Assessing the Safety of Central Counterparties*

Mark Paddrik†
H. Peyton Young‡

July 14, 2021

Abstract

We propose a general framework for empirically assessing a central counterparty's capacity to cope with severe financial stress. Using public disclosures data for global central counterparties (CCPs), we show how to estimate the probability that a CCP could cover any specified fraction of payment defaults by its members. This framework supplements conventional standards of risk management such as Cover 2 and provides a comparative and comprehensive approach to assessing risk protection across CCPs that is not predicated on a specific number of member defaults. We apply the approach to a wide range of CCPs in different geographical jurisdictions and asset classes and find that there are substantial differences in protection coverage. In particular, large European CCPs appear to be significantly safer than their counterparts in Asia-Pacific and North America. These differences are also reflected in supervisory data that provide CCP members' risk assessments of the CCPs to which they belong.

Keywords: central counterparty, default waterfall, guarantee fund, default probability
JEL Classification Numbers: G10, G23, G28, L14

*The authors thank Darrell Duffie, Stefano Giglio, David Johnson, Albert Menkveld, David Murphy, Paul Nahai-Williamson, Sriram Rajan, Robert Steigerwald and Stathis Tompaidis for their helpful comments on earlier drafts. Additionally, we would like to thank OFR's High Performance Computing, Data, and Legal teams for collecting and organizing the data necessary to make this project possible. The views expressed in the article are those of the authors and do not necessarily represent the views of the Office of Financial Research, or the U.S. Department of the Treasury.

†Office of Financial Research, U.S. Department of the Treasury, 714 14th St NW, Washington, DC 20220, United States; email: Mark.Paddrik@ofr.treasury.gov.

‡Department of Mathematics, London School of Economics, London WC2A 2AE, United Kingdom; University of Oxford, Oxford OX1 3UQ, United Kingdom; Office of Financial Research, U.S. Department of Treasury, Washington DC. 20220, United States; email: Hobart.Young@ofr.treasury.gov.

Since the financial crisis of 2008-09, central counterparties have assumed a major role in clearing over-the-counter derivatives (BCBS and IOSCO (2015)). Under central clearing, parties to a derivatives contract enter into two matched contracts with the central counterparty (CCP) that offset one another. There are many advantages to this arrangement: it facilitates greater transparency and standardization of contracts, there is greater potential for the netting of positions (Duffie and Zhu (2011)), and intermediation chains are shortened, which in principle can reduce contagion (Evanoff et al. (2006); Cont and Kokholm (2014)). A significant disadvantage is that central clearing concentrates risk in a few critical counterparties whose default could seriously disrupt financial markets.

It is therefore crucial to understand how well-protected CCPs are against sudden shocks that might cause multiple members to default on their payments or cause severe disruption to clearing operations. Although historically such events have been rare, there have been some close calls. For example, on September 10, 2018, Nasdaq Clearing nearly failed due to the inability of one member to meet margin calls on a very large position in electricity futures contracts. Due to a sudden change in circumstances, the one-day variation margin owed on these contracts greatly exceeded the member's initial margin. This forced the CCP to liquidate the member's positions and to call on its guarantee fund to pay the counterparties to these contracts.

In this paper, we introduce a novel concept for measuring the degree of stress on a CCP and show how to estimate it using publicly available data. Specifically, we define a CCP's *stress index* on a given day to be the total margin calls divided by the amount of prefunded resources available to pay them (initial margin, paid-in capital, and guarantee fund). If this ratio is greater than 1, the CCP does not have sufficient funds to meet all payment obligations from pre-funded resources in case none of the margin calls are met. We call such an event a *guarantee fund* (GF) *breach*. The concept of the stress index is therefore similar in spirit to that of liquidity coverage ratio, which compares a firm's liquid resources to its payment obligations over a given period of time, assuming that no incoming payments are realized. We propose a framework for empirically estimating the probability of a GF breach for CCPs operating in different geographical regions and specializing in different asset classes.

In the last section of the paper, we consider an alternative measure of CCP safety, namely the probability of default as estimated by the CCP's own members. These estimates are provided quarterly to the Federal Reserve as part of the annual Comprehensive Capital Analysis and Review (CCAR). Although this information is confidential, we are able to provide the members' estimated probabilities in aggregate form (by subgroups of CCPs). Taken together, the two approaches provide alternative ways of assessing the relative safety of CCPs in different jurisdictions and specializing in different markets. It turns out that, under both measures the largest European CCPs, as a group, appear to be less exposed to default risk than their counterparts in North America and Asia-Pacific.

1 CCP Default Waterfall Structure

We begin by briefly recalling the layers of protection that CCPs employ to guard against defaults by their members.[1] These multiple layers constitute the CCP's *default waterfall*. Although the detailed rules governing default waterfalls vary somewhat, their overall structures are similar and follow standard industry guidelines (ISDA (2013), ISDA (2015)). The stages of a typical default waterfall are depicted in Figure 1.

Most stages of the default waterfall are funded in advance by contributions from members and are available independently of any shocks that may subsequently arise. However, the final stage – members assessments – is called upon only if the prior pre-funded stages prove to be inadequate to meet members' defaults. The usefulness of this power is somewhat problematic, because a crisis could threaten the CCP's solvency in a very short time frame (e.g., one day), so that assessments could be difficult to carry out in a timely manner; moreover, the assessment power is useless against members who are already in default. For this reason we shall omit the assessment stage from our analysis and assume that the guarantee fund is the last line of defense.[2]

The first layer of protection against member default is the initial margin (IM) that the members post with the CCP. These funds are held in a segregated escrow account for each member and its clients, and can only be applied to the payment shortfalls of the account holder. The standards and factors for setting IM vary somewhat among CCPs, though they typically employ at minimum a Value-at-Risk (VaR) measure, such as 99.5 percent, in conjunction with a specified margin period of risk – usually two or five days. The distribution of variation margins (VM) payments used to

[1]See Cox and Steigerwald (2017) for an informative overview of CCP risk management structures.

[2]Other papers on CCP risk assessment make a similar assumption; see for example Murphy and Nahai-Williamson (2014).

