BANK OF ENGLAND

Financial Stability Paper No. 30 – October 2014

Dear Prudence, won't you come out to play? Approaches to the analysis of central counterparty default fund adequacy

David Murphy and Paul Nahai-Williamson

BANK OF ENGLAND

Financial Stability Paper No. 30 – October 2014

Dear Prudence, won't you come out to play? Approaches to the analysis of central counterparty default fund adequacy

David Murphy and Paul Nahai-Williamson

We have benefited hugely from comments from the paper's reviewers, Paul Bedford, Evangelos Benos, Joe Noss, Tomo Ota, Edwin Schooling Latter, Nick Vause and Graham Young.

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

paul.nahai-williamson@bankofengland.co.uk
Paul Nahai-Williamson, Bank of England, Threadneedle Street, London, EC2R8AH

ISSN 1754–4262

Contents

Dear Prudence, won't you come out to play? Approaches to the analysis of central counterparty default fund adequacy

David Murphy and Paul Nahai-Williamson

Central counterparties (CCPs) are a key feature of the post-crisis financial system, and it is vital that they are robust. Indeed, as Paul Tucker said, 'it is an understatement that it would be a disaster if a clearing house failed' (Tucker (2011)). Therefore the question of how safe CCPs are is an important one. A key regulatory standard for CCPs is 'cover 2': this states that systemically important clearing houses must have sufficient financial resources to 'cover', or be robust under the failure of, their two largest members in extreme but plausible circumstances. This is an unusual standard, in that it is independent of the number of members a CCP has. Therefore it is natural to ask how prudent the cover 2 standard is for different sizes of CCP. This is the question investigated in this paper.

We first use a simple model to quantify the likelihood of CCP failure. This model is used to produce stylised results showing how the probability of failure of a CCP that meets the cover 2 standard can be estimated. Second, we present a simple approach to explore how the distribution of risk among clearing members affects the prudence of the cover 2 standard. Our results give some reassurance in that we find that CCPs meeting the cover 2 standard are not highly risky provided that tail risks are not distributed too uniformly amongst CCP members. They do however suggest that CCPs and their supervisors should monitor this distribution as central clearing evolves.

1 Introduction

Central counterparties are pieces of market infrastructure which sit between market participants, guaranteeing their performance to each other on certain transactions (Murphy (2013)). Securities, repo and derivatives markets all have multiple CCPs, and indeed the use of CCPs is or will be mandated in some cases. For instance, both the European Union's (EU's) EMIR (European Union (2012)) and United States' Dodd-Frank Act (US Congress (2010)) require that large over-the-counter (OTC) derivatives market participants use CCPs for certain transactions.

CCPs play a key role in reducing contagion: they should make it far less likely that a failure of one market participant causes direct losses at another. They should therefore act as a buffer between their members by absorbing most shocks. They can only do this if they have sufficient resources to be robust under plausible stresses. The importance of this robustness has been recognised for some time. Indeed, as Ben Bernanke said in 2011 (Bernanke (2011)):

For more than a century, financial stability has depended on the resilience under stress of clearinghouses and other parts of the financial infrastructure. As we rely even more heavily on these institutions in the United States and around the world, we must do all that we can to ensure their resilience.

This paper contributes to the growing literature on CCP safety (Cumming and Noss (2013), Duffie and Zhu (2011), Jones and Pérignon (2008), Lin and Surti (2013), Murphy (2012, 2013)). In particular we examine how absorbent the buffer a CCP provides really is. This question is best approached by looking at how losses created by the failure of CCP members are allocated, so we turn to this question next.

1.1 The default waterfall

CCPs stand between their members. If one of these members fails, it may cost the CCP money to close out the defaulter's portfolio. CCPs have various protections against this risk:

- First, CCPs have standards for the parties they will admit as members;

- Second, CCPs adopt a 'defaulter pays' model: all members have to agree to both initial and variation margin arrangements whereby typically variation margin entirely covers the market value of the position at the time that margin requirements are estimated, and initial margin covers the likely change in value of the position over some margin period of risk to a high degree of confidence. These margin amounts are available to the CCP to absorb any costs of default.

- After this, CCPs often have some of their own capital at risk, although this is not a requirement in all jurisdictions.

- There is also a fund which is jointly provided by all members. This *default fund* (DF) forms a layer of mutualised capital which can absorb losses which exceed the defaulter's margin and CCP equity-at-risk. Typically first the defaulter's default fund contributions are at risk, followed by the rest of the fund.

- CCPs also usually have the ability to 'call' their members for further amounts of default fund.

- Then finally there may be other arrangements, such as haircutting variation margin, to allocate losses. More of the CCP's equity may be available to absorb losses too.

The order in which different elements of a CCP's financial resources are accessed is known as the CCP's *default waterfall*: the funded elements of the waterfall are known as the CCP's *financial resources*, and many large CCPs share the pattern set out above for these elements.

There is some variation from CCP to CCP about the later elements of the waterfall, but it is reasonable to suppose that there might be a loss of market confidence in any CCP which had to resort to these approaches, so the question of the adequacy of the first four funded elements of the waterfall is of particular interest.

1.2 Design choices and constraints in CCP financial resources

There are, in the abstract, a range of choices in the design of CCP default waterfalls. For instance, initial margin levels can be small, with correspondingly big default funds; or it can be much larger with less loss mutualisation (Murphy (2013), Nahai-Williamson *et al* (2013)). There are however regulatory constraints here. The internationally agreed Principles for Financial Markets Infrastructure (Committee on Payment and Settlement Systems and Technical Committee of the International Organization of Securities Commissions (2012)), for instance set a minimum size of initial margin in Principle 6:

A CCP should adopt initial margin models and parameters that are risk-based and generate margin requirements sufficient to cover its potential future exposure to participants in the interval between the last margin collection and the close out of positions following a participant default. Initial margin should meet an established single-tailed confidence level of at least 99% with respect to the estimated distribution of future exposure.

This requirement mandates that initial margin covers at least 99% of the distribution of losses based on the current portfolio and current market conditions.

Principle 4 requires large CCPs to have default funds that are large enough to be able to absorb the loss resulting from the simultaneous default of the two largest debtors to the CCP:

In addition, a CCP that is involved in activities with a more-complex risk profile or that is systemically important in multiple jurisdictions should maintain additional financial resources sufficient to cover a wide range of potential stress scenarios that should include, but not be limited to, the default of the two participants and their affiliates that would potentially cause the largest aggregate credit exposure to the CCP in extreme but plausible market conditions.

It is worth unpicking this requirement a little. Suppose that we are concerned with an OTC derivatives CCP, and that large losses can only occur on cleared portfolios (rather than, say, CCP investments or its liquidity arrangements). In that case initial margin will cover nearly all of the 'ordinary conditions' losses, while the Principle 4 requirement entails an obligation to have a big enough default fund such that the CCP can absorb the default of the two largest clearing members in stressed conditions from its financial resources. This requirement is known as 'cover 2' for short.

Note that cover 2 is a pure 'stressed conditions' requirement, whereas initial margin is primarily a 'current conditions' measure.(1)

1.3 Our approach

The cover 2 requirement is important because it is an internationally agreed minimum standard for the amount of financial resources that systemically important CCPs must have. Both EU (European Union (2012)) and US (Commodity Futures Trading Commission (2013)) regulatory frameworks contain a cover 2 requirement for these entities, for instance. Therefore there is practical interest in the question of how prudent the cover 2 standard is for different CCPs.

Cover 2 can be regarded as analogous to the 8% minimum capital ratio in Basel I, in that it is an internationally harmonised simple minimum requirement which supervisory authorities can supplement with their own (perhaps more nuanced) stipulations.

In this paper we examine two aspects of the prudence of the cover 2 standard. First we show how market data can be used to estimate the complete distribution of a CCP's counterparty credit risk during stressed conditions. This approach gives some insight into how much risk can be absorbed by a CCP which just meets the cover 2 minimum standard.

Mandatory clearing requirements will cause CCPs to grow over time. They may acquire more clearing members as clearing is extended to more types of market participant, more products, and more jurisdictions, and as liquidity concentrates in leading CCPs for certain product lines. Therefore our second contribution is to examine the prudence of cover 2 as a function of the number of clearing members at a CCP.

Our modelling approach here is deliberately simple: we use elementary models of multiple clearing member defaults and treat default separately from exposure, for instance. This allows us to illustrate the issues simply and starkly: more sophisticated models would show (somewhat) greater risks.

The rest of this paper is structured as follows. Section 2 examines the protection that CCPs have against the failure of their members. This allows us to build a simple model to quantify the likelihood of CCP failure in Section 3. We present some stylised results in this setting in Section 4, showing how the probability of failure of a CCP that meets the cover 2 standard can be quantified. A key issue here is the distribution of cleared risk amongst clearing members. Section 5 examines the prudence of the cover 2 standard as this distribution changes. Section 6 concludes.

2 Cover 2 and default fund sizing

The process by which CCPs determine the size of their default fund, and hence meet the cover 2 requirement, is worth reviewing in some detail.(2) In this section we first examine the calculation of cover 2, then discuss its implications both for the size of CCP default funds and for the safety of CCPs.

