Central Counterparty Exposure in Stressed Markets

Wenqian Huang [a], Albert J. Menkveld [b], Shihao Yu [c]

a Bank for International Settlements, 4051 Basel, Switzerland;
b School of Business and Economics, Vrije Universiteit Amsterdam, 1081 HV Amsterdam, Netherlands;
c Tinbergen Institute, 1082 MS Amsterdam, Netherlands
Contact: wenqian.huang@bis.org* Corresponding author

Published in Management Science (2024) 67 (6), 3596-3617. DOI: 10.1287/mnsc.2020.3601

Abstract: Time is valuable, particularly in stressed markets. Because central counterparties (CCPs) have become systemically important, we need to understand the dynamics of their exposure toward clearing members at high frequencies. We track such exposure and decompose it, yielding the following insights. The composition of CCP exposure is fundamentally different in the tails. At extreme levels or during rapid increases, there is elevated crowding. This is the result of clearing members all concentrating their positions on a single security or a particular portfolio, which is desirable if motivated by hedging but worrying if due to speculation.
Keywords: Central counterparties (CCPs), crowding risk, market stress

1. Introduction

Regulators are worried about central counterparty (CCP) risk management in fast markets. Sudden extreme price dislocations ("flash crashes")[1] coupled with superhuman trading speeds could have systemic consequences. If traders are unable to deliver on their trades, then CCPs become liable for their losses because a CCP effectively insures the counter parties in these trades. The margins posted by these de-faulting traders might not be sufficient to cover these losses. A recent example is the 2018 failure of a Nasdaq clearing member, when losses swallowed up two-thirds of the default fund.[2] A mutualized loss of this kind might itself trigger further defaults, in which case the event could become potentially systemic. State-of-the-art risk management of CCPs is, therefore, of first-order importance. CPMI-IOSCO (2017) emphasizes the need for monitoring intraday CCP exposure:

> Adverse price movements, as well as participants building larger positions through new trading (and settlement of maturing trades), can rapidly increase a CCP's exposures to its participants. This exposure can relate to intraday changes in both prices and positions. For the purposes of addressing these and other forms of risk that may arise intraday, a CCP should address and monitor onan ongoing basis. . . (CPMI-IOSCO 2017, p. 32)

In this paper, we propose a way for CCPs to monitor their exposure on an intraday basis, with a focus on stressed markets. In such markets, trading is likely to be fast-paced, and data therefore stream at extreme speeds. The approach should be able to cope with such "big data" challenges. More importantly, the monitoring should yield valuable economic insights that generate an understanding of "what just happened" and potentially guide interventions. We turn to the academic literature to formulate hypotheses to guide our high-frequency analysis of CCP exposure. Several studies have identified a fire-sale channel as the root cause of price dislocations. The narrative is as follows. During normal times, arbitrageurs smooth prices by trading against pricing errors (thereby essentially engaging in market making). Suppose that, at some point, a critical mass of traders crowds into a single risk factor. That is, their portfolio positions are very similar, say, having long positions of a book-to-market or size-based portfolio. 3If these positions suddenly experience a significant loss, then arbitrageurs face high variation margin calls (to mark-to-market their positions). If these arbitrageurs are capital-constrained, then they might be forced to free up capital by selling some of their positions. This selling pressure might trigger trades at fire-sale prices, thus leading to more losses, triggering further selling, etc. (Shleifer and Vishny 1997, Gromb and Vayanos 2002, Brunnermeier and Pedersen 2009). Perhaps the most prominent example of such dynamic is the "quant meltdown," where the arbitrageurs are hedge funds and the portfolios are indeed factor-based portfolios (Khandani and Lo 2007, 2011). This type of fire-sale channel implies that extreme exposures coincide with elevated crowding, price crashes, and volatility spikes. With these motivations in mind, let us now discuss in more detail how we proceed in the rest of this paper. We develop an approach for tracking and decomposing CCP exposure intradaily. The exposure

measure is based on the tail risk of losses in an oncoming period, aggregated across all clearing members (Duffie and Zhu 2011, Menkveld 2017).[4] The measure relies on analytical results that are all straightforward to compute. It further allows for decomposition across clearing members or securities.

We implement the approach on a sample of high-frequency CCP data to test three hypotheses on CCP exposure in stressed markets. We define such markets for a CCP as those where either the CCP exposures are at the highest *levels* or the CCP exposure *changes* are extremely high.[5] The three hypotheses pertain to the following questions:

1. Are extreme increases in CCP exposure driven by the same factors as regular exposure changes, or does one see, for example, elevated crowding?
2. Is the same true for extreme *levels* as opposed to extreme *changes*? Again, is crowding a larger part of it?
3. Finally, when comparing CCP exposure for these extreme levels relative to normal levels, is the relative contribution of clearing-member house accounts higher relative to client accounts? If so, then this is worrisome, as clearing members are typically highly leveraged financial intermediaries and therefore less able to absorb large shocks.

Our empirical analysis of these questions is based on a high-frequency 2009–2010 sample of a European CCP: the European Multilateral Clearing Facility (EMCF). The EMCF was the largest equity CCP in Europe and later merged with the Depository Trust & Clearing Corporation (DTCC) in the United States to become the world's largest equity CCP. Counterparty risk arises in equity trading because the settlement of a trade typically occurs three days after it is concluded. Therefore, a trade is like a three-day forward contract between a buyer and a seller. Counterparty risk then pertains to the possibility that one side might default in this period. Admittedly, analysis of a CCP that insures credit default swaps or interest rate swaps would have been more relevant in terms of systemic risk, but disaggregated CCP data are extremely hard to come by (see literature review here). Therefore, the application to actual CCP data could in and of itself be considered a contribution.[6]

There are several key findings. First, CCP exposure changes, on average, are almost entirely driven by changes in the positions of clearing members due to their trading. However, when zooming in on extreme exposure increases, security volatility and position crowding start to contribute substantially. For the top 100 increases, they collectively contribute 30%, where volatility contributes 13% and crowding 17%.

Second, we find a similar result when comparing the full sample with the subsample of high exposure *levels*—more crowding in the latter. More specifically, CCP exposure concentrates on a few clearing members and a smaller set of risk factors. For example, comparing the full sample with the top 1% subsample, the contribution of the largest five members increases from 28% to 47%. The contribution of the largest principal component across all risk factors increases from 7% to 42%.

Third, it is not true that, at high exposure levels, the CCP is relatively more exposed to house accounts. There is only a modest increase from 67% to 70% when comparing the full sample to the top 1%. However, we do find stronger concentration *within* the set of house accounts. A large share of the total house-account exposure originates from just a few clearing members.

In sum, the findings collectively suggest that stronger crowding/concentration characterizes CCP exposure both for large exposure increases and at extremely high exposure levels. There is, however, only a minor increase in the contribution of house accounts at these high levels.

One additional finding worth emphasizing is that idiosyncratic events can severely impact CCP exposure. For example, a disappointing earnings announcement by Nokia at noon on April 22, 2010 caused its share price to fall by about 15% in the minutes afterward, leading to an exposure increase of almost 16 times its average size. The decomposition of the exposure change shows that volatility is the largest component causing the jump. The exposure jump, however, was only a relatively small part of the extremely large CCP exposure increase that day. Most of this appears to be caused by clearing members increasing their Nokia position, either long or short, during heavy trading in the afternoon. (Note that this is a nontrivial finding, as the large volumes could have been due to traders reducing their positions after observing elevated volatility.[7]) The decomposition of exposure changes further reveals a substantial contribution of the crowding component that day. Members tilted their portfolio toward Nokia. Altogether, the Nokia example neatly illustrates the paper's main finding that volatility and crowding become important components of CCP exposure in the tails.

Our paper contributes to a rapidly expanding empirical literature on central clearing. CCP trade data disaggregated across members are scarce. Proprietary daily data have been used to compare CCP exposure to the margins collected (Jones and Perignon 2013, Lopez et al. 2017, Menkveld 2017). Duffie et al. (2015) analyze a snapshot of bilateral exposures on uncleared credit default swaps to assess the netting efficiency potential of central clearing. Event studies on CCP introductions yield insights into how trading is affected (Loon and Zhong 2014, 2016; Menkveld et al. 2015; Benos et al. 2020). We contribute to this literature by proposing an approach to monitoring CCP exposure intradaily along with an economically

motivated decomposition. The methodological contribution relative to Menkveld (2017) is that we decompose that author's exposure measure to diagnose the nature of exposure changes. Although Menkveld (2017) discusses how to decompose across clearing members, the author does not decompose exposure *changes* into changes in the various variables that enter the exposure computation (e.g., changes in volatility, changes in return correlations, or changes in positions). Such a decomposition is needed to test the first hypothesis.[8]

The paper contributes to a nascent literature on CCP systemic risk. Capponi et al. (2019) analyze the endogenous buildup of asset concentration due to central clearing. Amini et al. (2020) study the effects of central clearing on a network of inter-bank liabilities and show that systemic risk can be reduced under certain equity and guaranty fund policy of the CCP. In a similar setting, Amini et al. (2015) investigate partial netting that accounts for knock-on effects and asset liquidation effects. Glasserman et al. (2016) compare margining in dealer markets and a centrally cleared market. Menkveld (2016) endogenizes the fire-sale premium that a CCP will have to pay in the catastrophic state in which a critical mass of members default and liquidity supply is thus impaired.[9]

The rest of the paper is organized as follows. Section 2 formalizes and motivates the three overriding hypotheses. Section 3 presents the approach to monitoring and decomposing CCP exposure. Section 4 describes the data and discusses implementation issues. Section 5 presents the empirical results of testing the three hypotheses. Section 6 concludes.

2. Hypotheses

This section develops three hypotheses that will be taken to the data. Each hypothesis is stated formally and then followed by a motivation.

Hypothesis 1. *The drivers of CCP exposure changes are different in the (right) tail.*

CCP exposure changes can be driven by a variety of factors that are either price-related (e.g., volatility or correlation) or trade-related (i.e., trade causes member positions to change). We expect the latter to dominate CCP exposure changes in normal times. However, we conjecture that turbulent periods are characterized by elevated volatility and lots of trading. The strong positive correlation of volatility and trading volume is a well-known stylized fact in the microstructure literature (Jones et al. 1994).

The intense trading at times of extremely high volatility does *not* necessarily imply that CCP exposure increases rapidly. A sudden volatility increase might actually trigger traders to *reduce* their existing positions to contain risk. Such trading benefits a CCP because it reduces its exposure.

On the other hand, a volatility shock might lead to (more) speculation, in which case traders increase their positions. Heterogeneity in beliefs or in signals might generate such stronger position taking (Kim and Verrecchia 1994). Or, in a more recent paper, Crego (2019) proposes a channel by which risk-averse informed traders strategically wait to trade on their (idiosyncratic) signal until the arrival of a public signal which removes significant uncertainty. Either way, member positions would increase in magnitude, and CCP exposure rises as a result.

An even more worrisome channel that could cause high volatility and fast trading is a so-called self-reinforcing fire-sale channel. For example, financially constrained arbitrageurs (hedge funds, sell-side banks, high-frequency traders, etc.) hit by adverse price shocks might have to quickly liquidate their large positions and thereby cause transitory price shocks (Shleifer and Vishny 1997, Gromb and Vayanos 2002, Brunnermeier and Pedersen 2009).