Figure 1: Stages of CCP Default Waterfall

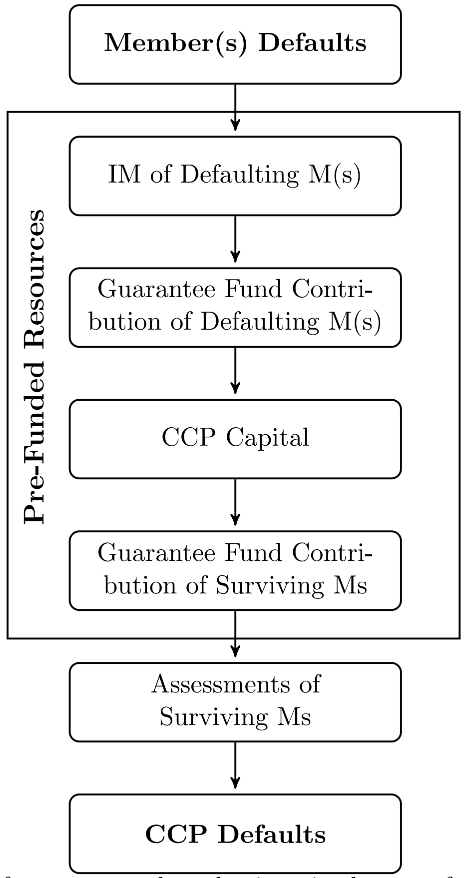

Note: The chart depicts the series of resources and mechanisms in the waterfall which will be accessed if previous ones are insufficient to cover total default losses in the event of a clearing member (M). The solid arrows depict the most common set of waterfall resource contingencies. A defaulting clearing member's obligation is first covered by their initial margin (IM). If the member's IM is insufficient to cover its obligations, the resources of the following stages will be used.
Source: Authors' creation.

determine the VaR is estimated from a look-back window of the member's previous VM payments over several prior years, adjusted for the current size and composition of the member's portfolio. Our assessment of the empirical frequency of IM breaches over the period July 2015 - March 2020, which includes a very significant tail event (the COVID 19-induced crisis in March 2020) suggests that the majority of CCPs do meet a standard in excess of 99.5 percent. As we shall see in Section 2, the frequency of IM breaches increased markedly during the COVID 19-induced crisis in March 2020, indicating that there is a substantial correlation between members' IM breach probabilities. Other research has documented the same phenomenon due to crowded positions in members' portfolios (Jones and Pérignon (2013); Menkveld (2015), (2017); Menkveld and Vuillemey (2020)).

The central question that we address in this paper is how much the initial margin plus paid-in capital and guarantee fund offer as a buffer against potential member defaults, when the probability of multiple member defaults is not known. In other words, how likely is it that aggregate IM breaches are larger than the CCP's paid-in capital and guarantee fund put together? We call this the *GF breach probability*, and its complement is the *comprehensive coverage probability*.

This standard is different, and potentially more stringent, than the so-called 'Cover 2 rule.' Under Cover 2, the CCP's guarantee fund is supposed to be capable of absorbing the simultaneous default of the two largest exceedances under 'extreme but plausible market conditions.'[3] There are several drawbacks to this approach. First, there is no good reason why we should consider the simultaneous default of just two members. Why not three, four, or five? Indeed, the more severe the financial shock, and the more members the CCP has, the more likely it is that a given number of members will default on their payments simultaneously.[4] Second, the Cover 2 rule is typically applied under 'extreme but plausible market conditions' that rely on a CCP's internal model of stress, rather than on externally verifiable assumptions. What is missing is the equivalent of a VaR standard that allows an empirical assessment of potential risk across different CCPs.

As we shall see, there is now sufficiently detailed time series data on CCPs to be able to estimate the probability distribution of margin calls and the likelihood of GF breaches empirically. This allows one to estimate the probability of covering any given percentage of payment defaults, not just the default of two members.

2 Estimating the Likelihood of Initial Margin Breaches

In this section and the next, we formally define the concepts of *IM breach* and *GF breach*, then show how to estimate them given the limitations of the publicly available data. Fix some CCP and let i designate the account of a member firm or a client that is sponsored by a member. Let IM_{it} be the initial margin posted by member i at the end of trading day t. Let the random variable VM_{it} represent the net variation margin that will be owed by i to the CCP on day t; if the CCP owes member i on day t we let $\mathrm{VM}_{it} = 0$. The random variable [IM_{it} - IM_{it-1}] is the top-up in IM

[3] Principle 4 of the Bank for International Settlements and International Organization of Securities Commissions (2012) Principles for Financial Market Infrastructures address this definition, though it leaves room for interpretation on what is to be included under an extreme but plausible stress scenario for the assessment of pre-funded resources.

[4] Default correlation can arise through common exogenous shocks as well as network spillover effects, which can increase the probability and severity of multiple defaults by members (Ghamami et al. (2020); Paddrik et al. (2020); Paddrik and Young (2021)).

demanded by the CCP on day t for account i. Thus the total margin call (MC)[5] payment of i on day t is

$$\text{MC}_{it} = \text{VM}_{it} + [\text{IM}_{it} - \text{IM}_{it-1}]. \tag{1}$$

We say that account i incurs an *IM breach* on day t if $\text{MC}_{it} > \text{IM}_{it-1}$. If the margin call is not met, the member will be declared in default and its positions liquidated, which will typically take several days. The CCP uses the member's IM on deposit to cover the unpaid VM plus the cost of sale, which will be higher the more volatile are current market conditions. We shall view the margin top-up as a provision reflecting these increased costs. Thus, if there is an initial margin breach the total expected loss incurred by the CCP will be the amount of the breach, that is, MC_{it} - IM_{it-1}.

The publicly available data on global CCPs come from CPSS-IOSCO Principles for Financial Market Infrastructure's Public Quantitative Disclosures (PQD). The data does not provide the amounts VM_{it} and IM_{it} for individual members on particular days, but it does report the number of times over the approximately 63 trading days in a given quarter that member accounts incur IM breaches. Averaged over the number of member accounts and the number of trading days we can therefore estimate the daily probability with which a typical account suffered an IM breach.

Suppose that the daily probability of not suffering a breach is p (which can be interpreted as a one-day VaR). Then the probability that the average account incurs at least one breach per quarter is $1-p^{63}$. Table 1 shows the quarterly IM breach probabilities, and corresponding daily VaR levels, for the whole sample of CCPs as well as for three geographical subgroups: Asia-Pacific, Europe, and North America. Note that all groups exhibit an average daily VaR in excess of 99.5 percent.[6]

Table 2 compares IM breach probabilities over the 18 periods 2015 Q3-2019 Q4 with the breach probabilities in 2020 Q1, which was an unusually stressed quarter. Note that for European and North American CCPs the probability of IM breaches increased more than three-fold in this quarter as compared to the previous period. For example, about 8.9 percent of the North American accounts

[5] Coupon payments on a specified settlement day are included indirectly because they are reflected in changes in VM from the prior day. We do not include changes in the contributions to the guarantee fund, as they are not reported our data.

[6] Similarly, Capponi et al. (2020) analyze individual portfolio level data for ICE Clear Credit and find that initial margin breaches are rare.

