RESERVE BANK
OF AUSTRALIA

Research Discussion Paper

Central Counterparty Loss Allocation and Transmission of Financial Stress

Alexandra Heath, Gerard Kelly and
Mark Manning

RDP 2015-02

The Discussion Paper series is intended to make the results of the current economic research within the Reserve Bank available to other economists. Its aim is to present preliminary results of research so as to encourage discussion and comment. Views expressed in this paper are those of the authors and not necessarily those of the Reserve Bank. Use of any results from this paper should clearly attribute the work to the authors and not to the Reserve Bank of Australia.

Enquiries:

Phone: +61 2 9551 9830
Facsimile: +61 2 9551 8033
Email: rbainfo@rba.gov.au
Website: http://www.rba.gov.au

ISSN 1448-5109 (Online)

Central Counterparty Loss Allocation and Transmission of Financial Stress

Alexandra Heath*, Gerard Kelly* and Mark Manning**

Research Discussion Paper
2015-02

March 2015

* Economic Group
** Payments Policy Department
Reserve Bank of Australia

The authors would like to thank Sheri Markose, Ali Rais Shaghaghi and Simone Giansante for valuable inputs to this work while collaborating on related research, including access to the genetic algorithm that has been applied to complete the bilateral position matrix described in Section 3. The authors would also like to thank Christopher Kent, John Simon, Jenny Hancock and Manmohan Singh for helpful comments, and Wing Hsieh for assistance in gathering relevant balance sheet data. The views expressed in this paper are our own and do not necessarily reflect those of the Reserve Bank of Australia. Any errors are our own.

Authors: heatha, kellyg and manningm at domain rba.gov.au

Media Office: rbainfo@rba.gov.au

Abstract

Among the reforms to over-the-counter (OTC) derivative markets since the global financial crisis is a commitment to collateralise counterparty exposures and to clear standardised contracts via central counterparties (CCPs). The reforms aim to reduce interconnectedness and improve counterparty risk management in these important markets. At the same time, however, the reforms necessarily concentrate risk in one or a few CCPs and also increase institutions' demand for high-quality assets to meet collateral requirements. This paper looks more closely at the implications of these reforms for the stability of the financial network. Following Heath, Kelly and Manning (2013), the paper examines liquidity and solvency risk under alternative clearing configurations, but extends the analysis in two main ways. First, rather than using simulated data, it uses actual data on the derivative positions of the 41 largest bank participants in global OTC derivative markets in 2012 (as previously used by the Bank for International Settlements' Macroeconomic Assessment Group on Derivatives). Second, it extends the methodology to consider in greater depth the implications of loss allocation by CCPs to meet obligations once pre-funded financial resources have been exhausted, and in particular the mechanism of variation margin gains haircutting. This mechanism is considered in international standard-setters' guidance on recovery planning for CCPs and has been adopted by some CCPs. The paper demonstrates that designing and operating CCPs in accordance with international standards can limit the potential for stress to propagate through the system, even in very extreme market conditions.

JEL Classification Numbers: E42, G17, G230

Keywords: clearing, netting, financial stability, central counterparty, derivatives, loss allocation, recovery and resolution

Table of Contents

Central Counterparty Loss Allocation and Transmission of Financial Stress

Alexandra Heath, Gerard Kelly and Mark Manning

1. Introduction

Since the global financial crisis, the G20 has overseen an ambitious program of regulatory reform in financial markets. One goal of the reform program is to 'make derivative markets safer' by reducing interconnectedness, improving counterparty risk management and increasing transparency. An important step towards meeting this objective is the 2009 commitment by G20 Leaders that 'all standardized OTC derivative contracts should be … cleared through central counterparties [CCPs]' (G20 2009). Further, policymakers have developed standards that require current and potential future counterparty exposures to be collateralised where contracts cannot be centrally cleared.

The resulting increase in the importance of CCPs in OTC derivative markets is well documented. It has been noted by many commentators that, given their central role, CCPs could be a channel for the transmission of financial shocks.[1] In this paper, we build on the conceptual framework of Heath *et al* (2013) to gain a better understanding of how the potential for transmission of stress can be mitigated by the risk management and loss allocation arrangements established by CCPs.

In contrast to Heath *et al* (2013), we use actual rather than simulated data on derivative positions, as well as banks' Tier 1 capital and liquid asset holdings. The analysis in Heath *et al* highlights the importance of considering a network that extends beyond the 'core' of the financial system. Reflecting this, our sample, which is based on the data collected for the Macroeconomic Assessment Group on Derivatives (MAGD) coordinated by the Bank for International Settlements, includes the 41 largest bank participants in global OTC derivative markets in the fourth quarter of 2012. This extends well beyond the core of 16 highly interconnected dealer banks to also include banks that have fewer counterparties. Using the available data, we consider market participants' positions in five OTC derivative asset classes (interest rates, credit, currency, commodities and equity)

1 See, for example, Pirrong (2011).

and explicitly model exposures and collateral requirements under scenarios with different clearing configurations.

We apply a variant of the methodology in Heath *et al* (2013) that simulates extreme changes to OTC derivative prices and directly traces the propagation of contagion through the system. This analysis supports the view that a well-designed CCP operating in accordance with international risk management standards can be a source of stability in the system, rather than a source of instability. Fundamental to the analysis is the observation that a CCP cannot generally be a trigger for initial stress in the system. A CCP does not typically take on discretionary risks, only assuming financial risks that arise from the positions it clears for its participants. A CCP seeks to maintain a balanced position at all times and is exposed to potential stress only if one or more of its participants default.[2]

Consistent with international standards, to limit the propagation of stress in the event of a participant default, a CCP collects initial margin from each participant to cover at least 99 per cent of potential price changes in the products that it clears. It also maintains an initial buffer of pooled resources to ensure that it could withstand the default of the participant (or, for larger CCPs, the two participants) to which it has the greatest credit exposure in extreme, but still plausible, market conditions.

If, however, either the CCP experienced multiple participant defaults or the market conditions prevailing at the time of these defaults were more extreme than the scenarios considered when calibrating its additional resources, the CCP's available financial resources could be exhausted. To ensure that it could still meet its obligations to non-defaulting participants in such a scenario, the CCP would allocate any uncovered losses to its participants. One way to do this would be to 'haircut' variation margin that participants were owed. While such loss allocation could be a channel for transmitting stress to participants, our analysis demonstrates that even in a range of very extreme scenarios, any such losses would be sufficiently widely dispersed that stress would be well contained.

After beginning with some background and relevant literature in Section 2, we turn in Section 3 to the key inputs to our analysis. In particular, we describe the dataset,

2 The notable exceptions to this general observation are the general business risks that a CCP assumes and the risks associated with its reinvestment of cash collateral. International standards place tight limits on a CCP's discretion in these activities.

the scenarios under consideration, and some of the key exposure metrics used in the analysis. Section 4 presents the analysis of contagion by modelling the design features of CCPs and the propagation of extreme shocks through the financial network. The results of this section show that, while introducing CCPs creates critical nodes in the financial network, if they are designed and operated in accordance with international standards, they can be expected to increase stability and reduce the propensity for contagion. Section 5 considers the policy implications, and Section 6 concludes.

2. Background and Relevant Literature

2.1 The Systemic Importance of Central Counterparties

A CCP assists institutions in the management of counterparty credit risk by interposing itself between the counterparties to trades in securities and derivatives markets – becoming the buyer to every seller, and the seller to every buyer. These arrangements support anonymous trading, deepen market liquidity, and generally maximise the netting of exposures across participants. They also deliver operational efficiencies and help to coordinate actions in the event of a market participant's default.

At the same time, however, they result in significant concentration of risk in the CCP. This risk can crystallise if a participant defaults on its obligations to the CCP, since the CCP must continue to meet its obligations to all of the non-defaulting participants. The CCP does this by replacing the trades of the defaulted participant, but may incur losses should the replacement trade be executed at an unfavourable price. This is known as replacement cost risk.

Importantly, a participant default on obligations arising in a CCP will typically be the result of financial difficulties that it experiences, or constraints that it faces, outside of the CCP. The CCP is then a potential channel for the transmission of stress to other participants and the financial system more widely, rather than an initial trigger for stress. Policymakers acknowledge that confidence in underlying markets could be severely tested if a financial (or indeed an operational) shock disrupted a CCP’s activities. These markets might then cease to function, leaving market participants unable to establish new positions or manage existing exposures.[3]

How widely stress could ultimately be transmitted will depend crucially on the CCP's design and its risk management arrangements: whether the CCP is sufficiently collateralised; how quickly it can liquidate non-cash collateral assets or re-invested cash collateral; how effectively it manages the default and closes out the risk associated with the defaulted participant's trades (e.g. by entering into a replacement contract with a new counterparty to rebalance its position); etc. Any shortcomings in the design or risk management framework of the CCP could, in the event of a shock, have spillover effects throughout the system (RBA 2014).

The need for sound risk management arrangements to address concentration risks is well recognised both by policymakers (Tucker 2011, 2014; Bailey 2014; Cœuré 2014; Powell 2014) and by industry participants (JP Morgan Chase 2014; ISDA 2015). Accordingly, international policymakers and standard-setters have focused increasingly on CCP resilience in recent years. The G20-led international initiative to expand the scope of CCP clearing to OTC derivative markets added impetus to these efforts (G20 2009).

In particular, new international standards have been developed for the design, operation and risk management arrangements of CCPs and other financial market infrastructures (FMIs). These Principles for Financial Market Infrastructures (PFMIs), developed by the Committee on Payments and Market Infrastructures (CPMI; formerly the Committee on Payment and Settlement Systems (CPSS)) and the International Organization of Securities Commissions (IOSCO), were published in 2012. Among other things, they establish minimum requirements in all areas of risk management: e.g. credit, liquidity, investment, business, legal and operational risks (CPSS-IOSCO 2012).

3 Wendt (2015) describes a range of contagion channels in the event of a shock – either to a participant of a CCP, or to the CCP itself – that arise from the connections that a CCP has with its 'ecosystem' (i.e. connections with the financial markets that a CCP serves, its participants, and linked CCPs and other financial market infrastructures).

2.2 Collateral and Netting

CCPs typically manage replacement cost risk through the use of variation and initial margin. Variation margin is typically exchanged at least daily – usually in cash – to reflect mark-to-market price changes on participants' outstanding positions. Initial margin is collected to cover, with a high probability, potential future exposure arising between the last variation margin payment and the closeout or replacement of a defaulted counterparty's trades. Initial margin requirements may be met either in cash or using high-quality non-cash assets that carry low credit, market and liquidity risk. Consistent with the PFMIs, initial margin is typically calibrated to at least a 99 per cent confidence interval. Only the defaulted participant's initial margin can be used in the event of a default.

A CCP's initial margin resources are supplemented with a pool of resources, typically pre-funded by contributions from all participants (along with a layer of CCP equity). This default fund is managed as a mutualised resource and sized to ensure that, in combination with the defaulted participants' margin, the CCP could withstand the default of its largest participant (*Cover 1*) or, in the case of CCPs that are systemically important in multiple jurisdictions, the largest two participants (*Cover 2*) in 'extreme but plausible' market conditions. There is not yet a consistent interpretation of 'extreme but plausible', but some CCPs target market stress equivalent to a 'once-in-30-years' price change.

In recent years, the use of margin has also become more commonplace in non-centrally cleared markets, although bilateral collateral agreements have to date typically covered only variation margin and not initial margin (ISDA 2014). This is set to change in light of new regulatory standards under which the exchange of both variation and initial margin will become mandatory between bilateral derivative market counterparties. These standards are due to be phased in from December 2015 (BCBS-IOSCO 2013). The BCBS-IOSCO standards will also establish a minimum level of initial margin coverage of 99 per cent.

The expansion of CCP clearing to OTC derivative markets, and margining of non-centrally cleared derivative transactions, will increase market participants' demand for high-quality assets and change how collateral markets operate (Singh 2013). There have been a number of attempts to estimate the magnitude of this increase in demand (Heller and Vause 2012; ISDA, IIF and AFME 2012; Levels and

Capel 2012; Sidanius and Zikes 2012; CGFS 2013; Duffie, Scheicher and Vuillemey 2014). These studies have delivered a wide range of estimates, which largely reflect assumptions about the underlying volatility of OTC contracts, the share of the market that is ultimately centrally cleared, and the netting efficiency of alternative clearing arrangements (Cheung, Manning and Moore 2014).

Netting efficiency depends on the product and counterparty scope of a given clearing arrangement, the profile of positions, and the margining methodology applied:

- Variation margin is calculated as a net payment/receipt, based on observed price changes across *all* products covered by the clearing arrangement. In the case of non-centrally cleared trades, separate variation margin payments/receipts are calculated vis-à-vis each bilateral counterparty. In the case of central clearing, variation margin payments/receipts are multilaterally netted across all counterparties.

