-margin.

Let $u_t = -min(r_t, 0)$ denote the portfolio losses at time t and let m_t be the daily margin requirement (expressed as a return) produced by a margin model M. We can define a cost function by translating into dollar terms any over-margining observed and adding it up over the subset U of days in the sample where no breaches have occurred:

$$C(M) = \frac{1}{N} \sum_{t \in U} P_t \times (e^{m_t} - e^{u_t}) \tag{5}$$

This of course is not an estimate of the real costs of posting collateral to the CCP, but only a way of comparing two models in terms of how much they over-margin.

6 Analysis and results

6.1 Setting the scene: the S&P500 volatility shocks

For the purposes of comparing the model outputs with the volatility observed in the markets, and to define what we consider as stressed periods, we will first estimate the conditional volatility of the series of returns. For this, we will use a normal GARCH(1,1) process,

$$X_t = \sigma_t z_t, \qquad z_t \sim \mathcal{N}(0,1) \tag{6}$$

$$\sigma^2_{t+1} = \omega + \alpha X^2_t + \beta \sigma^2_t \tag{7}$$

where the innovations z_t are independent and identically distributed (iid), and $\omega > 0$, $\alpha > 0$, and $\beta \geq 0$ are constants.[15]

As we can see in the lower panel in Figure 2, the volatility shock during March 2020 is sharper than that of 2008-2009, leading to volatility value which is 6.5 times the unconditional volatility.[16] The more intense nature of the 2020 shock is confirmed when we compare with the PT ratios and the 5-day increases in volatility observed during the stressed period of 2007-2012 (Table 1).

For the purposes of implementing two of the antiprocyclicality measures presented in Section 2, we need to have a definition of what constitutes a stressed period. Following Murphy, Vasios and Vause

[15] Other GARCH(p,q) specifications are possible, but GARCH(1,1) is commonly accepted as already providing a parsimonious good fit for financial series. We could also consider other frequently used variants of GARCH, like GJR-GARCH, but since the particular specification is not central for purpose of the paper, we stay with the simplest model.

[16] This is in line with the spike observed in the VIX Index, for example, where on March 16th the index reached 82.69, which is 5.9 times its value on January 3rd, 2020.

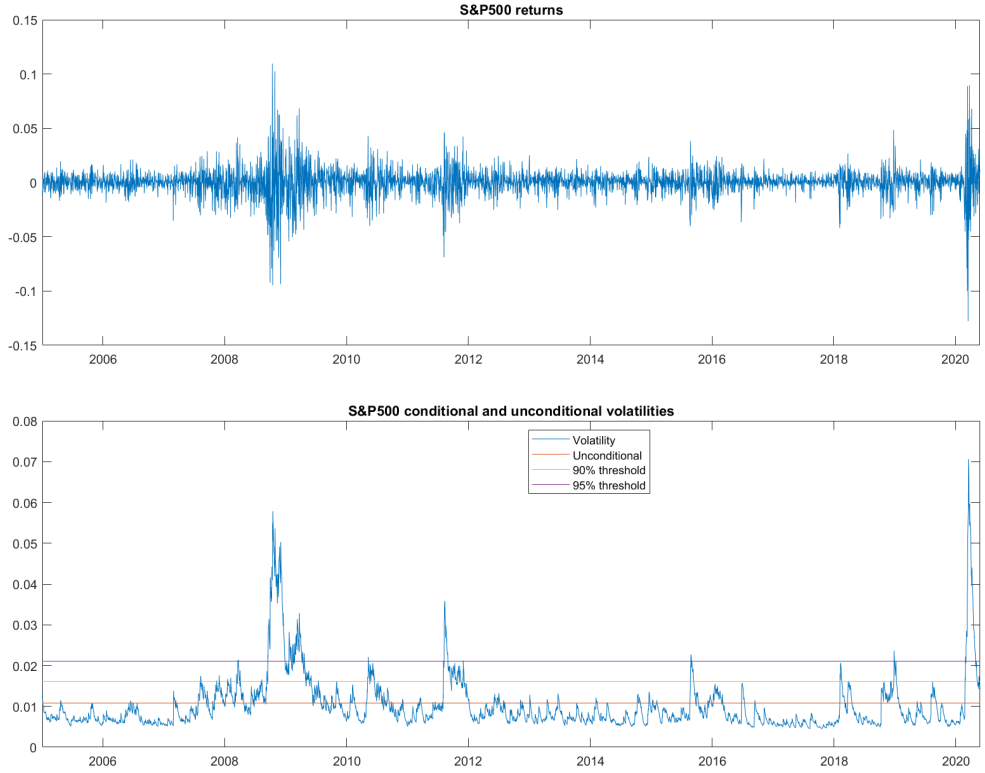

Figure 2: The top panel shows the log-returns of the S&P 500 during the period from January 3rd, 2005 to May 26th, 2020 (3,875 observations). On the lower panel we have fitted a normal GARCH(1,1) conditional variance model to the sample. The GARCH specification is $\sigma^2_{t+1} = 2.6211 \times 10^{-6} + 0.84488 r^2_t + 0.1327\sigma^2_t + \epsilon_t$. *The horizontal lines represent the unconditional volatility* $\sigma = 0.0108$, *and the volatility levels (0.0161 and 0.0211) corresponding to the 90% and 95% quantiles of the distribution of volatilities during the period. The spike in volatility observed on March 2020 reaches a maximum value* $\sigma = 0.0706$, *which is 6.5 times the unconditional volatility.*

(2016) we will define a stress period for our portfolio as the set of days t where conditional volatility σ_t exceeds the 90th percentile of the observed conditional volatilities across the cycle. Looking at the distribution of conditional volatilities, there have only been 332 days above the 90% percentile during the period from 2008 to 2019 (Figure 2).

6.2 Model calibration

In the case of HS models, the lookback period is the main factor determining the sensitivity of the model. In the case of FHS, it is the decay factor. Using the different performance measures we can select the models that provide required coverage and display lower procyclicality.