Such liquidations would not be a concern if these arbitrageurs had diverse positions (Wagner 2011). This is not the case, however, if these traders followed similar trading strategies and their portfolios thus crowd on a small set of risk factors/portfolios (Stein 2009). In such a scenario, there might not be enough cash-in-the-market to liquidate these positions, and markets have to clear at fire-sale prices. A prominent example is the "Quant Meltdown" of 2007, when quantitative equity market-neutral hedge funds crowded on similar trading strategies and made record losses (Khandani and Lo 2007, 2011). The risk of a CCP finding itself in such a scenario is particularly high when there is substantial crowding in its members' portfolios.

To test for such crowding in stressed markets, one needs to be able to decompose changes in CCP exposure into price- and trade-related components. One of the trade-related components should then be crowding across clearing members.

Hypothesis 2. *The structure of CCP exposure levels is different in the (right) tail.*

Hypothesis 2 restates Hypothesis 1 but this time in terms of *levels* instead of *changes*. The reason to also study whether there is, for example, elevated crowding for extreme exposure levels is derived from studies on historical CCP failures. Bignon and Vuillemey (2020) study the 1974 failure of the Paris Commodity Clearing House. They show that in a year starting from November 1973, the position of the largest clearing member rose from 9% of the total open position in sugar futures to 56% of it. Another example is the 1987 failure of the Hong Kong Futures Guarantee Corporation,

where at the point of failure the largest four members had accumulated 80% of the short position in all contracts (Cox 2015).

If crowding is prominent, understanding what causes the crowding requires one to be able to decompose the exposure level across clearing members and across risk factors. The reason is that strong crowding could occur when outstanding positions are held by only a few clearing members. The CCP failures in Paris and Hong Kong are examples of such crowding. There is, however, a more opaque way for there to be elevated crowding. In the extreme case, all clearing members contribute equally to CCP exposure, but they crowd on a single risk factor, such as a particular security or portfolio. The 2007 Quant Crisis is an example of such "risk-factor crowding." Let us turn to a simple example to make this distinction between the two types of crowding as clear as possible.

Suppose there are four clearing members and two assets with independently distributed payoffs. First consider the baseline case of member 1 and member 2 having traded one unit of asset A and therefore having open and opposite positions in this asset. Suppose the same holds for members 3 and 4 in asset B. In this baseline case, there is no crowding. Let us now consider the two polar cases of crowding. An example of perfect *member* crowding is when members 1 and 2 trade as in the baseline case, and members 3 and 4 refrain from trading. The reason is that there is concentration in CCP exposure, as only two members contribute. For an example of perfect risk-factor crowding, consider again the baseline case but now with members 3 and 4 also trading one unit of asset A. Note that in this case, all clearing members contribute equally to CCP exposure, yet there is perfect crowding. In both cases, there is a strong correlation in portfolio returns across clearing members, which increases the expected aggregate loss and therefore CCP exposure (Menkveld 2017, section 1.5).

When testing the second hypothesis, it is desirable to measure how much crowding contributes to CCP exposure and whether such crowding is member or risk-factor crowding. In Section 3.3, we discuss in detail how to measure crowding with the proposed CCP exposure measure.

Hypothesis 3. *The relative contribution of house accounts to CCP exposure increases in the (right) tail.*

The third hypothesis focuses on the two types of clearing-member accounts: house accounts and client accounts. House accounts capture the trades that clearing members do for their own books, whereas client accounts register their cleared trades on behalf of clients. It is worth decomposing CCP exposure across these two types of accounts, as one could argue that CCP exposure to house accounts carries more risk. Clearing members are often highly leveraged financial intermediaries whose trading is unlikely to be pure hedging. For example, they often engage in market making to absorb temporary order imbalances. Therefore, more exposure to house accounts at times of high CCP exposure is worrisome. Testing the third hypothesis will show whether this is the case.

3. Approach

This section presents an approach to monitoring CCP exposure intradaily. It is based on the framework proposed by Duffie and Zhu (2011) and extended by Menkveld (2017) to include tail risk and crowding. CCP exposure is essentially a measure that is based on the distribution of losses in clearing-member accounts for the oncoming period. We study the value at risk (VaR) for these losses following Menkveld (2017).[10] We first present the exposure measure in detail, then show how one could decompose exposure *change* needed for testing the first hypothesis, and finally present the decomposition of exposure *level* which is needed for testing the second and third hypotheses.

3.1. The CCP Exposure Measure: A VaR of Aggregate Loss

Consider the case of a single CCP, I securities, and J clearing members (or traders; the two terms will be used interchangeably). P_t is an $I \times 1$ vector consisting of current security prices. R_t is an $I \times 1$ vector that contains next period's security returns. R_t is assumed to be normally distributed:[11] $R_t \sim \mathbf{N}(0, \Omega_t)$, where Ω_t is the $I \times I$ covariance matrix of security returns. Let $n_{j,t}$ be the $I \times 1$ vector of member j's current positions expressed in euro. The portfolio return in euro for the member in the next period is then a scalar $X_{j,t}$, where $X_{j,t} = n'_{j,t} R_t$.

Collect all $n_{j,t}$ into an $I \times J$ matrix N_t, which thus becomes the (euro) position matrix of all members. Collect all $X_{j,t}$ into the $J \times 1$ vector X_t, which thus becomes the future return vector for all members, where $X_t = N'_t R_t$. Because X_t is linear in R_t, X_t is normally distributed: $X_t \sim \mathbf{N}(0, \Sigma_t)$, where $\Sigma_t = N'_t \Omega_t N_t$ is the $J \times J$ covariance matrix of (euro) portfolio returns.

As a CCP is exposed to losses, define

$$L_{j,t} = -\min(0, X_{j,t}) \tag{1}$$

as the loss in member j's portfolio. Then, aggregate loss A_t is as follows:

$$A_t = \sum_j L_{j,t}. \tag{2}$$

Duffie and Zhu (2011) propose to base CCP exposure on the *mean* aggregate loss,

$$\mathrm{E}(A_t), \tag{3}$$

and derive an analytical expression for it which suffices for their analysis of netting efficiency. Menkveld (2017) considers the VaR of aggregate loss a more appropriate measure for CCP exposure and refers to it as *ExpCCP*. Following standard practice and maintaining tractability, Menkveld (2017) uses the delta-normal method to compute the VaR:

$$ExpCCP_t \equiv \mathrm{VaR}(A_t) = \mathrm{E}(A_t) + \alpha \mathrm{var}(A_t)^{\frac{1}{2}}, \qquad (4)$$

where α is a parameter that needs to be calibrated. We follow Menkveld (2017) and use *ExpCCP* as our exposure measure. In Appendix A, we list all the results needed to compute *ExpCCP*.

3.2. Decomposition of CCP Exposure Change

The first hypothesis states that the drivers of CCP exposure change are different in the tail. As discussed in the hypothesis section, sudden extreme CCP exposure increases might be driven by volatility shocks and crowding in addition to position changes. To test such a hypothesis, one needs to decompose exposure changes and verify to what extent volatility and crowding contribute a larger part in the tail.

We propose to decompose exposure changes based on a relatively straightforward one-factor-at-a-time (OFAT) approach (Daniel 1973). The underlying factors will consist of price-related factors and trade-related factors. Price-related factors include security return volatility, correlation, and price level. Trade-related factors are member positions and crowding across members. The remainder of this subsection describes the approach in detail.

Let us start by writing $ExpCCP_t$ as defined in (4) as a function of the underlying variables:

$$ExpCCP_t = f(\Sigma_t). \qquad (5)$$

To arrive at a meaningful decomposition across factors, we use the following two insights:

1. Following the financial econometrics literature, we decompose covariance matrices into their diagonal and off-diagonal components (Bollerslev 1990, Engle 2002):

$$\Psi_t = D_{\Psi_t} R_{\Psi_t} D_{\Psi_t}, \qquad (6)$$

where D_{Ψ_t} is a diagonal matrix with $\psi_{ii,t}$ as the ith diagonal element, and R_{Ψ_t} is the correlation matrix associated with the covariance matrix Ψ_t. This decomposition will turn out to be useful for identifying correlation effects in security returns and crowding across members.

2. Σ_t is itself a function of "deeper" variables:

$$ExpCCP_t = f(\Sigma_t) = f\left(N_t \Omega_t N_t'\right) = f\left(\Omega_t, P_t, \tilde{N}_t\right), \qquad (7)$$

where the variables are the covariance matrix of security returns Ω_t, the price level P_t, and the member portfolio matrix $\tilde{N}_t$ expressed in terms of the number of securities (as opposed to N_t, which is expressed in euro). The reason for using $\tilde{N}_t$ instead of N_t is to be able to pull out a price-level effect when considering the change from N_{t-1} to N_t.

Combining (6) and (7) yields

$$\begin{aligned} ExpCCP_t &= f(\Sigma_t) = f\left(D_{\Sigma_t} R_{\Sigma_t} D_{\Sigma_t}\right) \\ &= f\left(D_{\Sigma_t}\left(D_{\Omega_t}, R_{\Omega_t}, P_t, \tilde{N}_t\right),\right. \\ &\quad \left. R_{\Sigma_t}\left(D_{\Omega_t}, R_{\Omega_t}, P_t, \tilde{N}_t\right)\right), \end{aligned} \qquad (8)$$

which expresses $ExpCCP_t$ in terms of price-related variables $(D_{\Omega_t}, R_{\Omega_t}, P_t)$ and trade-related variables $(\tilde{N}_t)$. The OFAT decomposition changes these variables sequentially from their value at $t-1$ to their value at t. The sequencing matters (as will be discussed in depth at the end of this subsection), and we pick the baseline sequencing motivated by the following principles:

- We first change price variables and then change trade variables. The reason for this sequencing is that it identifies a pure price effect. In other words, the price components communicate what the CCP exposure change would have been had members' portfolios not changed.
- Changes in idiosyncratic volatility precede changes in correlations. In other words, we first consider changes in the diagonal and then changes in the off-diagonal of a covariance matrix. This approach makes interpretation of the components straightforward: changes in variances become pure in the sense that they are evaluated while keeping correlations constant.

