

Crowded Trades: An Overlooked Systemic Risk for Central Clearing Counterparties¹

Albert J. Menkveld

First version: April 8, 2014

¹ Albert J. Menkveld, VU University Amsterdam, FEWEB, De Boelelaan 1105, 1081 HV, Amsterdam, Netherlands, tel +31 20 598 6130, albertjmenkveld@gmail.com and the Duisenberg School of Finance. I am grateful to EMCF for in-depth discussions about clearing.

Crowded Trades: An Overlooked Systemic Risk for Central Clearing Counterparties

Abstract

Counterparty default risk might hamper trade and trigger a financial crisis. The introduction of a central clearing counterparty (CCP) benefits trading but pushes systemic risk into CCP default. Standard risk management strategies at CCPs currently overlook a risk associated with crowded trades. This paper identifies it, measures it, and proposes a margin methodology that accounts for it. The application to actual CCP data illustrates that this hidden risk can become large, in particular at times of high CCP risk.

1 Introduction

The 2007-2008 financial crisis and the more recent European debt crisis renewed the debate on how to best measure and regulate systemic risk (see, e.g., [Bisias et al., 2012](#); [Brunnermeier and Oehmke, 2013](#)). Regulators advocate a central clearing counterparty (CCP) to mitigate the risk of market failure (see, e.g., Dodd-Frank and EMIR). A market might fail when agents are reluctant to trade as they fear counterparty default. The introduction of a CCP could solve such trade deadlock as the CCP effectively becomes the counterparty to each trade.

A CCP might reduce counterparty risk but it does not eliminate it. In fact, counterparty exposures concentrate in the CCP. Counterparty default becomes systemic in nature as CCP default affects all clearing members at the same time.¹ Regulators are aware as the [ESRB \(2012, p. 16\)](#) for example writes: “Mandatory clearing will turn CCPs into systemic nodes in the financial system.”

Risk management at the level of CCPs becomes a first-order concern. The Bank for International Settlements (BIS) and the International Organization of Securities Commissions (IOSCO) cooperated to develop recommendations for CCPs as a standalone entity ([BIS-IOSCO, 2004](#)) and as part of the “financial market infrastructure ([BIS-IOSCO, 2012](#)).” They emphasize that a CCP should measure its credit exposures to participants and limit them through margin requirements or other control mechanisms. A standard approach is to calculate (initial) margins on a member by member basis. This margin scales with the ‘product’ of the net positions in a member’s trade portfolio and the volatility of returns on these net positions (see, e.g., [Hedegaard, 2012](#)). An example of such approach is SPAN, the margin methodology developed by the Chicago Mercantile Exchange (CME) and implemented by 54 exchanges and clearing houses around the world.

This paper identifies and studies crowded trades as a risk to CCPs that is overlooked in current member-by-member margining methodologies. It explores the intuition that losses in members’ portfolios become correlated when their trades crowd on a single security or risk factor. If this

¹Example of CCP failures are Caisse de Liquidation, Paris (1974), the Kuala Lumpur Commodity Clearing House (1983), and the Hong Kong Futures Guarantee Corporation (1987) (see [Hills et al., 1999](#)).

factor experiences a large price shock multiple members experience large margin calls *simultaneously*. If unable to post new collateral, the CCP might have to cover losses not on a single portfolio but on many portfolios at the same time.

The paper develops four sets of results based on CCP “aggregate exposure,” the CCP risk measure proposed by [Duffie and Zhu \(2011\)](#). Aggregate exposure is defined as the summation over losses in members’ portfolios, disregarding profits. This loss is what the CCP is effectively exposed to when it inherits trade portfolios of members in default. The mean and variance of aggregate exposure are derived analytically based on absolute moments for normal distribution ([Nabeya, 1951](#)) and moments for the truncated normal ([Rosenbaum, 1961](#)). The first result is that CCP aggregate exposure becomes more volatile when its members’ trades crowd *ceteris paribus*. The right tail of the aggregate exposure distribution becomes ‘fatter,’ i.e., extreme losses become more likely. Crowded trades do turn out to be hidden risk.

Second, the extent to which member trades crowd can be represented by an index, CrowdIx. This index is constructed as the ratio of two volatilities. The numerator is the volatility of actual aggregate exposure. The denominator is the ‘maximized’ volatility obtained after re-allocating member trades to a single risk factor.²

Third, an alternative margin methodology is proposed to account for crowded risk: Margin(A). Let CCP risk be expressed by the value-at-risk (VaR) of aggregate exposure. A standard way to calculate VaR is to set it equal to the mean plus a constant times the standard deviation (the “delta-normal method” according to [Jorion, 2007](#)). This approach is particularly appealing in this setting as the mean and standard deviation of aggregate exposure are analytical results. The proposed new margin methodology Margin(A) simply sets the aggregate margin equal to this VaR measure. It has a couple of appealing features. First, homogeneity of degree one in member portfolio risk implies that Margin(A) decomposes naturally across members. Each member is margined according

²Note that the textbook case of all investors trading the market portfolio corresponds to maximum crowded risk for a CCP.

to the ‘shadow cost’ of its portfolio, i.e., the contribution of the portfolio to CCP risk. Second, sensitivity to risk factors is derived as an analytical result. Third, the computational cost is trivial as all results are analytical. This is particularly attractive when a CCP wants to monitor its exposure at high-frequency, has many members, and clears many securities. Finally, the standard VaR implementation follows industry practice and it is therefore more likely to be widely accepted.

Fourth, real-world data are used to illustrate that crowded risk matters. Proprietary CCP data was obtained for equity trades in Denmark, Finland, and Sweden. The sample runs from October 19, 2009 through September 9, 2010. The actual margins the CCP imposed are compared to the margin implied by Margin(A), both in terms of aggregate margin and member-by-member margin. The actual daily aggregate margin hovers around €200 million throughout the sample and peaks twice to more than €1 billion. CrowdIx reaches extreme levels on exactly these two days. Appropriately accounting for crowded risk would require margins that are hundreds of millions euros higher. The allocation of aggregate margin across members also differs substantially. On one of these days, one member was charged more than €150 million whereas the alternative measure would have charged less than €50 million. Another member was charged about €100 million but should have been charged more than €150 million. Inspection of their portfolios reveals that the latter member joined a crowded trade whereas the former did not.

The paper contributes to several literatures. Perhaps most closely related is the literature on how to measure systemic risk (reviewed in-depth in, e.g., [Bisias et al., 2012](#)). One approach is to gauge a financial institution’s contribution to systemic risk by measuring the left-tail co-variation of change in its market value with change in the value of the market as a whole (e.g., CovaR or Systemic Expected Shortfall proposed by [Adrian and Brunnermeier \(2011\)](#) and [Acharya et al. \(2010\)](#), respectively). An alternative forward-looking measure is to collect financial institutions’ exposures to a set of risk factors and identify risk through scenario analysis ([Duffie, 2012](#); [Brunnermeier, Gorton, and Krishnamurty, 2013](#)). This would be a natural task for the newly established Office of Financial Research. This exposure information could also be used to calculate

CrowdIx/Margin(A)-type measures to gauge to what extent taxpayers are exposed to the risk associated with crowded trades by financial institutions.

The paper also contributes to the literature on an important ‘follow-up’ question: now that systemic risk is measured, how should it be allocated? [Brunnermeier and Oehmke \(2013\)](#) review various allocation rules. The proportional rule allocates risk according to an institution’s individual risk divided by the sum of all institutions’ individual risk ([Urban et al., 2003](#)). Standard margin methodologies allocate risk this way. The with-and-without rule allocates it based on the difference in systemic risk when an institution is included or excluded ([Merton and Perold, 1993](#)). The Euler-or-gradient rule is the marginal version of this rule ([Patrik, Bernegger, and Rüegg, 1999](#); [Tasche, 2000](#)). Finally, the Shapley rule allocates according to Shapley value, a concept from cooperative game theory ([Tsanakas, 2009](#); [Tarashev, Borio, and Tsatsaronis, 2009](#)).

[Brunnermeier and Oehmke \(2013, p. 62-63\)](#) characterize a good allocation as follows:

“Ideally, the allocation should be such that (i) the sum of all risk contributions equals the total systemic risk and (ii) each risk contribution incentivizes financial institutions to (marginally) take on the appropriate amount of systemic risk. However, capturing both total and marginal risk contributions in one measure is a challenging task, because the relationship between the two may be non-linear. In fact, the marginal contribution of one institution may depend on the risks taken by other institutions.”