It has been argued that one of the reasons for the large margin calls observed during the March events is that margin models "seem to have underestimated market volatility, in part because they have relied on a short period of historical price movements from tranquil times. These CCPs had to catch up and increase margins at the wrong time, squeezing liquidity when it was most needed." This belief seems to be based on the idea that shorter lookback periods are "dominated by tranquil times and are therefore unable to capture the suddenly rising counterparty credit risk in times of stress" and that with a long look-back period, IM models "are more likely to include high levels of historical volatility – and volatility

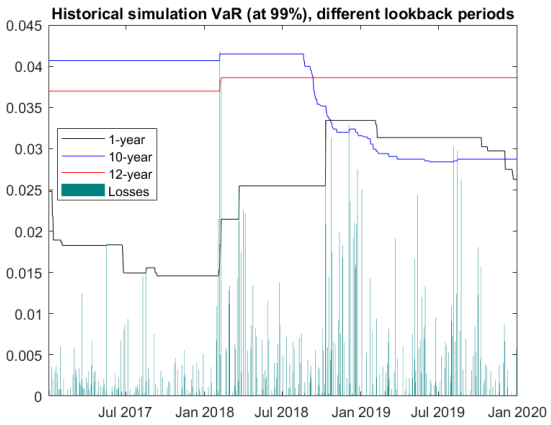

	1 year	10 years	12 years
Breaches			
Number	7	4	2
Agg. Size	0.06573	0.00510	0.00493
Max	0.02629	0.00183	0.00486
Avg	0.00009	0.00001	0.00001
ES Score	0.00181	0.00510	0.00601
Procyclicality			
PT ratio	2.61	1.59	1.51
5-day max	28%	7%	7%
1-day max	26%	5%	5%
Volatility	2%	1%	1%
Cost:			
C(M)	60.58	90.14	96.20

Figure 3: VaR estimates together with backtesting results (number of breaches, maximum breach, aggregate and average size of breaches), procyclicality measures (PT ratio, 5-day and 1-day maximal changes, and volatility of margin) and the cost metric. There are 754 observations, VaR is calibrated at 99% and the target coverage corresponds to 7.5 breaches. Note the ghost effect on the 10-year VaR, which is a consequence of the 2008 events falling out from the moving window.

spikes would be less likely to surprise" (Huang and Takats, 2020).[17]

Such arguments are flawed because they fail to distinguish between conditional and unconditional volatilities and to take into account the diluting effects of larger samples. In fact, the relation between the ability to capture a sudden rise in risk and the length of the lookback period tends to be precisely the other way around: shorter lookback periods are better to capture sudden spikes in risk. It is true that a longer lookback will be more likely to capture a more diverse set of conditions (including stress periods), leading to a more accurate approximation of the unconditional distribution. But at the same time, in a historical simulation model, the ability to capture a sudden increase in risk will tend to decrease as the sample increases, because new observations will have lower weight and the less valid the *iid* assumption will be (Pritsker, 2006). Since volatility changes through time, this implies less ability to capture conditional volatility.

This simple but crucial difference between the local (conditional) and the global (unconditional) properties is illustrated in Figure 3, where we compare the historical VaR estimates obtained using three lookback periods.

The variable relation between the parameters of the model and its output can already be seen in this simple case. It is clear, for example, that longer lookback periods do not necessarily imply more conservative (i.e. higher) margins, even if longer periods include stress events. This is the case in the first part of the sample, where the 10-year margin is above the 12-year. Only as stressed events fall out of the window - a typical ghost effect - the relation is inverted with the 10-year even falling below the 1-year. In fact, the 1-year model seems the only one reacting to the changes observed in January 2019.

It is also worth noting that with 99% target coverage, the expectation is to have 7.5 breaches. The best calibrated model seems to be the 1-year VaR, which also appears (not surprisingly) as the most procyclical. And, while the longer lookback periods are less procyclical, the additional margin required increases costs by around 50%.

One way of balancing these well-known trade-offs between longer and shorter periods consists in using a weighted combination of periods of different lengths or reflecting different market conditions. Another

[17] It is worth noting that the above implications (longer implies better and shorter implies worse) are not equivalent. Asserting that both of them hold suggests a stronger (but false) belief that the model has the ability to adequately capture risk if and only if the lookback period is long.

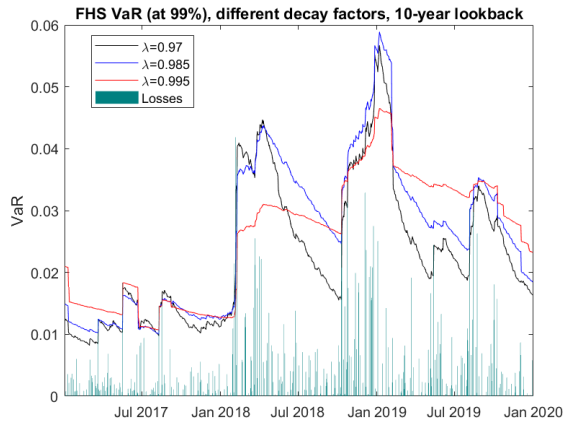

Decay	0.97	0.985	0.995
Breaches			
Number	14	12	12
Agg. Size	0.1088	0.1020	0.1066
Max	0.0249	0.0266	0.0258
Avg	0.0001	0.0001	0.0001
ES Score	0.0018	0.0016	0.0022
Procyclicality			
PT ratio	4.63	3.35	1.87
5-day max	84%	57%	16%
1-day max	52%	33%	11%
Volatility	3%	2%	1%
Cost:			
C(M)	55.47	59.04	58.64

Figure 4: FHS VaR estimates obtained using different decay factors, together with backtesting results (number of breaches, maximum breach, aggregate and average size of breaches), and procyclicality measures.

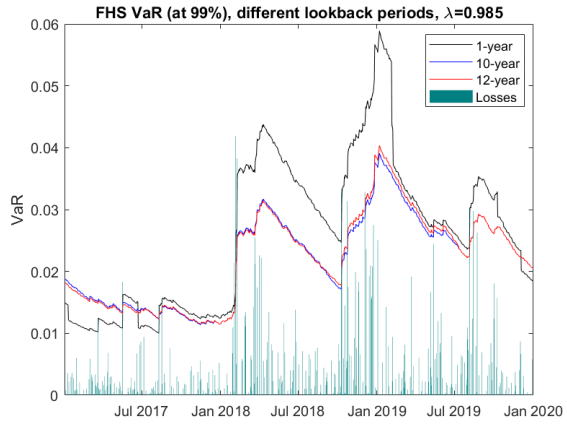

	1 year	10 years	12 years
Breaches			
Number	9	12	11
Agg. Size	0.0747	0.1020	0.0991
Max	0.0253	0.0266	0.0266
Avg	0.0001	0.0001	0.0001
ES Score	0.0022	0.0016	0.0016
Procyclicality			
PT ratio	6.06	3.35	3.43
5-day max	100%	57%	56%
1-day max	50%	33%	33%
Volatility	4%	2%	2%
Cost:			
C(M)	70.35	59.04	59.47

Figure 5: We show the FHS VaR estimates obtained with different lookback periods and a decay factor $\lambda = 0.985$, together with backtesting results (number of breaches, maximum breach, aggregate and average size of breaches)and procyclicality measures.

option is to keep a relatively long window and apply a weighting scheme where recent observations have more weight, like in the FHS approach.