2.1 Calculating the cover 2 requirement

The typical approach to determining cover 2 is as follows:

1. The CCP constructs a range of scenarios of possible market moves in conditions that would cause stress to its clearing members.(3) Typically these include both historical scenarios (such as the 2008 credit crisis or the events in late 1997 around the failure of Long Term Capital Management), and hypothetical scenarios.

2. The CCP determines the portfolios to be stressed. This list would include the house and client portfolios of all clearing members.(4)

3. The financial (P/L) impact of applying each scenario to each portfolio is determined. Profits are discarded.

(1) Some regulations, such as the EU's EMIR discussed further below, strengthen the standards for initial margin by requiring that some risk factor returns from a stressed period are included in the data used to calibrate CCPs' initial margin models. However the general point is clear: initial margin estimates are based on the return distribution from ordinary and perhaps stressed periods, not stressed periods alone.
(2) CCPs can have multiple services, with one default fund per service. For ease of explanation we assume a single service: our analysis can easily be extended to multiple services.
(3) This move should be a multi-day one, where the period chosen reflects the liquidation horizon of the products cleared by the CCP during the stressed markets which are likely to occur following a clearing member default.
(4) We have further simplified by assuming a single client omnibus account.

4. For each portfolio and each scenario, the relevant amount of initial margin (IM) is subtracted from the loss. This gives a matrix of 'stressed losses over initial margins (SLOIMs)', as depicted in **Table A**, where SLOIM(1C,S) is the stressed loss over IM that the first client account 1C would suffer in scenario S.

Table A Illustration of the calculation of SLOIMs

Account	Scenario 1	2	...	n
Clearing member 1 House	SLOIM(1H,1)	SLOIM(1H,2)	...	SLOIM(1H,n)
Clearing member 1 Client	SLOIM(1C,1)	SLOIM(1C,2)	...	SLOIM(1C,n)
Clearing member 2 House	SLOIM(2H,1)	SLOIM(2H,2)	...	SLOIM(2H,n)
Clearing member 2 Client	SLOIM(2C,1)	SLOIM(2C,2)	...	SLOIM(1C,n)
...	...	...	...	...

To make this concrete, suppose that the actual stressed losses over IM for an example CCP with four clearing members are as in **Table B**. This table shows three scenarios: real CCPs would of course use many more, but three suffices for explanatory purposes.

Table B Example SLOIMs by account, US$ millions

Account	Scenario 1	2	3
Clearing member 1 House	20	40	80
Clearing member 1 Client	30	50	500
Clearing member 2 House	160	70	200
Clearing member 2 Client	140	110	240
Clearing member 3 House	100	50	250
Clearing member 3 Client	15	30	280
Clearing member 4 House	0	12	250
Clearing member 4 Client	340	6	50

5. Next, the losses for each account at the same clearing member are added.(1) The effect on our example is shown in **Table C**.

Table C Example SLOIMs by clearing member, US$ millions

Account	Scenario 1	2	3
Clearing member 1	50	90	580
Clearing member 2	300	180	440
Clearing member 3	115	80	530
Clearing member 4	340	18	300

6. Each column is then sorted, so that the first entry is the largest loss in that scenario, the second entry is the second largest loss, and so on. Reading a given column down, then, gives an estimate of the stressed loss over IM if the clearing member with the largest exposure were to default after the scenario move; then the loss for the second largest clearing member, and so on. **Table D** shows the results of this for this example.

Table D Example sorted SLOIMs, US$ millions

Defaulter	Scenario 1	2	3
Biggest	340	180	580
2nd biggest	300	90	530
3rd biggest	115	80	440
4th biggest	50	18	300

7. For each scenario, the loss if the largest two clearing members were to default is determined. This is simply the sum of the first and second entries in that scenario's column.

8. The cover 2 requirement is the largest of these losses. Thus for instance in **Table D**, the losses are US$640 million, US$270 million and US$1.11 billion for scenarios 1, 2 and 3 respectively, so the cover 2 requirement for this example would be US$1.11 billion. The most stressful event considered is the simultaneous default of clearing members 1 and 3 in conditions given by scenario 3.

There are therefore two principal drivers of the SLOIM: the scenarios used; and the portfolios that they are applied to. The former should be revised when new vulnerabilities become apparent, but they often do not change from day to day. The latter however are highly variable, and hence CCPs must ensure that their stress tests are genuinely stressful for the portfolios they clear, and then perform these stressful stress tests daily on each of these portfolios.

In the sections that follow, we abstract away from this process, assuming that the scenarios chosen do indeed represent a full range of extreme but plausible events given the portfolios cleared.

2.2 Default fund sizing and cover 2

Some of the key regulatory requirements for systemically important CCPs can be summarised as in **Figure 1**.(2)

These apply from the bottom up, so first initial margin must be appropriately sized to meet the desired confidence interval

(1) For simplicity we have assumed here first that all clients are in an omnibus account, something that will not be true in most cases; and second that surpluses of house margin are not used to meet client deficiencies, something that is permitted in European regulation. We have moreover assumed that the CCP does not have more than one member from any group, and hence there is no need to aggregate exposures from affiliates.
(2) This is a high level summary and does not cover the detailed requirements in different jurisdictions, eg for 'wind-down' capital. We also do not discuss the issue of which parties provide these resources in what proportions.

Figure 1 Summary of sizing requirements for various elements of CCP financial resources

Rest of the default fund	Cover 2 requirement for total default fund
CCP skin in the game	Minimum 'skin in the game' capital requirement
Defaulter's DF	Risk-based allocation of total DF to defaulter
Defaulter's margin	Margin must be sufficient to cover 99% of the risk on each portfolio separately

(which is floored at 99% or more), then the CCP must have sufficient 'skin in the game' to meet its equity capital requirement (where such a requirement applies), then the total default fund must be larger than the stressed loss over IM caused by the default of the two largest clearing members in the worst scenario.

2.3 The unusual nature of cover 2 as a risk measure

It is clear from the description above that the cover 2 requirement only depends on the largest two stressed losses over IM in the worst scenario. The risks of the clearing members who are not one of the two largest do not contribute at all to the requirement. This makes cover 2 a rather unusual risk measure: typically, adding incrementally more risk to an entity increases its measured risk. Thus for instance if we add a new market risk position which is not negatively correlated with the existing portfolio, then the total value-at-risk goes up. With cover 2 however adding new clearing member risk to a CCP does not necessarily increase the risk to be covered by the default fund: it is only if this increase contributes one of the worst two losses that the risk measure changes. Implicitly, then, the cover 2 standard assumes that the largest two clearing members are the only source of counterparty credit risk to the CCP which are relevant for the sizing of default fund, or equivalently that no more than one additional clearing member can ever default before a CCP has replenished its default fund following an initial default.

A natural question that arises from this observation is how the prudence of cover 2 changes as more clearing members sign up to a CCP. Is the cover 2 standard always prudent enough for all CCPs as the number of members doubles or trebles, for instance?

We analyse this question in two different ways in subsequent sections: Sections 3 and 4 use a market-consistent approach, while Section 5 studies the problem analytically. First, though, we need to set some groundwork.

2.4 Combining defaults

The 'cover 2' measure assumes that precisely two clearing members will default and that these will be the largest two.(1) This is clearly arbitrary: in general any number of members can default and the ones that do default will not necessarily be the largest ones.

There are two opposing forces at work in the cover 2 definition:

- First, the probability of more than two failures in a collection of financial institutions over a fixed time goes up as the number of institutions in that collection goes up: three random failures out of 20 are more likely than three out of ten.

- Second, as the number of clearing members at a CCP goes up then the probability that any two independent defaults will be those of the biggest two members goes down. This simply reflects the fact that an independent draw of two random members from ten is more likely to give two particular members than the same draw from 20.

The first of these effects makes cover 2 less prudent as the number of clearing members increases. However the second acts in the opposite direction, making cover 2 more prudent as the number of clearing members increases. This is because the bad event that we are covering becomes less likely.

These effects can be quantified and combined with information about how much clearing member portfolios might lose between the last successful margin call and close-out, to gain insight into the prudence of cover 2. The next section examines the ingredients of this model, and compares the two effects mentioned above.

3 Model ingredients and first steps

Any model of CCP counterparty credit risk has to answer three key questions:

1. How likely is the default of each clearing member?

2. What is the relationship between clearing member defaults?

3. If a clearing member defaults, what is the resulting loss or distribution of possible losses for the CCP?

There are different ways to answer these questions. For instance, taken literally, the cover 2 standard implicitly answers them by assuming:

(1) It is should be noted that CCPs do not need to identify the largest two members and the size of their SLOIMs in *all* circumstances. Rather the aim of CCP default fund sizing stress tests is to identify the characteristic size of the losses in a wide range of extreme but plausible circumstances, and to use these SLOIMs to inform the default fund size. Put another way, CCPs do not need to accurately predict who will lose what in every plausible stressed circumstance, but they do need a prudent estimate of what a large SLOIM might be.

Box 1
Credit Copulas 101

If we have two clearing members, A and B, each of which have liquidly traded credit default swaps which reference them, then we can estimate the probability that each clearing member will default, p_A and p_B say.