Table 1: Quarterly IM Breach Probabilities of CCPs

	All	Asia-Pacific	Europe	North America
Quarterly Account IM Breach Probability 2015 Q3 - 2020 Q1 (%)	12.46	12.66	12.89	10.15
Corresponding Daily VaR (%)	99.79	99.78	99.78	99.83
CCP Sample	77	26	41	10

Source: CCPView Clarus Financial Technology; authors' analysis.

suffered at least one IM breach per quarter during the period 2015 Q3 through 2019 Q4, whereas 27 percent of the accounts suffered a breach in the March 2020 quarter. This highlights the fact that breaches tend to be correlated due to exogenous financial shocks that hit many member accounts simultaneously.

Table 2: Quarterly IM Breach Probabilities for the March 2020 vs. Previous Quarters' Averages

	All	Asia-Pacific	Europe	North America
Quarterly Account IM Breach Probability 2015 Q3 - 2019 Q4 (%)	8.23	12.22	5.55	8.88
Quarterly Account IM Breach Probability 2020 Q1 (%)	20.25	15.81	21.79	27.07
CCP Sample	77	26	41	10

Source: CCPView Clarus Financial Technology; authors' analysis.

3 Estimating the Likelihood of Guarantee Fund Breaches

Fix a CCP and let GF_t be the amount in the guarantee fund plus the paid-in CCP capital on day t. The *stress index* on day t is

$$\mathrm{S}_t = \frac{\sum \mathrm{MC}_{it}}{\mathrm{GF}_{t-1} + \sum_i \min(\mathrm{MC}_{it}, \mathrm{IM}_{it-1})} \tag{2}$$

where i ranges over the members of the CCP. The CCP incurs a *GF breach* on day t if the sum of the IM breaches exceeds the guarantee fund, that is,

$$\sum_i [\mathrm{MC}_{it} - \mathrm{IM}_{it-1}]^+ > \mathrm{GF}_{t-1}. \tag{3}$$

As noted earlier, a breach is not the same as a default. The CCP is under severe stress if there is a GF breach, but it does not actually default unless a sufficient number of members default on their payments when the GF breach occurs.[7] The left-hand side of Equation (3) is an estimate of how much the CCP would have to make up (in payments owed to members) if all the positions were liquidated and the members' IM were seized. If this exceeds the remaining amount of funded resources, then we say that there is a GF breach.

The *stress index* is a measure of the liquid pre-funded reserves available to pay members in the event of a default. In this sense it is similar to the concept of the liquidity coverage ratio. The greater the index value, the more stress is placed on the guarantee fund to meet potential payment obligations. In particular, if the index exceeds one, the GF will be insufficient to meet the obligations to the members in full (in case of complete default), and the amounts owed to members will be reduced pro rata (a procedure known as Variation Margin Gains Haircutting).

The probability of a GF breach is given by the expression

$$\beta_t = \mathrm{P}\left[\sum_i [\mathrm{MC}_{it} - \mathrm{IM}_{it-1}]^+ > \mathrm{GF}_{t-1}\right]. \tag{4}$$

The PQD data do not provide information about the amount of MC owed by each member or the amount by which MC exceeds the member's IM on any given day. Thus we cannot estimate the probability of a GF breach directly as given by expression (4). Nevertheless, we can exploit the fact that the data include the *aggregate* daily VM owed by members and their clients (averaged over each trading day in a quarter), as well as the *maximum aggregate* VM owed by members and clients on some day during the quarter. By *aggregate* VM owed we mean the total VM summed over all firms who owe money on net to the CCP on a given day.[8] Let

$$\mathrm{VM}_t = \sum_i \mathrm{VM}_{it}, \tag{5}$$

$$\mathrm{VM}^{\max} = \max_t \mathrm{VM}_t. \tag{6}$$

The PQD data also report the total amount of IM posted on an average day during the quarter:

[7] In times of severe stress the CCP may also experience operational difficulties that prevent payments from being made to members in a timely fashion.

[8] Due to the CCP's matched book, the net amount owed to the CCP summed over all firms is zero. The vulnerability of the CCP to delinquent payments is the sum over all firms that owe the CCP on net.

$$\text{IM}^{\text{avg}} = \frac{\sum_i \sum_t \text{IM}_{it}}{T}, \tag{7}$$

where T is the number of trading days in the quarter. In addition, the data give the maximum amount of IM top-up, IMT^{max}, that members must post on any given day during the quarter:

$$\text{IMT}^{\text{max}} = \max_t \left[\sum_i [\text{IM}_{it} - \text{IM}_{it-1}] \right]^+ . \tag{8}$$

Note that IM top-ups are assessed on both the buy and sell side of any contract, so on average half of the top-up will be assessed against those members that owe VM to the CCP. We shall therefore assume that the maximum payment owed to the CCP on any given day during the quarter is

$$\text{MC}^{\text{max}} = \text{VM}^{\text{max}} + \text{IMT}^{\text{max}}/2.^9 \tag{9}$$

We now describe our framework for estimating GF breach probabilities given the data limitations. The true GF breach probability on a given day t is defined as in (3). Note that we can bound this value from below as follows:

$$\beta_t^- = \text{P}\left(\sum_{i \in \text{S}_t} \text{MC}_{it} - \sum_{i \in \text{S}_t} \text{IM}_{it-1} > \text{GF}_{t-1} \right) \text{where } \text{S}_t = \{i \in \text{N} : \text{MC}_{it} > 0\}. \tag{10}$$

Here $\sum_{i \in S_t} \text{IM}_{it-1}$ is the amount of initial margin held in accounts that owe the CCP on day t. To the extent that MC_{it} is positive but less than $\text{IM}_{i,t-1}, \beta_t^-$ underestimates the true breach probability β_t.

Each realization of $\text{MC}_t = \sum_{i \in \text{S}_t} \text{MC}_{it}$ represents a one-day change in the value of the positions in S_t plus calls for additional IM. A standard assumption is that one-day gains and losses in the value of each set of positions are more or less symmetrically distributed about zero. Note that this assumption is consistent with correlated as well as independently distributed changes in value. Therefore, due to the CCP's matched book, the set N - S_t owes the CCP an amount VM_t with

[9] This is an approximation because we do not know whether the maximum top-up occurred on the same day as the maximum VM payment. Given CCP risk practices, however, we do know that intra-day top-ups are demanded on days with unusually high volatility, so this approximation is probably close to being correct.

about the same probability that S_t owes the same amount VM_t to the CCP. It follows that the *expected* amount of initial margin associated with the set of accounts that owe the CCP on any given day is approximately one-half of the total margin.[10] This leads to the estimate

$$\beta_t^- \approx \mathrm{P}\left(\mathrm{MC}_t > \mathrm{IM}_{t-1}/2 + \mathrm{GF}_{t-1}\right).^{11} \tag{11}$$