In the case of FHS models, the choice of the decay factor is the main driver of the model's sensitivity: the closer the factor is to 0, the more reactive the model is, and the closer to 1, the more it resembles a constant volatility process. Figure 4 illustrates this by considering FHS VaR (at a 99% confidence) estimated with $\lambda = 0.97, 0.985$ and 0.995.

The length of the lookback period has less impact, especially as λ gets smaller. With a decay factor $\lambda = 0.985$, Figure 5 illustrates the effect of different lookback periods on a FHS VaR and compares the models in terms of the performance measures.

In addition to the previous observations for the HS case, we see that the impact of the decay factor on the number of breaches is not linear and that the longer periods appear poorly calibrated (in terms of breaches). The 12- and 10-year models show similar n-day maximal increases, while the 1-year produces larger values. In general, the differences between 10 and 12-year seem too small to be significant.

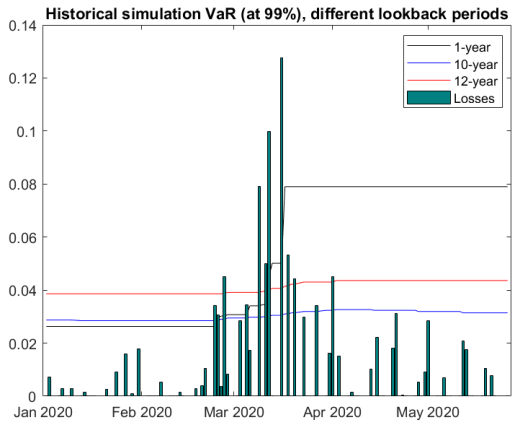

	1 year	10 years	12 years
Breaches			
Number	8	12	8
Agg. Size	0.2204	0.3153	0.2199
Max	0.0775	0.0972	0.0870
Avg	0.0022	0.0032	0.0022
ES Score	0.0113	0.0036	0.0065
Procyclicality			
PT ratio	2.70	1.35	1.36
5-day max	85%	15%	14%
1-day max	37%	9%	9%
Volatility	5%	3%	3%
Cost:			
C(M)	127.98	67.95	96.83

Figure 6: VaR 99% estimates on 2020 together with backtesting results (number of breaches, maximum breach, aggregate and average size of breaches) and procyclicality measures. There are 100 observations, which implies a target of one breach.

6.3 Responding to the rise in volatility in March 2020

We will now see how these models respond as the 2020 crisis unfolds. Figures 6 and 7 show the margin levels produced during 2020 by the HS and FHS models. It is striking to see how the differences observed during the calibration period appear dwarfed when the plot is rescaled to include the 2020 events. In the case of HS VaR models, it is remarkable to see how the models with longer lookback (10 and 12 years) produce risk estimates which are largely insensitive to the new information arriving, regardless of whether a stress period is included or not. The shock is so large and sudden that even the 1-year lookback, which increases in small steps, is unable to fully respond in time. The HS VaR with a 10-year lookback would have increased the cost of margin (clogging money that should circulate), and would probably have required an undesirable adjustment to the margin requirements.

In the case of FHS VaR estimates (Figure 7) we also include the floor defined as the 10-year HS VaR as a reference. Clearly, during the first weeks of 2020, the floor would have had a significant impact on any of the FHS margin estimates, producing margins which are at least 40% higher than those produced by the FHS models. But as the crisis unfolds, margin levels quickly increase above the floor level.

Looking at both HS and FHS examples, it is clear that the FHS models perform significantly better than HS, at least in terms of risk sensitivity. The models that are considered less procyclical (lower PT ratios and lower n-day measures) tend to fail to capture the increase in risk, leading to larger breaches. This confirms that, since shorter lookback periods are more reactive to new information arrival, they are better placed to capture the suddenly rising counterparty credit risk in times of stress, therefore reducing the possibility of the CCP being under-margined and forced to "catch-up".[18]

Finally, when looking at the procyclicality measures produced by any of these margin models and comparing with the actual changes in volatility (Table 1), none of the models appear to be overly pro-cyclical.[19]

[18] This is the opposite of what seems to be implied in Huang and Takats (2020).

[19] When looking at the actual margin increases observed during the crisis and comparing, for example, with the sharp increase on the VIX index, this also seems to be the case. For example, FIA (2020) reports that initial margin on the E-mini S&P 500 futures was $6,600 per contract in March 2nd, and by March 23rd it had been raised to $12,000 per contract (an 80% increase). This should be contrasted against the backdrop of the S&P500 conditional volatility increasing at least by a factor of three within the same period (see Figure 9).

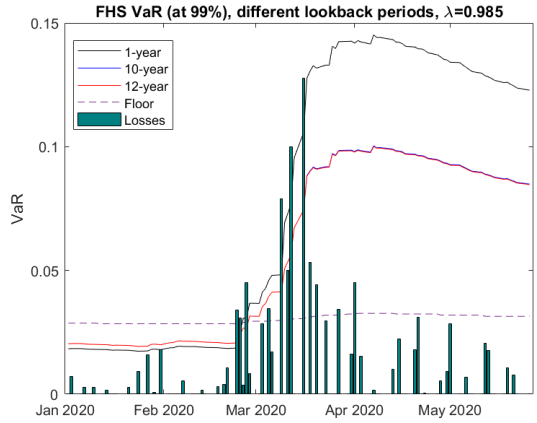

	1 year	10 years	12 years
Breaches			
Number	6	6	6
Agg. Size	0.1112	0.1730	0.1730
Max	0.0308	0.0532	0.0532
Avg	0.0011	0.0017	0.0017
ES Score	0.0202	0.0152	0.0146
Procyclicality			
PT ratio	6.61	4.21	4.20
5-day max	74%	45%	45%
1-day max	31%	21%	21%
Volatility	6%	4%	4%
Cost:			
C(M)	229.93	172.55	172.25

Figure 7: FHS VaR estimates at 99% confidence with a decay factor $\lambda = 0.985$, and lookback periods of 1, 10 and 12 years. The table shows backtesting results (number of breaches, maximum breach, aggregate and average size of breaches), together with procyclicality and cost measures. There are 100 observations, which implies a target of one breach. The floor is the 10-year HS VaR.