Margin(A) is an Euler-or-gradient rule that meets these desirables. In particular, an institution’s margin *does* depend on the risks that others take as it accounts for crowded trades. Admittedly, it seems particularly well-suited for linear products (e.g., equity) but requires linearization for non-linear products (e.g., options). An alternative non-linear model-based approach is suggested by [Brunnermeier and Cheridito \(2013\)](#).

A related contemporary paper is [Cruz Lopez et al. \(2014\)](#). The authors also identify crowded trades as a hidden source of credit risk for a CCP. They propose ‘CoMargin,’ a margin methodology that sets margins for a particular member at a level high enough to avoid losses *conditional* on one

or two other members being in financial distress. It is derived from the CoVaR concept introduced by [Adrian and Brunnermeier \(2011\)](#). The relative merit of Margin(A) over CoMargin is that it avoids conditioning and therefore any discussion on who the appropriate clearing members are to condition on. No clearing member would volunteer to be selected as, by construction, his portfolio would correlate highly (i.e., perfectly) with the benchmark portfolio. He would therefore be charged most. In addition, it is hard to take the aggregate portfolio across all clearing members as a benchmark as this portfolio is zero by definition (for every clearing member buying there is one selling).

The remainder of the manuscript is organized as follows. Section 2 presents the model, develops the crowded-risk multiplier (CrowdIx) and the new margin methodology implied by it (Margin(A)). Section 3 uses the Nordic sample as a case study to illustrate CrowdIx and Margin(A). Section 4 concludes.

2 Simple model to capture crowded risk

2.1 Model description

Central clearing counterparty (CCP) risk is analyzed in the framework proposed by [Duffie and Zhu \(2011\)](#). Let a period be defined by the frequency at which a CCP collects margin. For example, a period will be one day for our data sample. Consider a set of I securities with one-period returns that are normally distributed:

$$R \sim N(0, \Omega). \tag{1}$$

Let the vector n_j be the start-of-period yet-to-settle trade portfolio of member j . This is the portfolio that the CCP effectively insures against counterparty default. Once settlement has taken place for a particular trade, it will be removed from the trade portfolio as it no longer requires insurance. n_j will be referred to as trade portfolio or portfolio in the remainder of the manuscript. Mathemat-

ically, the aggregate exposure A associated with the trade portfolios of all member firms is

$$A = \sum_j E_j, \quad (2)$$

where E_j is the exposure to losses in member j 's portfolio

$$E_j = -\min(X_j, 0), \quad (3)$$

where X_j is the one-period profit on member j 's trade portfolio

$$X_j = n_j' R, \quad \text{hence } X \sim N(0, \Sigma), \quad \Sigma = N' \Omega N, \quad N = [n_1, \dots, n_J], \quad (4)$$

and, for the remainder of the paper, Σ denotes the covariance matrix of X .

As in [Duffie and Zhu \(2011\)](#), analysis of CCP risk centers on the aggregate exposure variable A . They state (p. 78):

“For given collateralization standards, the risk of loss caused by a counterparty default is typically increasing in average expected exposure.”

Note that this average expected exposure is defined based on the start-of-period trade portfolio. It therefore does not account for likely changes in this trade portfolio throughout the trading period. This intraperiod portfolio change risk is potentially important but out of scope here (and left for future work). The focus is on how to improve on existing margin methodologies (that also ignore intraperiod risk) by recognizing the systemic risk that is embedded in crowded trades.

The aggregate exposure mean and standard deviation are calculated based on two set of results: absolute moments of a bivariate normal distribution ([Nabeya, 1951](#)) and moments of a truncated bivariate normal distribution ([Rosenbaum, 1961](#)). All derivations are in appendix [A](#).

Aggregate exposure mean. The aggregate exposure mean is

$$E(A) = \sum_j \sqrt{\frac{1}{2\pi}} \sigma_j. \quad (5)$$

The mean increases linearly in the standard deviation of member j 's one-period profit.

Aggregate exposure standard deviation. The aggregate exposure standard deviation is

$$\text{std}(A) = \sqrt{\sum_{k,l} \left(\frac{\pi - 1}{2\pi} \right) \sigma_k \sigma_l M(\rho_{kl})}, \quad (6)$$

where ρ_{kl} is the correlation between the portfolio profit of member k and l and

$$M(\rho) = \frac{\left[\frac{1}{2}\pi + \arcsin(\rho) \right] \rho + \sqrt{1 - \rho^2} - 1}{\pi - 1}. \quad (7)$$

Note that $M(-1) = -1/(\pi - 1)$, $M(0) = 0$, and $M(1) = 1$.

[insert Figure 2 here]

The $M(\cdot)$ function is the engine of all crowded-risk analysis. It maps correlations in portfolio profits to correlations in portfolio exposures. Figure 2 plots the function and illustrates that it leaves positive correlations almost untouched (as it is close to a 45 degree line), but reduces the size of negative correlations. The intuition is that, unlike aggregate profit, aggregate exposure cannot benefit from a profit in one account to compensate for the loss in another. The clearing house cannot make up for the portfolio loss of members in default by collecting from its nondefaulted members who (by the zero-sum nature of trading) enjoy a portfolio profit. In some sense, negative correlations become less useful. Positive correlations do not die as a loss in one account becomes more likely when there is a loss in another account. In essence, the CCP insures ‘left-tail draws’ from the portfolio profit distribution³ but does not benefit in the cross-section from offsetting right-tail draws. It does suffer from other left-tail draws.

³Note that these correspond to right-tail draws from the portfolio *loss* distribution.

2.2 A measure of crowded risk

Definition 1 *CrowdIx for Σ (the covariance matrix of one-period profits on the J members' portfolios) is defined as*

$$\text{CrowdIx} = \text{std}(A)/\text{std}(\tilde{A}), \quad (8)$$

where $\text{std}(\tilde{A})$ is the aggregate exposure when all members' trades are re-allocated to a single risk factor based on the first-fit-descending binning algorithm.

CrowdIx compares the standard deviation of aggregate exposure of actual trade portfolios to the standard deviation for a benchmark of re-allocating member trades to a single risk factor (e.g., the market index). First off, note that CrowdIx is based on standard deviation only as the aggregate-exposure mean does not change when only re-allocating trades. More specifically, the benchmark portfolio is constructed such that each member's portfolio risk remains unchanged.

Consider for example the case of four clearing members, $J = 4$. All trades are equal in size and risk factors are independent identically distributed random variables. Suppose member one bought risk factor one from member two, and member three bought risk factor two from member four. The benchmark portfolio is easily created by re-allocating the latter trade to risk factor one. Such perfect re-allocation is not possible in general. Suppose that member one bought risk factor one from member two, member two bought risk factor two from member three, and member three bought risk factor three from member one. In this case it is impossible to re-allocate all trades to a single risk factor under the conditions that all members' individual portfolio risk is unchanged.

The general approach is to re-allocate all trades to a single risk factor to the maximum extent possible. Actual maximization is a non-trivial problem as the set of possible re-allocations is simply too big. It consists of 2^J possibilities as each member becomes either a buyer or seller of the single risk factor. In our data, for example, the size of this set would be $2^{55} \approx 3.6 * 10^{16}$. The approach taken here is to revert to the first-fit-descending (FFD) binning algorithm ([Coffman](#),

Garey, and Johnson, 1996). Members are sorted from largest to smallest portfolio risk. They are then sequentially assigned to be either buyer, seller, or land in a residual category. The objective is to maximize the total amount transacted in this risk factor. The residual category is ignored. It is expected to be small given a highly skewed distribution in member portfolio risk due to lots of small-risk portfolios. Appendix B describes the procedure in detail.

Lemma 1 *CrowdIx has a lower bound:*

$$CrowdIx \geq \sqrt{\frac{1}{\lfloor J/2 \rfloor}}, \quad (9)$$

where the function $\lfloor x \rfloor$ yields the largest integer smaller or equal to x .

The minimum is attained when all members' portfolios are equally risky and portfolio risk is spread perfectly across orthogonal risk factors. In this case, none of the off-diagonal elements in the portfolio risk covariance matrix Σ is positive. There simply is no positive correlation in profit among clearing members in such situation. If one is able to perfectly re-allocate all trades to a single risk factor then the sum across off-diagonals is unaffected (as aggregate profit remains zero). If effectively adds a mean-preserving spread to the off-diagonals. This raises aggregate exposure as the positive mass on off-diagonals gets multiplied by $M(1) = 1$ and negative mass is multiplied by $M(-1) = 1/(\pi - 1)$.⁴ The maximum factor by which aggregate exposure can be raised from perfect diversity to full crowding scales with the number of clearing members. This is not surprising given that more clearing members imply more opportunity for diversity of trades. A detailed proof is in Appendix B.