Given the data limitations, we do not know the values of the variables in Equation (11) on particular days t. Nevertheless, we can make a rough estimate of the GF breach probability as follows. Fix some CCP and a particular quarter (which we omit from the notation). Let $\mathrm{IM}^{\mathrm{avg}}$ be the average initial margin held by the CCP over the *previous* quarter.[12] Similarly, let $\mathrm{GF}^{\mathrm{avg}}$ be the average size of the guarantee fund during the previous quarter and let $\mathrm{MC}^{\mathrm{max}}$ be the maximum MC owed on some day during the quarter. We then obtain the following estimated lower bound on the GF breach probability

$$\beta^* = \mathrm{P}\left(\mathrm{MC}^{\mathrm{max}} > \mathrm{IM}^{\mathrm{avg}}/2 + \mathrm{GF}^{\mathrm{avg}}\right). \tag{12}$$

Similarly we can estimate a lower bound on the maximum stress index during the quarter by the expression

$$\frac{\mathrm{MC}^{\mathrm{max}}}{\mathrm{IM}^{\mathrm{avg}}/2 + \mathrm{GF}^{\mathrm{avg}}}. \tag{13}$$

All of the variables on the right-hand side of Equations (12) and (13) are reported for each CCP in the public data set. An alternative way of writing expression (12) is to divide through by $\mathrm{IM}^{\mathrm{avg}}$ and let $r = \mathrm{GF}^{\mathrm{avg}}/\mathrm{IM}^{\mathrm{avg}}$. We then obtain the expression

$$\beta^* = \mathrm{P}\left(\frac{\mathrm{MC}^{\mathrm{max}}}{\mathrm{IM}^{\mathrm{avg}}} > 1/2 + r\right). \tag{14}$$

[10] There are situations where the IM posted is asymmetric with respect to the Buy and Sell positions, degree of concentration, and degree of liquidity. It is not clear, however, that these asymmetries lead to more than half the IM being held against the aggregate VM due on any given day.

[11] In fact this expression underestimates β_t^-, because the tail of the Pareto distribution is convex, so by Jensen's inequality the expected tail probability is greater than the tail probability of the expectation.

[12] Some of the initial margin may be held in non-cash assets, but the initial margin reported in the data is its estimated cash value.

Equation (14) implies that each dollar of GF is twice as efficient as each dollar of IM in protecting the CCP against default, that is, the *mutualization factor* is at least two (Murphy (2017)). In fact expression (14) underestimates the comparative advantage of GF because we approximated the GF breach probability in expression (10) by ignoring the extent to which MC_{it} may be less than IM_{it-1} for some of the members who owe MC to the CCP on day t. If we let σ be the expected proportion of members' IM that is *dormant* (i.e., in excess of their margin calls) on any given day, then $\sigma \geq 0.5$, because on average half the members do not owe any VM to the CCP and their IM top-ups (if any) will typically be less than the amount of IM they already posted. For example, suppose that on a given day half the accounts owe money to the CCP, and of those, half owe more than they posted in IM, while the other half owe very little. Then three-quarters of the IM is dormant, that is $\sigma = 0.75$. Thus, we can bound the GF breach probability by the expression

$$\mathrm{P}(\mathrm{MC}^{\max}/\mathrm{IM}^{\mathrm{avg}} > 1 - \sigma + r), \text{ where } \sigma \geq 0.5, \tag{15}$$

and the corresponding mutualization factor is $1/(1-\sigma)$.

The PQD data do not provide enough information to estimate σ. In what follows we shall make the conservative assumption that $\sigma = 0.5$, and estimate the right-hand side of Equation (12) by fitting a distribution to the quarterly realizations of the random variable

$$\mathrm{X} = \frac{\mathrm{MC}^{\max}}{\mathrm{IM}^{\mathrm{avg}}/2 + \mathrm{GF}^{\mathrm{avg}}}. \tag{16}$$

X is a conservative estimate of the maximum level of stress experienced on some day during a given quarter. If the data allowed it, we would estimate expression (16) for each CCP separately. We only have 19 quarters of data, however, and some CCPs do not report in all quarters. Thus, for a given CCP, we do not have enough data to estimate the distribution of X (let alone the tail of the distribution) with much accuracy. Instead, we shall pool the data and treat the realizations of X as if they came from a single CCP. Later we shall refine the analysis by considering subgroups of similar CCPs defined by size and geographical location. For the pooled data we find a very good fit is achieved with a Frechet distribution,

$$\mathrm{F}(x) = \exp - \frac{s}{(x-m)^{\alpha}}, \tag{17}$$

where m is the location parameter, $s^{1/\alpha}$ is the scale parameter, and α is the shape parameter (see Figure 2).

Figure 2: Frechet Distribution fitted to Stress Index Realizations X

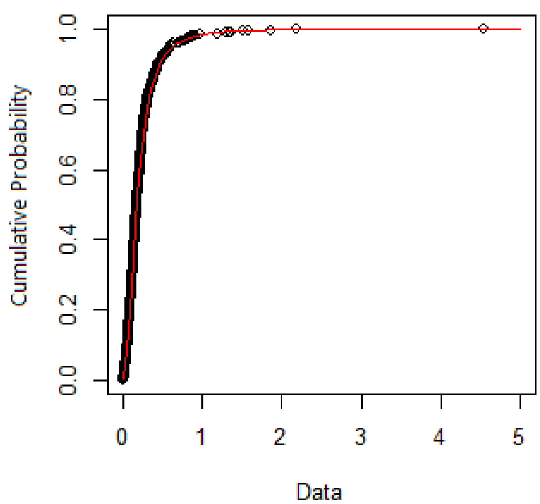

(a) Fitted CDF

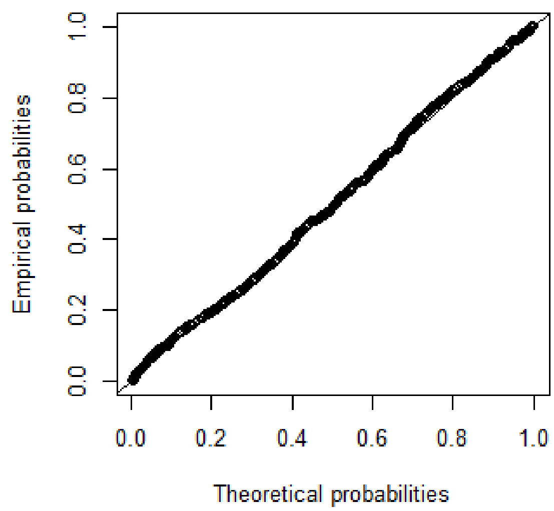

(b) P-P plot

Note: Frechet distribution fitted to the empirical distribution of X over all CCPs that provide data for at least 10 out of the 19 quarters 2015 Q3 - 2020 Q1. Here s = 0.3119, α= 3.0873, and m = -0.1837.
Source: CCPView Clarus Financial Technology; authors' analysis.