The costs of overmargining

A simple way of reducing procyclicality measurements without necessarily decreasing the ability of the model to capture an increase in risk is to floor the margin so that in tranquil times margin does not fall below a certain level. Such an approach is already in place in many jurisdictions and across CCPs. However, there is a limit to what can be achieved by increasing the floor, as increasing such requirements can have negative systemic risk implications by making central clearing uneconomical. Any increase in centrally cleared margins must also be examined together with its interaction with the framework for non-centrally cleared markets (Cominetta, 2019). In particular, a higher cost of central cleared trades could provide the wrong incentives and undermine the central clearing mandate, which was a fundamental pillar of the financial reforms after the Great Financial Crisis.[20]

To get a sense of how this relation between floors and costs plays out in our example, we used as baseline the FHS VaR model with a floor defined by 10-year VaR, in line with ESMA's APC rule, and we examined the impact that increasing the floor has on procyclicality and on the cost function $C(M)$. As we can see in Figure 8, while costs increase linearly with the floor level, the impact on procyclicality is slower. In the case of the 1-year lookback period, for example, the 5-day procyclicality measure only begins to significantly decrease after the floor is multiplied by a factor of 2.5 or more, implying a minimum cost increase of around 32%. We observe a similar pattern when considering a 10-year lookback window. In this case, there is no significant impact on the 1- or 5-days procyclicality unless floors increase by a factor around 2 and margin costs increase at least by 25%. In other words, the shock observed during March is so extreme that increasing the floor by a factor of 2 would still have little impact on the maximum speed of margin increase.

[20] An adequate balance between the costs of central versus non-central clearing cannot be taken for granted. For example, Robertson (2018) found that, when comparing margins on interest rate swaps obtained with the SIMM model that is used in the bilateral space with those obtained using a FHS VaR market model calibrated at 99.7% confidence level, the average ratio of SIMM margin to the FHS VaR margin was 1.72. However, for some portfolios the ratio was as low as 0.96, meaning that they attracted higher margin under the CCP's FHS VaR model.

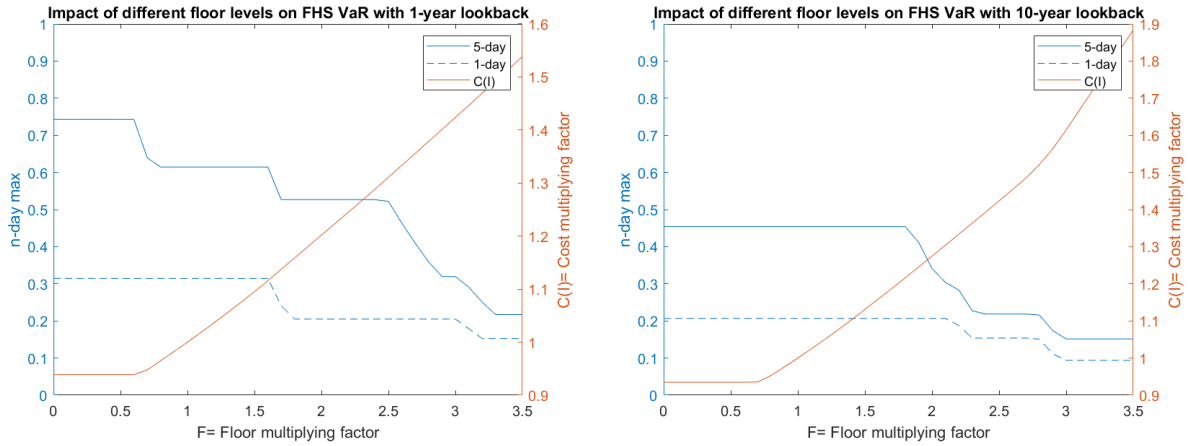

Figure 8: The graphs illustrate the impact on procylicality and costs of increasing/decreasing a floor. The baseline model is an FHS VaR with a floor defined by the 10-year HS VaR. The floor is then multiplied by a factor F that ranges between 0 and 3.5 to obtain different levels of mitigation. Procyclicality is measured using the 1- and 5-day maximum increases. The right axis represents the increase in costs, $I(C)$, expressed as a multiple of the costs of the baseline model. The point $F = 1$, $I(C) = 1$ corresponds to the baseline model. The left and right panels illustrate the 1-year lookback and 10-year lookback cases, respectively. The period is from January to May 2020 (100 observations).

6.4 A worst-case scenario analysis

To complete the picture of the different trade-offs between risk protection and margin anti-procyclicality we also need to examine the potential outcomes had the portfolio defaulted.

We will assume that the portfolio defaults on day T, with T being any day between February 13th and March 27th 2020, and will estimate the losses that the CCP would have to bear, considering each of the FHS models, with or without APC mitigation. Given that the 10- and 12-year lookback periods are almost identical during this period, we will only show the 10-year case, to avoid cluttering the plots.

The buffer, the stress and the floor

In the case the CCP opted to apply a 25% buffer, it will have to decide when to start releasing the buffer, at what speed, and when to replenish it. Murphy, Vasios and Vause (2016) suggest using the 90% quartile of volatility as the threshold to define a stress. Figure 9 shows the GARCH volatility estimates (using the specification defined previously), together with the 90% and 95% quartiles of the distribution of volatilities. In practice, given the speed of events in March, using the 90% or the 95% threshold would not have made much difference: in the first case the release date would have been February 26th, and two days later in the second case.

There is no indication in the regulation of how quickly the buffer should be consumed. If we assume that the buffer is consumed immediately it would only take five days (in the case of the 1-year model) or six days (in the case of the 10-year model) to be consumed. Under these assumptions, by March 9th the buffer would have been exhausted.

For the floor we will use the 10-year HS VaR, which would be consistent with EMIR, for example. Since the floor stays almost constant around a value of 0.028, by February 27th margin estimates would already be above the floor. Members would have avoided initial margin calls during the preceding days, but would have not avoided the large calls to come.