CrowdIx cannot be bounded from above. One reason is that perfect re-allocation might not be feasible. The earlier example of three trades among clearing members serves as a counter example. Conceptually one could simply define the maximum aggregate exposure as the maximum attained

⁴The function M is only evaluated at +1 and -1 as re-allocating to a single risk factor implies either perfect positive correlation (both members are buyers or both members are sellers) or perfect negative correlation (one member is buyer, the other is seller).

by portfolios in the set restricted by unchanged trade portfolio risk by each member. As noted, this set is simply too large to evaluate. In practice, the FFD procedure seems to always yield higher-crowding benchmark portfolios. In our data sample CrowdIx values stayed far below one. Its maximum is 0.83 (see Section 3).

The CrowdIx bounds' analysis illustrates that two sources contribute to crowded risk. First, it occurs when equal-risk members' trades crowd in a single risk factor. Second, it occurs when portfolio risk differs across members *irrespective* of the trades they implemented. This second source is perhaps best illustrated by taking it to the extreme. Suppose only two members trade then necessarily their trades crowd and CrowdIx is one.

2.3 Example to illustrate crowded risk

A simple example illustrates the concept of crowded risk. Consider two securities that are available for trade. Their one-period returns are distributed as two independent standard normal variables. Let four members implement two trades.

Crowded trades. Suppose member 1 sells a unit of security 1 to member 2, and member 3 sells a unit of the same security to member 4. In this case all members' trades crowd on the same security. The trade portfolio matrix therefore is

$$N = \begin{pmatrix} 1 & -1 & 1 & -1 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad \text{and} \quad \Sigma = N' \Omega N = \begin{pmatrix} 1 & -1 & 1 & -1 \\ -1 & 1 & -1 & 1 \\ 1 & -1 & 1 & -1 \\ -1 & 1 & -1 & 1 \end{pmatrix}, \quad \text{respectively,} \quad (10)$$

which implies the following mean and covariance for exposures

$$E(E) = \sqrt{\frac{1}{2\pi}} \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}, \quad \text{var}(E) = \frac{1}{2\pi} \begin{pmatrix} \pi - 1 & -1 & \pi - 1 & -1 \\ -1 & \pi - 1 & -1 & \pi - 1 \\ \pi - 1 & -1 & \pi - 1 & -1 \\ -1 & \pi - 1 & -1 & \pi - 1 \end{pmatrix}. \quad (11)$$

The mean and variance of aggregate exposure are therefore

$$E(A) = 4 \sqrt{\frac{1}{2\pi}} \approx 1.60 \quad \text{and} \quad \text{std}(A) = 2 \sqrt{\frac{\pi-2}{\pi}} \approx 1.21, \quad \text{respectively.} \quad (12)$$

CrowdIx is one in this case as $\text{std}(\tilde{A}) = \text{std}(A)$ since the benchmark portfolio equals the actual portfolio.

Noncrowded trades. Suppose that members 3 and 4 trade security 2 instead of security 1, i.e., both trades now exhibit perfect diversity. The trade portfolio matrix and the covariance matrix of portfolio profits are

$$N = \begin{pmatrix} 1 & -1 & 0 & 0 \\ 0 & 0 & 1 & -1 \end{pmatrix} \quad (13)$$

and the covariance matrix of portfolio profits is

$$\Sigma = \begin{pmatrix} 1 & -1 & 0 & 0 \\ -1 & 1 & 0 & 0 \\ 0 & 0 & 1 & -1 \\ 0 & 0 & -1 & 1 \end{pmatrix}. \quad (14)$$

The mean and covariance matrix of exposures is

$$E(E) = \sqrt{\frac{1}{2\pi}} \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}, \quad \text{var}(E) = \frac{1}{2\pi} \begin{pmatrix} \pi-1 & -1 & 0 & 0 \\ -1 & \pi-1 & 0 & 0 \\ 0 & 0 & \pi-1 & -1 \\ 0 & 0 & -1 & \pi-1 \end{pmatrix}. \quad (15)$$

The mean and variance of aggregate exposure ($A = \iota' E$) are therefore:

$$E(A) = 4 \sqrt{\frac{1}{2\pi}} \approx 1.60 \quad \text{and} \quad \text{std}(A) = 2 \sqrt{\frac{\pi-2}{2\pi}} \approx 0.85. \quad (16)$$

CrowdIx is $\sqrt{1/2} = 0.71$ in this case because $\text{std}(\tilde{A}) = 2 \sqrt{(\pi-2)/\pi}$. The index reaches the lower bound (see lemma 1).

[insert Figure 3 here]

In sum, the mean aggregate exposure for the CCP is the same for the two polar cases, but the standard deviation is 41% higher in the crowded trades. Higher moments also differ as is evident from the simulated distribution of aggregate exposure depicted in Figure 3. The 95% exposure percentile is in both cases (noncrowded and crowded) accurately approximated by $\mu + 1.96\sigma$ as the true percentile at this value is 95.2 in the crowded trade case and 95.9 in the noncrowded case.

2.4 Contribution to crowded risk

If one observes high values of CrowdIx, it is useful to identify its cause. This section presents two results that are helpful in this respect. First, a decomposition of CrowdIx across members reveals which members contribute most. Second, a partial derivative result helps identify on what security/risk factor the trades crowd.

Decomposition across members. The standard deviation of aggregate exposure (eqn. (6)) is homogeneous of degree one in member portfolios' risk. The exposure standard deviation is therefore naturally assigned to the various traders' portfolios by applying Euler's homogeneous function theorem:

$$\text{std}(A) = \sum_k \sigma_k \frac{\partial \text{std}(A)}{\partial \sigma_k} = \sum_k \sigma_k \sum_l \frac{1}{\text{std}(A)} \left(\frac{\pi - 1}{2\pi} \right) \sigma_l M(\rho_{kl}). \quad (17)$$

The contribution of member k is therefore naturally identified as $\sigma_k(\partial \text{std}(A)/\partial \sigma_k)$, i.e., the amount of risk in his trade portfolio times the marginal cost to a CCP of such risk. This marginal cost itself is naturally interpreted as the shadow cost of a trade portfolio in terms of CCP risk. It measures the extent to which a member to joined crowded trades. Mathematically, it is captured by the factor $M(\rho_{kl})$.

The CrowdIx methodology and the decomposition result can also be applied to a subset of agents, e.g., market intermediaries or high-frequency traders (HFTs). The aggregate exposure A defined in eqn. (2) is then redefined as the sum over only this subset of agents. It becomes the

aggregate exposure of the CCP to this subset of clearing members.

CrowdIx could even be adapted to measure the extent to which a subset of agents crowd *among themselves*. This is particularly relevant for a CCP who is worried about liquidation risk. If, for example, all intermediaries default at the same time, a CCP is likely to pay a fire sale premium when offloading the inherited portfolios to any remaining ‘slow moving capital’ (Menkveld, 2013). The natural benchmark portfolio for the adapted measure is the one where all intermediaries’ bets are on *one* side of a single risk factor. Off-diagonal of Σ become unconstrained (i.e., negative covariance is avoided) and therefore maximized by setting them equal to the product of the corresponding members’ portfolio risk (i.e., correlations become one among clearing members profit). The lower bound of such modified CrowdIx measure becomes $\sqrt{1/J}$ (cf. lemma 1).

Sensitivity to particular security/risk factor. To discover on what security/risk factor trades crowd, it is useful to measure how aggregate exposure changes if one more unit of risk is added to the risk factor of interest, say f (e.g., a particular security’s return, the market return, or a Fama-French factor). The partial derivative of $\text{std}(A)$ with respect to σ^f is derived in appendix A. The result is:

$$\frac{\partial}{\partial \sigma^f} \text{std}(A) = \left(\frac{\pi - 1}{4\pi} \right) \frac{\sigma^f}{\sigma^A} \sum_{k,l} \left[2M'(\rho_{kl})B_{kl} + \left(\frac{\sqrt{1 - \rho_{kl}^2} - 1}{\pi - 1} \right) \left(\frac{\sigma_l}{\sigma_k} B_{kk} + \frac{\sigma_k}{\sigma_l} B_{ll} \right) \right]. \quad (18)$$

with

$$M'(\rho) = \frac{\frac{1}{2} + \arcsin(\rho)}{\pi - 1}, \quad B_{kl} := n_k' \beta \beta' n_l, \quad \text{and} \quad \beta = \text{cov}(R, r^f) / \text{var}(r^f). \quad (19)$$

The expression leads to several observations. First, the expression is linear in σ^f , the risk of factor f . Second, a key factor is B_{kl} that captures not only how trade portfolio profits load on the factor (e.g., B_{kk} for member k ’s portfolio), but also on the extent to which they crowd on that factor ($B_{kl}, k \neq l$). Third, the effect that B_{kl} has on aggregate exposure is governed by the derivative of

the $M(\cdot)$ function. The slope of this function is higher in the positive domain than in the negative domain (see Figure 1), i.e., a risk change in the factor kicks in stronger in areas where portfolio risk crowds relative to areas where it does not.