- Initial margin is similarly calculated separately vis-à-vis each bilateral counterparty in non-centrally cleared arrangements, and multilaterally across all counterparties where positions are centrally cleared. There is, however, typically less scope for netting across products in calculating initial margin requirements. For CCPs, the PFMIs require that so-called 'margin offsets' are limited to combinations of products where prices are significantly and reliably correlated.

In general, netting efficiencies are likely to be greater if trades are centrally rather than non-centrally cleared. However, the netting advantage of central clearing will be smaller the more concentrated is activity across counterparties, the more fragmented is central clearing, and the more directional are participants' positions (Duffie and Zhu 2011; Heath *et al* 2013).

2.3 CCP Recovery and Loss Allocation

Reflecting the central and systemically important role that CCPs play, policymakers have also made progress on initiatives to enhance arrangements for the recovery and resolution of CCPs and other FMIs, including providing guidance on the requirement in the PFMIs that CCPs develop recovery plans.

In the event that the market conditions prevailing at the time of bank default were more extreme than the market scenarios the CCP considered when calibrating its additional resources, or that multiple banks were simultaneously in stress, the defaulted banks' initial margin and the CCP's default fund resources could be fully depleted. To deal with such scenarios, the PFMIs require that CCPs' recovery plans include arrangements to fully address any uncovered losses and liquidity shortfalls (FSB 2013; CPMI-IOSCO 2014).

One mechanism for uncovered loss allocation that has been widely debated, and in some cases adopted, is 'variation margin gains haircutting' (VMGH) – see Elliott (2013), Gibson (2013), ISDA (2013), CPMI-IOSCO (2014), and Duffie (2014). This involves writing down a CCP's variation margin outflows in proportion to the amount owed to each 'winning' participant, so as to fully allocate the loss.

One benefit of this approach, as Gibson (2013) demonstrates, is that VMGH mimics the allocation of losses to creditors that would otherwise arise in insolvency. This reflects that a CCP does not typically issue debt; rather, its obligations arise solely from clearing on behalf of its participants. At any point in time, therefore, the participants that are owed variation margin are the CCP's creditors. It is also comprehensive; to the extent that a CCP's only obligations are variation margin payments to winning participants, these can be fully met by uncapped VMGH (Singh 2014).[4] Finally, VMGH is reliable; since VMGH operates via a write-down of outgoing payments from the CCP, participants do not need to raise liquidity to meet their obligations in loss allocation.

However, as Duffie (2014) notes, VMGH may lead to 'unequal and unpredictable' loss allocation, since those who bear the loss are those that just 'happen to be' on the winning side of a trade on the day the CCP enters stress. Further, to the extent that some participants rely on the amounts written down in order to fund other obligations – e.g. hedges – such loss allocation could stress the solvency of participants. Those with highly directional positions vis-à-vis the CCP – including those hedging exposures outside of the CCP – are more likely to have net gains or

4 One qualification to this is that a CCP may incur a loss in closing out its exposures beyond the mark-to-market revaluation reflected in the variation margin obligation.

losses than those with more balanced positions. Accordingly, directional participants will be more exposed to loss allocation under VMGH.

An alternative loss allocation mechanism is *ex post* calls on participants – otherwise referred to as 'default fund top ups' or 'emergency assessments'. Such calls would typically be allocated proportionally with each participant's contribution to the pre-funded default fund, or its share of initial margin. Accordingly, applying this mechanism, losses may be allocated more widely, more equally and more predictably. If *ex post* calls were uncapped, this mechanism would also be comprehensive. However, particularly relative to VMGH, participants could face liquidity challenges in meeting their obligations.

3. Data and Exposure Analysis

The first step in exploring the question of how collateralisation and CCP clearing affect the stability of a financial system is to construct a matrix of bilateral positions between counterparties.

To do this, we use data on the total derivative assets and liabilities of 41 financial institutions across 5 categories of OTC derivatives at the end of 2012. These data were compiled for the Macroeconomic Assessment Group on Derivatives (MAGD 2013). The original data, and the transformations that we use to create a matrix of bilateral net notional positions, are described in Section 3.1.

The next step is to derive the bilateral exposure matrices, which are a function of the bilateral positions and the opportunities for netting. As noted, the scope for netting will depend on the extent to which transactions are cleared non-centrally, through a single CCP, or through separate CCPs for each asset class. The set of clearing arrangements that we consider is described in Section 3.2. To illustrate the dataset and establish some stylised facts that are relevant for the analysis in the remainder of the paper, we also present and discuss how exposures and collateral requirements change under each alternative clearing arrangement.

3.1 The Dataset and the Position Matrix

The data used in this analysis were compiled by the MAGD and consist of reported balance sheet data for 41 banks that are involved in OTC derivative trading. Of the

41 banks, 16 are widely recognised as forming the 'core' of the OTC derivative markets. The remaining 25 banks were chosen because they participate in OTC derivative markets, interact with CCPs, and/or are large regional banks (MAGD 2013, Table 3). These banks are typically smaller and are more likely to be involved in OTC derivative markets as part of their client business rather than as dealers with a market-making role. We have not included non-banks or any non-financial institutions (end users) in this network.

For B banks, we define the OTC derivative obligations owed by bank i to bank j in product-class k to be X_{ij}^{k}. Bank i's total derivative liabilities in product-class k will be given by the sum of its obligations to all other banks, $\Sigma_{j=1}^{B} X_{ij}^{k}$, and its total derivative assets will be given by $\Sigma_{j=1}^{B} X_{ji}^{k}$. The available data provide us with these aggregates, which can be thought of as the row and column sums of a matrix of bilateral gross market values – that is, current exposures arising from accumulated past price movements.

We infer the bilateral gross market values for each product class using a genetic algorithm that distributes the aggregate gross market asset and liability values across bilateral relationships. As in Markose, Giansante and Shaghaghi (2012) and Shaghaghi and Markose (2012), the algorithm does this in a way that minimises the errors in the relevant row and column sums, subject to the constraint that the bilateral relationships are consistent with a core-periphery structure. In particular, it uses 'connectivity priors' about the nature of relationships between counterparties that were used in the MAGD exercise. That is, the 16 core banks are assumed to have transactions with all the other banks in this group with 100 per cent probability; peripheral banks are assumed to have a 50 per cent probability of having a relationship with a core bank and a 25 per cent probability of having a relationship with another peripheral bank. These assumptions are similar in spirit to those used in Heath *et al* (2013).

The bilateral gross notional positions are estimated by multiplying the values in each row of the product matrices by the ratio of gross notional liabilities to gross market value liabilities.[5] In cases where gross notional liabilities are not reported

5 This exercise can also be done using the ratios of net notional assets to gross market value assets with the columns of the matrices. There is little difference in the resulting bilateral net notional positions.

(five banks in the sample), the average ratio for the remaining banks is used. The matrix of bilateral gross notional OTC derivative positions for product-class k is denoted $\mathbf{G}^k$. The matrix of bilateral *net* notional positions is then given by $\mathbf{N}^k = \mathbf{G}^k - \mathbf{G}^{k\prime}$, and is skew symmetric such that $N_{ij}^k = -N_{ji}^k$.

To support the analysis of stability and contagion, we supplement the OTC derivative position data with published 2012 balance sheet data for each bank on Tier 1 capital, cash and cash equivalents, and available-for-sale assets.

We acknowledge that the choice of dataset and the construction of the bilateral position matrix have three inherent limitations, although we do not believe that these materially affect the key policy messages arising from the analysis.

- First, since we use a static dataset of OTC derivative positions, compiled at a point in time under the prevailing market structure, we cannot capture the extent to which derivative positions are endogenous to the market structure. For instance, to the extent that central clearing drives netting and risk management efficiencies, a bank that was otherwise constrained by counterparty credit limits or other position limits might have an incentive to increase its derivative positions. In a similar vein, the dataset was compiled at a time of relatively benign market conditions. We are therefore unable to examine the endogeneity of positions to market conditions. That is, we cannot capture pre-crisis dynamics, such as the build-up of positions due to the under-pricing of risk.

- Data on Tier 1 capital and liquidity are similarly drawn from point-in-time observations under the prevailing market structure. Since alternative clearing arrangements will alter exposures and collateral requirements, banks' Tier 1 capital positions and liquid asset holdings would be expected to adjust accordingly. However, we have applied banks' observed point-in-time capital and liquidity positions across all modelled clearing scenarios.

- Finally, we do not directly observe the bilateral position matrix, but rather populate this matrix using the genetic algorithm described above. We believe that the algorithm is a good approximation of these underlying interconnections, but acknowledge the dependence on this assumption.

3.2 Clearing Scenarios and Netting

Having calculated the matrix of bilateral net notional positions, we estimate the exposures between two counterparties and the associated collateral demand arising from the need to pay initial margin under alternative assumptions about the way transactions are cleared.

As noted in Section 2.2, different clearing arrangements have different implications for the scope for netting of exposures. Prior to the financial crisis, most OTC derivative transactions were cleared directly between the transacting counterparties, with netting occurring across products within a given bilateral relationship. At the time of writing it has become common for some products – notably, interest rate and credit derivatives – to be cleared by product-specific CCPs, or for a single CCP to clear unrelated products via separate services that do not permit margin offsets and are supported by segregated default funds. Under such arrangements, netting occurs separately across counterparties for each product class. Generally, the greatest netting efficiency would arise where a single CCP cleared the full range of derivative products via a single service and allowed netting across both products and counterparties.

The scenarios considered are summarised in Table 1. To implement scenarios that involve central clearing, the matrices of bilateral net notional positions **N** are augmented by additional rows and columns representing CCPs to create new matrices **W**. For each bank i (within the population of B banks) that novates a proportion s^k of its net notional derivative positions N_{ij}^k to CCP c, bilateral net notional amounts outstanding with another bank j are given by $W_{ij}^k = \left(1 - s^k\right) N_{ij}^k$, and those with CCP c are given by $W_{i(B+c)}^k = \Sigma_{j=1}^B s^k N_{ij}^k$. It is also true that $W_{(B+c)j}^k = \Sigma_{i=1}^B s^k N_{ij}^k = -W_{j(B+c)}^k$.

Scenario 1 assumes that 75 per cent of interest rate derivatives, 50 per cent of credit positions, 20 per cent of commodity positions and 15 per cent of both equity and currency positions are cleared centrally through separate CCP services for each product class. The proportions are similar to the 'central' post-reform scenario used in the MAGD exercise for interest rates, credit and currency, but lower for commodity and equity derivatives. The lower penetration of clearing for commodity and equity derivatives acknowledges the slower-than-expected progress towards central clearing of these product classes since the MAGD exercise was undertaken.

Table 1: Clearing Structure Scenarios

Scenario	CCP service	Per cent centrally cleared, by product class
1	Product specific	75 per cent interest rate; 50 per cent credit; 20 per cent commodity; 15 per cent equity; 15 per cent currency
2	Single	As in Scenario 1
3	Product specific	100 per cent of each product class
4	Single	100 per cent of each product class

Scenario 2 assumes that the same proportions of each product class are cleared centrally, but that this is done through a single CCP service. In this case, $W^k_{i(B+1)} = \Sigma^B_{j=1} s^k N^k_{ij} = -W^k_{(B+1)i}$.

To provide an upper bound, and to isolate the stability implications of CCP clearing, we also consider a scenario that assumes s^k is 1; i.e. all bilateral trades are centrally cleared. We consider the case of separate CCPs – or segregated services – for each product (Scenario 3), as well as a single CCP or fully integrated services for the five products (Scenario 4).

3.3 Expected Exposures and Collateral Demand

The focus of our analysis is future exposures. It is already common practice for variation margin to be exchanged, not only on centrally cleared OTC derivative positions, but also on non-centrally cleared positions – at least for transactions between large banks (ISDA 2014). Accordingly, for the purposes of our analysis, the starting assumption is that all current exposures arising from observed price changes are already fully collateralised by the exchange of variation margin.[6]

6 Note that, since variation margin is typically exchanged in cash, these funds may be re-used immediately by the recipient. In contrast, initial margin posted to CCPs is segregated and not available for re-use unless there is a default event. Similarly, under the soon-to-be-implemented BCBS-IOSCO standards, initial margin for non-centrally cleared trades must be held in such a way as to protect the collateral receiver and there are significant limitations on re-use of initial margin.