We claim that the right tail of a Frechet distribution is very close to the tail of a Pareto distribution. Indeed, suppose that x is sufficiently large that the Frechet tail probability p is small, say

$$p = 1 - e^{-s/(x-m)^{\alpha}} < 0.1. \tag{18}$$

Note that $\ln(1-p)$ is very close to $-p$ when $0 < p < 0.1$. It follows that $p \approx s/(x-m)^{\alpha}$, that is, the tail of the Frechet is approximately Pareto distributed with the same parameters.

In what follows we shall fit a Pareto distribution to the tail of the empirical distribution (instead of fitting a Frechet to the whole distribution), since it is the far right tail that is relevant for our analysis of GF breach probabilities. Specifically, we shall fit a Pareto distribution $\mathrm{P}(\mathrm{X} > x) =$

$1 - s/x^{\alpha}$ to the observed empirical distribution of X.[13] The shape parameter α reflects the variance in MC payments relative to funded resources IM/2 + GF. For a given value of s, larger values of α lead to lower variance and hence to lower breach probabilities.[14]

Let $\mathrm{F}(x)$ be the empirical cumulative distribution function of X over all quarters in which the relevant variables are reported. In our sample we omit any CCPs that report less than ten quarters. We then estimate the parameters of the Pareto distribution from the linear regression

$$\ln(1 - \mathrm{F}(x)) = -\alpha \ln x + b + e, \tag{19}$$

where $b = \ln s$ and e is the error term which is assumed to be zero in expectation. The resulting estimate of $\mathrm{F}(x)$ is

$$\mathrm{F}(x) = 1 - \frac{s}{x^{\alpha}}, \tag{20}$$

and our estimated lower bound on the GF breach probability is $\beta^* = 1 - \mathrm{F}(1) = s$.

4 Breach Probabilities by Geographic Region and Size

We now apply this estimation methodology to CCPs that are grouped by size and geographical location. We include only those CCPs that report at least ten quarters of data over the period 2015 Q3 through 2020 Q1. Within each geographical group, we consider the ten largest CCPs as measured by total IM held. We omit the smaller CCPs from the analysis because their reported numbers tend to be quite erratic, and are not as important from a systemic risk perspective.[15]

Table 3 gives the empirical and model-estimated breach probabilities for the various subgroups. Note that the regression model defined in Equation (19) yields a very high R-squared for each subgroup. Figure 3 shows the model-estimated probability distribution of the stress index for each of the three regions. For each value of the index x in the range 0.4 - 3.0 the graph shows the

[13] If we fit a distribution of form $\mathrm{P}(\mathrm{X} > x) = s/(x - m)^{\alpha}$ with a location parameter $m > 0$ we find that the resulting fit is not significantly better than if $m = 0$.

[14] Andersen and Dickinson (2018) also estimate margin payments are Pareto distributed following from the empirical work done in Arnsdorf (2012).

[15] The Asia-Pacific group includes Japan, Singapore, Hong Kong, Australia and Russia. The group of North American CCPs does not include the largest one, CME, for which much of the data is missing in the public reports. If the CME is more conservatively managed than the average North American CCP, the estimated probabilities of GF breach in Tables 3 and 4 may be on the high side for this group.

probability that x is exceeded in a randomly drawn quarter. Note that the tails are very elongated, especially for the North American and Asia-Pacific CCPs. This suggests that extreme levels of stress can occur with non-negligible probability. Note also that the stress distribution for European CCPs is markedly lower than for the other two groups. This suggests that European CCPs may, on average, be safer than their counterparts in the other regions. This finding is consistent with another measure of CCP safety that we shall discuss in Section 6.

Table 3: Estimated GF Breach Probabilities Per CCP in Each Region

	All	Asia-Pacific	Europe	North America
Empirical Quarterly Frequency (%)	0.61	0.63	0.00	1.71
Estimated Quarterly Frequency (%)	0.81	0.92	0.12	1.71
Model Parameters	$\frac{0.0081}{x^{2.4900}}$	$\frac{0.0092}{x^{2.4536}}$	$\frac{0.0012}{x^{3.9676}}$	$\frac{0.0171}{x^{1.9527}}$
R^2	0.9930	0.9367	0.9744	0.9833
Quarters Sample	455	164	174	117
CCP Sample	27	10	10	7

Source: CCPView Clarus Financial Technology; authors' analysis.

Figure 3: Modelled Probability Distribution of the Stress Index by Region

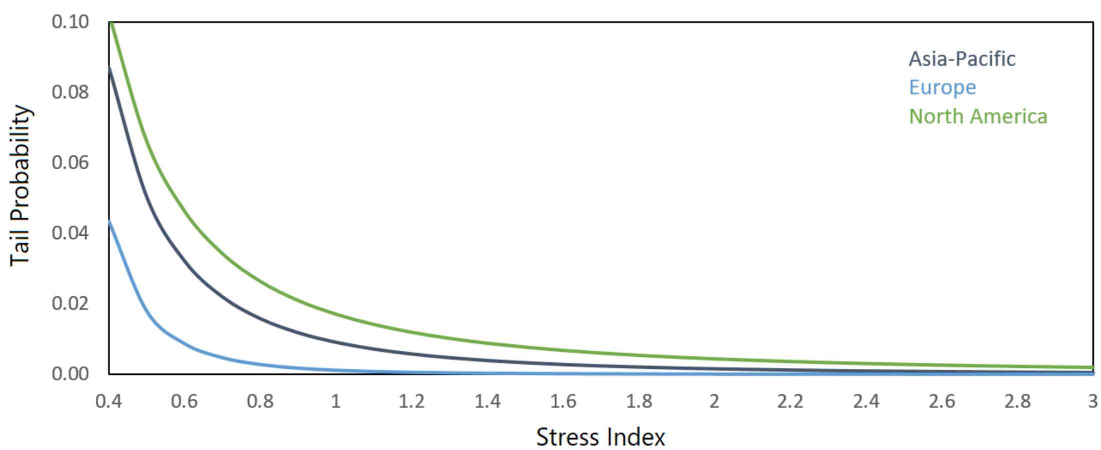

Source: CCPView Clarus Financial Technology; authors' analysis.