2.5 Margin implied by aggregate exposure, $\text{Margin}(A)$

CrowdIx can be used to set margins that assign CCP (exposure) risk appropriately across its clearing members. The standard SPAN methodology of the Chicago Mercantile Exchange sets margins solely based on risk in a member's portfolio irrespective of whether his trades crowd or not. The margin scales with the right tail of the portfolio loss distribution as it based on value at risk (often shortened to VaR). The CME reports that SPAN is used by 54 exchanges and clearing organizations world-wide.⁵

The margin methodology that EMCF, the CCP in our sample, uses is the in-house correlation haircut model (CoH). CoH is similar in spirit to SPAN, but remains proprietary and mostly opaque. The most detailed public description of CoH is in a report by the Dutch Central Bank (Zhu, 2011). It states that a margin calculation (i) takes into account the correlation between the various products that are part of the investor portfolio, (ii) determines the risk factors that have the greatest impact on the portfolio, (iii) shifts these components to find worst case scenario (maximum loss), and (iv) attribute back the contribution per product in the determined worst case.

If CCP risk were to be margined for instead of individual member portfolio risk, then a CrowdIx inspired margin would be preferred over both SPAN and CoH. A natural analog to the existing methodologies is to base margin on the right-tail of aggregate exposure. A particularly appealing value-at-risk measure is what Jorion (2007, p. 260) refers to as the “delta-normal method.” It calculates VaR for a portfolio return as the mean plus alpha times the standard deviation. This suggests defining a CrowdIx-based margin as

⁵See “SPAN’s Risk Based Margin Requirements Allow Effective Margin Coverage and Efficient Use of Capital SPAN’s Risk Based Margin Requirements Allow Effective Margin Coverage and Efficient Use of Capital,” available from the CME website.

$$\text{Margin}(A) := E(A) + \alpha \text{std}(A). \quad (20)$$

Margin(A) is appealing for several reason. First, it follows standard practice. It is based on one of the three standard VaR approaches advocated in [Jorion \(2007\)](#); the other two are “historical simulation” and “Monte Carlo simulation.” Second, the mean and standard deviation are analytical and therefore easily computed (particularly relevant in the case of many securities, many clearing members, and high-frequency monitoring required due to sub-millisecond trading). Third, the analytical decomposition and sensitivity results developed for CrowdIx carry over to Margin(A). This is particularly useful for real-time analysis. Which member is most ‘systemically risky’? And, what risk factor do members’ bets crowd on? The remainder of the section develops these decomposition and sensitivity results.⁶

The CrowdIx decomposition result based on eqn. (17) is useful to develop a similar result for Margin(A). The part of Margin(A) that can be attributed to member j is

$$\text{Margin}(A)_j = \sqrt{\frac{1}{2\pi}}\sigma_j + \alpha \sigma_j \frac{\partial \text{std}(A)}{\partial \sigma_j} = \sigma_j \left(\sqrt{\frac{1}{2\pi}} + \alpha \sum_l \frac{1}{\text{std}(A)} \left(\frac{\pi - 1}{2\pi} \right) \sigma_l M(\rho_{jl}) \right), \quad (21)$$

where $\text{Margin}(A)_j$ denotes the margin imposed on member j . It follows directly that the individual margins add up to the CCP aggregate margin.

The Margin(A) sensitivity result is

$$\frac{\partial}{\partial \sigma^f} \text{Margin}(A) = \frac{\partial}{\partial \sigma^f} E(A) + \alpha \frac{\partial}{\partial \sigma^f} \text{std}(A), \quad (22)$$

where the standard deviation partial derivative is given by eqn. (18) and

⁶Chebyshev’s inequality implies that a result based on the first two moments is useful beyond the assumption of normality. It yields an upper bound for VaR. Chebyshev says $P[|X - \mu| \geq k\sigma] \leq 1/k^2$ where μ and σ are the distribution’s mean and standard deviation, respectively. For example, in the later application α is set to seven and margin therefore corresponds to a VaR that is at most $1/7^2 \approx 0.02$ (normality implies a VaR of 1.3×10^{-12}).

$$\frac{\partial}{\partial \sigma^f} E(A) = \sum_j \sqrt{\frac{1}{2\pi}} \frac{\sigma_f}{\sigma_j} B_{jj}, \quad (23)$$

and B_{jj} as defined in eqn. (19).

3 Application of the crowding index, CrowdIx

3.1 Data and summary statistics

This section explores the crowding index, CrowdIx, by applying it to actual data. The sample available for analysis consists of about a year's worth of 'trade' reports filed by clearing members to their central clearing counterparty, EMCF. It also contains the margin EMCF required each member to post each day. The sample period is October 19, 2009 through September 10, 2010. The reports cover transactions in stocks that are listed in Denmark, Finland, or Sweden.

The data set includes trades across almost all exchanges that traded these stocks at the time: NASDAQ OMX, Chi-X, Bats, Burgundy, and Quote MTF. The only exchange not included is Turquoise which had a very small market share at the time. Turquoise trades were outside of the realm of EMCF.

Fifty-seven clearing members traded these stocks (see Appendix C). They span a variety of financial firms, e.g., international brokerage firms (e.g., Goldman Sachs and JPMorgan), local banks (e.g., Nordea and Swedbank), and high frequency traders (e.g., GETCO and Knight). Fifty-five of them were active in the sample period. Members are anonymized and referred to by a two digit code (random numbers between 0 and 100).

[insert Table 1 here]

Table 1 provides some summary statistics. The sample consists of 1.4 million reports in 242 securities across 228 trading days. This implies that on average a stock traded about 25 times per

day. Signed volume (buy volume gets a positive sign, sell volume gets a negative sign) sums to zero for each stock each day. This result serves as a straightforward data integrity check.

Overall statistics reported in panel A lead to a couple of observations. First, clearing members collectively filed about 6000 trade reports to the CCP each day. Trade size is heavily skewed to the right as the average number of shares traded is 25,600 (€287,600) whereas the median is only 2,600 (€36,100). The largest trade was for roughly 19 million shares worth €142 million.

Cross-sectional statistics show substantial difference in activity across clearing members and across stocks. Panel B reports how average activity is distributed across clearing members. The cross-sectional mean of the average daily number of reports filed by a clearing member is 114.5; the median is 64.9. The least active member filed a trade report less than once a month whereas the most active member filed 736.4 reports daily. Daily volume statistics show similar skewness. The cross-sectional statistics for stock activity also exhibit skewness to the right although to a lesser degree. Panel C shows that the mean and median for the average daily number of reports per stock are 26.0 and 20.6, respectively. The least active stock trades less than once a month on average whereas the most active stock trades 84.2 times per day.

3.2 Aggregate exposure distribution and the implied margin

Margin(A), the margin implied by aggregate exposure (see Section 2.5), is calculated for each day in the sample. A time-varying covariance matrix for returns is calculated as an exponentially weighted moving average (EWMA) of the outer product of historical daily returns. Such covariance matrix has become a standard in industry and is used by, for example, MSCI's RiskMetrics. The parameter used for the daily frequency is 0.94. It implies that the most recent week carries $(1-0.95^2)=27\%$ of total weight; the most recent month 71%. EMCF's proprietary methodology CoH also uses EWMA estimates of the covariance matrix. The CCP inherited CoH from Fortis Bank Global Clearing who wrote in the annual accounts (p. 21): "E.g. equity prices are stressed by seven standard deviations (EWMA or implied), giving a confidence level of 99.98%." The quote

further leads us to set the scaling factor used to calculate $\text{Margin}(A)$ to seven, i.e., $\alpha=7$.