Participant j's expected future exposure to participant i (either a bank or a CCP) is the expected value of j's losses in the event of i's default, after accounting for initial margin. This can be written as $R_{ij} = \mathbb{E}\left[\max\left(V_{ij} - C_{ij}, 0\right)\right]$, where:

- V_{ij} is equivalent to the variation margin that would have been paid by participant i to participant j, had participant i not defaulted. If we define Δp^k as the change in the price of product k since the last variation margin payment (assumed to be normally distributed around zero), then the next variation margin payment is given by $V_{ij}^k = W_{ij}^k \Delta p^k$, with $V_{ij} = \Sigma_{k=1}^5 V_{ij}^k$. $V_{ij} > 0$ denotes that participant j expects to receive a variation margin payment from participant i, while $V_{ij} < 0$ denotes that participant j is expected to pay variation margin to participant i. For participants i and j, the random variable for variation margin obligations over the margining period is $V_{ij} \sim N\left(0, \sigma_{w_{ij}}^2 W_{ij}^2\right)$, where $W_{ij} = \Sigma_k \left|W_{ij}^k\right|$ and $\sigma_{w_{ij}}^2 = \mathbf{w}'\mathbf{\Omega}\mathbf{w}$, for a 1 × 5 vector $\mathbf{w}' = \left(\frac{W_{ij}^1}{W_{ij}}, \ldots, \frac{W_{ij}^5}{W_{ij}}\right)$, and a 5 × 5 covariance matrix $\mathbf{\Omega}$ for price changes across the five derivative product classes.[7]

- C_{ij} is the collateral posted by participant i as initial margin against its derivative positions with participant j. Initial margin is calculated to cover with a high probability any variation margin that participant j would fail to receive in the event of the default of participant i, between the time of default and the time of closeout of the outstanding derivative exposure.[8] We scale up daily derivative price volatilities to cover closeout periods of five and ten days. Five days is the typical closeout period assumed in practice by CCPs to calibrate initial margin on OTC derivative products. The future regulatory minimum in non-centrally cleared settings is ten days, reflecting the likelihood that it will be more difficult to close out positions in a decentralised setting than via a CCP's coordinated default management process. Assuming that the distribution of expected price changes for a given product has a mean of zero, collateral to cover initial margin

---

7 The zero mean implies that derivatives are fairly priced, valued at zero at the time they are written, with a symmetric distribution of potential price movements such that both long and short sides of the position are as likely to pay as to receive variation margin.

8 Collateral posted by a bank to a CCP as initial margin is intended to cover any variation margin that the CCP would fail to receive in the event that the bank defaulted, since the CCP retains an obligation to pay variation margin to non-defaulted banks with mark-to-market gains.

is calculated as $C_{ij} = m\sigma_{w_{ij}} W_{ij}$, where: m is the number of standard deviations of the portfolio variance covered; $\sigma_{w_{ij}}$ is the per-unit portfolio standard deviation; and W_{ij} is the size of the portfolio position.[9] The portfolio standard deviation is, in turn, a function of the price volatility for each product class, which we take from the MAGD exercise (Table 2). The covariance between price changes across products is assumed to be zero. Note that in the case of product-specific CCPs, the price volatility of the portfolio is equivalent to the price volatility of the relevant product class.

Table 2: Derivative Price Volatilities over Alternative Closeout Periods

Product class	Daily volatility	5-day volatility	10-day volatility
Interest rates	0.068	0.152	0.215
Credit	0.119	0.266	0.376
Equity	0.635	1.420	2.008
Currency	0.068	0.152	0.215
Commodity	0.387	0.865	1.224

Note: Derivative price volatilities over closeout periods of longer than a day are estimated by multiplying the daily volatility by the square root of the number of days in the closeout period

As an indication of the magnitude of exposure if no initial margin was collected ($m = 0$ with $C_{ij} = 0$ for all i and j), Table 3 presents uncollateralised expected exposure over various margining periods. Several observations can be made.

First, exposures increase at a decreasing rate as the assumed time between default and closeout increases. This is as would be expected, given the assumption that prices move in a random walk. Also, as expected, exposures decrease as netting opportunities increase. Clearing all OTC derivative products through a single CCP service lowers exposures relative to the case of using separate CCP services for each product (Scenarios 2 and 4, relative to Scenarios 1 and 3, respectively). Centrally clearing a larger share (Scenarios 3 and 4, relative to Scenarios 1 and 2) of the OTC derivative portfolio also lowers exposures.

9 Therefore, for standard normal random variable v: $R_{ij} = \mathbb{E}\left[\max\left(V_{ij} - C_{ij}, 0\right)\right] = \sigma_{w_{ij}} W_{ij} \cdot \mathbb{E}\left[\max\left(v - m, 0\right)\right] = \sigma_{w_{ij}} W_{ij} \int_m^\infty (v - m)\phi(v)\,dv$ where $\phi(\cdot)$ is the standard normal probability density function. This gives $R_{ij} = \sigma_{w_{ij}} W_{ij} \left(\phi(m) + m\left(\Phi(m) - 1\right)\right)$ where $\Phi(\cdot)$ is the standard normal cumulative density function.

Table 3: Expected Exposures with Zero Collateral Coverage
$US billion

Scenario	1-day	5-day	10-day
1	64.13	128.65	176.99
2	60.35	123.74	171.25
3	38.89	50.43	59.09
4	25.78	33.43	39.17

The collateral required for initial margin assuming 99 per cent coverage of one-tailed price movements ($m = 2.33$), the minimum coverage level in the PFMIs and the BCBS-IOSCO margining standards for non-centrally cleared derivatives, is presented in Table 4.

Table 4: Initial Margin at 99 Per Cent Coverage
$US billion

Scenario	Total	Bank-to-bank	Bank-to-CCP	CCP-to-bank
1	942.10	892.88	49.22	0.00
2	930.25	892.88	37.37	0.00
3	121.82	0.00	121.82	0.00
4	80.76	0.00	80.76	0.00

Note: A 10-day closeout period is assumed for bank-to-bank margin, a 5-day closeout period is assumed for bank-to-CCP margin

Each participant holds initial margin against one direction of possible price movements. For either counterparty to a derivative position, uncovered price movements correspond to a single tail of the price-movement distribution, because a counterparty default only gives rise to a replacement cost loss if the default coincides with an adverse price movement. If the default coincides with a favourable price movement, there is no loss. Note that banks post margin against outstanding positions with CCPs, but that CCPs do not post margin with banks. Initial margin again increases with the assumed closeout period, and decreases as the scope for netting increases. In the case of a single CCP, the decline in initial margin requirements is substantial.

When the level of initial margin coverage is high, the remaining uncollateralised exposure is substantially reduced. In interpreting Table 5, it is important to note that the data reflect only the expected uncollateralised exposure that would crystallise in the event of a counterparty default and take no account of the probability of default. Given that CCPs are highly regulated, single-purpose institutions that have a specialist risk management function, the likelihood that a bank's exposure to a CCP crystallises is very low – even though the loss given default is sizeable. This observation will be examined further in Sections 4 and 5.

Table 5: Expected Exposures at 99 Per Cent Initial Margin Coverage

$US billion

Scenario	Total	Bank-to-bank	Bank-to-CCP	CCP-to-bank
1	13.86	1.30	11.94	0.62
2	10.84	1.30	9.06	0.47
3	31.09	0.00	29.54	1.05
4	20.61	0.00	19.58	1.03

Note: A 10-day closeout period is assumed for bank-to-bank margin, a 5-day closeout period is assumed for bank-to-CCP margin

3.4 Realised Exposures

For stability analysis, the tail of the distribution is more relevant than expected outcomes. In the analysis in Sections 4 and 5, therefore, we consider single extreme realisations of OTC derivative price changes and associated *ex post* exposures.

For each of the five products, let $W_{ij}^{k} = -W_{ji}^{k}$ be the net open position between i and j in product k and let σ_k be the standard deviation in the size of the price changes of product k. Then let v_k be the realised price change in product k in numbers of standard deviations. The variation margin flows from i to j will be:

$$V_{ij} = \Sigma_k v_k \sigma_k W_{ij}^{k}$$

Realised exposures are now defined as any net variation margin receipt that exceeds the size of the initial margin set aside for that position. The realised exposure of participant j to participant i will equal the positive variation margin obligation from participant i to participant j, less the initial margin on the position. We denote this as M_{ij}:

$$M_{ij} = \max\left(V_{ij} - C_{ij}, 0\right)$$

In what follows we consider an 'expected tail realisation', which is the expected price change conditional on that price change being larger than the price change on which initial margin was calibrated. This 'conditional expected future exposure' is one way to define a 'large' price change that isn't simply an arbitrary large number of standard deviations. Of course, since these calculations are based on a normal distribution, the expected tail realisation is only a fraction of a standard deviation above the point at which initial margin is set. Accordingly, in the analysis that follows in Section 4, we supplement this approach with additional tests that consider market outcomes further into the tail.[10]

Initial margin is set on the basis of m standard deviations, which corresponds to a realised price change of $\bar{v} = \mathbb{E}\left[v \middle| v > m\right] = \dfrac{\phi(m)}{1-\Phi(m)}$. For example, a one-tailed coverage level of 99 per cent would have a value of m of approximately 2.33 and a value of $\bar{v}$ of approximately 2.67. In this case, the 'realised exposure' for a single-product portfolio would be about 0.34 times the portfolio standard deviation (over the exposure period). A one-tailed coverage level of 50 per cent would have a value of $\bar{v}$ of approximately 0.40.[11]

A price change of $+\bar{v}$ standard deviations would represent a large positive price change (to the benefit of banks with net long positions and to the cost of banks with net short positions), and a price change of $-\bar{v}$ standard deviations would represent a large negative price change (to the benefit of banks with net short positions, and to the cost of banks with net long positions). For a positive W_{ij}^{k},

10 An alternative approach, which we leave to future research, would be to use price change distributions that exhibit 'fat tails'. Such distributional assumptions are common in the finance literature.

11 Zero initial margin would correspond to a coverage level of 50 per cent, because for either participant, one half of the price-movement distribution represents favourable price movements that would not give rise to any replacement cost loss in the event of counterparty default. With no initial margin held, possible counterparty losses follow a mixture distribution, with a 50 per cent probability of being zero, and a 50 per cent probability of following a half-normal distribution with expected value of $\sqrt{2/\pi}$. Accordingly, expected counterparty losses where zero initial margin is held would be $1/\sqrt{2\pi} \approx 0.4$ times the portfolio standard deviation.

which represents i being short and j being long, this upward movement in prices results in a variation margin payment from i to j.[12]

As an illustration of the magnitudes involved, Table 6 presents realised exposures based on the conditional expectation of a price movement that is beyond the 99th percentile of the price distribution (for an illustrative combination of positive and negative price changes across the product classes).

Table 6: Realised Exposure

$US billion

Scenario	Total	Bank-to-bank	Bank-to-CCP	CCP-to-bank
1	148.52	93.38	39.87	15.26
2	132.83	93.38	28.18	11.27
3	136.46	0.00	98.69	37.78
4	76.64	0.00	53.28	23.36

Notes: Based on the conditional expectation beyond 99 per cent initial margin coverage; assumes price changes of 2.67 standard deviations: positive for interest rates, currency and credit; negative for equities and commodities

4. Contagion Analysis

Applying the dataset and the methodologies for calculating price changes, initial margin requirements and realised exposures set out in Section 3, we interrogate the widely aired concern that a CCP could be a channel for the transmission of stress to the rest of the system in the event of a severe shock.

The methodology applied here and described in Section 4.1 traces the propagation of contagion through the system in response to a simulated extreme realised price change. Our results are presented and discussed in Section 4.2.

12 It should be noted that, to the extent that our estimates of net open interest are derived from current exposures arising from an unknown price history, the correspondence between signs and directions is essentially arbitrary; we can determine whether positions are directional and whether they are offsetting, but cannot determine whether directional positions are long or short.

4.1 Methodology

We apply two types of 'shock' to the system:

- *Pure price shock.* We shock the system with a vector of realised extreme derivative price changes that generate variation margin obligations for participants in the network, based on the price volatilities in the MAGD exercise (Table 2). In the event that a bank's variation margin obligations exceed a threshold level of its unencumbered liquid asset holdings, that bank is deemed to be unable to meet its variation margin requirements and is therefore in 'liquidity stress'. Depending on the magnitude of the shock, the degree of collateral coverage, and the capital positions of the stressed bank's counterparties, this could generate 'secondary solvency stress' for others in the system. This is how contagion propagates in our model.

- *Price and solvency shock.* We shock the system with a vector of realised extreme derivative price changes that generate variation margin obligations for participants in the network. This time, however, rather than allowing initial stress to arise via the liquidity channel, we consider sequential exogenous solvency shocks (in the spirit of Furfine (2003)) to banks in the network that cause them to default on their variation margin obligations. Again, depending on the magnitude of the shock, the degree of collateral coverage, and the capital positions of the stressed banks' counterparties, this may generate secondary solvency stress for others in the system.