Copulas are a way of answering the question 'what is the probability that *both* A and B will default?' Clearly if the two clearing members are closely related, perhaps because they are affiliates, then we might have the probability of both A and B defaulting be $\max(p_A, p_B)$. On the other hand, if A's default is truly independent from B's, then the probability of both members defaulting might be $p_A \times p_B$.

More typically, there will be some positive relationship between the two defaults, not least because both clearing members are banks, and likely have trading relationships with each other. The Gaussian copula expresses this relationship via a positive correlation. Specifically, in this setting we assume that each clearing members' default is driven by some random variable, and that we can express the relationship between defaults by correlating these variables.

This is a simple choice: more sophisticated copulas allow a more sophisticated relationship between defaults to be expressed (Burtschell, Gregory and Laurent (2009)) — and there are also non-parametric approaches to the same problem. Since we are primarily interested in a handful of defaults out of a collection of perhaps 50 clearing members, the Gaussian copula is an acceptable model:(1) its limitations become more apparent when looking at larger numbers of defaults.

(1) It would however be interesting to rerun this analysis using other copulas, or to impose a more nuanced comovement structure, for instance assigning a higher default correlation to clearing members from the same country.

- The biggest two members are certain to default, but no other defaults can occur; and

- The loss on default is the loss in the worst 'extreme but plausible' scenario.

3.1 Market-consistent default estimates

The market can provide an answer to question 1. We examine the quoted premia (or 'spreads') of credit default swaps referencing each clearing member, and infer a probability of default from these using an assumed recovery rate (Schönbucher (2003)).(1) There are issues in this approach (Huang and Huang (2012), Noss (2010)), not least the assumption that all of the credit spread is compensation for default risk (instead of, *inter alia*, liquidity risk) but nevertheless this approach is standard in the literature and we adopt a version of it. Specifically for ease of calculation we assume that all clearing members have an identical probability of default equal to the average probability of default inferred from the credit default swap (CDS) spreads of banks which are clearing members of global CCPs.

3.2 Market-consistent default associations

The relationship between defaults is more problematic. If the prices of the tranches of a collateralised debt obligation (CDO) whose underlying credits were the clearing members were available, then we could potentially use those to infer a market-implied default association: Box 1 explains the associated mathematical tool of a copula in a little more detail.

The problem is that the required prices are not available; at best we have reliable prices for the tranches of different CDOs, namely the liquidly traded credit derivatives indices (Amato and Gyntelberg (2005)). Moreover the 'right' mathematical model of default association is still an open question, and certainly no market consistent model of CDO tranches is entirely free from model risk and calibration issues.(2) Faced with this, we chose the simplest possible approach, modelling default association with a Gaussian copula (Li (2000)) informed by typical long-run average correlations, viz. 30%. There is model risk in this approach, as it will underestimate the probability of many banks failing in a small period of time. Hence we will (slightly) over-estimate the prudence of cover 2 when we use this model in Section 4.

3.3 Cover 2 as the number of clearing members varies

This setting, even without knowing the distribution of possible losses for the CCP, can be used to illustrate the two effects discussed in Section 2.4. We simply take our simple model of default and calculate the probability of experiencing two defaults in a fixed period, and the probability that those two defaults will be the worst two.(3)

Chart 1 shows the results of this calculation, while Box 2 discusses the question of what the fixed period should be.

(1) An even more direct approach would be to obtain a market quotation for a third-to-default CDS on the basket of clearing members. However, given the highly bespoke nature of this product, this quotation would likely incorporate a significant structuring premium.
(2) Either the model's calibration is not time-stable, or the model has so many free parameters that it runs the risk of 'over-fitting' the data, or both (Burtschell, Gregory and Laurent (2009), Cousin and Laurent (2008)).
(3) The use of our simple model implies in particular the assumption that all clearing members have the same default probability, and differing SLOIMs.

Box 2
The time period for defaults

Suppose that a CCP suffers a loss from one default. CCP rule books permit, and regulators will often require, that the CCP 'refill' the default fund within a fixed period. The CCP is 'on risk' during this period as another default could occur and further erode or even exhaust the default fund. However, for most leading CCPs, this period is rather short: calls for additional default fund often have to be met within a week or less.

This means that we may be interested in studying the probability of multiple defaults in rather short periods of time. Unfortunately the standard quantitative framework used for pricing credit derivatives is ill-suited to this problem as it is designed for products with much longer lives: five years is the most liquid tenor for many types of credit derivative, and one-week CDS on baskets or indices are extremely unusual. Thus there are no direct market-based estimates of the probability of multiple defaults in short periods of time.(1) This in turn means that indirect methods must be used to estimate the probability of situations such as one clearing member defaulting within a week of another.

We intend to use the probabilities of default and of loss inferred from our model to illustrate an approach to analysing CCP safety, so we use three-month probability of defaults. This is highly conservative in the light of the discussion above.

(1) Most credit derivatives models assume at least piecewise linear hazard rates. In fact, defaults show strong clustering. This issue is not material when pricing the tranches of a five-year CDO, but it does matter a great deal when, as we are, trying to price the mezzanine tranche of a one week structure.

Chart 1 Properties of cover 2 as the number of clearing members varies[(a)]

(a) Probability of two and more than two defaults

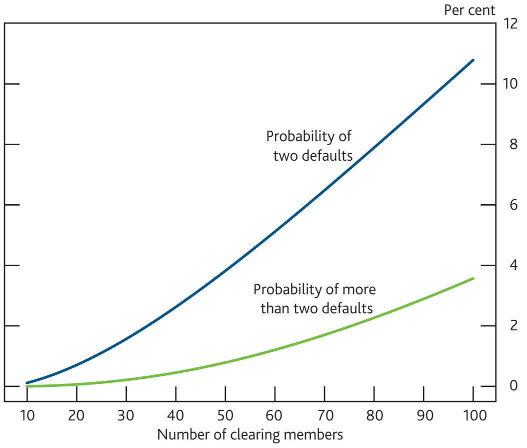

(b) Probability of two defaults being the worst two

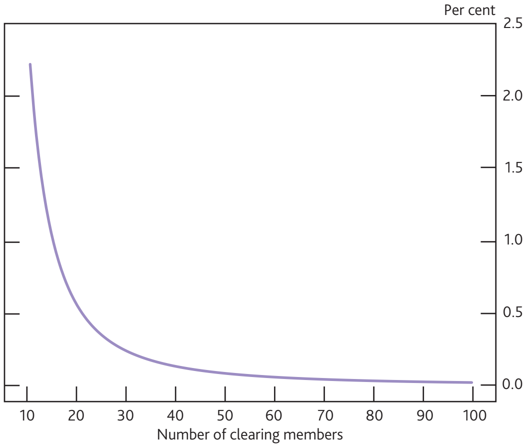

Source: Bank calculations.

(a) **Chart 1b** shows the three-month probability of exactly two clearing members failing (in blue) and more than two failing (in green); and **Chart 1b** shows the probability of two random defaults being the worst two, all as a function of the number of clearing members.

If clearing members have the same credit quality, their probability of failure is unrelated to the size of their risk at the CCP, and our model of default association is reasonable (clearly big and unjustified assumptions), then this analysis tells us two things: the probability of more than two defaults rises for CCPs with more clearing members; but the chances that the experienced defaults will be the most dangerous ones decreases substantially.

One obvious question here is 'how damaging are defaults that are not of the biggest two members?' In order to answer that, we need some notion of the distribution of risk amongst members, so we turn to this in the next two sections.

4 Default fund safety contingent on stress

This section considers the use of stress scenarios to determine the losses suffered in the event of clearing member default. The results of these tests will give estimates of the riskiness of the CCP due to counterparty credit risk contingent on the market being stressed.

The stress test data discussed in Section 2.1 can be visualised by plotting stressed loss over IM for each clearing member. For instance, scenarios 1, 2 and 3 of **Table D** give rise to the blue, green and purple data respectively in **Chart 2**.

4.1 Using the stress scenarios

In practice, large OTC derivatives CCPs have more than four members and use more than three stress scenarios when estimating cover 2. **Chart 3** illustrates the six most stressful scenarios for a particular clearing service at a large UK CCP. Here the data has been scaled by dividing each loss by cover 2, so it can be seen that the largest single loss is 71% of cover 2 in the green line scenario. The graph continues out to the

Chart 2 Example stressed losses over IM for three scenarios

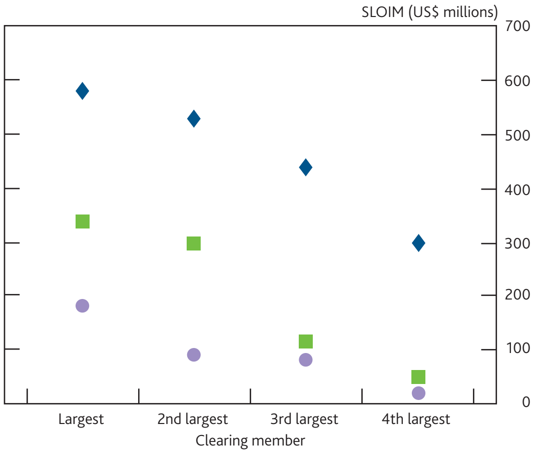

Source: Bank calculations.