These principles therefore suggest the following baseline OFAT decomposition:

$$\begin{aligned} &\Delta ExpCCP_t \\ &= f\left(D_\Sigma\left(\overset{1}{D_{\Omega_t}}, \overset{2}{R_{\Omega_t}}, \overset{3}{P_t}, \overset{4}{\tilde{N}_t}\right), R_\Sigma\left(\overset{1}{D_{\Omega_t}}, \overset{2}{R_{\Omega_t}}, \overset{3}{P_t}, \overset{5}{\tilde{N}_t}\right)\right) \\ &\quad - f\left(D_\Sigma\left(D_{\Omega_{t-1}}, R_{\Omega_{t-1}}, P_{t-1}, \tilde{N}_{t-1}\right),\right. \\ &\quad \left. R_\Sigma\left(D_{\Omega_{t-1}}, R_{\Omega_{t-1}}, P_{t-1}, \tilde{N}_{t-1}\right)\right), \end{aligned} \qquad (9)$$

where the sequencing is illustrated by the (red) numbers on top of the various variables. The decomposition yields five components. For example, the first component, $RetVola_t$, is computed as follows:[12]

$$\begin{aligned} RetVola_t &= f\left(D_\Sigma\left(\boldsymbol{D_{\Omega_t}}, R_{\Omega_{t-1}}, P_{t-1}, \tilde{N}_{t-1}\right),\right. \\ &\quad \left. R_\Sigma\left(\boldsymbol{D_{\Omega_t}}, R_{\Omega_{t-1}}, P_{t-1}, \tilde{N}_{t-1}\right)\right) \\ &\quad - f\left(D_\Sigma\left(D_{\Omega_{t-1}}, R_{\Omega_{t-1}}, P_{t-1}, \tilde{N}_{t-1}\right),\right. \\ &\quad \left. R_\Sigma\left(D_{\Omega_{t-1}}, R_{\Omega_{t-1}}, P_{t-1}, \tilde{N}_{t-1}\right)\right), \end{aligned} \qquad (10)$$

which captures the contribution of volatility change.

We list the five components here and discuss each of them in detail. Note that the numbering corresponds to the red numbers in (9).

3.2.1. Price Components.

1. *RetVola*: The impact of a change in return volatility on CCP exposure change. This effect captures the well-known empirical fact that volatility is time-varying (commonly referred to as generalized autoregressive conditional heteroskedasticity (GARCH) or stochastic volatility in the financial econometrics literature).
2. *RetCorr*: The additional impact of a change in the correlations of security returns on CCP exposure change. The time-varying nature of such correlations is another well-known empirical fact and can be identified, for instance, through a dynamic conditional correlation (DCC) model (Engle 2002).
3. *PrLevel*: The additional impact of a change in the price level of securities. This effect is entirely due to covariance matrices being defined in relative terms (i.e., they are based on relative returns as opposed to euro returns). For example, a covariance matrix might not have changed in the interval, but if price levels dropped, then CCP exposure dropped because the latter is defined in terms of the euro. Such effect is picked up by *PrLevel*.

3.2.2. Trade Components.

4. *TrPosition*: The additional impact of trades. These trades might expand or reduce members' existing positions. CCP exposure therefore does not necessarily increase after new trades. It declines if their overriding effect was to reduce members' outstanding positions.
5. *TrCrowding*: The additional impact due to changes in the correlations of member portfolio returns, beyond what is caused by changes in the correlations of security returns (as that change is captured by *RetCorr*). *TrCrowding* is therefore solely the result of position changes due to trading. If these portfolio correlations increase (in magnitude), then CCP exposure increases.[13]

In Appendix C, we illustrate the decomposition of exposure changes by presenting a simple example. We discuss how the various components change when changing either price- or trade-related variables.

3.2.3. In-Depth Discussion of Component Identification.

The identification of the components that drive exposure change in the tail ($\Delta ExpCCP$) deserves a more thorough discussion. Such identification is nontrivial for essentially two reasons. First, exposure is a *nonlinear* function of the various variables (e.g., security-return correlations, member positions). Therefore, a decomposition cannot assign changes uniquely to the various components. To illustrate this point, consider the following two simple functions: the linear function, $f(x,y) = x + y$, and the nonlinear one, $g(x,y) = xy$. For f, any change can be uniquely decomposed as $\Delta f = \Delta x + \Delta y$.

For g, however, any change is nontrivial to write in terms of Δx and Δy. The approach we picked is OFAT, which can decompose Δg in two ways. One can first change x and then y, yielding $[g(x+\Delta x,y)-g(x,y)] + [g(x+\Delta x,y+\Delta y)-g(x+\Delta x,y)] = [(\Delta x)y] + [(x+\Delta x)\Delta y)]$, whereby the terms in square brackets correspond to the contribution of x and y, respectively. Alternatively, one can first change y and then x, yielding $[g(x,y+\Delta y)-g(x,y)] - [g(x+\Delta x,y+\Delta y)-g(x,y+\Delta y)] = [x\Delta y] + [\Delta x(y+\Delta y)]$, where the first term gets assigned to y and the second to x. In summary, the OFAT decomposition in this case is given in Table 1.

Note that the interaction term $(\Delta x)(\Delta y)$ gets assigned to the component that is updated later. In economic applications, there might be reasonable arguments to pick a reasonable baseline sequencing (as in our case), but it is always useful to consider all possible sequences and report lower and upper bounds to the size of each component. The wedge between the two bounds tells the researcher to what extent the decomposition critically depends on the sequencing that the researcher picked.

Second, if the objective is to study how time series Δf and Δg are driven by Δx and Δy, the identification of components for both Δf and Δg suffers from nonzero correlations between the underlying variables. Consider the case of a perfect correlation between Δx and Δy in the time series. Then any function-value changes are driven by changes in *both* x and y simultaneously. The individual contribution of each variable therefore cannot be (statistically) determined. This is a genuine feature, not a flaw, of our method. The comovement in x and y is itself an important property of the system. We will return to this issue in Section 5.1, where we study covariation across exposure components.

3.3. Decomposition of CCP Exposure Level

Testing the second and third hypotheses requires a decomposition of CCP exposure levels (as opposed to exposure changes). To test whether there is more crowding at higher exposure levels, it is desirable to decompose CCP exposure across members and across securities. If one finds more concentration either across members (member crowding) or across securities (risk-factor crowding), then there is elevated crowding as discussed in Section 2.

Table 1. A Conceptual OFAT Decomposition

Sequencing	x component	y component
First x, then y	$(\Delta x)y$	$x\Delta y + (\Delta x)(\Delta y)$
First y, then x	$(\Delta x)y + (\Delta x)(\Delta y)$	$x\Delta y$

ExpCCP being homogeneous of degree one in member portfolio volatility and in security volatility suggests a natural decomposition. Let us focus on the decomposition across members to clarify (Menkveld 2017, section 1.5). As *ExpCCP* is homogeneous of degree one in member portfolio risk σ_j,[14] applying Euler's homogeneous function theorem yields the following:[15]

$$ExpCCP = \sum_j \sigma_j \left(\frac{\partial}{\partial \sigma_j} ExpCCP \right). \tag{11}$$

Therefore, the contribution of member j is as follows:[16]

$$\begin{aligned} ExpCCP_j &= \sigma_j \left(\frac{\partial}{\partial \sigma_j} ExpCCP \right) \\ &= \sqrt{\frac{1}{2\pi}} \sigma_j + \sum_{i \in \{J\}} \frac{\alpha}{\sigma_A} \left(\frac{\pi - 1}{2\pi} \right) \sigma_i \sigma_j M(\rho_{ij}), \end{aligned} \tag{12}$$

where σ_A is the standard deviation of aggregate loss, and M is defined in (15) in Appendix A. This result shows that member j's contribution to *ExpCCP* is equal to the portfolio risk σ_j times the (marginal) price of such risk in terms of CCP exposure, $\partial ExpCCP / \partial \sigma_j$. This type of decomposition is used when testing for elevated member crowding on high exposure levels (Hypothesis 2) and for verifying whether house accounts contribute more to *ExpCCP* in these conditions (Hypothesis 3).

A decomposition across securities is derived analogously, where the risk units are ω_k[17] instead of σ_j. A detailed derivation is included as Appendix D. This decomposition is used to test for elevated risk-factor crowding at high exposure levels (Hypothesis 2).

4. Application

This intermezzo section presents the data and discusses various implementation issues. These issues include normality of returns (needed for *ExpCCP*), estimation of the return covariance matrix, and setting the parameter α in the delta-normal VaR.

The data sample used for testing the hypotheses was made available by EMCF. EMCF, now merged with DTCC in the United States to become EuroCCP, is an equity CCP for Nordic equity markets, including Denmark, Finland, and Sweden. The sample consists of trade records with time stamp, transaction size, transaction price, an (anonymized) counterparty ID, and information on whether it was a house- or client-account trade. A trade done on a house account is for a clearing member's own book, whereas a client-account trade is done on behalf of its customers.[18] The sample runs from October 19, 2009, through September 10, 2010, and includes trades on almost all exchanges: NASDAQ-OMX, Chi-X, Bats, Burgundy, and Quote MTF. The only exchange with Nordic trades that it did not clear was Turquoise. Turquoise, however, had a market share of less than 1% at the time.

An equity CCP insures counterparty credit risk for equity trades in the period that starts when a trade is concluded and ends when it settles. When an exchange concludes a trade, the money and the securities are not immediately transferred. Such transfer happens three days later in our sample. Should one side to the trade default in this period, the CCP inherits its position, and the trade will follow through all the way to settlement.

A three-day deferred settlement is conceptually similar to a three-day forward contract between the two sides of the trade. To fix language, we therefore refer to yet-to-settle positions as "positions." Note that these positions change overnight in the absence of any trade. This change is simply due to settlement of legacy trades, and these trades are therefore removed from member positions. In other words, if a member does not trade for three consecutive days, his or her position in all equities becomes zero, as all of his or her earlier trades settled. Finally, we refer to a member's set of open positions at any point in time as his or her portfolio. We emphasize that this should not be confused with a member's portfolio in terms of its equity holdings. It simply refers to the yet-to-settle trades, as these are relevant for CCP exposure because it is for these open positions that the CCP insures counterparty risk.

4.1. Data

4.1.1. Summary Statistics. Table 2 introduces the sample by presenting various summary statistics. The sample captures trading in 242 stocks on 228 days. It contains 226 trading accounts, 87 of which are house accounts and the remaining 139 are client accounts.

The table shows that Nordic stocks are reasonably actively traded, leading to substantial variation in account positions. On average, stocks are traded 1,180 times per day, generating an average volume of €9 million. The standard deviation in account positions is €1.5 million. The corresponding within-account standard deviation is relatively modest: €0.6 million. In other words, most variation in positions is across accounts. Separating between house and client accounts shows that house-account positions tend to be larger in magnitude. Their standard deviation is €1.9 million, whereas it is €1.1 million for client accounts. In addition, the average of end-of-day positions is zero because, for every buyer, there is a seller.

4.2. Implementation Issues

4.2.1. Volume Clock to Recover Normally Distributed Returns. It is well known that financial returns are not normally distributed when sampled using the

Table 2. Summary Statistics

Panel A: General information	
Number of trading days	228
Number of stocks	242
Number of accounts	
House accounts	87
Client accounts	139
Total	226

Panel B: Trade information across stocks	Mean	Standard deviation	Median
Mean of daily number of trades	1,180	2.112	204
Mean of daily volume (shares)	797,844	2,148,216	66,286
Mean of daily volume (euro)	9,042,586	19,535,974	743,348

Panel C: Trade information across clearing members (by account type)	All accounts	House accounts	Client accounts
Mean of end-of-day position (euro)	0	–11,237	11,567
Standard deviation of end-of-day position (euro)	1,535,067	1,880,137	1,069,244
Within-member standard deviation end-of-day position (euro)	619,105	988,685	387,785

Notes. This table presents summary statistics for the CCP data sample. Trades on house accounts are done for a clearing member's own book. Trades on client accounts are done for clients.

wall clock. Returns exhibit negative skewness and excess kurtosis, especially at high frequencies. However, the financial econometrics/microstructure literature has shown that normality of security returns can be recovered when time is measured on a volume clock as opposed to the wall clock (Clark 1973, Ané and Geman 2000, Easley et al. 2012). When using a volume clock, security prices are sampled each time a prespecified amount of volume has been traded. It turns out returns based on such prices are much closer to being normally distributed with less negative skewness and less excess kurtosis.