For the purposes of this analysis, any failure to receive variation margin owed by a bilateral counterparty on a derivative position that is not covered by initial margin is not merely an opportunity loss, but rather a realised loss. This is equivalent to assuming that OTC derivative positions are not speculative, but rather are entered into to hedge other balance sheet exposures. To maintain its hedge, a bank would immediately seek to replace the position; in the event of an adverse price movement in excess of margin coverage, the bank would realise a loss. Similarly, any loss allocation by a CCP – either via VMGH or an *ex post* call (see Section 4.1.2) – constitutes a realised loss.

To see this more concretely, consider a given set of derivative price movements that creates a set of realised variation margin obligations, V_{ij}, which represent the variation margin paid by i to j. Bank i's total outgoing variation margin payments are $V_i^+ = \Sigma_j \max\left(V_{ij}, 0\right)$, summing over all positive net obligations to counterparties j. These outgoing variation margin payments are compared with the banks' liquid assets. In the pure price shock scenario, outgoing variation margin payments could trigger stress for a bank if they represent a large share of a bank's unencumbered liquid asset resources. Conservatively, we do not account for incoming variation margin payments, since to do so would assume that those banks from which variation margin inflows were due were themselves not in stress. In the model, we do not allow for partial payment of variation margin by a bank in stress. If a bank cannot meet its variation margin obligation in full, no payment is made, and the defaulted bank's initial margin is retained by its counterparty.

Importantly, we do not allow outgoing variation margin obligations to be a direct *initial* source of stress in the case of a CCP. Since a CCP maintains balanced positions, outgoing obligations are always fully funded in the absence of a participant default. Only in the event that a participant enters stress and cannot meet its obligations to the CCP will the CCP be exposed to potential stress. In the case of a CCP, therefore, we focus solely on the propensity for *secondary* stress.

4.1.1 Contagion

We denote bank i's liquid assets as L_i. In this section, we define liquid resources to be $L_i = L_i^A - C_i - F_i$ where L_i^A is bank i's cash, cash equivalents and available-for-sale assets, C_i is the total initial margin posted by bank i and F_i is bank i's contribution to the CCP's default fund.[13]

Bank i's available capital is defined as $K_i = K_i^{T1} - F_i$, where K_i^{T1} is total Tier 1 capital. The full deduction of F_i from capital represents a conservative application of new international bank capital requirements for exposures to CCPs.[14] Again, note that data are from a point in time and do not capture changes to a bank's capital position and liquid asset holdings that might be expected to occur in response to changes in the clearing structure.

13 The sum of cash, cash equivalents and available-for-sale assets is used as an imperfect proxy for a bank's liquid assets. It is acknowledged that 'available for sale' represents an accounting convention that does not necessarily capture the underlying liquidity of the asset.

14 Capital requirements for bank exposures to CCPs were revised in 2014 and will take effect from January 2017 (BCBS 2014).

A CCP's resources comprise initial margin and its default fund – calibrated to Cover 2 (as defined in Section 2.2). K_c represents CCP c's default fund, which contains the resources available to absorb uncovered losses due to the non-receipt of variation margin from a participant in excess of that participant's initial margin. We assume that the default fund is calibrated such that, taking into account defaulters' initial margin, the CCP could withstand the default of the two largest participants in extreme but plausible market conditions.[15] For product-specific CCPs, we assume a single product price change at 99.987 per cent (one-tailed) of the price distribution (equivalent to a once-in-30-years event). For a multi-product CCP, we assume a portfolio price change at 99.987 per cent (one-tailed) of the joint price distribution, with zero covariance between price changes across products.

In the pure price shock scenario we assume bank i will be in liquidity stress if its outgoing variation margin payments exceed a proportion, $0 < \rho_L < 1$, of its liquid assets. More specifically, bank i will be in liquidity stress after the derivative price movement if $V_i^+ / L_i > \rho_L$.

Let $\mathbf{u}^0$ be an indicator vector where u_j^0 (the j-th entry of $\mathbf{u}^0$) is equal to 1 if j is stressed, and 0 if j is not stressed. As noted, CCPs cannot enter stress at this initial stage.

We assume a bank will be in secondary solvency stress if the incoming variation margin payments that it fails to receive exceed a proportion $0 < \rho_K < 1$ of its capital. As before, bank i's 'residual exposure' (i.e. exposure after initial margin) to bank j is given by M_{ji}.

Bank i's losses due to counterparty default are then $D_i = \Sigma_j M_{ji} u_j^0$.

If $D_i / K_i > \rho_K$, bank i is deemed to be in secondary solvency stress. The threshold for secondary solvency stress need not be the same as that for initial liquidity stress, although in the results that we present in Section 4.2 we do keep these thresholds the same. A separate threshold of ρ_C may be applied for a CCP's default fund resources to determine when a CCP becomes stressed. In our analysis we

15 The 'largest' participants in this context are the participants to which the CCP has the largest exposures.

assume that $\rho_C = 1$. Where losses exceed these thresholds, they define the entries of another indicator vector $\mathbf{u}^1$ and the process iterates.

4.1.2 CCP loss allocation

Should the pre-funded financial resources from either defaulted participants' initial margin or the pre-funded default fund be insufficient to meet a CCP's obligations to non-defaulted participants, the CCP must proceed to loss allocation in order to avoid its own insolvency. This is the channel by which a CCP can transmit secondary solvency (or liquidity) stress back to the system in our model.

As discussed in Section 2.3, one mechanism for loss allocation that has been widely debated, and in some cases adopted, is VMGH. We model this approach by writing down a CCP's variation margin outflows in proportion to the amount owed to each winning participant, so as to fully allocate any uncovered loss.[16] Assuming that participants rely on the amounts written down to fund other obligations, VMGH could be a channel for solvency stress. In particular, should the capital of a participant facing a write-down thereby fall below the assumed stress threshold, $\rho_K K$, that participant is deemed to be in solvency stress.

Our central results in Section 4.2 are based on VMGH as the CCP loss allocation mechanism. By way of comparison, in Section 4.2.4, we consider the alternative loss allocation mechanism of *ex post* calls on participants. Under this method of loss allocation, any shortfall in the CCP's default fund arising from stressed participants' failures to pay variation margin is distributed by the CCP to all non-defaulted participants in the form of calls for additional contributions. These are made in proportion to the size of these participants' original positions. This differs from VMGH in that the allocation of losses is not only to those with outstanding

---

16 VMGH therefore allows for partial payment of variation margin obligations by a CCP. In the non-centrally cleared case, by contrast, if a bank cannot meet its variation margin payment in full, it pays nothing at all. This might be regarded as automatically creating a bias in the results against non-central clearing. In unpublished results, therefore, we also consider a loss allocation approach in which the CCP withholds all outgoing variation margin payments when it faces an uncovered loss. This is clearly a very extreme point of comparison, since the CCP withholds variation margin payments many multiples larger than is strictly necessary to meet the uncovered loss. It nevertheless acts as a robustness test on our central results. As would be expected, this approach generates a much higher incidence of stress than in the central case results presented in Section 4.2.

positions that are in the money. Further, in contrast to VMGH, additional contributions must be paid out of participants' liquid assets. Accordingly, they can contribute to liquidity stress more directly.

Our analysis does not consider replenishment of a CCP's default fund resources once they have been depleted; that is, the injection of funds to return the default fund to its previous level so that the CCP can continue in business after all losses have been allocated. However, in practice, a CCP's rules would require participants to make new contributions to the default fund to return it to a size consistent with the regulatory minimum. This is essential to the CCP's continued capacity to provide a credible replacement cost risk guarantee to market participants. Since such contributions remain an asset of the participant until they are drawn down by the CCP to meet a loss, they do not constitute a charge on capital. However, in practice replenishment could be a trigger for liquidity stress to the extent that the participant had insufficient liquid assets to meet the obligation at short notice.

4.2 Results

The range of parameter combinations that could be examined is extremely large. We therefore present only a small number of combinations sufficient to illustrate the key mechanisms at work, drawing out a few key messages. We examine stresses that could arise from very extreme price changes, unanticipated co-movement of prices across products, lengthy closeout periods and multiple participant defaults. In Sections 4.2.1 and 4.2.2, we consider cases in which the initial stress arises from a pure price shock that tests banks' liquidity capacity. In Section 4.2.3, we go on to consider price shocks combined with exogenous solvency shocks. Finally, in Section 4.2.4, we compare the central results with outcomes that arise when a CCP allocates losses via *ex post* calls on participants.

4.2.1 Pure price shocks

With five product classes, there are many possible combinations of extreme price moves. In this section, we consider a positive price change of a given magnitude for interest rates, currency and credit derivatives, combined with a negative price change of the same magnitude for equity and commodity derivatives. In unreported results, we have considered a number of other combinations of extreme price moves. While there are some differences in the details of these results, overall they are qualitatively similar.[17]

The combinations of price shocks and stress thresholds that are considered in this subsection are set out in Table 7. For each combination, we examine stress transmission under each of the four clearing structure scenarios introduced in Table 1, Section 3.2.

Table 7: Parameter Combinations for Contagion Analysis – Pure Price Shocks

	(i)	(ii)	(iii)	(iv)
Price change (standard deviations)	2.67	3.89	3.89	6.00
Initial liquidity stress threshold (banks, ρ_L)	10%	10%	20%	10%
Secondary solvency stress threshold (banks, ρ_K)	10%	10%	20%	10%
Secondary solvency stress threshold (CCP, ρ_C)	100%	100%	100%	100%

The first combination corresponds to the price shock considered in Section 3. That is, we consider a realised price change for each product class (in the directions described above) equivalent to the expected price change conditional on its being greater than 2.33 standard deviations. This is equivalent to the expected price change beyond the 99th percentile of the distribution.

We then consider more extreme price moves, as well as different assumptions about the liquidity and secondary solvency stress thresholds. First, we apply a standard deviation that equates to the expected price change conditional on its being greater than 3.65 standard deviations (which is the magnitude of price change assumed in calibrating the default fund). The conditional expected price change is therefore 3.89 standard deviations. We combine a price change of this magnitude, first with stress thresholds for bank capital and liquidity of 10 per cent, and then with stress thresholds of 20 per cent. As a point of comparison far out in the tail, we also present results in which the realised price change for each product is six standard deviations.

17 These results are available from the authors on request.

It should be noted that the simultaneous incidence of tail price movements in all products is a very extreme case, designed to expose the channels for contagion. For instance, in Scenarios 2 and 4, the portfolio-based initial margin for both bilateral counterparties and multi-product CCPs is calculated assuming zero covariance between products. Simultaneous incidence of tail movements therefore constitutes a realised joint price shock that is far more extreme than the incidence of a large price shock for each product individually.

To test the resilience of the system further, especially where the CCP is a central node, realised exposures are based on 10-day price changes for all clearing arrangements in all scenarios. This equates to an assumed 10-day closeout period, which is a more conservative closeout assumption than is used in calibrating initial margin and other financial resources of the CCP, increasing the probability that these resources are challenged.

In our experiments, we allow initial margin coverage on both non-centrally and CCP cleared positions to vary, demonstrating how this affects the stability of the system. It should be noted that, even at low initial margin coverage levels, a CCP remains protected by its pre-funded default fund resources. In calibrating the default fund to Cover 2, the CCP takes into account the initial margin of the two participants to which it has the largest exposures; the less margin that is provided by these participants, the more is required in the default fund, and vice versa.[18]

A sample of results for each parameter combination is presented in the following sections.

4.2.1.1 Combination (i) – price change beyond margin coverage, low stress threshold

We first examine stress propagation in the case of a price change of 2.67 standard deviations, where the thresholds for initial liquidity and secondary solvency stress are both 10 per cent. The results are presented in Figure 1. The figure shows the number of banks experiencing solvency stress (red line) and the number

---

18 In unreported results, available from the authors on request, we consider a range of alternative coverage assumptions for default fund resources. Unsurprisingly, when shocks are very large and default fund coverage is lower, uncovered losses are higher and there is greater recourse to loss allocation. In such circumstances, there is also an increased likelihood of contagion.

experiencing liquidity stress (blue line) under each of the four scenarios. The total number of banks in stress, the sum of the two, is depicted by the dashed line.