[insert Figure 4 here]

Figure 4 illustrates the time variation in aggregate margin both for margin implied by aggregate exposure, $\text{Margin}(A)$, and the actual margin collected by the CCP. For the remainder of the analysis, $\text{Margin}(A)$ was scaled up by 6.13 so as to make both margins have the same time average. The focus is on comparing the time variation as well as the cross-sectional variation in both margins, not on the level *per se*.

Panel A illustrates that through time both aggregate margin series are highly correlated. Their empirical distributions are both heavily skewed with two extreme peaks occurring on April 26, 2010 and May 10, 2010. Margins typically hover around €150 million, but shoot up to more than €500 million for these two days. The second peak corresponds to ‘macro shock’ as it followed an announcement that the Eurozone countries and IMF would bailout Greece. The first one is more of a surprise as it was triggered by an idiosyncratic event, i.e., Nokia reported first-quarter earnings that were far below analyst expectations.

The more striking finding revealed in Panel A is that when margins shoot up $\text{Margin}(A)$, the one that accounts for crowded risk, shoots up substantially more. For example, the total ‘Greek-bailout’ margin collected by the CCP was €511 million, whereas $\text{Margin}(A)$ would amount to €747 million for that day, i.e., €236 million more. The total ‘Nokia’ margin collected by the CCP was €495 million where $\text{Margin}(A)$ amounts to €644 million, a difference of €149 million.

Panel B illustrates that crowding is what drives the wedge between the two margin methodologies. It zooms in on the month surrounding the two peak margin events and overlays the crowding index, CrowdIx . CrowdIx is 72% for the Nokia peak and 62% for the Greek-bailout peak. The CrowdIx sample mean is 47% with a standard deviation of 8%. A level of 62% or higher therefore indicates high crowded-trade risk. As a matter of fact, the Nokia-day level is the second highest level in the sample.

The disproportionately large jump in $\text{Margin}(A)$ relative to CCP-imposed margin deserves further study. The new margin methodology accounts for crowding and CrowdIx does indeed reach extreme levels in this period. As $\text{Margin}(A)$ is based on an easily computed analytical mean and standard deviation, it is worth studying the actual distribution to verify the existence of a ‘fat’ right tail due to crowding. The approach taken here is to simulate returns from the standard normal, calculate the aggregate exposure realization for each simulation, and plot its distribution. As benchmarks, the same procedure was carried out for the median and minimum CrowdIx day. The distribution for these benchmark days are rescaled so as to match the mean of the day of interest.

[insert Figure 5 here]

Figure 5 reveals that the distribution for the sample day with largest margin does indeed exhibit a fat right tail. The rescaled distributions for the median and minimum CrowdIx benchmark days show considerable less mass in the right tail. The 90% quantile for the Greek-bailout day is €267 million compared to €196 and €175 million for the median and minimum CrowdIx benchmark days, respectively. The 99.9% quantile (corresponding to an extreme event once every four years) is €510 million for the Greek-bailout day relative to €360 million and €277 million for the two benchmark days. Tail risk defined this way appears to be about a third to two-thirds higher due to crowding.

3.3 Aggregate margin decomposition across members

The $\text{Margin}(A)$ decomposition result of eqn. (21) is used to analyze which members contribute most to CCP risk. In particular, is a member’s contribution primarily due to his portfolio risk? Or does it also reflect the extent to which the member participated in crowded trades? In the latter case, CrowdIx -based margins allocate aggregate CCP risk differently across members and are therefore a meaningful innovation relative to existing margin methodologies.

[insert Figure 6 here]

Figure 6 illustrates the issue with three scatterplots of *actual* margin posted by a member relative to what they would have to post based on Margin(A). Panel A depicts the scatterplot for July 29, 2010. It is the median-CrowdIx sample day (CrowdIx=0.46) and could therefore be considered a ‘representative.’ If both margin methodologies agree on how to allocate aggregate CCP risk across members, then all points in the scatterplot should be on the 45 degree line. The plot illustrates that re-allocation of CCP risk is substantially different under Margin(A). Member 6 is perhaps the best example. The CCP required him to post more than €15 million whereas based on Margin(A), he would need to post only €5 million, less than one third of what was posted.

The Greek-bailout scatterplot in Panel B identifies a few members who contribute strongly to the high level of crowding on this day (CrowdIx=0.62). Member 41 for example who posted less than €150 million but should have posted €250 million. His trade portfolio characterized in the exhibit below the plot shows that its ten largest positions are short positions. Given the macro nature of the event, market return is the likely candidate for crowded trades (Section 3.4 will show that it is). Other members appear to be closer to the 45 degree line. The margin required for member 12 for example is about the same for the two margin methodologies. His portfolio is also shown and unlike member 41 its ten largest positions are neither fully long nor fully short in the market portfolio.

One interpretation of the Greek-bailout observations is that member 41 is arbitraging across the index futures market and the cash market. One could say that it is therefore not systemically risky. In this particular case the cash CCP is different from the index futures CCP (in case there is one). Any potential for offsetting gains for member 41 are therefore out of reach for the cash CCP. Economically, any further margin calls due to quick recovery of the market (requiring member 41 to post ‘variation margin’ to mark his position to market) should in this case not be a problem for member 41 as such recovery would free up margin on the long position in the index futures market. The key point is the cash CCP is simply not aware of offsetting positions and therefore the large Margin(A) contribution member 41 has to post reflects the enormous shadow cost of his position

to the cash CCP.

The Nokia day depicted in Panel C suggests that many members were contributing to the high crowded risk ($CrowdIx=0.72$), yet some large portfolio risk members stayed away from the crowded bet. Member 12 for example posted €60 million but would be owed less than €10 million based on $Margin(A)$. Not surprisingly, the Nokia stock is not a top-ten position in his portfolio. Member 41 on the other hand contributed most to CCP risk as he would be charged more than €80 million based on $Margin(A)$ but in reality only posted less than half of it. His portfolio does show significant Nokia exposure, 20.7% in terms of absolute net position.

3.4 Aggregate margin sensitivity to risk factors

The $Margin(A)$ sensitivity result of eqn. (22) reveals what risk factor(s) the trades of clearing members crowd on. This partial derivative result reveals how aggregate margin changes when a particular risk factor becomes one unit more risky (measured in terms of standard deviation units). Perhaps more informative are elasticities which are easily calculated by multiplying the partial derivative with the ratio of the standard deviation of the candidate risk factor and $Margin(A)$.

To illustrate sensitivity analysis $Margin(A)$ is analyzed for same three days as used in the decomposition analysis: the median- $CrowdIx$ day, the Greek-bailout day, and the Nokia day. The three factors considered are: the market return as measured by the index $STOXXNordic30$, the European telecom sector return $STOXXTelecom$, and the Nokia stock return. The telecom sector was chosen as the Nordic countries seem to be telecom-heavy through firms like Tele2, Ericson, and Nokia.

[insert Table 2 here]

The sensitivity results presented in Table 2 lead to a couple of observations. First, the strongest finding are elasticities larger than one. Perhaps the most interesting one is the 1.05 witnessed for the Nokia day with respect to the Nokia risk factor. If daily Nokia stock volatility would increase

by one percentage point, the aggregate margin would increase by €147 million. Elasticities for the market and for telecom are an order of magnitude smaller, i.e., 0.19 and 0 respectively. This illustrates the idiosyncratic nature of the event. A shock to a single stock's volatility demands more aggregate collateral to be posted than a shock to the market. This is extraordinary as on the 'representative' median-CrowdIx day the elasticity for the market is 0.91, for the telecom sector it is 0.46, and for Nokia it is 0.15.

Second, for the Greek-bailout day Margin(A) is particularly sensitive to the market return and the telecom sector return. Elasticities are 0.98 and 0.83 respectively. This is not surprising given the macro nature of the event.

4 Conclusion

Crowded trades constitute a risk to a CCP that is overlooked in standard margin methodologies. CCP risk is assumed to scale with the aggregate loss in its members' portfolios; potential offsetting profit on other members' portfolios is ignored. This 'aggregate exposure' distribution become more skewed to the right if member portfolios crowd. A crowding index, CrowdIx, is developed as a summary statistic of the level of crowding. The paper further proposes an alternative margin methodology, Margin(A), to accounts for the additional risk due to crowded trades.

CrowdIx has several attractive properties. First, it is easily computed as it relies on an analytical result. It does not require computationally burdensome simulations which are particularly costly when there are many clearing members, many securities, and a need for high-frequency monitoring. Second, it decomposes naturally across clearing members so that it is easy to identify who contributes most to the crowded risk. Third, an analytical partial derivative result enables one to quickly identify the set of securities/risk factor(s) on which member trades crowd.