As normally distributed portfolio returns are needed for computing *ExpCCP*,[19] we use a volume clock in our empirical analysis inspired by Easley et al. (2012). We set the average number of volume bins per day to 34, which corresponds to a 15-minute frequency on the wall clock, as the market is open from 09:00 to 17:30. The bin size, therefore, is picked to be the average daily euro volume divided by 34, yielding 6,770 *ExpCCP* observations. The choice for a 15-minute frequency is common in the microstructure literature, as it strikes a balance between sample size and microstructure noise (Hansen and Lunde 2006). As a robustness check, we consider other frequencies as well (see Section 5.1 and Appendix E.2).

Our implementation follows the volume-clock literature except for two notable differences. First, instead of creating the clock security-by-security based on security-specific volume, we group all securities together and create the clock based on market volume. Suppose the clock starts now, then the latest prices are stacked into a vector. If the volume bin is one million euros, we wait until one million euros have been traded across all securities, and at that moment we again stack the latest prices of all securities into a vector. Returns are then computed based on the standard log difference. The benefit of this approach is that we have a marketwide volume clock that allows *ExpCCP* to be calculated in volume time. Moreover, the wall clock is not completely ignored, as we reset the volume clock at market open. This way, the analysis avoids mixing in overnight effects and thus focuses on intraday exposures only.[20]

To assess whether volume-clock returns are indeed closer to normal than wall-clock returns, we compare various statistics for member portfolio returns. Wall-clock returns are based on a 15-minute sampling frequency. Table 3 presents skewness, excess kurtosis, and the Jarque-Bera statistic, which includes both skewness and kurtosis. Under the null of normality, these statistics are zero in expectation. These statistics are reported for the five largest clearing members, for all five pooled, and for all members pooled, respectively.

The results show strong evidence in favor of the volume clock when returns are required to be normal. All three statistics are substantially closer to zero for each of the five largest members. When pooled, skewness drops from 1.01 to 0.03, kurtosis drops from 46.79 to 3.20, and the Jarque-Bera statistic drops from 91.38 to 0.43. Similar patterns hold when considering all members instead of the five largest only.

Table 3. Statistics on Member Portfolio Returns: Wall vs. Volume Clock

	Skewness		Kurtosis		Jarque-Bera	
Member	Wall clock	Volume clock	Wall clock	Volume clock	Wall clock	Volume clock
First largest	−0.47	−0.05	15.51	1.97	10.06	0.16
Second largest	1.96	0.19	46.60	3.15	91.12	0.42
Third largest	1.50	0.01	30.66	3.16	39.54	0.42
Fourth largest	1.96	0.27	109.55	3.96	500.69	0.66
Fifth largest	−0.29	−0.24	8.69	3.42	3.16	0.50
Largest five pooled	1.01	0.03	46.79	3.20	91.38	0.43
All pooled	−0.62	−0.19	205.47	18.46	1,759.19	14.20

Notes. This table presents various statistics based on realized euro returns for member portfolios. These statistics are presented for wall-clock and volume-clock returns to assess to what extent the returns are normally distributed. The statistics include skewness, excess kurtosis, and the Jarque-Bera statistic. The latter combines the former two and is computed as $(S^2 + K^2/4)/6$, where S is skewness and K is excess kurtosis. The clock runs in 15-minute intervals for the wall clock and for a bin size that, on average, makes a volume bin last 15 minutes. Statistics are presented for the largest five members in terms of volume, for all five pooled, and for all members pooled.

These statistics suggest that, consistent with the literature, nonnormality is indeed much less of an issue for volume-clock returns.[21]

4.2.2. Estimation of Time-Varying Return Covariance. To account for time-varying volatility in returns, we estimate Ω_t as the exponentially weighted moving average (EWMA) of the outer product of returns. This approach is in line with standard practice (e.g., RiskMetrics and EMCF) and corresponds to estimating an IGARCH(1,1).

What remains is to pick the EWMA decay parameter. RiskMetrics uses 0.94 for its highest frequency: daily returns. Because round-the-clock variance is 38 times larger than the intraday 15-minute variance, we pick the decay parameter to be 0.9984 (because $0.9984^{38} = 0.94$). Ω_t is therefore calculated recursively as follows:[22]

$$\Omega_t = (1 - 0.9984)R_{t-1}R'_{t-1} + 0.9984\Omega_{t-1}. \qquad (13)$$

The sample used for our analysis starts on December 7, 2009, but we use data as of October 19, 2009, to have a burn-in period for Ω_t. We start off the recursion with the zero matrix, but given that 0.94 corresponds to a half-life of 11 days, the effect of this choice is negligible by the time we arrive at December 7, 2009.

4.2.3. Pick α to Make *ExpCCP* a 1% VaR. CPMI-IOSCO (2012) recommends that a CCP use a 1% VaR to set margins. We follow this lead and calibrate the alpha parameter in our delta-normal VaR to 2.5 to achieve an exceedance rate of 1%.[23]

5. Results

In this section, we first present the time series of CCP exposure. Several salient spikes will be discussed. In the three subsequent subsections, we test the three hypotheses in Section 2.

Figure 1 plots the time series of CCP exposure: $ExpCCP_t$. Panel (a) of Figure 1 plots exposure levels and shows one particularly large spike in May 2010. This turns out to be the peak month of the Greek sovereign debt crisis.[24] *ExpCCP* reached €5 million that month, which means that the 1% VaR of losses across all members in the oncoming volume bin is €5 million. Although this is about triple the average level, it is still a relatively moderate amount and would not cause a systemic crisis in and of itself. As stated in the introduction, equity CCPs are unlikely to be systemic, but as CCP data are extremely scarce, we are privileged to have access to such data. We believe that it is interesting to study *ExpCCP* dynamics (which is what we do in the remainder of the section) to test several hypotheses.

This result provides new evidence, from the angle of CCP exposure, on the spillover effects of the Greek sovereign debt crisis that have been extensively discussed in the literature. For example, Mink and De Haan (2013) find that news about the Greek bailout generally led to abnormal stock returns for European (including Nordic) banks: positive returns for regulatory initiatives that favor banks, negative returns otherwise. Bhanot et al. (2014) find that Greek yield spread increases are associated with negative abnormal returns on financial stocks throughout Europe. Beetsma et al. (2013) document spillover effects from the Greek yield spread to those of other European countries, and Candelon et al. (2011) find similar evidence when studying credit default swaps on sovereign debt. We will revisit the Greek crisis when decomposing *ExpCCP* in Section 5.2.

Panel (b) of Figure 1 plots exposure *changes* instead of levels. It shows that periods with high levels do not

Figure 1. (Color online) CCP Exposures, Both Levels and Changes

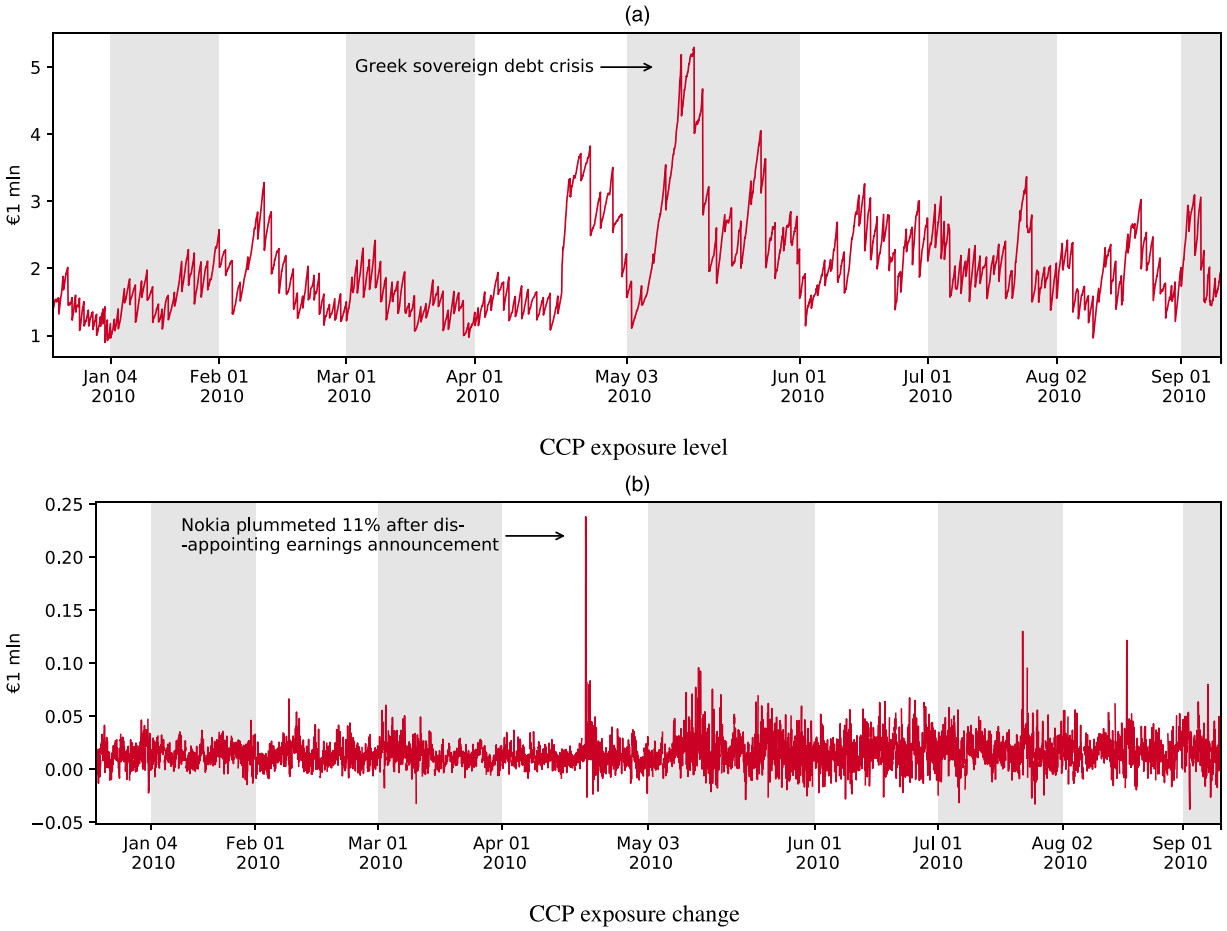

Notes. This figure plots CCP exposure $ExpCCP_t$. Panel (a) plots exposure levels and panel (b) plots exposure changes. Each shaded area corresponds to one month. A wider area indicates a higher monthly volume as the clock runs in volume time.

necessarily correspond to periods with extreme intraday increases. It is the latter that CPMI-IOSCO (2017) is particularly worried about when presenting its latest guidance on CCP risk management. The largest peak corresponds to the idiosyncratic event when Nokia announced earnings that were far below analyst expectations at noon on April 22, 2010. Its share price dropped by about 15% in subsequent minutes. Volume jumped and remained high throughout the afternoon, 400% above what volume was in the morning of that day. We revisit the Nokia event when decomposing $\Delta ExpCCP$ in Section 5.1.