Figure 1: 2.67σ Price Change, 10 Per Cent Stress Threshold

Number of stressed banks

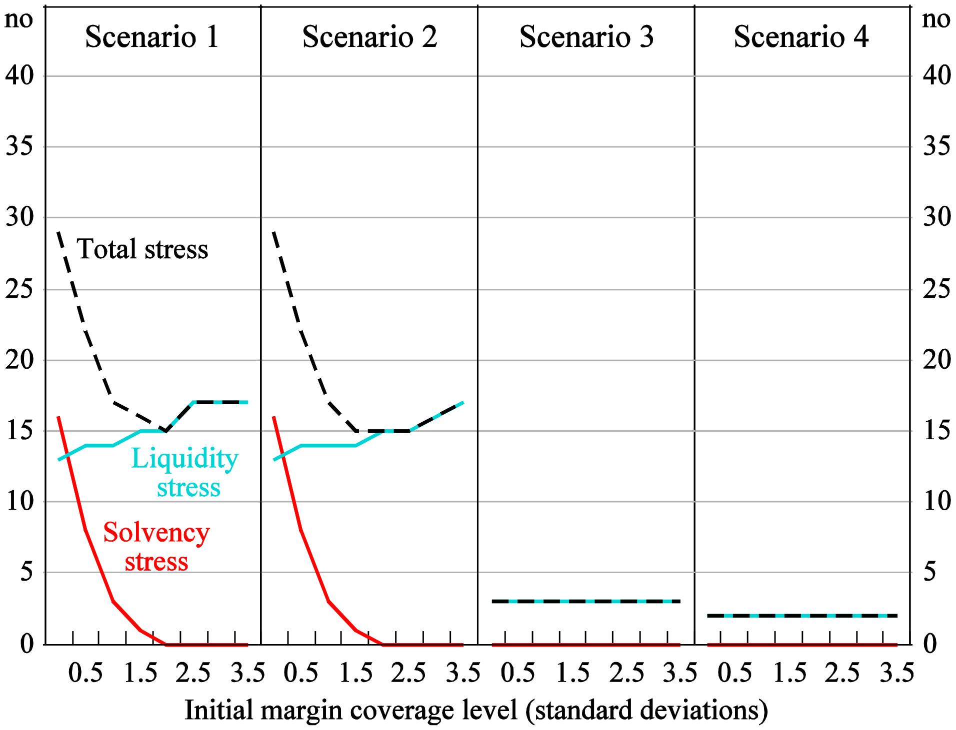

Some observations may be made.

First, under non-central clearing – which persists in Scenarios 1 and 2 – variation margin inflows and outflows vis-à-vis different counterparties cannot be offset. Accordingly, in these scenarios, a large number of banks face extremely high variation margin obligations, which many have insufficient unencumbered liquid assets to meet (based on their 2012 liquid asset holdings). At low levels of initial margin coverage, the counterparties of these banks have large uncollateralised exposures and the non-receipt of variation margin results in a direct charge against capital. With a low solvency threshold, many fall into solvency stress.

Second, the dynamic of stress in the system shifts as initial margin coverage increases, from secondary solvency stress to initial liquidity stress. Higher initial margin coverage encumbers a larger share of banks' high-quality liquid assets, leaving some banks more vulnerable to liquidity stress. At the same time, however, with higher initial margin coverage levels, the counterparties to these banks are better protected against non-receipt of variation margin. Indeed, at initial margin coverage of 99 per cent and beyond, there is no incidence of solvency stress under any scenario. These results are consistent with the u-shaped trade-off described in Heath *et al* (2013). There it was shown that as collateral coverage increased in less netting-efficient clearing structures, solvency stress declined substantially but liquidity stress increased.

Finally, at this magnitude of price change, the scenarios involving universal CCP clearing are generally ‘safer’ than those without. In both Scenarios 3 and 4, a small number of banks become liquidity constrained, but the system absorbs this stress with no flow-on secondary stress. This is true at all initial margin coverage levels. In these scenarios, the system is protected by the CCPs’ default fund resources, which are calibrated to ensure that the Cover 2 requirement is met. In Scenarios 1 and 2, by contrast, any losses beyond initial margin are directly charged against capital where positions are non-centrally cleared.

4.2.1.2 Combination (ii) – very extreme price change, low stress threshold

In this case, we take 3.89 standard deviations as the size of the realised extreme price move, being the conditional expectation beyond the 99.987 per cent price move used to size the CCP’s default fund. Price changes of this magnitude may therefore be regarded as ‘extreme but *im*plausible’, particularly coupled with the assumed realised level of co-movement and the longer closeout period. The outcome is summarised in Figure 2.

The most notable observation arising from this analysis is that there is a much higher level of stress across the system than in Figure 1. At very low levels of initial margin coverage, almost the entire system falls into stress in Scenarios 1 and 2. The incidence of secondary solvency stress again falls sharply as initial margin coverage increases; there is again evidence of the u-shaped trade-off between solvency and liquidity stress at high levels of initial margin coverage.

In Scenarios 3 and 4, with full CCP clearing, the system again absorbs stress with no flow-on secondary stress.

Figure 2: 3.89σ Price Change, 10 Per Cent Stress Threshold

Number of stressed banks

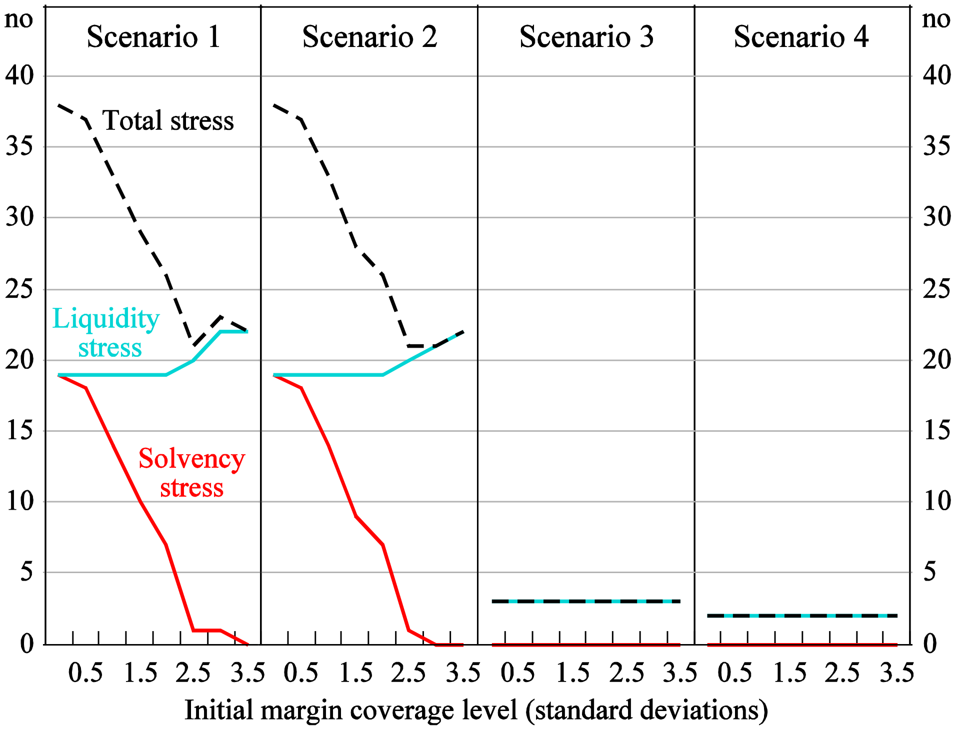

4.2.1.3 Combination (iii) – very extreme price change, higher stress threshold

The high incidence of stress observed in Figure 2, particularly in Scenarios 1 and 2, in part reflects the relatively low stress threshold for both unencumbered liquid assets and Tier 1 capital. We therefore also examine the implications of assuming a higher threshold for stress. In particular, we allow 20 per cent of either liquid assets or Tier 1 capital to be absorbed before a bank is assumed to be in stress. We continue to assume that the CCP can absorb all of its pre-funded resources before it is deemed to be in stress. As would be expected, the incidence of stress across all four scenarios is significantly lower and there is no spillover secondary stress in any scenario beyond initial margin coverage of one standard deviation (Figure 3).[19]

19 With an even higher stress threshold of 50 per cent, there is no incidence of either liquidity or solvency stress in Scenarios 3 and 4.

Figure 3: 3.89σ Price Change, 20 Per Cent Stress Threshold

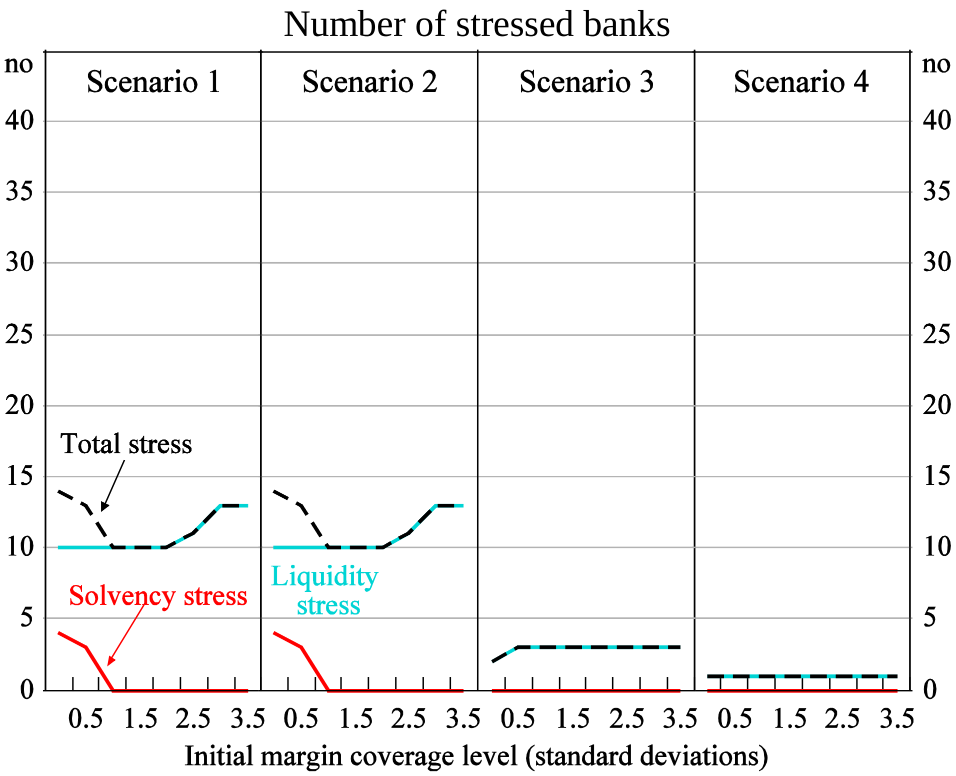

4.2.1.4 Combination (iv) – six standard deviation price change, low stress threshold

As a final experiment, we go even further into the tail of the distribution and test the capacity of the system to withstand a six standard deviation price change. We also lower the stress threshold to 10 per cent once more. The results are presented in Figure 4.

Unsurprisingly, in the less netting-efficient clearing structures in Scenarios 1 and 2, the magnitude of the shock tests the liquidity capacity of a large number of participants to meet their variation margin obligations. Across the full range of initial margin coverage levels considered in Figure 4, this gives rise to a residual loss which must be charged against capital. At the 10 per cent stress threshold, the loss exceeds available capital for a number of banks. The number of solvency-stressed banks declines as coverage increases, but does not fall to zero.

In the more netting-efficient Scenarios 3 and 4, the incidence of initial liquidity stress is again much lower. In Scenarios 3 and 4, four and two banks experience liquidity stress, respectively. Given the magnitude of the shock, the CCPs are unable to absorb the consequent losses within pre-funded resources. In Scenario 3, the consequent allocation of losses via VMGH ultimately triggers secondary solvency stress for one bank.

Figure 4: 6σ Price Change, 10 Per Cent Stress Threshold

Number of stressed banks

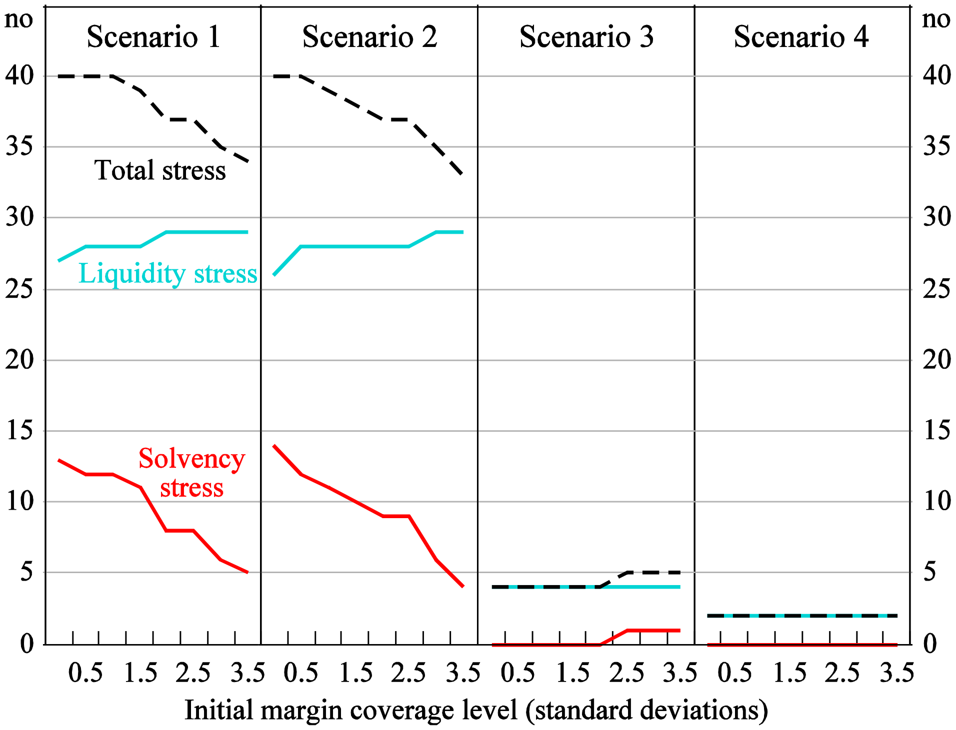

4.2.2 Loss allocation under VMGH

While there is little evidence of secondary solvency stress in the scenarios that involve universal CCP clearing in Section 4.2.1, it is instructive to consider in more detail the extent to which uncovered losses nevertheless arise, and to examine how these losses are then allocated via VMGH.