We now turn to the stress APC tool, which requires to apply at least a 25% weight to stressed observations. We interpret this as a weighted average of VaR estimations. In line with the previous

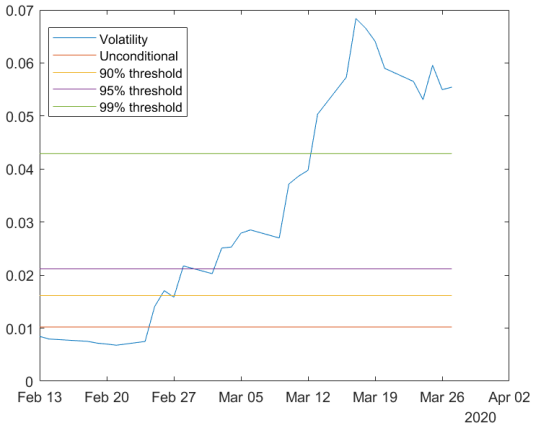

Figure 9: Estimates of S&P500 conditional volatility using the GARCH(1,1) specification. The horizontal lines reflect unconditional volatility together with the 90% and 95% quantile of the volatilities observed during the calibration period. Volatility increased above the 90% and 95% thresholds on February 27th and March 1st, respectively. From March 2nd to March 17th volatility increases by a factor of 3.4.

definition of a stressed period as a period where volatility is above its 90% quartile, we will use the 183 worse days of the 2008 crisis, from September 10th, 2008 to June 3rd, 2009 (183 days), to produce the stressed VaR component, $\text{VaR}^{\text{stressed}} = 0.0889$. We then consider the weighted average

$$\text{SVaR} = 0.25 \times \text{VaR}^{\text{stressed}} + 0.75 \times \text{FHSVaR} \tag{8}$$

where FHSVaR represents filtered historical VaR with either the 1- or 10-year lookback.

In Figure 10 we show how the different models (with and without APC) would behave during the most volatile days of 2020.

With the benefits of hindsight, one could dismiss the 1-year model with the argument that it overreacts after March 16th. But assessments must be made conditional on the information available at each point in time. On March 9th, with the information available on that day, one could have also thought the 1-year model was overreacting, and it turned out that was not the case.

The costs of underestimating risk

We will now assume that the portfolio defaults at day T. At day $T-1$, with the information available at that point time, the CCP cannot know for sure whether the model is over- or under-reacting. We can quantify the losses derived from a default happening any day within the sample.

Let m_t denote the risk estimate for day t conditional on information at day $t-1$. Let P_t and r_t the price and return observed at time t. In line with usual the margin call and collection timeline discussed in Section 4, we assume that margin collected in the morning of day t has been estimated with information available at $t-1$. In monetary terms, the margin collected at time t is $M_t = -P_{t-1}(e^{-m_{t-1}} - 1)$.

If a default happens on day T, it means that the CCP could not collect the margin that was due on that day, and will have to bear the potential losses of the portfolio. These losses will arise from the difference between the observed move in the portfolio (at the end of day $T-1$) and the last margin collected. We assume there were no intraday calls and therefore the last margin collected corresponds to the previous day. In other words, the CCP shortfall on day T assuming the portfolio defaults on day T would be

$$L_T = -\max\left[P_{T-2}(e^{r_{T-2}} - 1) - P_{T-3}(e^{-m_{T-3}} - 1), 0\right] \tag{9}$$

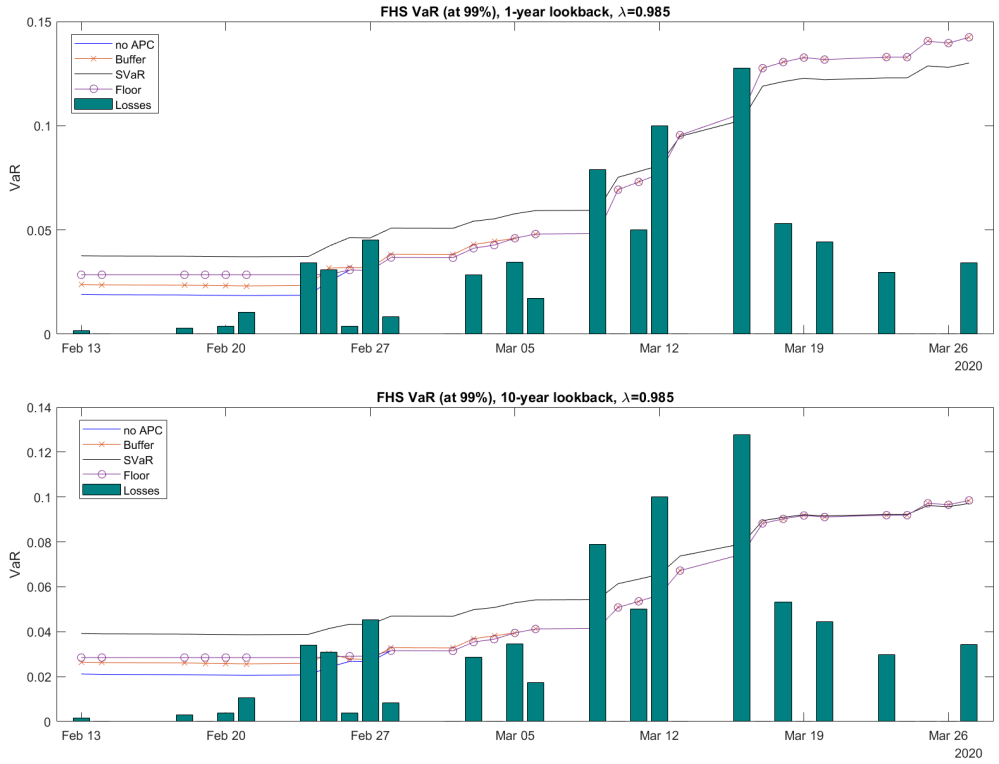

Figure 10: The graphs illustrate the performance of FHS models with and without APC tools from February 13th to March 27th. The decay factor is $\lambda = 0.985$ and there are 31 observations.