Chart 3 Stressed losses over IM as a fraction of cover 2 for six stressful scenarios, where clearing members are ordered from largest to smallest contribution to cover 2

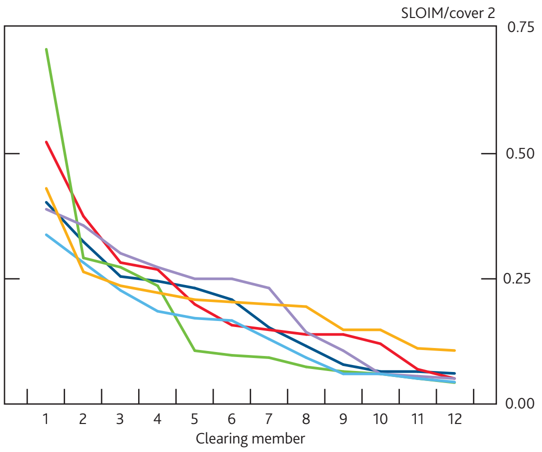

Sources: UK CCP and Bank calculations.

right for the remaining clearing members of the CCP: we have truncated at the twelfth clearing member because even for large CCPs, twelve or more defaults in a short period are vanishingly unlikely both in our copula and in reality absent an extreme systemic crisis.

These data can be used to derive the loss distribution for the CCP under the assumption of stressed markets and assuming that the default probability of a clearing member is independent of the size of its cleared exposure.(1) We combine scenario losses with the market consistent copula discussed in the previous section. The copula tells us how likely n losses are for n ranging from one to the total number of clearing members, while the scenario tells us how much will be lost for a given number of defaults.

There are actually a variety of choices here. The first relates to the choice of scenario:

- We could pick a single scenario on some basis. For instance, the scenario plotted in red is the one which generates cover 2, so that is one reasonable choice; or

- We could construct a new set of losses by taking the envelope of the scenarios. For our example we would use the green line scenario for the first loss, the red one for the next one, then the purple one for defaults three to seven, then the orange line scenario for losses beyond seven defaults. This would give some measure of robustness to situations where a handful or more clearing members are exposed to the same risk, and hence cover 5, say, comes from a different scenario to cover 1 or 2.

The second choice concerns the loss at default:

- We could assume that the first default is always of the member with the largest exposure, the next default that of the second largest, and so on; or

- We could randomly sample from the distribution of exposures.

We will examine just one possibility here: the next section illustrates how CCP counterparty credit risk can be analysed if we assume that clearing members' risk follows the scenario generating cover 2, and that the worst possible set of defaults happens. This example is chosen not because we think it is particularly realistic — it isn't — but rather as an illustration of how scenario losses and a copula model of the clearing members can be combined.

4.2 Worst defaults in the cover 2 scenario

Charts 4a and **b** illustrates the two pieces of the puzzle: loss data and probabilities of default. (**Chart 4b** uses log probabilities simply to display a wider range of values clearly.)

We can combine these to infer the probability of losses over the default fund, contingent on being in stress and contingent on the clearing members to which the CCP has the largest exposure defaulting. **Chart 5** illustrates the results for two to ten defaults. Two defaults is the most probable of the cases illustrated, and this results in no loss over cover 2, as we would expect. Three or more defaults can occur, though, and when they happen, they generate losses for this particular clearing service which are not fully absorbed by the default fund if the worst scenario is being experienced. The most probable of these loss-causing events is three defaults, occurring with probability 0.37%: the loss here is 27% of the cover 2 amount. Four, five etc — defaults are less likely but cause larger losses in the worst scenario if they do occur.

(1) The assumption that having more cleared risk does not increase the probability of default of the clearing member is reasonable given that by far the largest risks run by most clearing members are in their banking books.

Chart 4 Stressed losses over IM as a fraction of cover 2 for the worst scenario, and the probability of *n* defaults in a three-month period for a fixed clearing service

(a) SLOIM as a fraction of cover 2

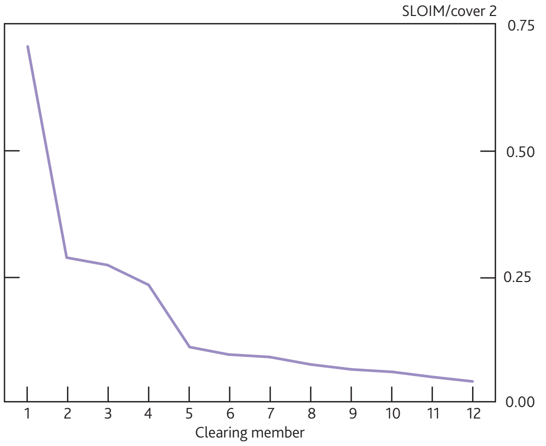

(b) The probability of *n* defaults

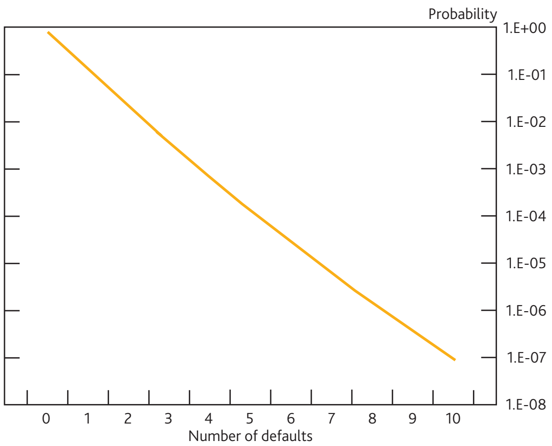

Source: Bank calculations.

Chart 5 The contingent probability of various losses in excess of cover 2, expressed as a fraction of cover 2

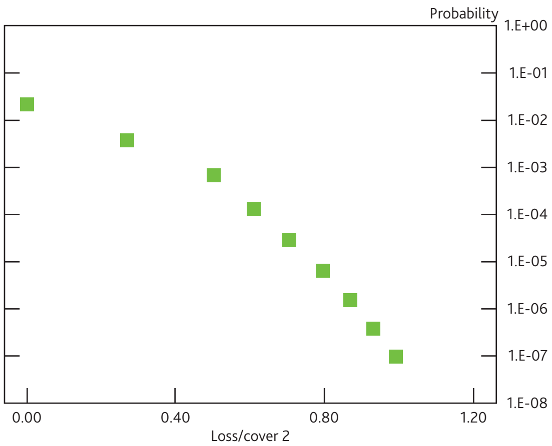

Source: Bank calculations.

4.3 Discussion

It should be noted here that some CCPs require additional initial margin from any member whose SLOIM exceeds a certain percentage of the total default fund. Therefore if one member consistently accounts for a majority of total SLOIM, a CCP might choose to have a smaller default fund than implied by the cover 2 number and make up the shortfall with initial margin collected from that member. This would result in a default fund that is less robust to member defaults, but would more equitably distribute the costs of mitigating the risks brought by different members.

The approach that we have taken up to this point depends on the distribution of SLOIMs among clearing members. Different distributions will lead to different levels of CCP safety, all other things being equal. Therefore it makes sense to examine some stylised distributions here: we turn to this next.

5. Default fund safety as the distribution of CCP risk varies

The previous section discussed how we can combine CCPs' SLOIM distributions with default probabilities and default correlations to produce a metric through which different CCPs' levels of counterparty credit risk can be compared. In this section, we explore how the relative robustness of different CCPs' cover 2 measures depends on the nature of their membership. One obvious factor that will affect a CCP's robustness is the credit quality of its members; for the purposes of this paper, we assume that all the CCPs' members are of the same credit quality in order to focus on how the distribution of risk alone impacts on the safety of the cover 2 standard. In this setting there are two key aspects of the CCP membership that are relevant to the safety of the cover 2 standard:

1. The way in which activity (and thus risk) is distributed among clearing members; and

2. The total number of members.

This section examines each of these aspects in turn. We begin by introducing three stylised distributions of risk.(1) Then we study the impact of the risk distribution on the safety of the cover 2 standard by simulating losses for each of these distributions. Finally, the effect of changing the number of clearing members is considered.

(1) Our intent is to illustrate how the distribution of risk (as measured by SLOIMs) among a CCP's members can impact on its resilience: it is not intended to represent the actual level of risk faced by current CCPs.

5.1 Three risk distributions

In considering the distribution of risk among clearing members, we will consider three cases:

a. risk is concentrated in a small subset of clearing members as is often found in practice;
b. all members contribute the same amount of risk; and
c. a single large clearing member accounts for a significant proportion of the total SLOIM.

The latter two distributions are not commonplace, but they are helpful in exploring the range of outcomes possible when applying the cover 2 standard.

Recall that the SLOIM per clearing member is a function of both their house and client portfolios. In what follows we will study, strictly, changing numbers of SLOIMs rather than changing numbers of clearing members: a clearing member that generates no SLOIM adds no extra risk to the CCP (at least under the 'lens' of stress testing), so arbitrarily many of these can be added without changing the risk measure. Therefore for this rest of this section, 'clearing member' should be read as 'clearing member with a positive SLOIM'.