For each of the alternative parameter combinations, and at initial margin coverage of 99 per cent, Table 8 presents the scale of uncovered losses under Scenarios 3 and 4, and the magnitude of the resultant haircut under VMGH. The results reveal that, except in the case of a six standard deviation price move, uncovered losses are absorbed with relatively small haircuts (between 6 and 8 per cent).

Table 8: Uncovered Losses and Variation Margin Gains Haircutting

Parameter combinations (severity of price change, stress threshold)

	(i)		(ii)		(iii)		(iv)	
Price change (standard deviations)	2.67		3.89		3.89		6	
Stress threshold	10%		10%		20%		10%	
Scenario	3	4	3	4	3	4	3	4
Uncovered losses (US$b)	0	0	8.9	6.2	8.9	0	44.0	20.5
Variation margin haircut (%)	0	0	6.2	8.0	6.2	0	19.8	17.1

Examining these results more closely, the following observations may be made:

- The absence of uncovered losses under combination (i) is to be expected since there is little evidence of initial liquidity stress in either scenario with price changes of 2.67 standard deviations. Furthermore, the CCPs' default fund resources are calibrated to withstand large participant defaults in much more extreme market conditions.

- With price changes of 3.89 standard deviations and a stress threshold of 10 per cent (combination (ii)), uncovered losses totalling US$8.9 and US$6.2 billion, in Scenarios 3 and 4 respectively, are allocated to CCP participants via VMGH. The incidence of uncovered losses is unsurprising, since the magnitude of the price shock exceeds the 99.987 per cent stress considered in calibrating the CCPs' pre-funded resources. Furthermore, at least two banks enter stress in each scenario. Dispersed among participants, however, the relatively small variation margin haircuts (6.2 and 8 per cent in Scenarios 3 and 4, respectively) do not challenge any participant's solvency threshold. With a higher stress threshold (combination (iii)), only one bank experiences initial liquidity stress in Scenario 4 and therefore losses are absorbed within the CCP's pre-funded resources. The magnitude of uncovered losses and the allocation of those losses are unchanged in Scenario 3, where three participants again experience liquidity stress.

- With extreme tail price changes of six standard deviations (combination (iv)), uncovered losses extend to US$44 billion under Scenario 3, and more than US$20 billion under Scenario 4. These losses must be allocated via VMGH, resulting in haircuts of between 17 and 20 per cent. This does not trigger any spillover secondary solvency stress in Scenario 4. However, in Scenario 3, the size of the loss allocation is sufficient to trigger solvency stress at one capital-constrained bank.

4.2.3 Price and solvency shocks

We now simulate the sequential exogenous default of banks, combined with extreme price changes across products.

In this analysis, we base our simulations on parameter combination (ii) in Section 4.2.1 (i.e. price changes of 3.89 standard deviations in each product and a solvency stress threshold for banks of 10 per cent), and restrict our attention to the empirically relevant case of 99 per cent initial margin coverage. As before, we assume that extreme price changes occur simultaneously across products (with these occurring in the same directions as assumed in Section 4.2.1) and that there is a 10-day closeout period. In our experiment, the ‘trigger’ bank defaults occur sequentially in order of the size of the CCP’s exposure to each (based on the single CCP service in Scenario 4), beginning with the largest.

The analysis is a useful complement to that in Section 4.2.1, since it does not rely on initial stress being endogenously determined by liquidity positions that prevailed in 2012 under a different clearing structure. The results are presented in Figure 5.

In the scenarios in which some non-central clearing persists, stress is well contained until four banks are initially stressed. Prior to this point, initial margin coverage at 99 per cent is sufficient to limit contagion. However, once this point is reached, sizeable uncovered losses on bilateral positions begin to draw down capital and Scenarios 1 and 2 appear much less stable than Scenarios 3 and 4. Stress transmission becomes particularly marked in the least netting-efficient clearing structure of Scenario 1, extending to ten banks when seven banks are initially stressed.

In Scenarios 3 and 4, by contrast, the system is better able to withstand multiple participant defaults due to the added safeguards of: multilateral netting of variation margin obligations; a mutualised default fund calibrated to 3.65 standard deviations; and wider dispersion of uncovered losses via VMGH. There is nevertheless some incidence of secondary stress in these scenarios, although this extends only to one bank in Scenario 4 (after 11 defaults), and ultimately two banks in Scenario 3 (after 10 defaults).[20]

Figure 5: Price and Solvency Shocks with VMGH

Number of banks in secondary stress

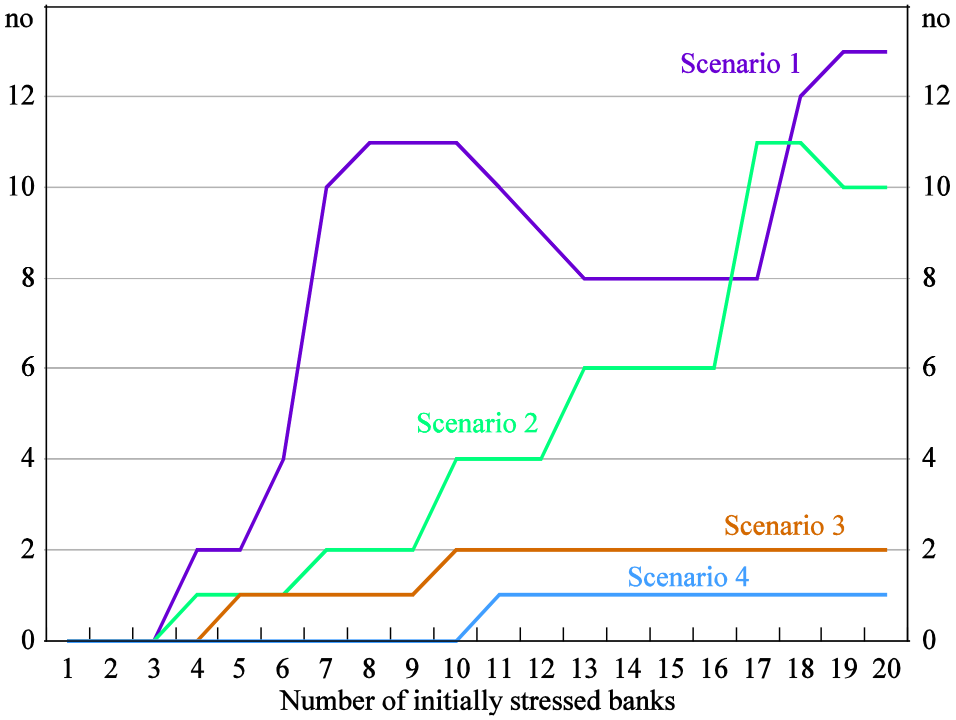

Note: 3.89σ price change, 10 per cent stress threshold

4.2.4 Alternative loss allocation mechanisms

The foregoing analysis has been based on CCPs allocating losses using VMGH. In this subsection, for the purposes of comparison, we assume that losses are allocated using *ex post* calls on all non-defaulted participants.

Figure 6 presents the results for parameter combination (ii) used in Section 4.2.1. Comparing this with Figure 2, allocating losses via *ex post* calls rather than VMGH generates broadly similar outcomes in each scenario, but with more evidence of liquidity stress in some scenarios. This is due to the fact that participants have to source liquidity to meet their *ex post* calls, which can be an additional channel for

20 With a 20 per cent stress threshold, there is no contagion in Scenario 4, and contagion to just one bank in Scenario 3, after 10 defaults. In the case of a 2.67σ price movement, there is no secondary stress under either Scenario 3 or 4, even with a 10 per cent stress threshold.

secondary stress where banks are liquidity constrained (particularly at high initial margin coverage levels where a larger proportion of liquid assets is encumbered).

Figure 6: Price Shocks with *Ex Post* Calls

Number of stressed banks

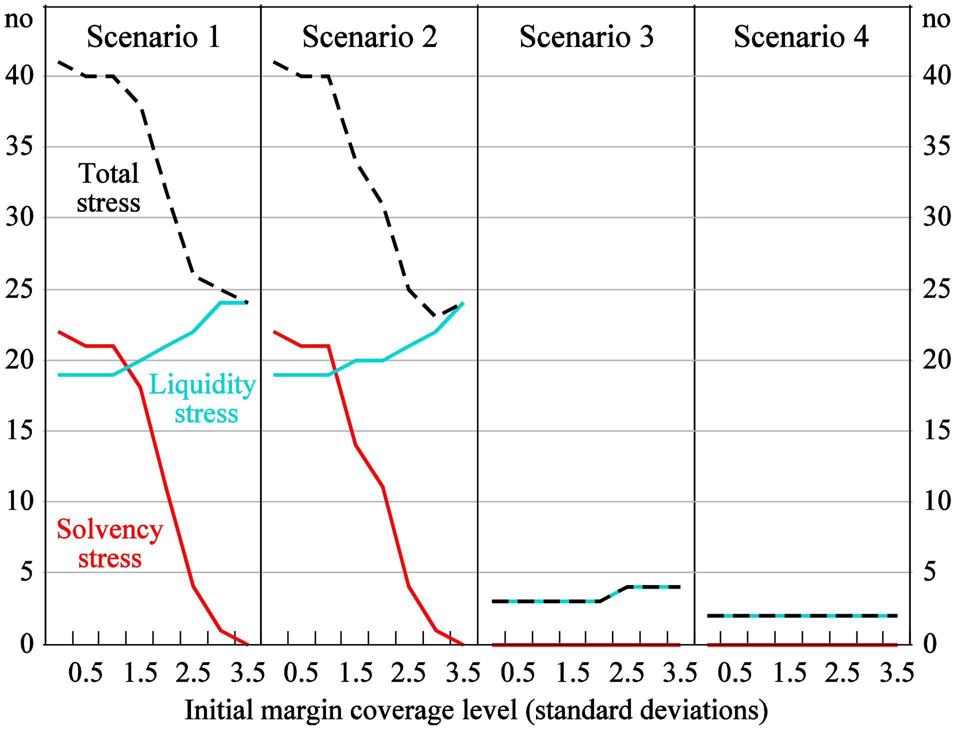

Note: 3.89σ price change, 10 per cent stress threshold

Ex post calls not only have different liquidity implications for participants relative to VMGH, but also alter the distribution of losses across participants. While under VMGH, only those with variation margin gains face losses, *ex post* calls disperse losses among all participants of the CCP. Accordingly, *ex post* calls may be expected to disperse losses not only more widely, but also more evenly than VMGH (see Figures A1 and A2). This is likely to be particularly true where some participants have large directional exposures.