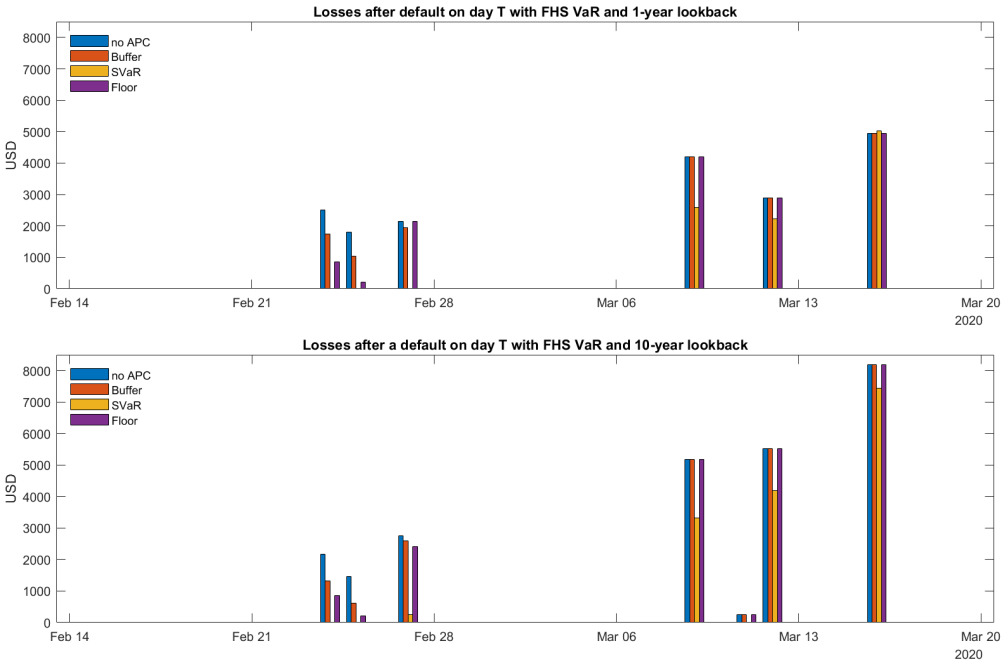

Figure 11: The charts illustrate the losses suffered assuming a default happening at day T, under the different FHS VaR models and assuming a position on one S&P500 E-mini futures contract. The upper panel corresponds to 1-year lookback periods and the lower panel uses a 10-year lookback period. The models are considered without any antiprocyclicality measure (no APC), when applying the 25% buffer (Buffer), when including weighted stressed observations (SVaR), and with the 10-year non-filtered VaR as floor (Floor).

Figure 11 illustrates how losses would have been as we consider different values of T. While in the last days of February the 1- and 10-year do not produce significant differences, on the critical days of March the losses under the longer period almost double those derived from the use of the 1-year. The CCP would have been consistently (and significantly) worse under the longer lookback periods.

To see in greater detail the trade-off between procyclicality and losses, in Table 2 we compare the models in terms of their procyclicality measures and we also compare them in terms of the shortfall derived from the hypothetical default happening at the opening after March 12th or 16th, which appear to be two of the worst days.

Table 2: The table reports the margin shortfalls under the 1-year and 10-year FHS VaR models assuming a default on day T. The models are tested including with APC tools (i.e. a 10-year HS VaR floor, a 25% buffer and the stressed VaR). The first four columns report the largest daily increase in margin and the procyclicality measures for the period from February 13th to March 27th. The Margin column reports the margin at time $T-1$. The Loss column shows the net loss for the CCP from S&P500 losses observed in March 12th or 16th had the portfolio defaulted at the opening of the next day T. CCP losses are calculated as the difference between the portfolio loss on day $T-1$ and the last margin collected. All estimates are based on holding one S&P E-mini futures contract.

		Max. daily				12 March		16 March	
		increase	**PT**	**5-day**	**1-day**	**Margin**	**Loss**	**Margin**	**Loss**
1-year									
	No APC	2315.08	5.66	74%	31%	10150.95	2886.05	11296.46	4948.05
	w/buffer	2315.08	3.82	61%	31%	10150.95	2886.05	11296.46	4948.05
	w/SVaR	1994.07	2.63	35%	18%	10820.30	2216.70	11228.34	5016.17
	w/floor	2315.08	4.54	61%	31%	10150.95	2886.05	11296.46	4948.05
10-year									
	No APC	1666.92	3.60	45%	21%	7510.41	5526.59	8063.27	8181.23
	w/buffer	1666.92	2.76	45%	21%	7510.41	5526.59	8063.27	8181.23
	w/SVaR	1502.99	1.92	27%	17%	8854.61	4182.39	8810.58	7433.93
	w/floor	1666.92	2.90	45%	21%	7510.41	5526.59	8063.27	8181.23

Not surprisingly, the longer period leads to lower procyclicality values, but at a much higher cost. A CCP using the 10-year models would have to face losses between 48% and 91% higher compared with the 1-year versions. On the other hand, the 1-year models produce maximum daily increases in margin between 33% and 39% larger than those produced by the 10-year models. A CCP using the 10-year model would effectively be subsidizing the risk of the portfolio, providing an incentive for the member to assume more risk.

In terms of the effectiveness of the APC measures in reducing procyclicality, the three APC tested reduce the PT ratios but the stressed FHS VaR seems to be the more effective, although the differences are small.

Notably, none of the APC tools is able to prevent having a large margin increase, a result that coincides with those recently produced by the Financial Stability Board based on the analysis of actual market data. In their report, the FSB established that "In the March market turmoil, IM increased sharply after a few days, which suggests that APC tools were able to dampen or slow down the IM increase only for a short time period" (FSB, 2020). Moreover, while the SVaR reduces losses in the case of 1-year model on the March 12th, it does the opposite on March 16th, a situation which illustrates the concerns expressed in Murphy, Vasios and Vause (2016) about the relative benefits of SVaR depending on how present shocks compare with the previous ones.

It is worth stressing that this is just a particular instance of a set of random variables, and we cannot conclude that the above relations will hold under different circumstances.

The conclusion of this example is not that the APC measures selected are ineffective, they do contribute to reduce procyclicality; but they can do so only to a limited extent. The problem is that there is only so much that the calibration of the margin model can do in terms of procyclicality, unless we are willing to renounce the goals of protecting the CCP and keeping central clearing economically viable.

7 Complexity

So the question remains, what else we need to do to remediate fragilities in the system that may lead to adverse liquidity feedback loops, like the one observed during March 2020?

In this last section we will emphasize the importance of looking at procyclicality from a systemic perspective, more akin to the approaches used in engineering to assess the safety of complex systems like nuclear plants or airport traffic, and consistent with the view of the financial system as complex adaptive network.

There is large amount of literature on the topic of complexity in financial systems.[21] The financial system displays many of the characteristics that are common to complex systems (Kastens et al., 2009): for example, it has multiple and heterogeneous interacting agents; there is a propensity to feedback loops, where the change in a variable may result in an amplification of that change; there are non-linear behaviours, where small changes can lead to large alterations; and it is hierarchical, in the sense that properties of the system at a higher level may not necessarily be explained in terms of elements in lower levels.

In this context, procyclicality is only one of the possible mechanisms that may contribute to the propensity of the system to generate adverse feedback loops which may amplify a stress. Its effects depend both on the state of the system and on its interactions. Despite its systemic nature, much of the debate around procyclicality tends to focus on CCPs and on the IM models, in isolation from the system. After the March 2020 events, some commentators characterized IM models as being highly procyclical, without any reference to the market conditions, as if procyclicality was a property that existed in isolation and not in relation to the system states and interactions at a given point in time. This sounds like saying that a thermometer is too procyclical because today it reported 10^oC and yesterday 6^oC, without verifying how the weather actually changed.