5.1.1 The exponential distribution

The analysis of actual CCP stress test loss distributions in the previous section showed that, in general, risk is not uniformly distributed among clearing members. For OTC derivatives CCPs for example, a small number of members often account for the majority of the total exposure. The SLOIM exposure distribution found in practice is often roughly exponential, as illustrated in **Chart 6**.(1) This shows the actual distribution of stressed losses observed in the most stressful scenario at one UK CCP, with the largest SLOIM normalised to one: a simple exponential fit to this data which approximately describes the distribution is also illustrated.(2)

Chart 6 Observed distributions of SLOIM among clearing members for a large UK CCP, with the best fit exponential distribution

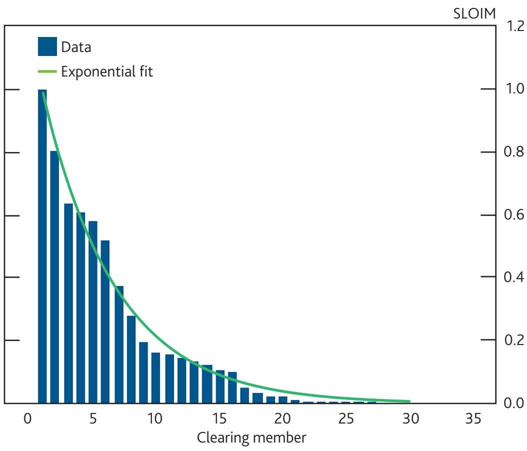

Sources: UK CCP and Bank calculations.

The shape of the SLOIM distribution for a given CCP is influenced by the evolution of the structure of the market it clears. Often, a small number of major liquidity providers emerge in a market; these parties then find it relatively easy to cement their position as their superior order flow provides informational and other advantages; the ability of these dealers to provide superior execution, especially for large orders, then attracts further clients. In fact power law distributions of market risk are widespread in financial markets (Gopikrishnan *et al* (2000)), to the point where *not* finding a power law is more noteworthy than finding one.

5.1.2 The uniform distribution

In order to consider the impact of the distribution of risk among members on the robustness of the cover 2 standard, it is useful to consider possible alternative distributions. One obvious edge case to consider is one in which all members of a CCP present approximately the same amount of risk; in other words, the CCP would face the same loss for any individual member that defaults. **Chart 7** illustrates this *uniform* SLOIM exposure distribution for a CCP with 30 members.

Chart 7 A stylised uniform distribution of SLOIM among clearing members

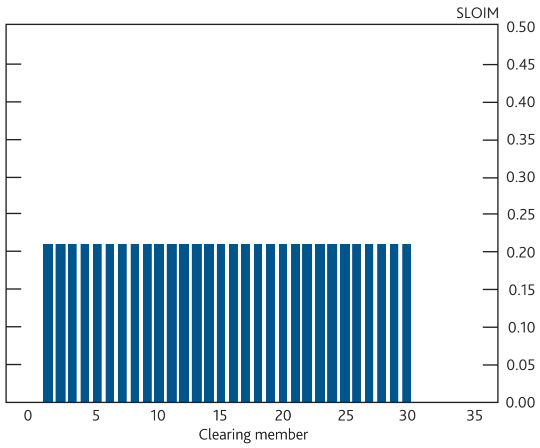

Source: Bank calculations.

5.1.3 The whale distribution

We can also consider the stylised case in which a single member accounts for a large proportion of the CCP's exposure. **Chart 8** shows a stylised distribution in which one member accounts for 40% of the total SLOIM exposure,(3) and

(1) An exponential distribution of SLOIMs is fairly typical of the loss distributions observed for many other scenarios used by UK CCPs, although both the number of members with non-zero SLOIMs, and the extent to which risk is more or less concentrated in a small number of members can vary significantly between CCPs and between different scenarios at the same CCP.
(2) The SLOIMs are for the combined losses of house and client accounts (where both house and client losses are individually floored at zero). The SLOIM distribution shown in **Chart 6** is taken from a CCP's stressed loss calculations from a single day, and the analysis below assumes that the SLOIM of each member in the distribution represents the loss that will be realised ex post should that member default.
(3) Note that while in a derivatives CCP no clearing member can have more than 50% of the total position, as the CCP must be balanced, a clearing member can have any fraction of the total SLOIM.

Chart 8 A stylised 'whale' distribution of SLOIM among clearing members

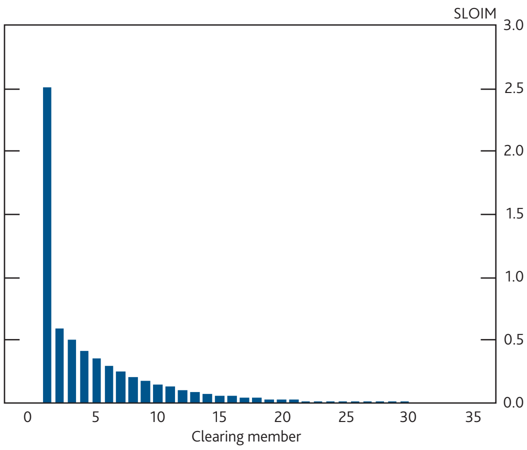

Source: Bank calculations.

the SLOIMs of the remaining 29 members follow an exponential distribution. This stylised distribution is henceforth referred to as the *whale* distribution.

In **Charts 6**, **7** and **8** the CCP's total SLOIM aggregated across members has been held constant: it is only the distribution of this exposure among members that has been changed.

5.1.3 Cover 2 in each of the three distributions

For the three cases considered above, the cover 2 requirements differ significantly. The cover 2 numbers as a proportion of total SLOIM are:

Table E The cover 2 requirement for three different exposure distributions

Distribution of exposures	Cover 2 requirement
Exponential	29% of total SLOIM
Uniform	7% of total SLOIM
Whale	49% of total SLOIM

From the above stylised examples, it is obvious that the size of a default fund produced by the cover 2 standard is highly dependent on the exposure profile of the CCP's membership, even when total SLOIM is held constant.

5.2 The prudence of cover 2 in each distribution

We can now explore how the distribution of risk among clearing members affects the robustness of the cover 2 standard in terms of the realised losses that a CCP might face in a market stress. To do this, we assign default probabilities to each member, and perform numerical simulations in which a member defaults with probability *p*, and survives with probability (1-*p*), using the same default correlation assumption as in the previous section. If a member defaults, the CCP realises a loss equal to that member's SLOIM, where the distribution of SLOIMs is given by one of the three distributions discussed in the previous section. We run a large number of simulations to produce a realised loss distribution for the CCP under each SLOIM distribution. The expected and unexpected losses associated with this distribution are then compared with the cover 2 amount.

For the purposes of this stylised analysis, an unrealistically high clearing member default probability of 5% is chosen. This allows us to more clearly illustrate differences in the CCP's realised loss distribution resulting from the different membership structures. As in Section 4, all members are assigned the same default probability.

5.2.1 The simulation of realised losses

Chart 9 shows the realised loss distributions produced by the simulation for the three SLOIM distributions as (unnormalised) histograms. The y-axis is truncated to make the loss distribution clearer; arrows indicate the cover 2 amount for each SLOIM distribution.

Chart 9 CCP realised loss distributions for the three exposure distributions(a)

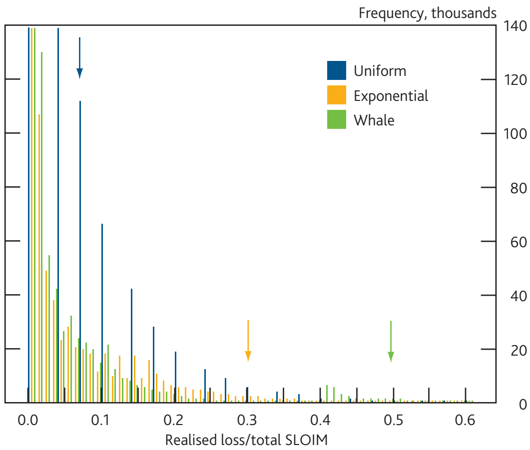

Source: Bank calculations.

(a) The respective cover 2 values are indicated with arrows. The axes are truncated to show the tail of the distribution more clearly.

Several observations can be made from **Chart 9**. First, where a CCP has similar exposures to all of its members as in the case of the uniform distribution (and where those members are of similar credit quality), the lower value of cover 2 means that there is a higher likelihood that it will face losses beyond cover 2, compared to the cases where a CCP has minimal exposure to many of its members.

Second, heterogeneity of a CCP's exposure to its members — represented by the exponential and whale SLOIM distributions — naturally leads to a longer tail, due to events in which it realises significant losses upon the default of one (or more) of the larger clearing members.

Third, where one clearing member accounts for a significant fraction of the total exposure (represented by the whale distribution), the CCP's counterparty credit risk profile will be bi-modal. That is, there will be a second peak in the loss distribution corresponding to occasions on which that clearing member defaults. It follows that if a CCP has two or three clearing members accounting for, say, 80% of the SLOIM in the market it clears, the realised loss distribution will have extra peaks far into the tail corresponding to the rare occasions on which each combination of these clearing members defaults.(1)

5.2.1 Cover 2 compared with simulated expected loss

The simulation of the three stylised loss distributions allows us to consider different aspects of the robustness of the cover 2 standard. One obvious metric is how the cover 2 standard compares with a CCP's *expected loss* conditional on being in stressed market conditions. This expected loss is simply the average realised loss observed in the numerical simulations. Although the three loss distributions presented above are quite different, they all produce a similar expected loss of 5%. Since the cover 2 size for each distribution varies significantly, so too does the multiple of expected loss that cover 2 represents (**Table F**):

Table F Cover 2 as a multiple of expected loss for three different exposure distributions

Distribution of exposures	Cover 2/expected loss
Exponential	579%
Uniform	133%
Whale	990%

The actual values above are not representative of the risks faced by CCPs, since we have used an unnaturally high default probability of 5% for each clearing member; the relative values for each loss distribution however are informative.