The relevance and strength of these results is illustration by an application to real data. CrowdIx and Margin(A) were calculated for a 2009-2010 sample of trades in Nordic stocks by members of

a large European central clearing counterparty.

Appendix

A Aggregate exposure mean and standard deviation

The mean and standard deviation of aggregate exposure is derived in this appendix. It is derived for the case where the distribution of member portfolio profit is standard normal, i.e., all diagonal elements of Σ are equal to one. The more general case is easily derived by scaling these results appropriately.

The mean and standard deviation of individual exposures follow immediately from results on the absolute moments of a bivariate normal ([Nabeya, 1951](#)):

$$E(E_j) = \frac{1}{2}E(|Z|) = \sqrt{\frac{1}{2\pi}}, \quad \text{and} \quad \text{std}(E_j) = \sqrt{\frac{1}{2}E(Z^2) - (E(E_j))^2} = \sqrt{\frac{\pi - 1}{2\pi}}, \quad (\text{A1})$$

where Z is a standard normal random variable.

The mean and standard deviation of aggregate exposure $A = \sum_j E_j$ further requires calculation of how correlation among portfolio profits X_j translates into correlation among exposures $E_j = -\min(X_j, 0)$. This result is derived in the next few steps.

1. [Rosenbaum \(1961, eqn. \(5\)\)](#) implies the following result for a truncated bivariate normal distribution constructed from two correlated standard normals, Z_1 and Z_2 , and truncation points set equal to zero:

$$L(0, 0; \rho)E(\tilde{Z}_1 \tilde{Z}_2) = \rho L(0, 0; \rho) + \frac{\sqrt{1 - \rho^2}}{2\pi} \quad (\text{A2})$$

where ρ is correlation among Z_1 and Z_2 , $L(h, k; \rho)$ is the probability mass in the top-right quadrant defined by truncation points h and k , and $\tilde{Z}_1$ and $\tilde{Z}_2$ represent the truncated random

variables.

2. $L(h, k; \rho)$ is not a analytical expression in general, but for the special case $h = k = 0$ it can be calculated explicitly (see, e.g., [Balakrishnan and Lai, 2009](#), p. 495)

$$L(0, 0; \rho) = \frac{1}{4} + \frac{1}{2\pi} \arcsin(\rho). \quad (\text{A3})$$

3. The exposure correlation now follows from the following steps.

- (a) The factor $L(0, 0; \rho)$ on the LHS of eqn. (A2) is deleted as the corresponding moment for exposures does not involve truncation. The product of exposures is zero in the other three quadrants.
- (b) The expression for $L(0, 0; \rho)$ in eqn. (A3) is inserted into eqn. (A2).
- (c) The covariance between exposures is then derived by subtracting the squared exposure mean (i.e., $1/2\pi$) from the exposure moment.
- (d) The correlation is derived by dividing this covariance by the exposure variance which equals $(\pi - 1)/2\pi$.

The result of these steps is the M function that maps portfolio profit correlations into exposure correlations, see eqn. (7).

Sensitivity to a particular security/risk factor. The partial derivative of the mean of aggregate exposure to risk factor f is:

$$\frac{\partial}{\partial \sigma^f} E(A) = \sum_j \sqrt{\frac{1}{2\pi}} \frac{\partial}{\partial \sigma^f} \sigma_j = \sum_j \sqrt{\frac{1}{2\pi}} \frac{\partial}{\partial \sigma^f} \sqrt{n_j' \beta \beta' n_j \sigma_f^2 + n_j' \tilde{\Sigma} n_j} \quad (\text{A4})$$

where $\tilde{\Sigma}$ is the covariance matrix of the residual after regressing a member's profit on the factor f . Eqn. (23) follows from this expression.

The partial derivative of $\text{std}(A)$ with respect to σ^f is:

$$\frac{\partial}{\partial \sigma^f} \text{std}(A) = \left(\frac{\pi - 1}{4\pi} \right) \frac{\sigma^f}{\sigma^A} \sum_{k,l} \left[2M'(\rho_{kl})B_{kl} + (M(\rho_{kl}) - \rho_{kl}M'(\rho_{kl})) \left(\frac{\sigma_l}{\sigma_k} B_{kk} + \frac{\sigma_k}{\sigma_l} B_{ll} \right) \right] \quad (\text{A5})$$

with

$$M'(\rho) = \frac{\frac{1}{2} + \arcsin(\rho)}{\pi - 1}, \quad B_{kl} := n_k' \beta \beta' n_l, \quad \text{and} \quad \beta = \text{cov}(R, r^f) / \text{var}(r^f). \quad (\text{A6})$$

Straightforward algebra yields eqn. (18).

B Construction and analysis of CrowdIx

This section describes how CrowdIx is constructed and discusses its properties. Let the total trade portfolio risk be defined as:

$$T = \text{tr}(\Sigma). \quad (\text{A7})$$

The crowding index, CrowdIx, is defined as the ratio of two standard deviations. The numerator is the standard deviation of aggregate exposure for actual member portfolios. The denominator is the standard deviation of a benchmark case where total risk T is unchanged, individual member portfolio risk is unchanged, but trades are re-allocated to the maximum extent possible to a single risk factor. The first-fit-descending binning algorithm is used to develop the maximum crowding benchmark case (Coffman, Garey, and Johnson, 1996). The procedure consists of the following steps.

1. Set the bin size equal to $T/2$.
2. Sort clearing members according to portfolio profit risk in descending order.
3. Assign members to bins sequentially. Start with the first bin and add member to the bin if it fits, otherwise try to assign it to the next bin. If unsuccessful, open a new bin and assign the

member to that bin (note that this is always possible since a member's risk cannot be larger than half total risk).

4. Make clearing members in first bin buyers and members in the second bin sellers.
5. Visit the remainder bin and re-assign its member's portfolio risk partially to bin one or bin two depending on which one is least full. Note that there cannot be more than one member in this bin as otherwise the sum across members' portfolio risk would exceed T . Re-assign the member's portfolio risk only so much as to make the bins equally full. This ensures that trade re-allocation is internally consistent; for every buyer there is a seller.

In the best case this procedure can reassign member portfolio risk perfectly so that the first two bins are filled perfectly (as is the case in the noncrowded case example discussed in section 2.3). In reality the bins will not be filled entirely. The space that remains is likely to be small as there is considerable variation in member risk and the largest members are assigned first.

Lower bound CrowdIx (lemma 1). Assume that the number of clearing members J is even; the case of odd J will be dealt with later. Total risk T achieves lowest aggregate exposure when distributed evenly across all clearing members and members trade in pairs in risk factors that are orthogonal. This creates a portfolio profit covariance matrix Σ with the following structure:

$$\Sigma = \begin{bmatrix} Q & -Q & 0 & 0 & \dots \\ -Q & Q & 0 & 0 & \\ 0 & 0 & Q & -Q & \\ 0 & 0 & -Q & Q & \\ \vdots & & & & \ddots \end{bmatrix}. \quad (\text{A8})$$

The standard deviation of aggregate exposure associated with Σ is calculated with the help of eqn. (6). Each block of the block-diagonal matrix Σ contributes

$$2 \left(1 - \frac{1}{\pi - 1} \right) Q, \quad (\text{A9})$$

because $M(-1) = 1/(\pi - 1)$. Therefore

$$\text{std}(A) = \sqrt{\left(\frac{J}{2}\right)\left(\frac{\pi - 1}{2\pi}\right)^2 \left(\frac{\pi - 2}{\pi - 1}\right) \frac{\text{tr}(\Sigma)}{J}}, \quad (\text{A10})$$

where A is the aggregate exposure associated with Σ . This expression simplifies to eqn. (8). The case of an odd number of clearing members is treated the same way as the case of $J - 1$ clearing members. The last two members are merged into a single member.