Not surprisingly, GF breaches tend to occur in bunches during periods of severe financial stress. This phenomenon can be illustrated by comparing the estimated breach frequencies in the relatively normal period 2015-2019 with the first quarter of 2020. Table 4 shows that GF breach frequencies increased dramatically in the first quarter of 2020. To be specific, our estimates suggest that two out of the seven largest CCPs in North America (28.57 percent) suffered a GF breach on at least one day during the March 2020 quarter. During the preceding 18 quarters, by contrast, there were no GF breaches. These null results are not inconsistent with the model estimates. Consider, for example, the North American model estimated from the full sample of 7 CCPs and 19 quarters. Assuming independent realizations, the probability of no GF breaches in the first 18 quarters is $(1-0.0171)^{7*18} = 0.114$. The corresponding probabilities for Asia-Pacific and Europe are 0.189 and 0.806 respectively. The probabilities of no GF breach will be higher if the realizations are positively correlated. Thus the null outcomes in the quarters prior to 2020 Q1 are not inconsistent with the model-estimated probabilities.

Table 4: Estimated GF Breaches in 2020 Q1 vs. Prior Quarters

	All	Asia-Pacific	Europe	North America
Estimated Number of GF Breaches				
2015 Q3 - 2019 Q4:	0	0	0	0
2020 Q1:	3	1	0	2
Estimated Quarterly Frequency Per CCP (%)				
Model w/o 2020 Q1:	0.24	0.36	0.08	0.39
Model w/ 2020 Q1:	0.81	0.92	0.12	1.71
Quarters Sample	455	164	174	117
CCP Sample	27	10	10	7

Note: Empirical and model-estimated quarterly GF breach frequencies in the first quarter of 2020 compared to the average quarterly GF breach frequencies over the 18 prior quarters of top 10 CCPs by region.
Source: CCPView Clarus Financial Technology; authors' analysis.

5 Comprehensive Versus Partial Coverage

The GF breach probability measures the potential vulnerability of a CCP to default, not the probability of default per se. Unlike Cover 2, for example, it is not predicated on the simultaneous default of a particular number of members; it holds for the default of any subset of members. Clearly this concept is more demanding than Cover 2: for a given level of initial margin, there is a higher probability that a given quantity of guarantee fund will be sufficient to cover the default of just two members as opposed to default by all of the members. In this section, we show how the concept can be modified to accommodate any desired level of *partial* coverage and we show how demanding different levels of coverage are. We also show that there is a simple way to evaluate the tradeoff between the level of coverage and the amount in the guarantee fund.

First, let us estimate how much guarantee fund a CCP needs to meet a specified probability of protection against GF breaches. If we fix the ratio $r = \text{GF}/\text{IM}$, we can approximate the probability of a GF breach by a function of form

$$\beta^*(r) = 1 - \frac{s}{(1+2r)^\alpha}, \tag{21}$$

for some constant s and scale factor α. Given some desired level of protection β^* against GF breaches, such as $\beta^* = 0.10$, 0.05 or 0.01, we can solve for the required ratio r as a function of β^*:

$$r(\beta^*) = \frac{(s/\beta^*)^{\frac{1}{\alpha}} - 1}{2} \tag{22}$$

For North American CCPs the current annual level of β^* is approximately 0.065, the current ratio of GF to IM is about 0.37, and $\alpha = 1.95$. The corresponding scale factor is $s = 0.191$.[16] This results in the function

$$r(\beta^*) = 0.5 \left[\frac{0.191}{\beta^*}\right]^{0.513} - 0.5. \tag{23}$$

We can view expression (23) as a rough measure of the risk-reward trade-off in sizing the GF. For a given level of IM, higher r means that more GF must be held, which is costly; the benefit is the decreased probability β^* that a GF breach will occur.

Another way of thinking about the trade-off is to estimate the amount of GF that would be needed to achieve a given amount of partial coverage with a specified probability, similar to a VaR. Suppose, for example, that we want to know how likely it is that some proportion $\lambda \in [0, 1]$ of the MC payments would be covered by the GF in case some of the members default. The *partial coverage breach probability* can be bounded by the following variant of Equation (14)

$$\beta^*(\lambda, r) = \text{P}\left(\text{MC}/\text{IM} > 0.5 + \frac{r}{\lambda}\right) = \frac{s}{(0.5 + r/\lambda)^\alpha}. \tag{24}$$

Thus we see that there is a linear relationship between r and λ: if we wish to cover only half of the payment defaults with some probability β^*, then we need only half as much guarantee fund. The trade-off between more coverage (λ) and more risk (β^*) is illustrated in Figure 4 for $r = 0.3$. In the

[16] This differs from the scale factor in Table 3 because we have annualized the breach probability and have assumed a specific value of r for purposes of illustration.

example figure, the CCP can cover a default of half the MC ($\lambda = 0.5$) with probability $1-\beta^* = 95.9$ percent, whereas it can cover the default of the entire amount ($\lambda = 1$) with probability $1-\beta^* = 92.4$ percent.

Figure 4: Trade-off between Annualized Breach (β^*) Probability and Coverage ($\alpha = 1.95$, $s = 0.191$, $r = 0.3$)

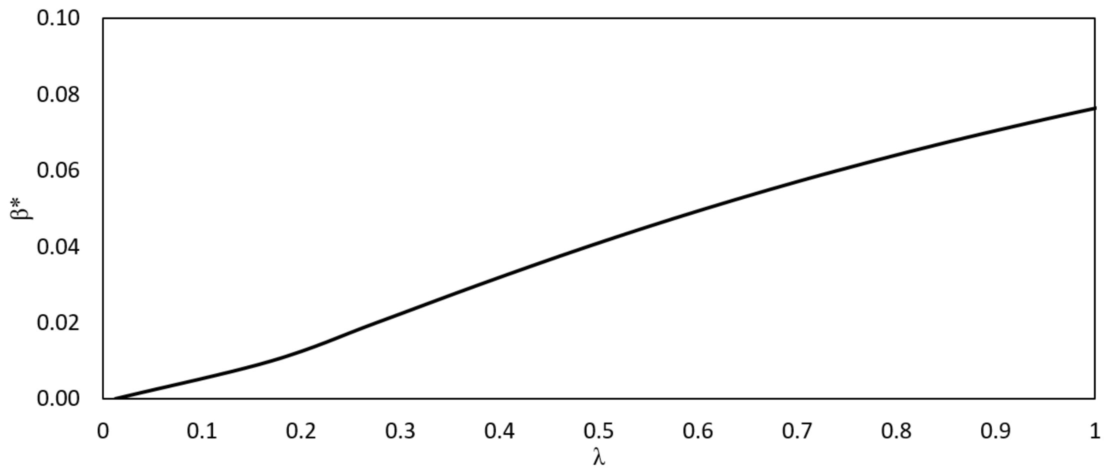

Source: CCPView Clarus Financial Technology; authors' analysis.