Turning to the analysis of simultaneous price and solvency shocks, there is more evidence of contagion in the case of *ex post* calls (Figure 7). In Scenarios 1 and 2, the incidence of secondary stress is both much greater, and occurs after fewer banks enter initial stress. The results for Scenarios 3 and 4, while similar to those for VMGH in Figure 5, reveal the perhaps surprising feature that a structure with separate product-specific CCP services is at least initially better able to withstand stress than one with a single multi-product CCP service. In some circumstances building firewalls between products can add resilience.[21]

Figure 7: Price and Solvency Shocks with *Ex Post* Calls

Number of banks in secondary stress

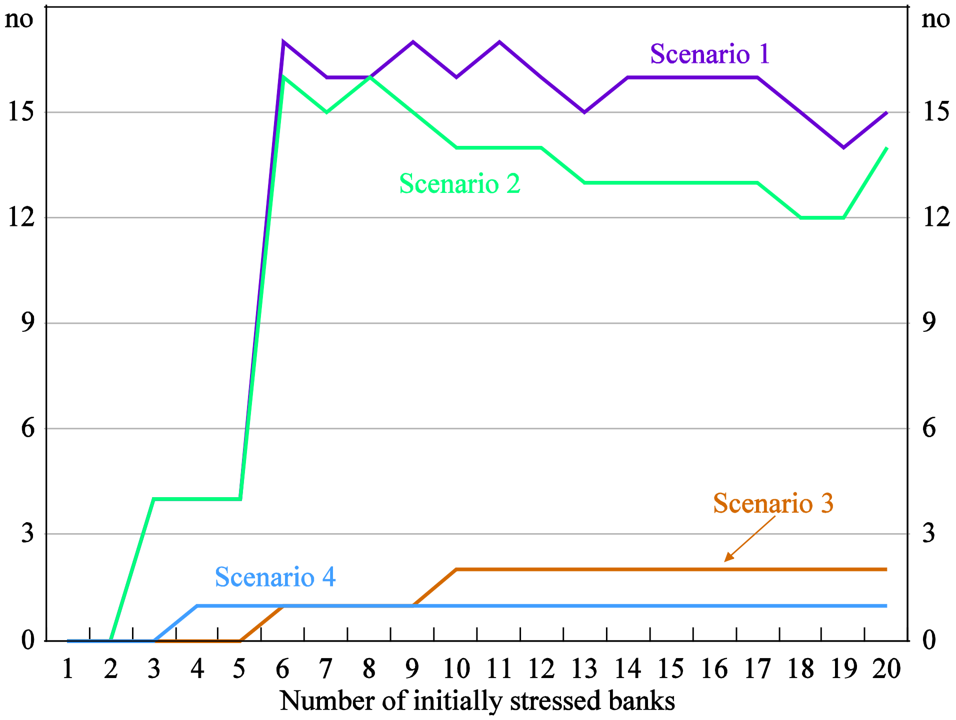

Note: 3.89σ price change, 10 per cent stress threshold

5. Policy Implications

The analysis in Section 4 gives rise to policy messages in three broad areas: (i) the trade-off between liquidity and solvency risk as collateral coverage increases; (ii) CCP default fund resources and loss allocation mechanisms; and (iii) network

21 The dynamics at play in the example with *ex post* calls are as follows. In Scenario 4, losses are allocated based on participants' outstanding positions across all product classes. In Scenario 3, by contrast, losses are allocated only in proportion to banks' positions in the single product class cleared by the relevant CCP. In the experiment depicted in Figure 7, by far the largest uncovered loss arises for interest rate derivatives. The bank that initially experiences secondary stress happens to have relatively small positions in interest rate derivatives, but large positions in some other asset classes. Accordingly, under Scenario 4, where losses are allocated based on all positions, the bank is required to absorb a higher proportion of the uncovered loss on interest rate derivatives than under Scenario 3. This is sufficient to push it into solvency stress after four exogenous defaults. In Scenario 3, this does not occur until six exogenous defaults have arisen.

analysis. These policy conclusions acknowledge the limitations of the model identified in Section 3.1 – i.e. that we do not capture the extent to which derivative positions, liquidity holdings and capital positions are endogenous to the clearing market structure; and that we do not directly observe the bilateral matrix of positions.

5.1 Liquidity and Solvency Stress – the Trade-off

Using actual data on banks' OTC derivatives positions, the analysis in this paper confirms the finding in Heath *et al* (2013) that there is a trade-off between liquidity risk and solvency risk. Particularly in Scenarios 1 and 2, where non-central clearing persists, we present evidence consistent with the u-shaped relationship identified in that paper: the incidence of solvency stress declines sharply as initial margin coverage increases, but at the same time the incidence of liquidity stress steadily rises. The results for Scenarios 1 and 2 are particularly relevant to the extent that progress towards CCP clearing of OTC derivatives has been slower than anticipated and some product classes may continue to be predominantly non-centrally cleared for some time.

Netting efficiency is a critical determinant of the shape of this trade-off. The analysis confirms that multilateral netting via CCPs can substantially reduce banks' variation margin obligations, even in extreme market conditions, making liquidity stress less likely (see Section 4.2.1). It can also lower the collateral requirements associated with banks' OTC derivative positions.

Nevertheless, even with CCP clearing, extreme price changes can give rise to high variation margin obligations that trigger liquidity stress among some participants. In interpreting this result, it is important to reiterate that the initial liquidity positions assumed in the analysis in Section 4 prevailed under a different clearing structure. Banks' liquidity holdings would be expected to increase materially in an environment with higher collateralisation, and to vary according to the clearing structure. The key message is that robust liquidity regulation is crucial, such as has been introduced in the form of the Liquidity Coverage Ratio under Basel 3.

5.2 CCP Default Fund Resources and Loss Allocation

We recognise in our framework that while CCP clearing concentrates risk in a single node in the network, that node cannot generally be a direct source of stress in the system. A CCP does not generally assume financial risks other than those arising from the positions that it clears for its participants. Accordingly, typically the only circumstance in which a CCP may experience stress is if one or more of its participants defaults. If this arises, the adequacy of the CCP's financial safeguards is critical to ensuring that any stress is contained. The results in Section 4 support the view that a CCP designed and operated in accordance with the PFMIs can be expected to promote stability in the financial network.

As discussed, the PFMIs specify that initial margin should cover at least 99 per cent of potential future price changes and that a (large, internationally systemically important) CCP should maintain additional pre-funded default fund resources to meet the Cover 2 standard. This nevertheless leaves the possibility either that more than two participants enter stress and/or the market conditions prevailing at the time of participant default are more extreme than those considered in the CCP's stress tests. As noted in Section 4.2, realised market conditions could be more extreme not only in terms of the magnitude of the price move across products relative to that assumed in calibrating default fund resources, but also in terms of the assumed co-movement between products and the assumed closeout period. In such circumstances, the CCP's pre-funded financial resources could be exhausted.

To ensure that it did not then become insolvent and cease its provision of critical infrastructure services, the CCP would, in accordance with the PFMIs, allocate any uncovered losses to its participants. In our analysis, allocating uncovered losses back to participants is the principal channel by which a CCP could transmit stress back to the wider system.

In the results presented in Section 4.2, there is little evidence of such flow-on solvency stress arising from loss allocation. There could nevertheless be circumstances in which stress transmission did occur – for instance, this is observed in our analysis when we consider a six standard deviation price change or multiple sequential participant defaults. Even in these circumstances, our analysis suggests that losses would be sufficiently widely dispersed that stress would be well contained.

Since our analysis is focused on 41 large banks, however, it does not capture the extent to which the allocation of losses via VMGH could impose stress on non-banks, such as investment funds and other end users of derivatives who may have more directional positions. Equally, however, extending the network beyond large banks could, by dispersing uncovered losses more widely, potentially leave the system even better able to absorb stress.

Precisely how stress would transmit in the event of an extreme shock is highly dependent on the particular scenario at hand: the particular loss allocation mechanism applied; the distribution and direction of positions across participants; the magnitude of price changes across product classes and their co-movement; and the financial position of participants at the time of the shock.

Given the multi-dimensionality of the problem, it is inherently difficult for CCP participants to estimate their contingent liability when uncovered losses arise. It is nevertheless important that – subject to confidentiality constraints – CCPs provide sufficient transparency about their exposures, risk models and frameworks to assist participants in modelling and managing their potential obligations in the event of loss allocation.

To this end, CPMI and IOSCO have developed a public quantitative disclosure framework for CCPs (CPMI-IOSCO 2015). This includes required disclosures around margin models and coverage of pooled financial resources. Such transparency is a very welcome development. CCPs could, however, perhaps go beyond the disclosure framework to make additional information available specifically to assist participants in their understanding of 'tail-of-tail' risks. Additional transparency would be particularly useful in the areas of stress testing and model validation.

5.2.1 Stress testing

Stress testing is at the core of a CCP's risk framework, and central to any analysis of the adequacy of a CCP's pre-funded resources and its capacity to absorb rather than transmit stress. There are two ways in which the comparability and interpretation of outputs from CCPs' stress tests could be further improved:

- The PFMIs appropriately allow CCPs discretion in establishing what constitutes an extreme but plausible scenario to be used in calibrating stress tests and pooled financial resources. This allows stress tests to be tailored to a CCP's particular product and participant profile. It also allows for innovation in stress-testing techniques over time. It is crucial that participants and regulators understand the range of scenarios used by CCPs for similar products to facilitate analysis and allow better comparison of resilience and loss-absorbing capacity across CCPs.

- To promote transparency and comparability further, regulators could consider the feasibility of regulatory stress tests of CCPs' exposures on similar lines to those that have been carried out in the United States and Europe for banks. This could be challenging, given that CCPs face very different market and operating environments and have different product and participant profiles. Nevertheless, where feasible, periodic regulatory stress tests could be a useful tool for benchmarking by both regulators and participants, and for the exercise of market discipline (see also, Bailey (2014)). Consistent with local oversight arrangements, however, CCPs should retain discretion to tailor the stress tests used in their risk management processes.

5.2.2 Model validation – reverse stress testing and sensitivity analysis

Since the circumstances in which stress is transmitted back into the system are those in which a CCP's pre-funded financial resources have been exhausted, participants should be made aware of scenarios in which this could arise. The PFMIs require that CCPs carry out 'reverse' stress tests to gauge the circumstances in which pre-funded resources could be exhausted. Transparency around the outcomes of such tests would assist participants (and regulators) in understanding the potential range of circumstances in which uncovered losses could arise and loss allocation mechanisms could be invoked.

CCPs are also required under the PFMIs to analyse the sensitivity of their margin models to key assumptions, such as closeout periods, the sample period used to estimate the price distribution, or any floors applied in the margin-setting process. As an example of sensitivity to key model assumptions, our analysis considers simultaneous extreme price movements across all five derivative product classes. It is assumed, however, in calibrating initial margin and default fund resources that the covariance of price changes is zero. Such extreme realisations of co-movement are a potential trigger for stress in scenarios involving multi-product CCPs, which reveals the importance of prudent recognition of price change covariances between products in calibrating a CCP's financial protections. This is reflected in the PFMIs, which require that portfolio margin offsets be applied only where 'the risk of one product is significantly and reliably correlated with the risk of the other product' (CPSS-IOSCO 2012).

5.3 Network Analysis

The contagion analysis in this paper emphasises the importance of understanding how, in the event a shock did arise, stress could be transmitted through the system. As trade repositories deliver more detailed information to support such data-intensive analysis, and as methodologies and techniques in this area are further refined (perhaps building on the techniques in this paper), regulators could consider complementing regulatory stress tests with system-wide 'big data' network and contagion analysis. Wendt (2015) makes a similar recommendation. Metrics, such as the eigen-pair method described in Markose (2012), provide a good first approximation of the stability of the network and which particular institutions may warrant closer attention. This could be a useful tool for regulators.

6. Conclusions

Applying an adaptation of the methodology developed in Heath *et al* (2013) and using actual data on banks' OTC derivatives positions, this paper has examined further the implications for the stability of the financial network of introducing central clearing and collateralisation of OTC derivatives trades.

At the time of writing, CCPs' risk management arrangements have been the subject of considerable policy debate. We have therefore modelled a series of extreme 'tail-of-tail' scenarios to examine the circumstances in which CCPs could feasibly be a channel for contagion in the event of wider stress in the financial system. We conclude that, while CCPs are central nodes in the financial network, maintaining CCP financial resources in accordance with international standards is consistent with maintaining system stability.

The analysis demonstrates that there could nevertheless be circumstances in which, having exhausted its pre-funded resources, a CCP transmits stress back to its participants by haircutting their variation margin gains. Although in our analysis there is little evidence of contagion, the precise scope for spillover stress is dependent on a number of factors, including the loss allocation mechanism applied, the distribution and direction of positions among participants, the magnitude of price changes across product classes and their co-movement, and the financial position of participants at the time of the shock.

Our analysis underscores the importance of understanding the level of stress that CCPs' pre-funded financial resources are designed to withstand, and also the channels by which losses could be transmitted back into the system in the event of a more extreme shock that depleted these resources. In considering the level of stress implied by these scenarios, it is important that attention is paid not only to assumed extreme price changes, but also the assumed co-movement between products and the closeout periods. Transparency of CCPs' own internal stress tests and other elements of the risk modelling framework would help to improve market participants' understanding and awareness of their exposure to such extreme tail events. To the extent feasible, this could usefully be combined with regulatory stress tests of CCPs' exposures and network analysis using individual banks' position-level exposure data.