Unsurprisingly, a default fund sized at cover 2 is a significantly weaker safety net for CCPs in which the distribution of exposure among members is more uniform than for those in which it is more heterogeneous. This is solely a function of the way in which the cover 2 standard has been defined: cover 2 is insensitive to whether the two members on which the default fund is being sized account for, say, 5% or 50% of the total risk.

Chart 1b showed that for an increasing number of members, the probability of two defaulters being those with the largest SLOIMs decreases rapidly. This implicitly assumes however that not all members pose the same risk to the CCP. For a uniform distribution, we cannot rely on this relationship to support the prudence of the cover 2 standard: if all members have similar SLOIMs, then any two defaults will produce similar losses and so the probability of two defaulters being the worst is roughly constant as the number of clearing members varies. With 30 members then, it is unsurprising that the cover 2 standard for the uniform distribution is relatively less robust. The impact of the number of clearing members on the robustness of the cover standard is discussed further in Section 5.3.

Considering only the expected loss (in stressed market conditions) obscures important information about the tail of the realised loss distribution. The three realised loss distributions shown above for example have very different tail risks: **Table G** shows the 99th percentile of the realised loss distribution (a common measure of unexpected loss), and how this compares with the cover 2 requirement.

Table G Cover 2 as a multiple of unexpected loss for three different exposure distributions

Distribution of exposures	99th percentile unexpected loss (as proportion of total SLOIM)	Cover 2/unexpected loss
Exponential	42%	69%
Uniform	37%	18%
Whale	55%	90%

Events lying in the tail here are *extremely* unlikely to occur: in considering SLOIMs, we are already assuming that members are defaulting during a market stress that is itself beyond the 99th percentile of the distribution of market outturns. The 99th percentile of the realised loss distribution is thus in the tail of the tail.

Again we see that a default fund sized at cover 2 for a uniform member loss distribution offers less safety than for a service with a less homogeneous member loss distribution.

Finally, we consider the expected shortfall beyond cover 2, as a measure of the average risk to the CCP should cover 2 prove insufficient (**Table H**):

Table H Expected shortfall for losses in excess of cover 2 for three different exposure distributions

Distribution of exposures	Expected shortfall over cover 2 as proportion of total SLOIM
Exponential	11%
Uniform	11%
Whale	9%

The expected shortfall is the average loss conditional on losses exceeding cover 2; as such, it is sensitive to realised loss events far in the tail. While the probability of exhausting cover 2 is larger for the uniform distribution than for the other two, the average incremental losses once it is exhausted are

(1) Whether the joint default of two or more of such members would be visible in the loss distribution would depend on the magnitude of the joint default probability.

similar. This is because the shapes of the tails of all three distributions above cover 2 are broadly similar.

5.2.2 The impact of default correlation

The estimates in **Tables G** and **H** are a measure of tail risk and so are sensitive to the default correlation assumptions used. If we assume that defaults are independent for example, the expected shortfall over cover 2 is larger for the uniform distribution than for the other distributions. This is an important point: in practice, higher default correlations do not materially increase expected loss for these distributions, and so have little impact on the robustness of cover 2 in relation to this metric; but increasing default correlations extends the tail of the loss distribution such that the robustness of cover 2 against unexpected losses reduces. **Charts 10a** and **b** show the loss distributions for the exponential SLOIM distribution where correlations are zero and 30% respectively. **Chart 10b** has more weight in the tail, illustrating that increasing the correlation between clearing members' defaults increases the CCP's potential worst case losses.

5.2.3 Discussion

The analysis above illustrates the importance of the metric used to evaluate the robustness of CCPs. If we are interested in a CCP's expected loss conditional on stressed market conditions, the distribution of risk among clearing members may not be that important, since we find the same expected loss for all three distributions. If instead we are interested in a CCP's unexpected loss conditional on stressed market conditions, the distribution of risk matters more, as the unexpected loss is larger where risk is concentrated in a subset of clearing members (**Table G**). In contrast, the cover 2 standard, by construction, is as robust or more robust when risk is concentrated among a relatively small subset of clearing members (**Tables F** and **G**).

It is also important to note that the stylised results above come with a significant health warning when considering the realities of central clearing, and default management in particular. While the default fund may be relatively bigger if one member dominates risk-taking than if risk is more evenly distributed, managing the default of that member is likely in practice to be significantly more difficult than managing the default of members with smaller exposures, and entails significant concentration and liquidity risks. In such a case, the extra prudence gained in the sizing of the default fund may be undermined by other risks, unless these are separately addressed (for instance by the use of extra margin requirements accounting for the bigger risks entailed in liquidating large positions in stressed markets).

5.3 Increasing the number of clearing members

It is intuitively obvious that if more members join a CCP and bring additional risk to it, then the robustness of the cover 2 standard will be diluted. However, we must be careful in

Chart 10 Unnormalised histograms of the loss distributions for a CCP with SLOIMs following an exponential distribution(a)

(a) Loss distribution with zero default correlation

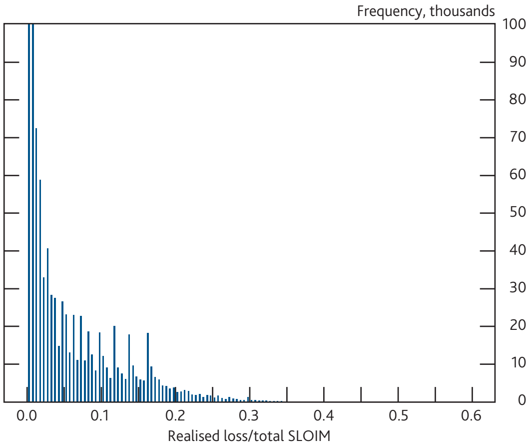

(b) Loss distribution with 30% default correlation

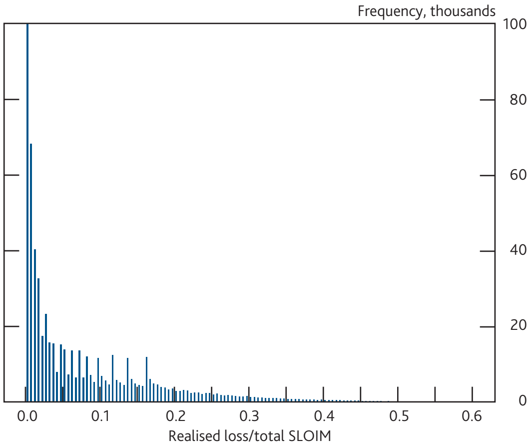

Source: Bank calculations.

(a) **Chart 10a** with zero default correlations. **Chart 10b** with default correlations of 30%. The axes are truncated to show the tails of the distribution more clearly.

assuming that more members imply more risk. Sometimes the main impact of new clearing members joining a CCP will be to increase competition, with the same amount of risk being shared among more members. On other occasions however, new clearing members will increase the total amount of risk at the CCP (while of course giving more potential contributors towards default management).

In the next two subsections we consider the impact of the number of direct clearing members on the robustness of the cover 2 standard. As in the previous section, we use an unrealistically high default probability of 5% to explore the tail of the distribution. First we consider the case in which new members increase the spread of existing risk at the CCP. Then we test a case in which new members add risk to the CCP. In order to retain comparability with previous sections, we add new members to the SLOIM distribution in a way that preserves the functional form of the distribution — in other words, the distributions do not change shape when we add members, but rather scale down (when new members spread existing risk), or extend outwards (when new member bring more risk). This approach is unlikely to reflect reality. If the addition of new members increases risk at a CCP for example, this is likely to impact on other members' SLOIMs and thus impact on the shape of the SLOIM distribution.

5.3.1 Increasing the number of clearing members 1: the same risk, spread among more members

Suppose that new members diversify activity at the CCP. To model this, we simply hold the total exposure in our system constant and scale our distributions in or out as the number of members increases or decreases.

This approach could be seen as a stylised description of a mature cleared market where the entry of a new member causes a fixed amount of risk to be shared out among participants. The exponential distribution models the situation found in many OTC markets, but other distributions closer to the uniform one might be found in situations with a different market structure.

The charts below plot how the robustness metrics considered in Section 5.2 evolve as the number of clearing members increases.

Chart 11 shows how cover 2 depends on the number of clearing members under each of the three loss distributions studied. For the exponential distribution, the cover 2 number reaches a plateau as the number of members increases as each new member has only a marginal impact on the losses of the largest two members. Having a 'whale' in the membership simply exacerbates this behaviour, as more of the weight of the distribution is concentrated in a small number of members. For the uniform distribution by contrast, the cover 2 default fund size continues to decrease as the number of members increases.