The portfolio profit covariance matrix for the benchmark portfolio of maximum crowded risk becomes:

$$\tilde{\Sigma} = \begin{bmatrix} Q & -Q & Q & -Q & \cdots \\ -Q & Q & -Q & Q & \\ Q & -Q & Q & -Q & \\ -Q & Q & -Q & Q & \\ \vdots & & & & \ddots \end{bmatrix}. \quad (\text{A11})$$

The standard deviation of the aggregate exposure associated with $\tilde{\Sigma}$ is calculated in the same way as the standard deviation for the actual portfolio. Each block

$$\begin{bmatrix} Q & -Q \\ -Q & Q \end{bmatrix} \quad (\text{A12})$$

contributes

$$2\left(1 - \frac{1}{\pi - 1}\right)Q, \quad (\text{A13})$$

because $M(-1) = 1/(\pi - 1)$. As there are $(J/2)^2$ such blocks relative to the $J/2$ blocks in the benchmark case of eqn. (A8), the result follows immediately, i.e.,

$$\text{std}(A)/\text{std}(\tilde{A}) = \sqrt{\frac{2}{J}}. \quad (\text{A14})$$

C Clearing members EMCF (December 2010)

ABN AMRO Clearing Bank N.V.	Numis Securities Ltd
BNP Paribas Securities Services S.A.	UBS Ltd
Bank of America Merrill Lynch	Barclays Capital Securities Ltd.
Citibank Global Markets and Citibank International	Alandsbanken Abp
JPMorgan Securities Ltd.	Alandsbanken Sverige AB
Goldman Sachs International	Amagarbanken A/S
Skandinaviska Enskilda Banken	Arbejdernes Landsbank A/S
KAS BANK N.V.	Avanza Bank AB
Parel S.A.	Carnegie Bank A/S
Deutsche Bank AG	Dexia Securities France
Citigroup	E-Trade Bank
MF Global UK Ltd	Eik Bank A/S
CACEIS Bank Deutschland	EQ Bank Ltd.
Danske Bank	Evli Bank Plc
ABG Sundal Coller Norge	FIM Bank Ltd.
DnB NOR Bank	GETCO Ltd.
Deutsche Bank (London Branch)	Handelsbanken
HSBC Trinkaus & Burkhardt	Jefferies International Ltd.
Istituto Centrale delle Banche Popolari Italiane SpA	Knight Capital Markets
Interactive Brokers	Lan & Spar Bank A/S
KBC Bank N.V.	Nordnet Bank AB
Nordea	Nomura International Plc
Swedbank	Nykredit A/S
Credit Agricole Cheuvreux	Pohjola Bank
Credit Suisse Securities (europe) Ltd	RBC Capital Markets
Morgan Stanley International Plc	Saxo Bank A/S
RBS Bank N.V.	Spar Nord Bank A/S
Instinet europe Ltd.	Sparekassen Kronjylland A/S
Morgan Stanley Securities Ltd.	

Source: [Zhu \(2011\)](#)

References

- Acharya, Viral V., Lasse H. Pedersen, Thomas Philippon, and Matthew P. Richardson. 2010. "Measuring Systemic Risk." manuscript, New York University.
- Adrian, Tobias and Markus K. Brunnermeier. 2011. "CoVaR." Manuscript, Princeton.
- Balakrishnan, Narayanaswamy and Chin-Diew Lai. 2009. *Continuous Bivariate Distributions*. New York: Springer.

- BIS-IOSCO. 2004. “Recommendations for Central Counterparties.” Manuscript, Bank for International Settlements and International Organization of Securities Commissions.
- . 2012. “Principles for Financial Market Infrastructures.” Manuscript, Bank for International Settlements and International Organization of Securities Commissions.
- Bisias, Dimitrios, Mark Flood, Andrew W. Lo, and Stavros Valavanis. 2012. “A Survey of Systemic Risk Analytics.” Manuscript, MIT.
- Brunnermeier, Markus K. and Patrick Cheridito. 2013. “Measuring and Allocating Systemic Risk.” manuscript, Princeton University.
- Brunnermeier, Markus K., Gary Gorton, and Arvind Krishnamurty. 2013. “Liquidity Mismatch.” In *Risk Topography: Systemic Risk and Macro Modeling*, edited by Markus K. Brunnermeier and Arvind Krishnamurty. Chicago: University of Chicago Press (forthcoming).
- Brunnermeier, Markus K. and Martin Oehmke. 2013. “Bubbles, Financial Crises, and Systemic Risk.” In *Handbook of the Economics of Finance, Volume 2, Part B*, edited by George M. Constantinides, Milton Harris, and Rene M. Stulz. Amsterdam: Elsevier.
- Coffman, Edward G., Michael R. Garey, and David S. Johnson. 1996. “Apprimation Algorithms for Bin Packing: A Survey.” In *Approximation Algorithms for NP-Hard Problems*, edited by Dorit S. Hochbaum. Boston: PWS Publishing Co.
- Duffie, Darrell. 2012. “Systemic Risk Exposures: A 10-by-10-by-10 Approach.” In *Risk Topography: Systemic Risk and Macro Modeling*, edited by Markus K. Brunnermeier and Arvind Krishnamurty. Chicago: University of Chicago Press (forthcoming).
- Duffie, Darrell and Haoxiang Zhu. 2011. “Does a Central Clearing Counterparty Reduce Counterparty Risk?” *Review of Asset Pricing Studies* 1:74–95.
- ESRB. 2012. “Annual Report.” Manuscript, European Systemic Risk Board.
- Hedegaard, Esben. 2012. “How Margins are Set and Affect Prices.” Manuscript, NYU.
- Hills, Bob, David Rule, Sarah Parkinson, and Chris Young. 1999. “Central Counterparty Clearing Houses and Financial Stability.” Manuscript, Bank of England Financial Stability Review 6: 122-134.
- Jorion, Philippe. 2007. *Value at Risk: The New Benchmark for Managing Financial Risk*. New York: McGraw-Hill.
- Cruz Lopez, Jorge A., Jeffrey H. Harris, Christophe Hurlin, and Christophe Pérignon. 2014. “Co-Margin: A System to Enhance Financial Stability.” Manuscript, HEC Paris.
- Menkveld, Albert J. 2013. “Systemic Liquidation Risk: Centralized Clearing, Margins, and the Default Fund.” Manuscript, VU University Amsterdam.

- Merton, Robert C. and André F. Perold. 1993. “Theory of Risk Capital in Financial Firms.” *Journal of Applied Corporate Finance* 6:16–32.
- Nabeya, Seiji. 1951. “Absolute Moments in 2-Dimensional Normal Distribution.” *Annals of the Institute of Statistical Mathematics* 3:2–6.
- Patrik, Gary, Stefan Bernegger, and Marcel Beat Rüegg. 1999. “The Use of Risk Adjusted Capital to Support Business Decision-Making.” manuscript, Swiss Reinsurance Company.
- Rosenbaum, S. 1961. “Moments of a Truncated Bivariate Normal Distribution.” *Journal of the Royal Statistical Society Series B (Methodological)* 23:405–408.
- Tarashev, Nikola, Claudio Borio, and Kostas Tsatsaronis. 2009. “The Systemic Importance of Financial Institutions.” manuscript, BIS Quarterly Review.
- Tasche, Dirk. 2000. “Risk Contributions and Performance Measurement.” manuscript, TU München.
- Tsanakas, Andreas. 2009. “To Split or Not to Split: Capital Allocation with Convex Risk Measures.” *Insurance: Mathematics and Economics* 44:268–277.
- Urban, Michael, Jörg Dittrich, Claudia Klüppelberg, and Rolf Stölting. 2003. “Allocation of Risk Capital to Insurance Portfolios.” *Bätter der DGVM* 26:389–406.
- Zhu, Siyi. 2011. “Is There a “Race to the Bottom” in Central Counterparties Competition?” Manuscript, De Nederlandsche Bank Occasional Studies 9/6.

Table 1: Summary statistics, overall and cross-sectional

This table presents summary statistics based on 1,434,946 Nordic ‘trade’ reports sent to the clearing house by 55 clearing members. The sample covers 242 stocks listed in Denmark, Finland, or Sweden. Each report contains a time stamp to the second, an anonymized clearing member ID, the symbol of a stock, price, a buy or sell indicator, and the size of the transaction in terms of shares. The sample period consists of 228 trading days. It starts on October 19, 2009 and ends at September 9, 2010. The sum of signed volume is zero for each stock.