5.1. Hypothesis 1: The Drivers of CCP Exposure Changes Are Different in the (Right) Tail

Hypothesis 1 essentially states that extremely large increases of CCP exposure are different in nature compared with regular changes. As discussed in the hypothesis development section (Section 2), they are likely to reflect a jump in volatility and elevated trading. There might also be crowding if all members tilt their portfolios to the single risk factor at the heart of the turbulence (in the later part of the section, we explore the Nokia event to illustrate). These stand in contrast to "average" changes in CCP exposure that we conjecture to mostly reflect member position changes due to trade.

To test the first hypothesis, we decompose CCP exposure changes into various components for three samples: the full sample and subsamples with the 100 largest and 10 largest increases. Table 4 presents the decomposition results and yields the following insights. First, for the full sample, exposure changes seem to be driven only by member position changes: *TrPosition* dominates all other components.

Second, when zooming in on the top 100 and top 10 $\Delta ExpCCP$, a different picture emerges. Whereas *TrPosition* drops to 59.5% and 33.2%, respectively, two other components, volatility and crowding, become much more important. Although the volatility component makes up only 1.8% of exposure changes for the full sample, it jumps to 17.1% and 58.3% for the top 100 and top 10 increases, respectively. The crowding component is only 3.0% of exposure changes

Table 4. Decomposition of Changes in CCP Exposure

	Full sample	Top 100 Δ*ExpCCP*	Top 10 Δ*ExpCCP*
Panel A: CCP exposure change decomposition in euro			
RetVola	272	10,949	69,311
RetCorr	113	3,555	−89
PrLevel	−133	3,195	−5,324
TrPosition	14,255	38,002	39,445
TrCrowding	443	8,186	15,571
Δ*ExpCCP*	14,949	63,887	118,914
Panel B: CCP exposure change decomposition in percentage			
RetVola	1.8	17.1	58.3
RetCorr	0.8	5.6	−0.1
PrLevel	−0.9	5.0	−4.5
TrPosition	95.4	59.5	33.2
TrCrowding	3.0	12.8	13.1
Δ*ExpCCP*	100.0	100.0	100.0

Notes. This table presents the decomposition of CCP exposure change for the full sample, the top 100, and the top 10 increases. Panel A presents the decomposition in euro. Panel B presents the same decomposition but in percentage. The five components capture changes in security-return volatility (*RetVola*), security-return correlations (*RetCorr*), the pricing level (*PrLevel*), members' outstanding positions (*TrPosition*), and the crowding measure of member positions (*TrCrowding*).

for the full sample but jumps to 12.8% and to 13.1% for the top 100 and top 10 increases, respectively.

Third, the price and correlation components remain small in the two subsamples and sometimes turn negative. Note that the components can be either positive or negative as they can either increase or reduce CCP exposure; and when scaled by the total positive exposure change (Δ*ExpCCP*), the negative components lead to negative percentages. Importantly, the relative contributions of the various components add up to 100%.

Interestingly, the price component contributes −4.5% in the top 10 Δ*ExpCCP*. We speculate that this result must be driven by price crash events that are typically accompanied by high volume and volatility spikes. A lower price level per se reduces CCP exposure simply because CCP exposure is denominated in euro. Suppose the 1% VaR is to lose 20 cents on one euro. *Ceteris paribus*, the 1% VaR would be to lose 10 cents on 50 euro cents (i.e., with price level halved). The "exposure" dropped from 20 cents to 10 cents in this example due to a lower price level.

However, one should not conclude that price crashes serve to reduce exposure. This goes back to our discussion of the identification of components. Price crashes might be *coupled* with volatility and positions in the sense that *conditioning* on extreme price drops, tail events for volatility, or position changes become more likely. We further scrutinize the identification of the components at the end of this section with the help of CoVaR (Tobias and Brunnermeier 2016).

Overall, all of these findings support the hypothesis that extreme increases in CCP exposure are different in nature than overall changes. Specifically, although CCP exposure changes on average are almost entirely driven by member position changes, extreme ones exhibit substantial contributions from volatility changes and changes in crowding.

In Appendix E, we show that our findings are robust to changing the estimate of the time-varying return covariance and changing the sampling frequencies. One notable result worth mentioning here is that, for lower frequencies, the differences between the full sample and the top 10 subsample are attenuated. This highlights the importance of monitoring changes in CCP exposure at high frequencies.[25]

5.1.1. Identification of the Components. As discussed in Section 3.1, the identification of the various components could depend on the sequencing of the components and, in the time series, on the comovements between components. In terms of the sequencing, in Appendix E we show that our findings are robust to changing the sequencing of the components. The analysis of covariation in the component series, however, not only is a robustness check but also yields some economic insights and is therefore presented in the remainder of this section.

To study to what extent the components covary, we pick the baseline sequencing and simply compute correlations across components in the time series. Table 5 presents the results. The three strongest correlations appear in boldface: 0.57 between crowding and positions, −0.25 between price level and volatility, and −0.12 between price level and positions. The high correlation between crowding and positions is simply due to the fact that both are trade-related components and thus are driven by member position changes. The negative correlation between price level and volatility is consistent with the well-documented leverage effect in the financial economics literature: negative price shocks coincide with disproportional volatility increases.

These correlations could be driven by left-tail realizations, by right-tail realizations, or by both. The most

Table 5. Correlation Across Δ*ExpCCP* Components

	RetVola	*RetCorr*	*PrLevel*	*TrPosition*
RetCorr	0.04***			
PrLevel	**−0.25*****	0.01		
TrPosition	0.00	0.05***	**−0.12*****	
TrCrowding	0.08***	0.06***	−0.06***	**0.57*****

Notes. This table reports correlations between the decomposition components (using the baseline sequencing). Correlations that are larger in magnitude than 0.10 are in bold.

*10% significance level; **5% significance level; ***1% significance level.

worrying pattern to a CCP is if correlations are driven by tail realizations that increase CCP exposure, such as left-tail realizations of the price component (price crashes) and right-tail realizations of the volatility component (volatility spikes). To zoom in on adversarial tail events, we turn to CoVaR, which has been developed by Tobias and Brunnermeier (2016) to analyze systemic risk in a similar vein.

The extent to which adversarial shocks coincide can be measured by CoVaR, which is defined as follows:

$$\Pr\left(X^j \geq \mathrm{CoVaR}_q^{j|i} | X^i = \mathrm{VaR}_q^i\right) = q$$
$$\text{with} \quad \Pr\left(X^i \geq \mathrm{VaR}_q^i\right) = q,$$

with $q = 0.01$ for a 1% VaR, and $\mathrm{CoVaR}_q^{j|i}$ is the 1% VaR of variable j conditional on variable i being at its 1% VaR level. We compare the CoVaR derived directly from our sample (i.e., sample CoVaR) to the CoVaR that is implied by the two variables being Gaussian (i.e., Gaussian CoVaR). The difference in magnitude between the two serves to measure the tail dependence between the variables, benchmarked against the normal distribution.

Note that, contrary to the definition of standard VaR, we use $\geq$ (instead of $\leq$) because exposures are defined in terms of losses, which are positive numbers. Thus, we focus on the right-tail realizations of the components. The only exception is *PrLevel*, for which we use $\leq$ as an adversarial event, as price crashes are left-tail realizations of the price component. Thus, the tail events are large *increases* of volatility, return correlations, positions, and crowding, and large *decreases* of the price level.

Table 6 presents the sample CoVaR, Gaussian CoVaR, and the difference in magnitude between the two. Note that for the price component, the CoVaRs are negative because its left-tail realizations are adversarial to the CCP. Highlighted in bold are the sample CoVaRs with largest absolute value in each column. This focuses attention on the strongest tail dependence between the variables. The highlighting shows that, in five of six cases, it is the large price drop that, when conditioned on, leads to the strongest tail dependence in the other components. Nonetheless, for price level itself, it is large volatility increases that lead to its strongest tail dependence.

For each column, the largest difference in magnitude between sample CoVaR and Gaussian CoVaR is italicized. Again, in most cases (four of six), the wedge is largest when a price drop is conditioned on. Price crashes seem to "trigger" the strongest adversarial covariation with the other components. This finding itself is consistent with the fire-sales dynamics, which

Table 6. CoVaR Across *ExpCCP* Components

	RetVola	*RetCorr*	*PrLevel*	*TrPosition*	*TrCrowding*
RetVola					
Gaussian CoVaR		4,101	–12,744	22,471	9,323
Sample CoVaR		12,409	**–24,062**	52,143	12,159
Difference in magnitude		*8,308*	*11,318*	29,672	2,835
RetCorr					
Gaussian CoVaR	12,856		–12,419	23,463	9,201
Sample CoVaR	9,767		–17,925	48,447	10,390
Difference in magnitude	–3,089		5,506	24,984	1,190
PrLevel					
Gaussian CoVaR	17,446	5,423		26,254	9,967
Sample CoVaR	**47,799**	**12,806**		**63,408**	**22,320**
Difference in magnitude	*30,353*	7,383		*37,154*	*12,353*
TrPosition					
Gaussian CoVaR	12,454	4,148	–10,629		12,042
Sample CoVaR	11,114	9,003	–21,671		17,128
Difference in magnitude	–1,341	4,855	11,042		5,086
TrCrowding					
Gaussian CoVaR	13,372	4,210	–10,442	31,161	
Sample CoVaR	13,621	10,448	–21,829	52,301	
Difference in magnitude	250	6,238	11,387	21,140	

Notes. This table reports the CoVaR between the components. $\mathrm{CoVaR}_q^{j|i}$ is reported in row i and column j (i.e., the row variables are conditioned on). It measures the extent to which components exhibit coupling effects in the sense that tail events co-occur. We report the sample CoVaR along with a Gaussian CoVaR, which is what CoVaR would have been had the distribution been normal (calibrated to the mean and covariance of the pair). We also report the difference in magnitude between the two, defined as Sign(Sample CoVaR)(Sample CoVaR – Gaussian CoVaR). Highlighted in bold are the largest sample CoVaRs in absolute value within a column. The largest differences in magnitudes when comparing the sample CoVaR with the Gaussian CoVaR (within a column) appear in italics.

was one of the channels that motivate the first hypothesis. Although the findings are suggestive of causality, we caution that one cannot make any causal statements.

To illustrate these general findings, Figure 2 zooms in on the Nokia event. Panel (a) decomposes the exposure jump of €0.24 million immediately following Nokia's disappointing announcement. A couple of features stand out. First, return volatility is by far the largest component: €0.30 million. Its effect is moderated by the negative price-level component: €−0.06 million. In other words, Nokia volatility spikes due to a large negative return of about −15%, but relative volatility applies at a lower price level because of the negative return. Finally, the trade components are both positive, implying that, on average, traders expand their positions in a way that leads to more crowding. These trade components are, however, dwarfed by the volatility component.