The combined effect of these factors is that the robustness of cover 2, as measured by both cover 2 over expected loss and cover 2 over unexpected loss, is significantly lower for the case in which risk is distributed evenly among members. **Chart 12** illustrates the latter effect.

This behaviour produces the result that for member distributions representative of those observed in practice, if we believe that new members act primarily to diversify existing risk at a CCP, then the dependence of the robustness of cover 2 on the number of members is significant only for a relatively low number of members.

We can also see that for the exponential and whale distributions, the robustness of cover 2 (as measured by the ratio of cover 2 to total losses) falls rapidly as membership

Chart 11 Cover 2 as a proportion of total SLOIM for three exposure distributions as the number of clearing members varies

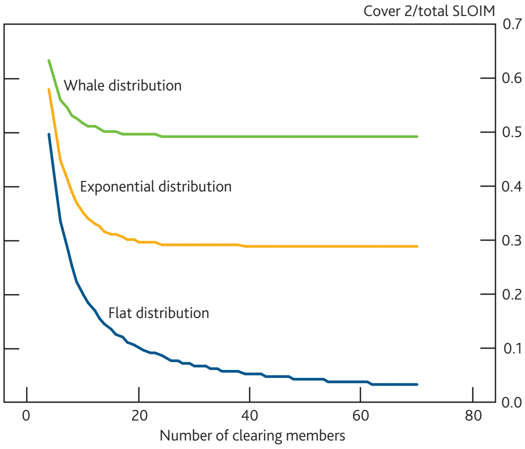

Source: Bank calculations.

Chart 12 Cover 2 as a proportion of the 99% unexpected loss for three exposure distributions as the number of clearing members varies

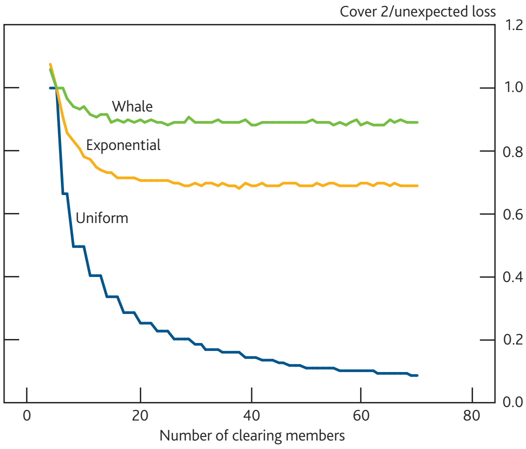

Source: Bank calculations.

increases from four members to around fifteen members; but thereafter is little changed as the membership increases. For a uniform distribution of member activity, the robustness of cover 2 continues to be eroded as the membership grows.

Finally, we show in **Chart 13** the expected shortfall over cover 2; that is, the loss that the CCP might face on average, conditional on losses exceeding the cover 2 default fund size. In this case, the results for the uniform distribution provide some reassurance; when a CCP has thirty members or more, the conditional loss over cover 2 is similar to or less than that associated with the heterogeneous distributions. So as discussed above, while cover 2 covers less of the expected loss, if it is breached, we expect similar losses to those seen in the other distributions.

Chart 13 The expected shortfall over cover 2 as a proportion of total SLOIM for three exposure distributions as the number of clearing members varies

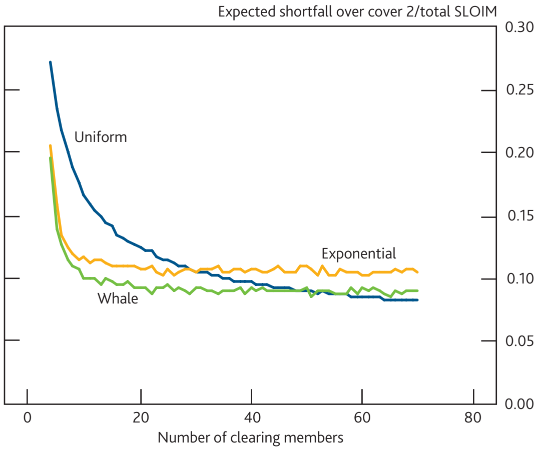

Source: Bank calculations.

5.3.2 Increasing the number of clearing members 2: when more members means more risk

We now consider the case in which increasing the number of clearing members increases the total risk to the CCP, based on the three SLOIM distributions we have considered thus far.

The question is then how new members should fit into these distributions. We choose to let new members extend the loss distribution according to its functional form. In other words, each new member has a loss that is dependent on the CCP's 'intrinsic' loss function:

- In the uniform distribution, each new member has the same SLOIM as existing members;
- For the exponential distribution however, new members are 'tacked on' to the end of the distribution and so each new member is assigned a lower SLOIM than the last; and
- For the whale distribution, we assume that the whale always accounts for 40% of exposure; and add new members on to the end of the exponential distribution that describes the remaining members' SLOIMs.

The approach chosen here clearly lays bare the potential weaknesses of the cover 2 standard when the distribution of risk among clearing members is uniform: each new member raises the total exposure in the system by the same amount; expected loss increases linearly; and the robustness of cover 2 falls, as the blue lines in **Chart 14** illustrate. For the exponential and whale distributions shown in orange and green respectively, the additional exposures associated with extending an already large membership are small enough to be irrelevant, leading to broadly the same stability in the robustness of cover 2 observed in the previous section (this of course being a result of our choice to add new members by simply extending the loss distribution).

In the case that additional members extend the loss distribution of the CCP, the expected shortfall over cover 2 also increases roughly linearly for the uniform distribution;(1) where new members can bring new risk to a CCP, and the market is structured in such a way that new members bring similar additional risk, the cover 2 standard is particularly vulnerable.

The analysis in this section is sensitive to the way in which we have chosen to extend the SLOIM distributions to account for new members — different approaches could produce qualitatively different results. This section serves to highlight the importance not only of how risk is distributed among a CCP's clearing members; but also of the way in which the addition of new clearing members affects this distribution of risk.

This is an area in which analysis is scarce. With the introduction of mandatory clearing and the expansion of CCP services in some OTC derivatives markets, it is becoming increasingly important that we have a clear view of how risks are distributed among a CCP's members, and how the robustness of the cover 2 standard will fare as cleared activity increases.

6 Policy implications and further work

We have presented two related approaches for analysing the safety of the cover 2 standard in a particular clearing service: one based on actual SLOIMs and a market consistent copula; and one based on theoretical loss distributions. Both approaches suggest that cover 2 is a prudent standard for the risk distributions likely to be found in practice, but both also suggest that it would be sensible for CCPs and their supervisors to monitor the whole of the loss distribution. Certain distributions of risk among clearing members, such as the uniform distribution studied in Section 5, give rise to situations where cover 2 is less prudent for CCPs with many clearing members. If these are found in practice, higher levels of financial resources may be needed to ensure clearing house robustness.

It may also be appropriate for this weakness in the cover 2 standard to be considered in any future revision of the international standards for systemically important CCPs. Perhaps a simple backstop to cover 2 could be considered, such as demanding that the default fund in addition meets the requirement that it is larger than some fixed percentage of the 'cover all' requirement. One basic requirement for calibrating this percentage would be knowledge of the ratio of cover 2 to cover all, so a reasonable first step would be the disclosure of both measures by all systemically important CCPs.

(1) Where default correlations are stronger, expected shortfall will increase further. Similarly, if member defaults are uncorrelated, the expected shortfall will be significantly lower (by a factor of around three for the results presented here).

Chart 14 Loss metrics as a function of number of clearing members, when more members bring more risk[(a)]