In a reductionist approach, properties of a system can be fully decomposed and reduced to the properties of its individual components. If the system fails, one can identify a chain of events where the failure can be fully explained as a direct consequence of the failure or malfunctioning of individual components. In complex systems this approach does not work because the safety of the system is an emerging property; that is, a property that is only observable at a level above that of the components involved. It is a consequence of the interaction between those components and cannot be traced back to their individual properties in isolation. This has several consequences in terms of preventing system failures: for example, while there are multiple contributors to a failure, each of these is insufficient in itself to create a failure. It is the linking of causes that creates the circumstances required for a failure.

It is also important to consider the interaction with incentives and behaviours. Financial networks are adaptive systems (Haldane, 2009), meaning that, in addition to complexity, agents change their behaviour as they respond to changes in the network. On the other hand, changes in incentives can also reshape the network. It is therefore critical to identify how incentives and behaviours could change as a consequence of new policies. For example, regardless of whether the APC measures mitigate the peaks of the original margin model or its troughs, the resulting model collects more margin per unit of risk in good times that it does in bad times, compared with the non-APC version. As observed by Raykov (2018), in a situation of high volatility plus insufficiently risk aversion such risk shifting may be enough to induce moral hazard in CCP members and lead to additional risk taking. In other words, as margin

[21] See, for example, Johnson (2003), for general references.

requirements are increasingly smoothed, or if limits on the size of margin calls are imposed, there may be more incentives for higher risk taking.

In a complex systems approach, the question is not about individual components but about the interdependencies that could lead to unsafe situations. To ensure safety, multiple layers of defence are needed. Going back to the procyclicality question, the procyclicality mitigation tools that CCPs have put in place are one layer of defence; but other layers are needed, including market participants ensuring their liquidity management strategies take account of the possibility that margin calls and requirements may rise significantly during periods of market stress (Bank of England, 2020), or liquidity-focused macroprudential stress tests to help managing systemic liquidity risk (King et al., 2020), for example.

8 Conclusions

Choosing and calibrating IM models so that they provide the expected coverage but with the least possible procyclical behaviour is undoubtedly an effective contribution to financial stability. We all agree that the IM model should be calibrated to address procyclicality to the extent that it is prudent and practical and, if it is not doing that appropriately, it should be recalibrated. But, if at the end of the day, fragilities in the system remain that contribute to adverse liquidity feedback loops, as we saw during the recent March 2020 events, what else we need to do?

In this paper we have argued that the answer to this question cannot simply be to impose further calibration constraints into the IM models. First, because IM is only one component of margin calls, and margin calls are mostly about VM, not IM. Second, because a calibration approach to procyclicality is limited by the fact that models need to be risk sensitive to measure risk exposures as accurately as possible, ensuring that the CCP remains adequately collateralized at all times. And it is also limited by the need to keep central clearing economically efficient. Third, because the same parametrization, under different initial conditions, may lead to different outcomes, given the stochastic nature of the process we are dealing with.

To illustrate these limitations, we empirically examined how some standard IM models would have reacted to the S&P 500 volatility increase observed during March 2020. During that period, the conditional volatility of the S&P500 increased by a factor of five (similar to the increase observed in the VIX index). The FHS margin models we tested reacted to this large increase, but the less procyclical ones failed to capture the risk and would have led to much larger losses had the hypothetical member defaulted. Similarly, in the real world, while reported margin increases were undoubtedly large, they should be assessed against the much higher increase in volatility levels and the risk of the CCP not being adequately collateralized. As a systemic phenomenon, procyclicality has to be measured and reported against underlying market conditions.

Moreover, while it seems tempting to call for less reactive margin models, one should not forget that CCPs also play a role in enforcing financial discipline. Margins provide incentives for participants to ensure that they are adequately managing their risk. Constraining the ability of the CCP to set prudent but adequate margins has the double negative effect of incentivising risk-taking while curtailing the ability of the CCP to correctly collateralize its exposures. If, for example, the CCP were to set a limit to the maximum rate of margin change over a defined period of time, wouldn't that constitute an incentive for its members to take additional risk?

While the tensions and trade-offs we have discussed will not go away, they would be better addressed by acknowledging that procyclicality is a systemic property and, as such, that it requires solutions that consider the system as a whole, identifying interdependencies, incentives, and behaviours that could lead to unsafe situations in the system.

References

Abruzzo, Nicole and Yang-Ho Park (2014). An Empirical Analysis of Futures Margin Changes: Determinants and Policy Implications. Finance and Economics Discussion Series 2014-86, Board of Governors of the Federal Reserve System (U.S.), revised 02 Sep 2014.

Bank of England (2020). What role did margin play during the COVID-19 shock? Bank Overground, 10 June 2020. https://www.bankofengland.co.uk/bank-overground/2020/what-role-did-margin-play-during-the-covid-19-shock

Bardoscia, Marco, Gerardo Ferrara, Nicholas Vause, and Michael Yoganayagam (2019). Simulating liquidity stress in the derivatives market. Bank of England Staff Working Paper no. 838.

Benos, Evangelos, Pedro Gurrola-Perez, and Michael Wood (2016). Managing market liquidity risk in central counterparties. *Journal of Financial Market Infrastructures* 5 (4), 105-125.

Biais, Bruno, Florian Heider, and Marie Hoerova (2016). Risk-sharing or risk-taking? Counterparty risk, incentives and margins. *Journal of Finance* 71, 1669–1698. https://doi.org/10.1111/jofi.12396.

Brunnermeier, Marcus K. and Lasse Heje Pedersen (2009). Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6), 2201-2238.

Campbell, Sean D. (2005). A review of backtesting and backtesting procedures. Finance and Economics Discussion Series from the Board of Governors of the Federal Reserve System (U.S.), No 2005-21.

Carter, Louise and Megan Garner (2015). Skin in the game - Central Counterparty Risk Controls and incentives. Bank of Australia Bulleting, June Quarter 2015.

Cerezetti, Fernando, Jorge Cruz Lopez, Mark Manning, and David Murphy (2019). Who pays? Who gains? central counterparty resource provision in the post-Pittsburgh world. *Journal of Financial Market Infrastructures* 7(3), 21-44. DOI: 10.21314/JFMI.2018.108.