Panel (b) zooms out and shows how exposure built up throughout the day. Its most salient feature is that, although the volatility spike dominates exposure change in the volume bin right after the event, it is only about a fifth of the exposure increase for that day. The reason is that volatility is only a major component in the bin just after the event; trade components dominate subsequent bins. Therefore, the high volume in the afternoon turns out to be due to traders expanding their positions, not reducing them. There is also elevated crowding, but its contribution is only

Figure 2. (Color online) Decomposition of the Largest CCP Exposure Increase: The Nokia Event

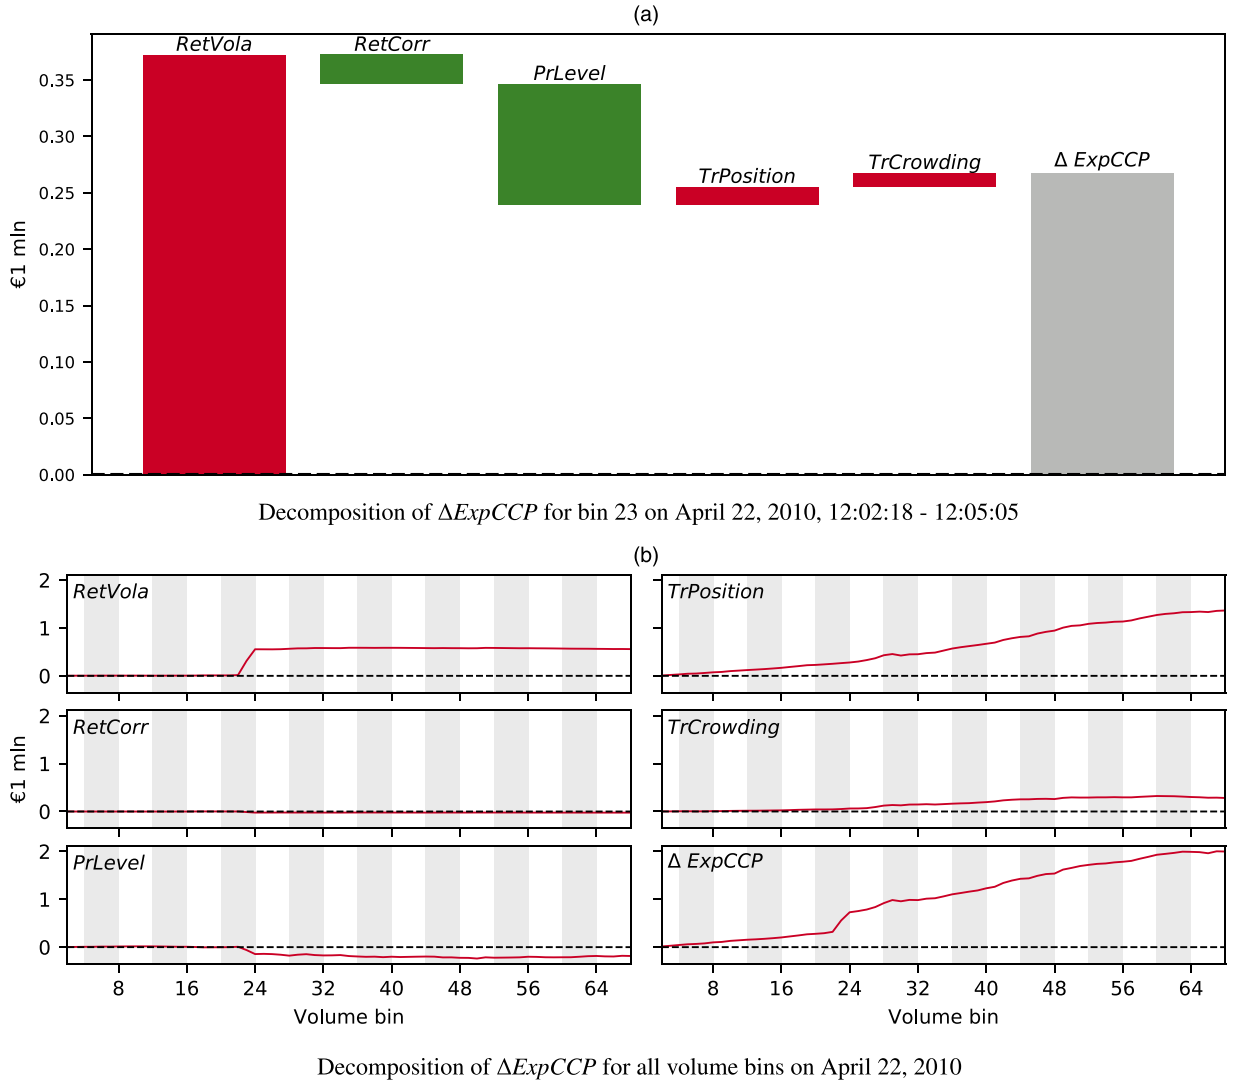

Notes. At noon on April 22, Nokia announced disappointing earnings, which caused a large price drop of 15% in a few minutes and a sharp increase in trading volume. CCP exposure rose steeply in the volume bin subsequent to the announcement. Panel (a) decomposes this exposure increase into five components: security-return volatility (*RetVola*), security-return correlation (*RetCorr*), price level (*PrLevel*), position changes (*TrPosition*), and the extent of crowding in member positions (*TrCrowding*). Panel (b) zooms out and cumulates these components for the full day.

about 20% of the total contribution of trade components. Finally, traders do not seem to take substantial positions *ahead* of the Nokia announcement, as all components only start to contribute substantially in the afternoon.

Perhaps the most important message of these Nokia results is that firm-specific shocks can have a systemic impact through heightened CCP exposure. News that strikes like lightning causes volatility to spike and, more importantly, makes traders expand their positions in ways that lead to more concentration in their portfolios (i.e., crowding).

5.2. Hypothesis 2: The Structure of CCP Exposure Levels Is Different in the (Right) Tail

The second hypothesis focuses on the highest exposure *levels* as opposed to the largest changes. Does one see evidence of elevated exposure concentration (i.e., crowding) across members, across (a combination of) stocks, or across both? Such a finding would raise concerns about market conditions that are potentially prone to fire-sale dynamics.

To verify whether the structure of CCP exposure is different in the tail, we decompose exposure for the full sample and for the subsamples of the top 10% and the top 1% CCP exposure levels. The reason for picking the top 10% here instead of the top 100 used in the previous subsection is that CCP exposure levels are very persistent as compared with exposure changes. The top 100 subsample is smaller than the top 10% sample, and therefore, when used in the level analysis, it would essentially point to the same period of time. The same argument applies to picking the top 1% instead of the top 10. The decomposition is conducted both across members and across stocks. We then compute the Herfindahl-Hirschman Index (HHI) along with shares of the largest 1, 5, and 10 contributors to measure the concentration level.

Table 7 reveals that concentration is elevated in the tail for members, but not for individual stocks. The shares of the top 1, 5, and 10 members increase substantially from the full sample to the top 1% subsample. For example, the share of the top five members increases from 27.8% in the full sample to 34.9% in the top 10% subsample and to 46.8% in the top 1%. The HHI shows a similar trend and increases from 0.030 to 0.046 and 0.085, respectively. There is no such trend for the decomposition across stocks. The share of the top five stocks, for example, stays rather flat. It changes from 43.3% in the full sample to 48.9% and 41.1% in the top 10% and top 1%, respectively.

The unchanged concentration for the decomposition across stocks does not preclude crowding in a particular *portfolio* of stocks. To study whether this is the case, we apply principal component analysis (PCA) on member portfolio returns for the full sample and for both subsamples. Table 8 shows that there does appear to be elevated crowding in the subsamples. It is strongest for the first principal component (PC1), for which the share in total variance increases from 7.8% in the full sample to 20.8% in the top 10% subsample and to 37.6% in the top 1% subsample. As the largest CCP exposures occur mostly in the Greek crisis period, it is likely that this component captures a market effect. To verify, we compute the correlation of PC1 with the local market index and indeed find the expected pattern: 0.43 for the full sample, 0.86 for the top 10% subsample, and 0.98 for the top 1% subsample.

Finally, to illustrate these results graphically, Figure 3 plots the HHI for both the across-member and across-stock decompositions (see Table 7). Panel (a)

Table 7. Decomposition of CCP Exposure Across Members and Across Stocks

	Full sample	Top 10% *ExpCCP*	Top 1% *ExpCCP*
Panel A: Decomposition of CCP exposure across traders			
Top 1 member (%)	9.3	14.4	25.5
Top 5 members (%)	27.8	34.9	46.8
Top 10 members (%)	41.7	48.2	57.3
HHI	0.030	0.046	0.085
Panel B: Decomposition of CCP exposure across stocks			
Top 1 stock (%)	18.7	28.0	16.1
Top 5 stocks (%)	43.3	48.9	41.1
Top 10 stocks (%)	59.3	62.6	57.3
HHI	0.080	0.176	0.053

Notes. This table presents the results of decomposing CCP exposure levels across members and across stocks for the full sample and for the top 10% and top 1% subsamples. Various concentration measures are reported: the share of the member/stock with the largest contribution, the five largest contributors, and the 10 largest contributors. The table further reports the HHI.

Table 8. Principal Component Analysis of Member Portfolio Returns

	Full sample (%)	Top 10% *ExpCCP* (%)	Top 1% *ExpCCP* (%)
PC1	7.8	20.8	37.6
PC2	5.2	8.9	10.8
PC3	2.7	6.4	6.2
PC1+PC2+PC3	15.7	36.0	54.7

Notes. This table uses principal component analysis to characterize the commonality in member portfolio returns for the full sample and for subsamples where CCP exposure levels are large. It reports the explained variation of the first, second, and third principal component along with the sum of these three.

plots the across-member HHI in solid red and overlays the CCP exposure level in dashed blue (using the second y axis). It illustrates that high concentration occurs mostly in the Greek crisis period. Panel (b) plots the across-stock HHI and, as expected, shows that it stays rather flat at times when CCP exposure peaks. This does not mean that concentration remains at the same level throughout. It does show a large peak around the Nokia event when exposure *increase* is the largest as analyzed in the previous subsection. Upon further inspection, we unsurprisingly find that the concentration occurs in the stock of Nokia.

5.3. Hypothesis 3: The Relative Contribution of House Accounts Increases in the (Right) Tail

The third hypothesis states that the relative contribution of house accounts is higher for extreme CCP exposure levels. This is potentially worrisome, as clearing members are highly leveraged financial institutions.