(a) Expected loss/total 30 member SLOIM

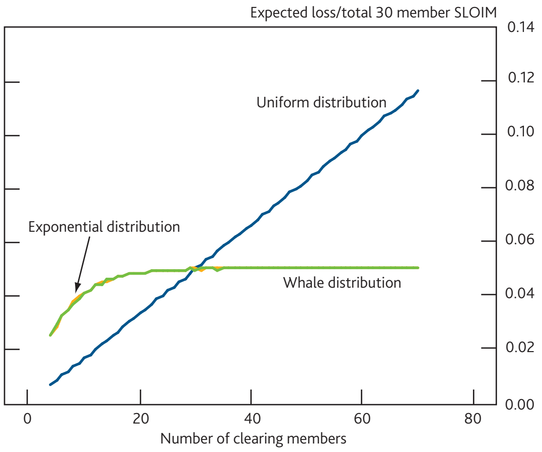

(b) Cover 2/total 30 member SLOIM

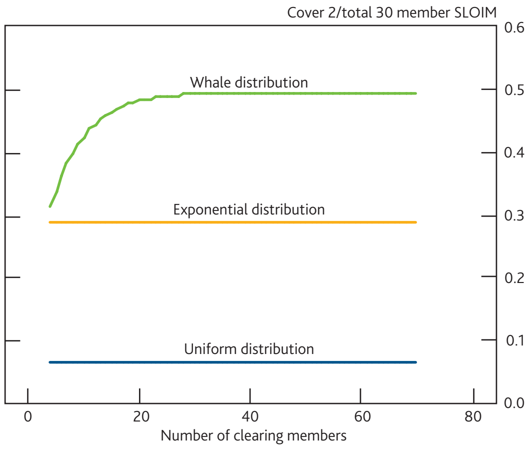

(c) Cover 2/expected loss

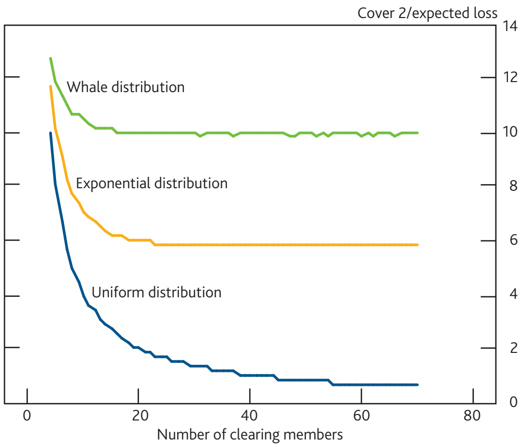

(d) Cover 2/unexpected loss

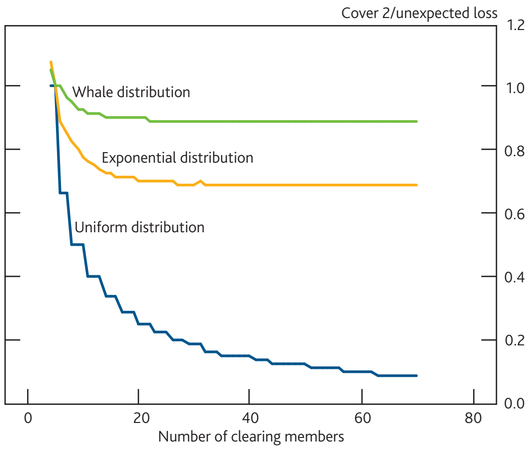

(e) Expected shortfall over cover 2/total 30 member SLOIM

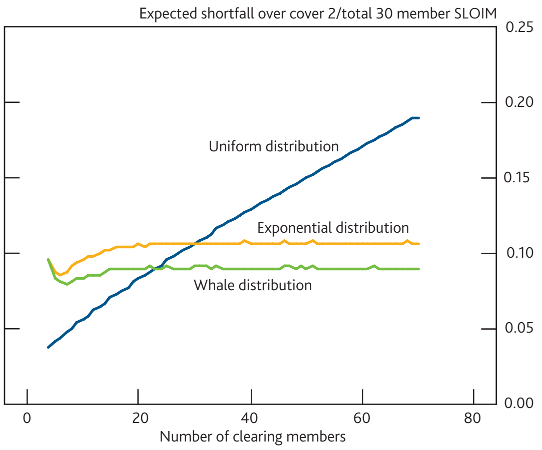

Source: Bank calculations.

(a) Cover 2 and measures of its safety for the three exposures distributions when more clearing members bring more risk to the CCP. Expected loss, cover 2 and expected shortfall in **Charts (a)**, **(b)** and **(e)** are normalised to the fixed total SLOIM used in the previous section (ie the total SLOIM when the CCP has 30 members) for comparability with previous results.

The question of how risk is distributed among clearing members is a timely one, as the central clearing landscape is changing: the imposition of mandatory clearing for some transactions and the increase in client clearing will both have an impact on the distribution of risk among clearing members and clients. It is plausible for example that client clearing may be concentrated in a small group of large banks; and that clients, at least in some markets, will tend to have directional positions leading to large potential SLOIMs for those banks.(1) The impact this would have on the robustness of cover 2 is an issue that merits further consideration.

Throughout this paper we have used SLOIM distributions based on a CCP's calculations from a single day. We have not considered the question of whether the functional forms of the SLOIM distributions are stable through time. A key question for future work is whether this assumption holds well enough to draw firm conclusions about the relative robustness of different CCPs based on the type of analysis presented herein.

The assumption that the scenarios used to determine cover 2 are extreme but plausible is key to our analysis. If the scenarios used are not sufficiently stressful, then the prudence of cover 2 may be undermined. Therefore both CCPs and their supervisors should regularly review whether the scenarios used continue to be appropriate.(2)

It is sometimes the case that a risk measurement framework offers not just the benefit of useful measures, but also requires for its implementation infrastructure and reporting that is independently helpful. This may be the case here: monitoring the CDS spreads of clearing members and constructing a plausible default copula for them; ensuring that CCP default fund scenarios are sufficiently stressful without being implausible; and reviewing changes to the CCP counterparty credit risk distribution as business changes are all potentially insightful. Indeed, without sufficiently stressful scenarios, cover 2 is a potentially imprudent standard. Supervisors may therefore wish to consider reporting requirements for CCPs which would allow them to gather data like this on a regular basis, and to consider using it to construct various measures of CCP robustness. Our work here is intended as one contribution in what we believe will be an on-going debate about the most useful risk measures for CCP users, CCPs themselves, and their supervisors.

(1) If a clearing member has a number of clients in the omnibus account with large directional positions, then it may be that the net position, and hence the SLOIM, is small. This situation would not pose significant risk to the CCP, but it does emphasise the importance of rules which ensure that the net close-out price is fair to both longs and shorts.

(2) Indeed, the requirement that 'On at least a monthly basis, a CCP should perform a comprehensive and thorough analysis of stress-testing scenarios, models, and underlying parameters and assumptions used to ensure they are appropriate' is part of the internationally agreed Principles for Financial Market Infrastructures (Committee on Payment and Settlement Systems and Technical Committee of the International Organization of Securities Commissions (2012)).

References

Amato, J and Gyntelberg, J (2005), 'CDS index tranches and the pricing of credit risk correlations', *BIS Quarterly Review*, March, pages 73–87.

Bernanke, B (2011), 'Clearinghouses, financial stability and financial reform', speech at the 2011 Financial Markets Conference, April, available at www.federalreserve.gov/newsevents/speech/bernanke20110404a.htm.

Burtschell, X, Gregory, J and Laurent, J-P (2009), 'A comparative analysis of CDO pricing models', *The Journal of Derivatives*, Vol. 16 No. 4, pages 9–37.

Commodity Futures Trading Commission (2013), 'Enhanced risk management standards for systemically important derivatives clearing organizations', Final Rule 17 CFR Part 39, August.

Committee on Payment and Settlement Systems and Technical Committee of the International Organization of Securities Commissions (2012), 'Principles for financial markets infrastructures', April, available at http://www.bis.org/cpmi/publ/d101a.pdf.

Cousin, A and Laurent, J-P (2008), 'An overview of factor models for pricing CDO tranches', in *Frontiers in Quantitative Finance: Credit Risk and Volatility Modeling*, John Wiley & Sons.

Cumming, F and Noss, J (2013), 'Assessing the adequacy of CCPs' default resources', *Bank of England Financial Stability Paper No. 26*, available at www.bankofengland.co.uk/research/Documents/fspapers/fs_paper26.pdf.

Duffie, D and Zhu, H (2011), 'Does a central clearing counterparty reduce counterparty risk?', *The Review of Asset Pricing Studies*, Vol. 1, No. 1, pages 74–95.

European Union (2012), 'Regulation (EU) No. 648/2012 of the European Parliament and of the Council, OTC Derivatives, Central Counterparties and Trade Repositories', July.

Gopikrishnan, P, Plerou, V, Gabaix, X and Stanley, H (2000), 'Statistical properties of share volume traded in financial markets', *Physical Review E*, Vol. 62, No. 4.

Huang, J-Z and Huang, M (2012), 'How much of the corporate-treasury yield spread is due to credit risk?', *The Review of Asset Pricing Studies*, Vol. 2, No. 2, pages 153–202.

Jones, R and Pérignon, C (2008), 'Default risk on derivatives exchanges: evidence from clearing-house data', paper presented at the Finance International Meeting AFFI – EUROFIDAI, December, Paris, France, available at www.en.affi.asso.fr/uploads/Externe/c1/CTR_FICHIER_335_1226348872.pdf.

Li, D (2000), 'On default correlation: a copula function approach', *Journal of Fixed Income*, Vol. 9, No. 4, pages 43–54.

Lin, L and Surti, J (2013), 'Capital Requirements for over-the-counter derivatives central counterparties', *IMF Working Paper No. 13/3*.

Murphy, D (2012), 'The systemic risks of OTC derivatives central clearing', *Journal of Risk Management in Financial Institutions*, Vol. 5, No. 3, pages 319–34.

Murphy, D (2013), 'OTC derivatives: bilateral trading and central clearing', Palgrave Macmillan.

Nahai-Williamson, P, Ota, T, Vital, M and Wetherilt, A (2013), 'Central counterparties and their financial resources — a numerical approach', *Bank of England Financial Stability Paper No. 19*, available at www.bankofengland.co.uk/research/Documents/fspapers/fs_paper19.pdf.

Noss, J (2010), 'Extracting information from structured credit markets', *Bank of England Working Paper No. 407*, available at www.bankofengland.co.uk/research/Documents/workingpapers/2010/wp407.pdf.

Schönbucher, P (2003), 'Credit Derivatives Pricing Models', Wiley.

Tucker, P (2011), 'Clearing houses as system risk managers', available at www.bankofengland.co.uk/publications/Documents/speeches/2011/speech501.pdf.

US Congress (2010), 'Dodd-Frank Wall Street Reform and Consumer Protection Act'.