We leave to future research further refinement of analytical techniques to deepen the analysis of how CCPs could transmit stress under alternative loss allocation mechanisms once pre-funded resources have been depleted. Alternative distributional assumptions for price changes could, for instance, be considered. Other topics for future research may include further analysis of the transmission of liquidity risk, and other channels for contagion, such as links between CCPs.

More work would also be useful on the implications of alternative loss allocation mechanisms for participant incentives. This could consider, for instance, the risk that some participants 'walk away' from a CCP in stress to avoid future obligations in loss allocation. There could also be implications for the use of CCPs by those with more directional positions – including perhaps investment funds and other ‘end users’ – that might be more exposed to mechanisms such as VMGH. Finally, the analysis in this paper has taken banks’ OTC derivative positions, liquidity holdings and capital positions as given. In practice, it is likely that these would all change endogenously in response to alternative market structures.

Appendix A: Distribution of Losses under Alternative Loss Allocation Mechanisms

Figures A1 and A2, below, compare the distribution of losses under VMGH and *ex post* calls on participants, for Scenarios 3 and 4. The parameter combinations we use are those considered in Section 4.2.4. The most striking features of the figures are that losses under VMGH are more uneven and more narrowly dispersed than under *ex post* calls; under VMGH, losses are concentrated among a small number of participants, each of which generally faces a larger loss than any individual participant under *ex post* calls.

Figure A1: Distribution of Loss Allocation with VMGH

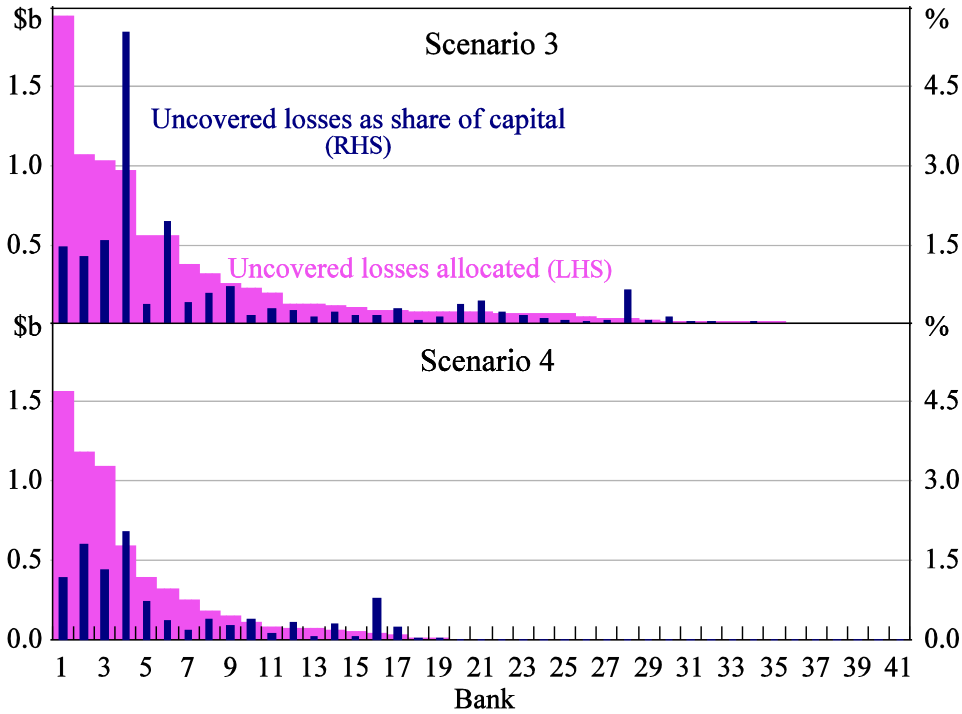

Notes: 3.89σ price change, 10 per cent stress threshold; banks ordered by size of losses allocated

Figure A2: Distribution of Loss Allocation with *Ex Post* Calls

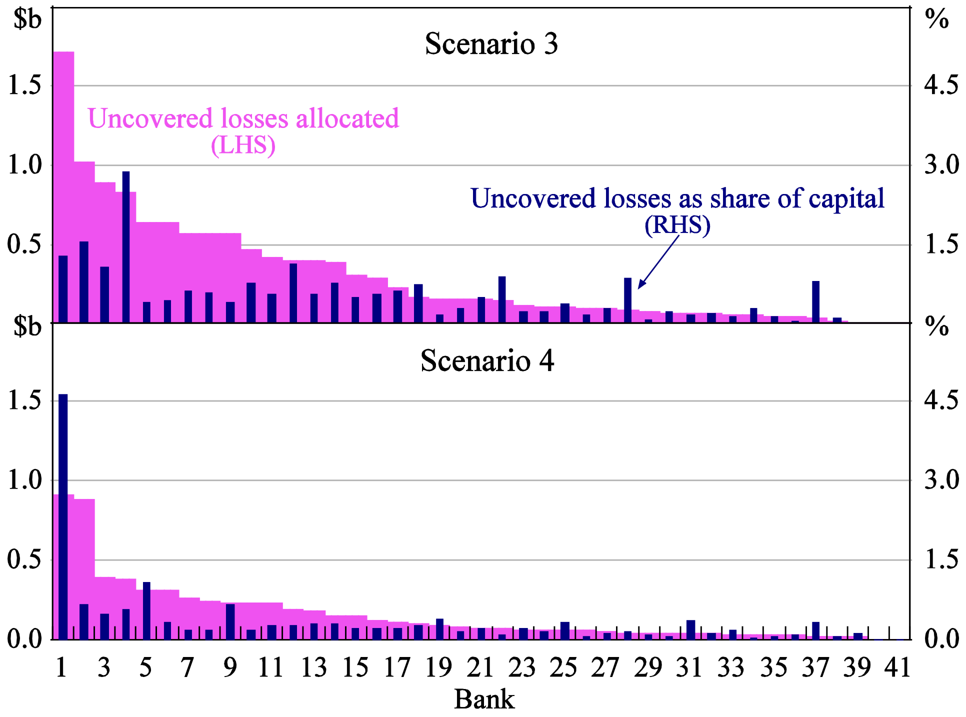

Notes: 3.89σ price change, 10 per cent stress threshold; banks ordered by size of losses allocated

References

Bailey D (2014), 'The Bank of England's Perspective on CCP Risk Management, Recovery and Resolution Arrangements', Speech given at Exchange of Ideas #2 'Central Clearing – Guarantee of Stability or New Moral Hazard', Hosted by Eurex, Deutsche Börse Group, London, 24 November.

BCBS (Basel Committee on Banking Supervision) (2014), *Capital Requirements for Bank Exposures to Central Counterparties*, Bank for International Settlements, Basel.

BCBS-IOSCO (Board of the International Organization of Securities Commissions) (2013), *Margin Requirements for Non-Centrally Cleared Derivatives*, Bank for International Settlements, Basel.

CGFS (Committee on the Global Financial System) (2013), *Asset Encumbrance, Financial Reform and the Demand for Collateral Assets*, CGFS Papers No 49, Bank for International Settlements, Basel.

Cheung B, M Manning and A Moore (2014), 'The Effective Supply of Collateral in Australia', RBA *Bulletin*, September, pp 53–66.

Cœuré B (2014), 'The Known Unknowns of Central Clearing', Speech given at the meeting on 'Global Economy and Financial System', Hosted by the University of Chicago Booth School of Business Initiative on Global Markets, Coral Gables, 29 March.

CPMI (Committee on Payments and Market Infrastructures)-IOSCO (2014), *Recovery of Financial Market Infrastructures*, Bank for International Settlements, Basel.

CPMI-IOSCO (2015), *Public Quantitative Disclosure Standards for Central Counterparties*, Bank for International Settlements, Basel.

CPSS-IOSCO (Committee on Payment and Settlement Systems and Technical Committee of the International Organization of Securities Commissions) (2012), *Principles for Financial Market Infrastructures*, Bank for International Settlements, Basel.

Duffie D (2014), 'Resolution of Failing Central Counterparties', Graduate School of Business, Stanford University, Working Paper No 3256.

Duffie D, M Scheicher and G Vuillemey (2014), 'Central Clearing and Collateral Demand', NBER Working Paper No 19890.

Duffie D and H Zhu (2011), 'Does a Central Clearing Counterparty Reduce Counterparty Risk?', *The Review of Asset Pricing Studies* 1(1), pp 74–95.

Elliott D (2013), 'Central Counterparty Loss-Allocation Rules', Bank of England Financial Stability Paper No 20.

FSB (Financial Stability Board) (2013), 'Application of the Key Attributes of Effective Resolution Regimes to Non-Bank Financial Institutions', October.

Furfine CH (2003), 'Interbank Exposures: Quantifying the Risk of Contagion', *Journal of Money, Credit, and Banking*, 35(1), pp 111–128.

G20 (Group of Twenty) (2009), 'Leaders' Statement: The Pittsburgh Summit', 24–25 September.

Gibson M (2013), 'Recovery and Resolution of Central Counterparties', RBA *Bulletin*, December, pp 39–48.

Heath A, G Kelly and M Manning (2013), 'OTC Derivatives Reform: Netting and Networks', in A Heath, M Lilley and M Manning (eds), *Liquidity and Funding Markets*, Proceedings of a Conference, Reserve Bank of Australia, Sydney, pp 33–73.

Heller D and N Vause (2012), 'Collateral Requirements for Mandatory Central Clearing of Over-The-Counter Derivatives', BIS Working Paper No 373.

ISDA (International Swaps and Derivatives Association) (2013), 'CCP Loss Allocation at the End of the Waterfall', Risk Management Technical Paper, August.

ISDA (2014), 'ISDA Margin Survey 2014', 10 April.

ISDA (2015), 'CCP Default Management, Recovery and Continuity: A Proposed Recovery Framework', Risk Management Position Paper, January.

ISDA, IIF and AFME (ISDA, the Institute of International Finance and the Association for Financial Markets in Europe) (2012), 'BCBS-IOSCO Proposal on Margin Requirements for Non-Centrally-Cleared Derivatives', Letter to BCBS and IOSCO, 12 December.

JP Morgan Chase (2014), 'What is the Resolution Plan for CCPs?', *Perspectives*, September.

Levels A and J Capel (2012), 'Is Collateral Becoming Scarce? Evidence for the Euro Area', *Journal of Financial Market Infrastructures*, 1(1), pp 29–53.

MAGD (Macroeconomic Assessment Group on Derivatives) (2013), *Macroeconomic Impact Assessment of OTC Derivatives Regulatory Reforms*, Bank for International Settlements, Basel.

Markose SM (2012), 'Systemic Risk from Global Financial Derivatives: A Network Analysis of Contagion and its Mitigation with Super-Spreader Tax', IMF Working Paper No WP/12/282.

Markose SM, S Giansante and AR Shaghagi (2012), '"Too Interconnected to Fail" Financial Network of US CDS Market: Topological Fragility and Systemic Risk', *Journal of Economic Behavior & Organization*, 83(3), pp 627–646.

Pirrong C (2011), 'The Economics of Central Clearing: Theory and Practice', ISDA Discussion Paper No 1.

Powell JH (2014), ‘A Financial System Perspective on Central Clearing of Derivatives’, Keynote Address at Seventeenth Annual International Banking Conference ‘The New International Financial System: Analyzing the Cumulative Impact of Regulatory Reform’, Held by the Federal Reserve Bank of Chicago and the Bank of England, Chicago, 6 November.

RBA (Reserve Bank of Australia) (2014), *Submission to the Financial System Inquiry*, Reserve Bank of Australia, Sydney, March.

Shaghaghi AR and SM Markose (2012) ‘Empirical Calibration of Global Financial Derivatives Network: Genetic Algorithm for Multi-Objective Optimization of Data-based Constraints’, University of Essex, Economics Department and Centre for Computational Finance and Economic Agents, unpublished manuscript.

Sidanius C and F Zikes (2012), ‘OTC Derivatives Reform and Collateral Demand Impact’, Bank of England Financial Stability Paper No 18.

Singh M (2013), ‘The Changing Collateral Space’, IMF Working Paper No WP/13/25.

Singh M (2014), ‘Limiting Taxpayer “Puts”—An Example from Central Counterparties’, IMF Working Paper No WP/14/203.

Tucker P (2011), ‘Clearing Houses as System Risk Managers’, Keynote Address at the DTCC-CSFI (Depository Trust & Clearing Corporation and Centre for the Study of Financial Innovation) Post-Trade Fellowship Launch, London, 1 June.

Tucker P (2014), ‘Regulatory Reform, Stability, and Central Banking’, Hutchins Center on Fiscal & Monetary Policy at Brookings, Hutchins Center Working Papers, 16 January.

Wendt F (2015), ‘Central Counterparties: Addressing their Too Important to Fail Nature’, IMF Working Paper No WP/15/21.

RESERVE BANK OF AUSTRALIA