Figure 3. (Color online) CCP Exposure Concentration Across Members and Across Stocks

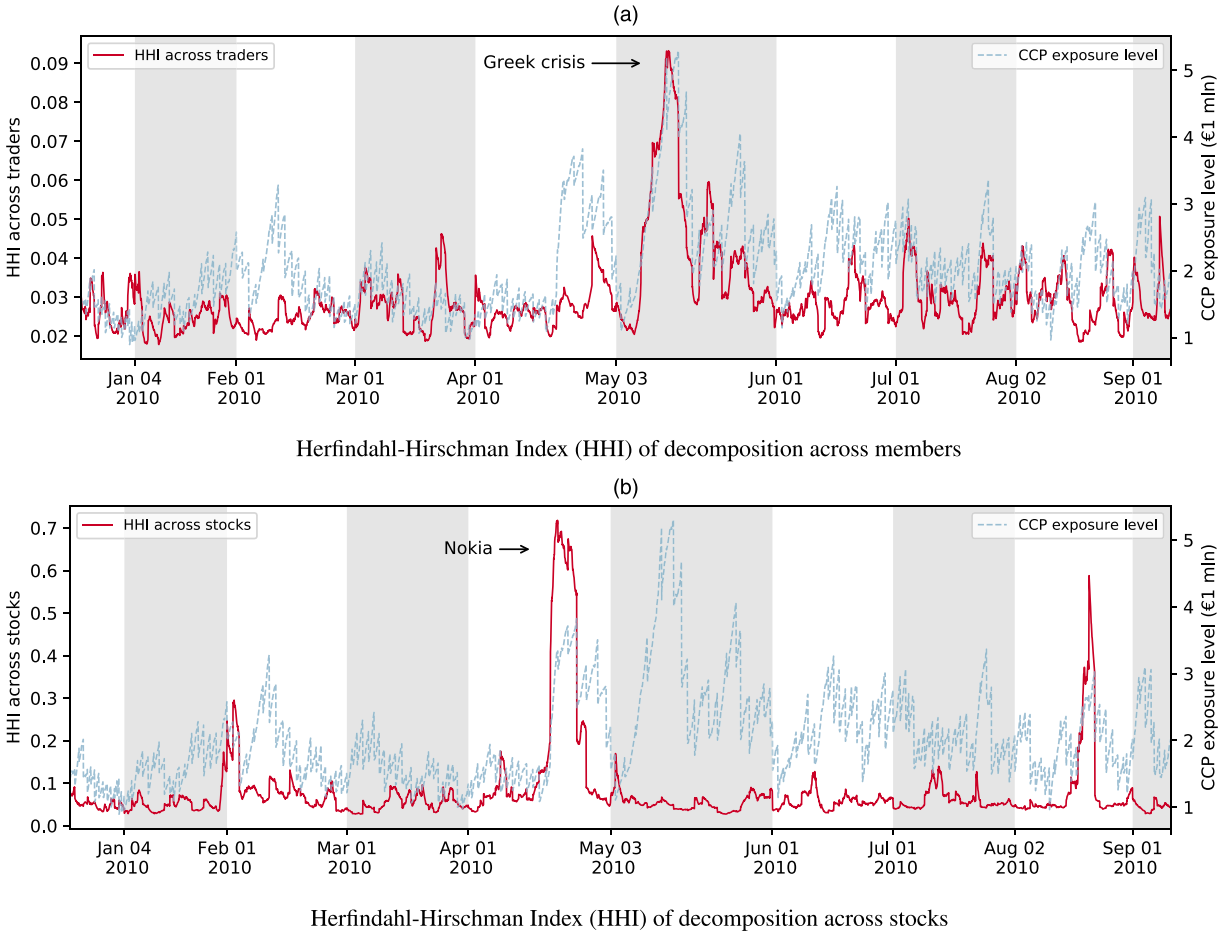

Notes. This figure plots the HHI of CCP exposure decomposed across members in panel (a) and across stocks in panel (b) in solid lines (left y axis). The plots also show the level of CCP exposure in dashed lines (right y axis).

Table 9 presents evidence largely rejecting the third hypothesis. The decomposition across house and client accounts in Panel A shows that house accounts contribute 66.8% to CCP exposure in the full sample. This contribution, however, hardly changes when measured for the top 10% subsample (66.0%) and increases only mildly to 69.7% in the top 1% subsample.

Panel B shows that in spite of the relative contribution of all house accounts combined being rather flat across subsamples, there is concentration *within* house accounts. The HHI computed based on each member's contribution to the total of house-account contributions increases from 0.051 for the full sample to 0.083 for the top 10% and to 0.160 for the top 1%. The results suggest that, in stressed markets, the positions in the books of some clearing members expand, whereas the positions of others shrink. This causes their total contribution to CCP exposure to remain unchanged, yet there is more concentration within house accounts.

There appears to be no such pattern for client accounts with a collective contribution that remains flat across the three samples and a within-client concentration that remains largely unchanged. The HHI is 0.068 for the full sample, 0.071 for the top 10% subsample, and 0.081 for the top 1% subsample.

In sum, the significantly higher concentration within house accounts is potentially worrisome. Most clearing members are highly leveraged sell-side banks, which, if trading for speculative reasons, might default on their position if they turn out to be on the wrong side of the bet. Given that they seem to crowd into the same (set of) risk factors, there might be multiple members who are heavily under water on their bets at the same time. Admittedly, it is unlikely that they would default on their equity trades, but if the pattern were to carry over to CCPs that clear interest rate derivatives or credit default swaps, then such a pattern would become a systemic worry.

6. Conclusion

In summary, we test three hypotheses about the exposure a CCP has vis-à-vis with its clearing members. All three hypotheses focus on tail events and whether the nature of CCP exposure changes in such cases. The academic literature has emphasized elevated concentration (i.e., crowding) in such stressed markets with a risk of fire-sale price dynamics.

We develop an approach for monitoring CCP exposure whereby both exposure *levels* and exposure *changes* can be decomposed to identify the relative contribution of various factors. The empirical results support the hypothesis that the nature of exposure levels or exposure changes is different in the tail: there is indeed more crowding in stressed markets. The hypothesized larger contribution of house accounts to total exposure in such conditions is not supported by the data. However, *within* house accounts, there is more concentration with a small number of clearing members contributing a disproportionate amount to total house-account exposure.

Our findings suggest that CCP executives and regulators should monitor at high frequencies with a particular focus on tail events. Whether contingency planning is needed and, if so, what form it should take is a topic for future research. We do, however, believe that the approach we have developed could be useful for monitoring CCP exposure at high frequencies. The proposed decomposition of exposure changes could help CCPs and regulators diagnose sudden large jumps in exposure. As all results are analytical (thus avoiding heavy-duty simulations), the approach can be implemented in real time. We believe this is an asset in today's extremely fast markets.

Acknowledgments

The authors are grateful for helpful comments from Evangelos Benos, Markus Brunnermeier, Jorge Cruz Lopez, Jean-Edouard Colliard, Tobias Dieler, Gerardo Ferrara, Pedro Gurrola-Perez, Mark Manning, David Murphy, Michalis Vasios, Andrew Vivian, Guillaume Vuillemey, Xiao Xiao, and Marius Zoican; seminar participants at the University of Hong Kong; and conference participants at the 2015 Australasian Finance and Banking Conference, 2015 Conference on Theories and Practices of Securities and Financial Markets, 2015 SYRTO Conference on Systemic Risk, 2016 Federal Reserve International Banking Conference, 2016 Eastern Financial Annual Meeting, 2019 Annual Conference of Western Economic

Table 9. Decomposition of CCP Exposure into House and Client Accounts

	Full sample	Top 10% *ExpCCP*	Top 1% *ExpCCP*
Panel A: Contribution to CCP exposure by account type			
Contribution by house accounts (%)	66.8	66.0	69.7
Contribution by client accounts (%)	33.2	34.0	30.3
Panel B: HHI within account type			
Herfindahl-Hirschman Index (HHI) within house accounts	0.051	0.083	0.160
Herfindahl-Hirschman Index (HHI) within client accounts	0.068	0.071	0.081

Notes. Panel A decomposes CCP exposure into house and client accounts. Panel B shows the concentration of CCP exposure *within* each account type by means of the HHI. Both panels consider the full sample and subsamples of the top 10% and top 1% CCP exposure levels.

Association International, 2019 Annual Conference of the Asia-Pacific Association of Derivatives (APAD), 2019 Central Bank Conference on Market Microstructure, 2019 Australian Finance and Banking Conference, and 2020 ECB Macroprudential Stress Testing Conference. The paper won the outstanding paper prize from the 2019 APAD conference. The authors thank SURFsara for support when using the Lisa Compute Cluster. Albert J. Menkveld and Shihao Yu gratefully acknowledge Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO) for a Vici and a Research Talent grant, respectively. Menkveld further thanks the Bank of Canada and the Bank of England for a week-long visit to refine his views on CCPs. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank for International Settlements.

Appendix A. Results Needed to Compute *ExpCCP*

Let L_t be the $J \times 1$ vector that stacks all $L_{j,t}$. Because $A_t = \sum_j L_{j,t}$, one needs to compute $\mathrm{E}(L_t)$ and $\mathrm{var}(L_t)$ to evaluate (4). Following Menkveld (2017, proposition 1) yields the following two results:

$$\begin{aligned} \mathrm{E}(L_t) &= \mu_t, \quad \mu_{j,t} = \sqrt{\frac{1}{2\pi}}\sigma_{j,t}, \\ \mathrm{var}(L_t) &= \Psi_t, \ \ \psi_{ij,t} = \frac{\pi - 1}{2\pi}\sigma_{i,t}\sigma_{j,t}M\big(\rho_{ij,t}\big), \end{aligned} \tag{A.1}$$

where $\sigma_{ij,t}$ is the (i,j)th element of the covariance matrix of member portfolio returns Σ_t, $\sigma_{i,t}$ is short for $\sigma^{\frac{1}{2}}_{ii,t}$, and $\rho_{ij,t} = \sigma_{ij,t}/\sigma_{i,t}\sigma_{j,t}$. The function

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

maps portfolio return correlations into portfolio loss correlations. Menkveld (2017) presents detailed proofs. *ExpCCP* can now be written explicitly as

$$\begin{aligned} &ExpCCP_t \\ &= \sum_j \sqrt{\frac{1}{2\pi}}\sigma_{j,t} + \alpha\left(\sum_i \sum_j \frac{\pi - 1}{2\pi}\sigma_{i,t}\sigma_{j,t}M\big(\rho_{ij,t}\big)\right)^{\frac{1}{2}}. \end{aligned} \tag{A.3}$$

Appendix B. Decomposition of CCP Exposure Change

This section presents the various components that add up to CCP exposure change:

$$\begin{aligned} \Delta ExpCCP_t = &\underbrace{RetVola_t + RetCorr_t + PrLevel_t}_{\text{Price components}} \\ &+ \underbrace{TrPosition_t + TrCrowding_t}_{\text{Trade components}}. \end{aligned} \tag{B.1}$$