The concept of partial coverage is similar in spirit to that of Cover 2, but it is more general. With sufficiently granular data we could compute the equivalent amount of partial coverage that is afforded under the Cover 2 standard for individual CCPs, but this is not possible with the public data. However, for each CCP the data does report the amount of IM posted by the top five members (and their clients) as a proportion ϕ of the total IM in each quarter. These proportions range from 10 to 80 percent with a median of 50 percent. Suppose that the (random) amount of MC owed by a member on a given day is proportional to its size as measured by its IM. Suppose also that the probability is fifty-fifty that any given member owes the CCP or is owed by the CCP on a given day. Then the proportion of MC owed by the two members with the largest exceedances on a given day can be approximated by $\lambda = 0.8\phi$.[17] Thus for the median CCP the equivalent amount of partial coverage afforded by Cover 2 is approximately $\lambda = 0.4$. It follows from the preceding discussion that comprehensive coverage ($\lambda = 1$) requires about $1/0.4 = 2.5$ times as much GF per dollar of IM to achieve the same level of protection.

Another way of studying the trade-off between full and partial coverage is to compare the level

[17] On average, the top two exceedances will be realized by two of the top five members, who represent $(2/5)(\phi/2) = 0.8\phi$ of the total IM posted by members who owe the CCP.

of protection in the two cases. Let P(λ, α, r) denote the probability of protecting against a loss of λ for any given parameters α and r. From Equation (24) it follows that for some scalar s,

$$\mathrm{P}(\lambda, \alpha, r) = 1 - \frac{s}{(0.5 + r/\lambda)^{\alpha}}. \tag{25}$$

Hence the level of protection against a total loss (relative to the level of protection against a partial loss) can be expressed as follows:

$$\mathrm{P}(1, \alpha, r) = 1 - (1 - \mathrm{P}(\lambda, \alpha, r)) \left[\frac{0.5 + r/\lambda}{0.5 + r} \right]^{\alpha} \tag{26}$$

This relationship is illustrated in Table 5, which shows how the VaR for full coverage varies with α and r when the VaR for 40 percent protection is set at 99 percent. Note that full coverage is achieved at quite a high level, even though it is a more demanding standard. This has to do with the heavy-tailed nature of the MC distribution, together with the fact that most CCPs are fairly highly concentrated in a few large members.

Table 5: Comprehensive coverage VaR when 40% coverage VaR is 99%.

		α		
		2	3	4
	0.3	97.5	96.2	94.0
r	0.4	97.2	95.4	92.3
	0.5	96.9	94.6	90.6

Source: CCPView Clarus Financial Technology; authors' analysis.

6 CCP Guarantee Fund Breach versus Default

The preceding sections have laid out a framework for assessing the ability of a CCP to withstand payment defaults by some or all of its members. This approach provides a useful tool for comparing the adequacy of CCP waterfall resources across a wide range of markets and geographical jurisdictions. Although high levels of stress may be associated with elevated probabilities of default, the approach we have described does not provide a measure of default probability per se. Indeed, a CCP can default for a variety of reasons including i) a breakdown in clearing operations; ii) cyberattacks, and iii) inability to pay due to defaults by members that owe the CCP (World

Federation of Exchanges (2020)). Our measure of stress is related only to the last of these three factors.

In this section we consider an alternative measure of CCP safety that is based on estimates of CCP default by the members themselves. The Federal Reserve collects quarterly counterparty risk estimates from each U.S. globally systemic important bank (GSIB), as part of the annual Comprehensive Capital Analysis and Review. Although we do not know how each GSIB conducts its estimates, we do know that they exhibit a significant amount of agreement for any given CCP, and they vary from quarter to quarter depending on financial market conditions. Thus, it appears that the members' default estimates are not ad hoc, but are based on some form of risk modelling. These estimates are provided by GSIBs for a wide range of CCPs in different markets and jurisdictions, so they are very useful as a comparative metric.

Figure 5 shows the average estimate of CCP default probability in each quarter for the top ten CCPs in each of the three geographical groups. These probabilities are obtained as follows. In a given quarter each GSIB member of a given CCP estimates a 5-year CDS spread on that CCP. We convert this to an annual default probability of the CCP and take the average over the CCP's members. The default probability for each group in a given quarter is the average over the CCPs in that group. Note that the estimated default probabilities are quite stable over the period 2016-2019, but jump up sharply during the first quarter of 2020. This suggests that the members' estimates do take into account current market conditions. Moreover, the members' estimates are reasonably consistent with one another: the ratio of the standard deviation to the mean estimate is less than 0.5 for the CCPs in each of the three groups (see Table 6).

Figure 5 shows that members consistently estimate European CCPs as being somewhat less prone to default than the other two groups. Our results on stress are broadly consistent with this comparison (see Table 3). However, Table 3 also suggests that, as a group, North American CCPs tend to be more prone to stress than Asia-Pacific CCPs, whereas the member estimates suggest that the latter are more prone to default. Thus, there is no one-to-one correspondence between GF breach probabilities and default probabilities. This is hardly surprising given that stress is only one of several factors that enter into estimates of CCP default.

Figure 5: U.S. GSIBs Estimates of Annual Default Probabilities of the Top 10 CCPs by Region

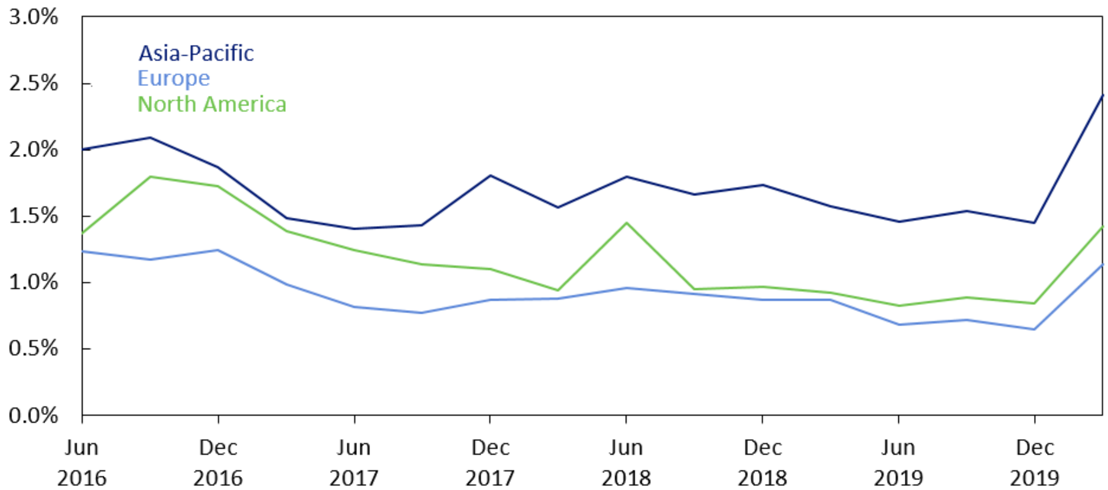

Note: Each line represents the average default probability estimate by each CCAR reporter across the top ten CCPs of each region matched against the top ten largest CCPs in the public quarterly data sample.
Source: Federal Reserve Y14 Q Schedule L.