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

Cominetta, Matteo, Michael Grill, and Audrius Jukonis (2019). Investigating initial margin procyclicality and corrective tools using EMIR. European Central Bank, Macroprudential Bulletin, 29 October 2019. https://www.ecb.europa.eu/pub/financial-stability/macroprudential-bulletin/html/ecb.mpbu201910_5~6c579ba94e.en.html#toc1

CPMI: Committee on Payment and Settlement Systems (2012) Principles for Financial Market Infrastructures. http://www.bis.org/cpmi/publ/d101a.pdf.

Cunliffe, Jon (2020). Financial System Resilience: Lessons from a real stress. Speech given at the Investment Association Webinar by the Bank of England Deputy Governor Financial Stability, London, Tuesday 9 June 2020.

Domanski, Dietrich, Leonardo Gambacorta, and Cristina Picillo (2015). Central clearing: trends and current issues. BIS Quarterly Review, December 2015. https://www.bis.org/publ/qtrpdf/r_qt1512g.pdf.

Dowd, Kevin (2005). Measuring market risk, 2nd Edition, Wiley & Sons.

ESMA: European Securities and Markets Authority (2013). European Markets Infrastructure Regulation (EMIR), Regulatory technical standards on requirements for central counterparties. https://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2013:052:0041:0074:EN:PDF.

ESMA: European Securities and Markets Authority (2018). Guidelines on EMIR Anti-Procyclicality Margin Measures for Central Counterparties. https://www.esma.europa.eu/sites/default/files/library/esma70-151-1496_guidelines_on_ccp_apc_margin_measures.pdf.

ESRB: European Systemic Risk Board (2020). Mitigating the procyclicality of margins and haircuts in derivatives markets and securities financing transactions, ESRB Report, January 2020.

Faruqui, Umar, Wenqian Huang, and Elöd Takáts (2018). Clearing risks in OTC derivatives markets: the CCP-bank nexus. BIS Quarterly Review, December 2018.

FIA: Futures Industry Association (2020). Revisiting Procyclicality: The Impact of the COVID Crisis on CCP Margin Requirements, Futures Industry Association, October 2020.

FSB: Financial Stability Board (2020). Holistic review of the March market turmoil. Financial Stability Board, 2020. https://www.fsb.org/wp-content/uploads/P171120-2.pdf

Glasserman, Paul and Qi Wu (2016). Persistence and Procyclicality in Margin Requirements. Office of Financial Research, Working Paper 17-01.

Gurrola-Perez, Pedro and David Murphy (2015). Filtered historical simulation Value-at-Risk models and their competitors. *Bank of England Staff Working Paper* No. 525. March 2015.

Gurrola-Perez, Pedro (2018). The validation of filtered historical Value-at-risk models. Journal of Risk Model Validation 12 (1), 85-112.

Haldane, Andrew G. (2009). Rethinking the Financial Network. Speech given by Andrew G. Haldane, Executive Director, Financial Stability, Bank of England.

Houllier, Melanie and David Murphy (2017). Borderline: judging the adequacy of return distribution estimation techniques in initial margin models. Bank of England Staff Working Paper No. 673. https://www.bankofengland.co.uk/-/media/boe/files/working-paper/2017/borderlinejudgingtheadequacyofreturndistributionestimationtechniquesininitialmarginmodels.pdf

Houllier, Melanie and David Murphy (2017). Initial margin model sensitivity analysis and volatility estimation. *Journal of Financial Market Infrastructures* 5(4), 77–103. DOI: 10.21314/JFMI.2017.078

Huang, Wenqian and Elöd Takáts (2020). The CCP-bank nexus in the time of Covid-19 , BIS Bulletin no. 13.

Hull, John and Alan White (1998). Incorporating volatility updating into the historical simulation method for Value-at-Risk. *Journal of Risk*, Volume 1, 1998.

Johnson, Neil F., Paul Jeffries, and Pak Ming Hui (2003). Financial Market Complexity: What Physics can tell us About Market Behavior. Oxford: Oxford University Press.

King, Thomas, Travis D. Nesmith, Anna Paulson, and Todd Prono (2020). *Central Clearing and Systemic Liquidity Risk*, Finance and Economics Discussion Series 2020-009. Washington: Board of Governors of the Federal Reserve System. https://doi.org/10.17016/FEDS.2020.009.

Laurent, Jean Paul and Hassan Omidi Firouzi (2017). Market Risk and Volatility Weighted Historical Simulation After Basel III. (December 28, 2017). http://dx.doi.org/10.2139/ssrn.3094463.

Leveson, Nancy G. (2011). Engineering a safer world. Systems thinking applied to safety. The MIT Press, Cambridge Massachusetts.

Maruyama, Atsushi and Fernando Cerezetti (2019). Central counterparty anti-procyclicality tools: a closer assessment. *Journal of Financial Market Infrastructures*, 7 (4), June 2019, 1-25. DOI: 10.21314/JFMI.2018.110 https://www.risk.net/journal-of-financial-market-infrastructures/6785891/central-counterparty-anti-procyclicality-tools-a-closer-assessment.

Murphy, David, Michalis Vasios, and Nicholas Vause (2014). An investigation into the procyclicality of risk-based initial margin models. *Bank of England Financial Stability Paper* No. 29.

Murphy, David (2015). I've got you under my skin: large central counterparty financial resources and the incentives they create. *Journal of Financial Market Infrastructures* 5(3), 57-74. DOI: 10.21314/JFMI.2017.073

Murphy, David, Michalis Vasios, and Nicholas Vause (2016). A comparative analysis of tools to limit the procyclicality of initial margin requirements. Bank of England Staff Working Paper No. 597.

O'Neill, Cyan and Nicholas Vause (2018). Macroprudential margins: a new countercyclical tool? Bank of England Staff Working Paper no. 765.

Pritsker, Matthew (2006). The Hidden Dangers of Historical Simulation. *The Journal of Banking and Finance*, Volume 30.

Raykov, Radoslav S. (2018). Reducing margin procyclicality at central counterparties. *Journal of Financial Market Infrastructures* 7(2), 43–59 DOI: 10.21314/JFMI.2018.106.

Robertson, Michael (2018). Cleared and Uncleared Margin Comparison for Interest Rate Swaps. CFTC Working Paper. https://www.cftc.gov/sites/default/files/idc/groups/public/%40economicanalysis/documents/file/dcr_cleared_uncleared_margin.pdf.

Wendt, Froukelien (2006). Intraday margining of central counterparties: EU practice and a theoretical evaluation of benefits and costs. De Nederlandsche Bank Working Paper no 107, July 2006.