	Mean	Std	Min	Median	Max
<i>Panel A: Overall summary statistics</i>					
Daily number of reports	6,293.6	699.0	1,135.0	6,426.5	7,663.0
Daily volume (in mln shares)	160.9	42.1	8.1	155.5	342.4
Daily volume (in mln euro)	1,809.8	475.1	272.4	1,762.3	3,649.6
Volume per report (in 1000 shares)	25.6	114.1	0.0	2.6	18,631.8
Volume per report (in 1000 euro)	287.6	1,067.6	0.0	36.1	142,271.3
<i>Panel B: Cross-sectional summary statistics, based on clearing-member averages</i>					
Daily number of reports	114.4	143.7	0.0	64.9	736.4
Daily volume (in mln shares)	2.9	4.2	0.0	0.7	20.8
Daily volume (in mln euro)	32.9	46.9	0.0	7.8	222.4
<i>Panel C: Cross-sectional summary statistics, based on stock averages</i>					
Daily number of reports	26.0	21.9	0.0	20.6	84.2
Daily volume (in mln shares)	0.7	1.6	0.0	0.1	14.2
Daily volume (in mln euro)	7.5	14.6	0.0	0.9	124.0

Table 2: CCP risk sensitivity to security/risk factor


This table shows how sensitive CCP risk is to a particular security or risk factor. CCP risk is measured by the aggregate margin that it would collect when accounting for crowded risk, i.e., Margin(A). Sensitivity could be used to identify securities or risk factors that members crowded on. Sensitivity is reported in two ways. The table shows what the change in Margin(A) is when one percentage point is added to the daily volatility of a particular risk factor. The further reports the elasticity of Margin(A) to change in the risk in the risk factor. Three days were picked from the sample: the median-CrowdIx day and the two days for which the CCP charged highest aggregate margin. The risk factors considered are the market return (based on the STOXXNordic30), the Nokia stock return, and the telecom sector return (based on STOXXTelecom).

	Date	CrowdIx	Risk factor	Margin(A) (million euro)	Δ Margin(A) on $\Delta\sigma_f=0.01$ (million euro)	Elasticity
Median CrowdIx day	Jul 29, 2010	0.46	Market	128	81	0.91
			Nokia	128	11	0.15
			Telecom	128	46	0.46
Greek bailout	May 10, 2010	0.62	Market	747	307	0.98
			Nokia	747	27	0.14
			Telecom	747	298	0.83
Nokia reports Q1	Apr 26, 2010	0.72	Market	644	116	0.19
			Nokia	644	147	1.05
			Telecom	644	-2	-0.00

Figure 1: Crowded vs. noncrowded risk

This figure illustrates two extreme cases: crowded trades in panel A and noncrowded trades in panel B. The n_j vector represents the position of member j after trading. The graphs depict the result after two trades: member 1 traded with member 2 and member 3 traded with member 4.

Panel A: Crowded trades



Panel B: Noncrowded trades




Figure 2: M function to map profit correlation into exposure correlation

This figure illustrates $M(\cdot)$ that maps correlations in member portfolio profits (X_j) to correlations in member portfolio exposures ($E_j = -\min(0, X_j)$).




Figure 3: Density aggregate exposure with four members ($J = 4$)

This figure graphs the density estimate (based on 100,000 draws) of the aggregate exposure when trades crowd and when they do not. The estimate is based on a simple setting. Two securities are available for trade. Their one-period profits are distributed as two independent standard normal variables. Four members implement two trades. In the case of crowded trades member 1 sells security 1 to member 2 and member 3 sells the same security to member 4. The aggregate exposure is defined as $A = \sum_j -\min(X_j, 0)$, where X_j is the profit on member j 's trade portfolio. The noncrowded setting is the same except that members 3 and 4 trade security 2 instead of 1.





Figure 4: Aggregate daily margin: actual vs. model-implied Margin(A)

This figure plots both the sum of actual margins and the model-implied Margin(A). The latter accounts for the CCP risk associated with crowded trades. Panel A graphs these series for the full sample period. Panel B zooms into the period with highest realized aggregate margin and overlays CrowdIx, a measure of crowded risk.

Panel A: Full sample period



Panel B: Zoomed into mid-April through mid-May 2010





Figure 5: CCP aggregate exposure distribution

This figure plots the simulated probability density function of CCP aggregate exposure for May 10, 2010. On this day the CCP collected most margin of all days in the sample. It was the period right after the first Greek bailout. The crowding index was particularly high on that day, i.e., $CrowdIx=0.62$. The aggregate exposure distribution is based on 100,000 simulations of security returns that are each assumed to be normally distributed. To illustrate the enhanced right-tail risk due to crowding, the plot also contains the distributions for the median- and minimum-CrowdIx days as a benchmark. The aggregate exposure for the benchmark days was multiplied by the ratio of aggregate CCP margin on the ‘Greek-bailout’ day and the benchmark day in order to make them comparable in terms of the shape of the distribution. The exhibit below the graph reports the 90%, 99%, and the 99.9% quantile for each distribution.




	Date	CrowdIx	Q(0.90) (million euro)	Q(0.99) (million euro)	Q(0.999) (million euro)
Greek bailout	May 10, 2010	0.62	267	405	510
Median CrowdIx day benchmark	Jul 29, 2010	0.46	196	289	360
Min CrowdIx day benchmark	Nov 12, 2009	0.31	175	233	277

Figure 6: Scatter plot of actual margin vs. model-implied margin


This figure contains three scatterplots of the margin that members actually posted versus the model-implied margin, Margin(A). The plots correspond to three days in the sample: the median-CrowdIx day and the two days for which the CCP charged highest aggregate margin. The exhibit below the scatterplots contains the ten largest positions in the trade portfolio of a member in the top-left corner and a member in the bottom-right corner.

Panel A: A representative day, i.e., CrowdIx at median level (July 29, 2010)




Clearing member 41				Clearing member 6			
Stock	NetPos (mln €)	AbsNetPos (mln €)	AbsNetPos (%)	Stock	NetPos (mln €)	AbsNetPos (mln €)	AbsNetPos (%)
ER	23.1	23.1	13.7	FSPAA	17.1	17.1	8.3
SHBA	14.5	14.5	8.6	ASSAB	-12.7	12.7	6.2
NOVNB	-11.0	11.0	6.5	SEBA	12.4	12.4	6.1
NBH	10.1	10.1	6.0	NBH	10.6	10.6	5.2
HMB	8.9	8.9	5.3	VWS	10.4	10.4	5.1
FSPAA	-7.2	7.2	4.3	SSABA	-7.5	7.5	3.6
SAND	6.5	6.5	3.8	MEO1V	-6.4	6.4	3.1
VOLB	5.6	5.6	3.3	FUM1V	6.2	6.2	3.0
BOLI	5.3	5.3	3.2	SAND	-5.6	5.6	2.7
ASSAB	-4.5	4.5	2.7	STERV	-5.3	5.3	2.6

Panel B: Greek bailout (May 10, 2010)



Clearing member 41				Clearing member 12			
Stock	NetPos (mln €)	AbsNetPos (mln €)	AbsNetPos (%)	Stock	NetPos (mln €)	AbsNetPos (mln €)	AbsNetPos (%)
ER	-99.9	99.9	13.5	SKFB	22.7	22.7	5.4
NOKI	-65.7	65.7	8.9	NOKI	20.7	20.7	4.9
HMB	-48.3	48.3	6.5	NOVNB	18.7	18.7	4.5
NBH	-35.6	35.6	4.8	NBH	18.3	18.3	4.4
ATCOA	-31.7	31.7	4.3	SHBA	15.5	15.5	3.7
SAND	-29.4	29.4	4.0	GETIN	15.4	15.4	3.7
TLS1V	-28.8	28.8	3.9	ER	15.4	15.4	3.7
FSPAA	-28.3	28.3	3.8	ABBN	-14.7	14.7	3.5
SHBA	-22.7	22.7	3.1	ZEN	-14.1	14.1	3.4
ZEN	-20.6	20.6	2.8	SAND	13.8	13.8	3.3

Panel C: Nokia reports Q1 (April 26, 2010)



Clearing member 41				Clearing member 12			
Stock	NetPos (mln €)	AbsNetPos (mln €)	AbsNetPos (%)	Stock	NetPos (mln €)	AbsNetPos (mln €)	AbsNetPos (%)
NOKI	-84.7	84.7	20.7	VOLB	35.7	35.7	12.6
ER	64.8	64.8	15.8	TLS1V	-17.4	17.4	6.2
FUM1V	-39.2	39.2	9.6	MAERS	-15.2	15.2	5.4
NDA1V	-31.7	31.7	7.7	ABBN	-13.2	13.2	4.7
VOLB	16.2	16.2	4.0	ALFA	-9.7	9.7	3.4
HMB	15.5	15.5	3.8	VWS	-9.2	9.2	3.2
STERV	15.3	15.3	3.7	TRELB	-9.0	9.0	3.2
TLS1V	9.8	9.8	2.4	TEL2B	-8.7	8.7	3.1
OUT1V	-8.9	8.9	2.2	ASSAB	6.8	6.8	2.4
SEN	-8.3	8.3	2.0	BOLI	6.3	6.3	2.2