Table 6: Annual Default Probabilities Statistics.

	All	Asia-Pacific	Europe	North America
Default Probability: Mean (%)	2.47	3.12	2.53	1.34
Default Probability: Std Dev (%)	1.91	1.73	2.27	1.10
Member Coefficient of Variation:	0.4678	0.4691	0.4638	0.4690
Quarter Sample	1042	465	286	291
CCP Sample	71	31	20	20
Default Probability of Top 10: Mean (%)	1.26	1.70	0.88	1.19
Default Probability of Top 5: Mean (%)	1.25	1.79	0.79	1.18

Note: Annual default probabilities as estimated by CCP members, averaged over the 15 quarters from 2016 Q2 - 2020 Q1 and averaged over the CCPs in each group. The standard deviation (Std Dev) refers to the standard deviation of the mean estimate over all quarters for each given subgroup of CCPs. The member coefficient of variation is a normalized standard deviation of member estimates in cases where more than one member provided an estimate for a given CCP.
Source: Federal Reserve Y14 Q Schedule L; authors' analysis.

7 Conclusion

The sizing of a CCP's waterfall determines its ability to fulfill its payment obligations under stress. The two main components of the waterfall are the initial margin posted against individual members' positions and the guarantee fund, which covers joint losses over and above the initial margin. There are well-established standards for setting the level of initial margin (BCBS and IOSCO (2015)): with probability at least 99.5 percent a member's IM should be able to cover its same-day variation margin obligation to the CCP. Unfortunately, no comparable standard exists for the sizing of the guarantee fund. As noted earlier, the Cover 2 rule offers no clear VaR standard for violation probabilities, and it is very sensitive to the number and size distribution of the CCP's membership.

In this paper, we have argued that the probability of a GF breach would be a clearer and more appropriate measure of a CCP's risk exposure. The measure can be used to set the appropriate level of a CCP's guarantee fund by specifying a VaR that should not be exceeded, either empirically or under stress scenarios specified by the regulators. Alternatively, it can be used as a metric of risk exposure that complements other approaches, such as Cover 2, and that can be estimated from public data across a wide variety of CCPs.A very useful addition to the public disclosures would be to have each CCP report, for each quarter, the maximum amount of payments owed to the CCP in excess of IM held against these payments, relative to the total amount of pre-funded resources.

Although our estimation methodology is restricted by the limitations of the data, it nevertheless provides a compact formulation that is easy to implement and allows one to study the trade-off between initial margin, guarantee fund size, and probability of a breach. Access to more complete data over a longer time horizon would allow the estimation of stress levels and GF breach probabilities for individual CCPs, whereas here we have been forced to pool the outcomes into subgroups of related CCPs. We conjecture that an individual-level analysis will show qualitatively similar results, namely, a heavy-tailed distribution of MC payments and a significant degree of correlation among MC payments by individual members. These results could be compared with members' confidential estimates of the default probabilities of individual CCPs. One could then study the relationship between default probability, breach probability, level of guarantee fund, nature of the membership, and various other explanatory variables. We leave such a study to future work.

References

Andersen, L. B. and Dickinson, A. S. (2018). Funding and credit risk with locally elliptical portfolio processes: An application to ccps. *Available at SSRN 3161154*.

Arnsdorf, M. (2012). Quantification of central counterparty risk. *Journal of Risk Management in Financial Institutions*, 5(3):273–287.

BCBS and IOSCO (2015). Margin requirements for non-centrally-cleared derivatives. Technical report, BIS and OICU-IOSCO, Basil, Switzerland.

Capponi, A., Cheng, W. A., Giglio, S., and Haynes, R. (2020). The collateral rule: Evidence from the credit default swap market. *Working Paper*.

Cont, R. and Kokholm, T. (2014). Central clearing of OTC derivatives: Bilateral vs multilateral netting. *Statistics & Risk Modeling*, 31(1):3–22.

Cox, R. T. and Steigerwald, R. S. (2017). A CCPis a CCP is a CCP. *Federal Reserve Bank of Chicago Policy Discussion Papers*.

CPMI-IOSCO (2012). Principles for financial market infrastructures. Technical Report. Bank of International Settlements, Basel, Switzerland.

Duffie, D. and Zhu, H. (2011). Does a central clearing counterparty reduce counterparty risk? *Review of Asset Pricing Studies*, 1(1):74–95.

Evanoff, D. D., Russo, D., Steigerwald, R., et al. (2006). Policymakers, researchers, and practitioners discuss the role of central counterparties. *Economic Perspectives*, (Q IV):2–21.

Ghamami, S., Paddrik, M., and Zhang, S. (2020). Central counterparty default waterfalls and systemic loss. *Office of Financial Research Working Paper*, 20-04.

ISDA (2013). CCP loss allocation at the end of the waterfall. Technical Report. New York, NY.

ISDA (2015). CCP default management, recovery and continuity: A proposed recovery framework. Technical Report. New York, NY.

Jones, R. A. and Pérignon, C. (2013). Derivatives clearing, default risk, and insurance. *Journal of Risk and Insurance*, 80(2):373–400.

Menkveld, A. J. (2015). Systemic risk in central clearing: Should crowded trades be avoided? *Working Paper*.

Menkveld, A. J. (2017). Crowded positions: An overlooked systemic risk for central clearing parties. *The Review of Asset Pricing Studies*, 7(2):209–242.

Menkveld, A. J. and Vuillemey, G. (2020). The economics of central clearing. *Annual Review of Financial Economics*. forthcoming.

Murphy, D. (2017). I've got you under my skin: Large central counterparty financial resources and the incentives they create. *Journal of Financial Market Infrastructures*, 5(3):57–74.

Murphy, D. and Nahai-Williamson, P. (2014). Dear prudence, won't you come out to play? Approaches to the analysis of central counterparty default fund adequacy. *Bank of England Financial Stability Paper No*, 30.

Paddrik, M. and Young, H. P. (2021). How safe are central counterparties in credit default swap markets? *Mathematics and Financial Economics*, 15(1):41–57.

Paddrik, M. E., Rajan, S., and Young, H. P. (2020). Contagion in derivatives markets. *Management Science*, 66(6):3125–3161.

World Federation of Exchanges (2020). Guidance on non-default loss. Technical report. Technical Paper.