B.1. Price Components

The three price components are as follows:

$$\begin{aligned} RetVola_t = &f\big(D\big(\boldsymbol{D_{\Omega_t}}, R_{\Omega_{t-1}}, P_{t-1}, \tilde{N}_{t-1}\big), \\ &R\big(\boldsymbol{D_{\Omega_t}}, R_{\Omega_{t-1}}, P_{t-1}, \tilde{N}_{t-1}\big)\big) \\ &- f\big(D\big(D_{\Omega_{t-1}}, R_{\Omega_{t-1}}, P_{t-1}, \tilde{N}_{t-1}\big), \\ &R\big(D_{\Omega_{t-1}}, R_{\Omega_{t-1}}, P_{t-1}, \tilde{N}_{t-1}\big)\big), \end{aligned} \tag{B.2}$$

$$\begin{aligned} RetCorr_t = &f\big(D\big(D_{\Omega_t}, \boldsymbol{R_{\Omega_t}}, P_{t-1}, \tilde{N}_{t-1}\big), \\ &R\big(D_{\Omega_t}, \boldsymbol{R_{\Omega_t}}, P_{t-1}, \tilde{N}_{t-1}\big)\big) \\ &- f\big(D\big(D_{\Omega_t}, R_{\Omega_{t-1}}, P_{t-1}, \tilde{N}_{t-1}\big), \\ &R\big(D_{\Omega_t}, R_{\Omega_{t-1}}, P_{t-1}, \tilde{N}_{t-1}\big)\big), \text{ and} \end{aligned} \tag{B.3}$$

$$\begin{aligned} PrLevel_t = &f\big(D\big(D_{\Omega_t}, R_{\Omega_t}, \boldsymbol{P_t}, \tilde{N}_{t-1}\big), \\ &R\big(D_{\Omega_t}, R_{\Omega_t}, \boldsymbol{P_t}, \tilde{N}_{t-1}\big)\big) \\ &- f\big(D\big(D_{\Omega_t}, R_{\Omega_t}, P_{t-1}, \tilde{N}_{t-1}\big), \\ &R\big(D_{\Omega_t}, R_{\Omega_t}, P_{t-1}, \tilde{N}_{t-1}\big)\big). \end{aligned} \tag{B.4}$$

B.2. Trade Components

The two trade components are as follows:

$$\begin{aligned} TrPosition_t = &f\big(D\big(D_{\Omega_t}, R_{\Omega_t}, P_t, \ \boldsymbol{\tilde{N}_t}\big), \\ &R\big(D_{\Omega_t}, R_{\Omega_t}, P_t, \tilde{N}_{t-1}\big)\big) \\ &- f\big(D\big(D_{\Omega_t}, R_{\Omega_t}, P_t, \tilde{N}_{t-1}\big), \\ &R\big(D_{\Omega_t}, R_{\Omega_t}, P_t, \tilde{N}_{t-1}\big)\big) \text{ and} \end{aligned} \tag{B.5}$$

$$\begin{aligned} TrCrowding_t = &f\big(D\big(D_{\Omega_t}, R_{\Omega_t}, P_t, \tilde{N}_t\big), \\ &R\big(D_{\Omega_t}, R_{\Omega_t}, P_t, \ \boldsymbol{\tilde{N}_t}\big)\big) \\ &- f\big(D\big(D_{\Omega_t}, R_{\Omega_t}, P_t, \tilde{N}_t\big), \\ &R\big(D_{\Omega_t}, R_{\Omega_t}, P_t, \tilde{N}_{t-1}\big)\big). \end{aligned} \tag{B.6}$$

Appendix C. Example of CCP Exposure Change Analysis

Table C.1 presents a simple example to illustrate the insights that one can get from a decomposition of CCP exposure changes. Suppose there are four agents ($A1$, $A2$, $A3$, $A4$) and two securities ($S1$ and $S2$) that cost €1 and have returns that are standard normal and mutually independent at least at the beginning of time. All agents start with a zero position in the securities. To illustrate real-time CCP exposure monitoring, we consider a particular sequence of events. We compute CCP exposure change after each event and present its decomposition. This controlled setting serves to familiarize with the approach before implementing it on real-world data.

The first two columns of Table C.1 describe the sequence of events. CCP exposure is computed after each event based on the loss distribution for the oncoming period. In some cases, events are illustrated by horizontal arrows that correspond to positions in the first security. Arrows that point right denote long positions. Left arrows denote short positions. Vertical arrows correspond to positions in the second security. Up arrows denote long positions. Down arrows denote short positions. The remaining columns show CCP exposure, its change, and the decomposition of this change into the five factors. These changes and decompositions are discussed here.

- $t = 0$. CCP exposure is 0 for the simple reason that none of the agents have a position.
- $t = 1$. $A1$ entered a long position of one unit on $S1$ and $A2$ is on the opposite side of that trade. CCP exposure becomes €2.3. The decomposition shows that €2.9 is due to expanded positions (*TrPosition*), and the crowding

Table C.1. Simple Example to Illustrate the Decomposition of CCP Exposure Changes

t	Trades/changes	$ExpCCP_t$	$\Delta ExpCCP_t = RetVola_t + RetCorr_t + PrLevel_t + TrPosition_t + TrCrowding_t$					
0	$\sigma_1 = \sigma_2 = 1$, $\rho = 0$, $p_1 = p_2 = 1$.	0.0	0.0	0.0	0.0	0.0	0.0	0.0
1	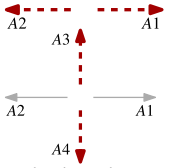	2.3	2.3	0.0	0.0	0.0	2.9	−0.6
2		3.7	1.4	0.0	0.0	0.0	1.8	−0.4
3	Volatility changes from $\sigma_1 = \sigma_2 = 1$ to $\sigma_1 = 2, \sigma_2 = 1$.	5.7	2.0	2.0	0.0	0.0	0.0	0.0
4	Return correlation changes from $\rho = 0$ to $\rho = 0.5$.	6.0	0.3	0.0	0.3	0.0	0.0	0.0
5	Price level changes from $p_1 = p_2 = 1$ to $p_1 = 0.5, p_2 = 1$.	3.9	−2.1	0.0	0.0	−2.1	0.0	0.0
6	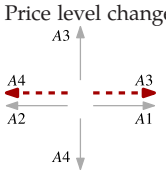	6.7	2.8	0.0	0.0	0.0	1.5	1.3
7	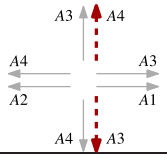	4.6	−2.1	0.0	0.0	0.0	−2.0	−0.1

Notes. This example illustrates how the one-factor-at-a-time decomposition approach identifies the different components in CCP exposure changes. There are four agents ($A1$, $A2$, $A3$, $A4$) and two securities ($S1$, $S2$). Arrows denote positions in these securities. Right and left arrows illustrate long and short positions in $S1$, respectively; up and down arrows illustrate long and short positions in $S2$, respectively. Red dashed arrows correspond to new trades. CCP exposures are computed with $\alpha = 2.5$, which is the calibrated value based on our real-world sample (see Section 4.2).

component is €–0.6 (*TrCrowding*). The reason for this negative crowding term is simply that, in this case, the members have taken the opposite side of the same trade and their portfolio returns are therefore perfectly negatively correlated.

- $t = 2$. $A3$ entered a long position of one unit in $S2$ with $A4$ taking the short side. CCP exposure increases by €1.4 to €3.7. The decomposition shows a positive *TrPosition* of €1.8 and a negative *TrCrowding* of €–0.4. The positive position risk is due to the new trade leading to larger positions. Furthermore, the new trade between $A3$ and $A4$ is in $S2$ and is therefore orthogonal to the positions between $A1$ and $A2$. In other words, the new trade between $A3$ and $A4$ lowers the correlations between member portfolio returns. Hence, there is less crowding now than before.
- $t = 3$. The return volatility of $S1$ increases from 1 to 2. CCP exposure increases by €2.0 to €5.7. The decomposition indeed attributes it to the volatility component (*RetVola*).
- $t = 4$. The correlation between the returns of $S1$ and $S2$ increases from 0 to 0.5. CCP exposure increases by €0.3 to €6.0. The decomposition assigns it to the correlations component (*RetCorr*).
- $t = 5$. The price of $S1$ drops from €1 to €0.5. CCP exposure drops by €2.1, which is completely assigned to the price level (*PrLevel*). This is simply the result of volatility being defined in relative terms. If it does not change but the price level drops, then the VaR, which is expressed in euro, drops.
- $t = 6$. $A3$ trades again with $A4$ but this time enters a one-unit-long position in $S1$, where $A4$ takes the short side. CCP exposure increases by €2.8 to €6.7. Positions now crowd on the risk factor $S1$. The decomposition assigns €1.3 of the increase to *TrCrowding* and the remaining €1.5 to *TrPosition.*
- $t = 7$. $A3$ and $A4$ effectively undo their first trade by entering a reverse trade. In this reverse trade, $A3$ is long one unit of $S2$ and $A4$ is short one unit. CCP exposure declines by €2.1 to €4.6. The decomposition shows that most of the decrease is due to a reduction in outstanding (net) positions (i.e., the drop is largely assigned to *TrPosition*). This event shows that trade does not necessarily imply more exposure; it could *reduce* exposure when, after the trade, positions shrink. Note that, combining $t = 6$ and $t = 7$, the *size* of trade positions have not changed—members are long or short the same amount of risk—but CCP exposure has increased due to crowding.

In summary, the decompositions of CCP exposure changes generate insight into the drivers of these changes. *TrPosition* picks up whether new trades extend or reverse legacy positions. *TrCrowding* captures the correlation of member portfolio returns. *RetVola*, *RetCorr*, and *PrLevel* identify exposure changes due to changes in the volatility of returns, their correlations, and price levels, respectively.

Appendix D. Decomposition of CCP Exposure Across Securities

ExpCCP being homogeneous of degree one in ω_k yields the following:[26]

$$ExpCCP = \sum_i \omega_k \left(\frac{\partial}{\partial \omega_k} ExpCCP \right). \tag{D.1}$$

Therefore, the contribution of security k is as follows:

$$\begin{aligned} ExpCCP_k &= \sum_{i,j} \omega_k \left(\frac{\partial}{\partial \omega_k} ExpCCP \right) \\ &= \sum_j \sqrt{\frac{1}{2\pi}} \frac{B_{jj}}{2\sigma_j} + \\ &\quad + \frac{\alpha}{2\sigma_A} \sum_{i,j} \left(\frac{\pi - 1}{2\pi} \right) \Bigg(M'(\rho_{ij}) B_{ij} \\ &\quad + \frac{\sqrt{1 - \rho_{ij}^2} - 1}{\pi - 1} \left(\frac{\sigma_j}{2\sigma_i} B_{ii} + \frac{\sigma_i}{2\sigma_j} B_{jj} \right) \Bigg), \end{aligned} \tag{D.2}$$

where

$$B_{ij} = n_i' \frac{\partial \Omega}{\partial \omega_k} n_j, \quad M'(\rho_{ij}) = \frac{\frac{1}{2}\pi + \arcsin(\rho_{ij})}{\pi - 1}. \tag{D.3}$$

Appendix E. Robustness Checks

E.1. Moving-Window Return Covariance Estimate

The exposure change decomposition analysis presented in Table 4 relies on an EWMA estimate of the covariance matrix of returns. To verify whether the results are robust, we redo the analysis with a rolling-window estimate of return covariance. For the length of the window, we picked the burn-in period used for EWMA (i.e., 50 days). We have considered other alternatives, such as parametric estimation of the time-varying covariance matrix. One natural approach is to estimate a multivariate GARCH, but implementation is infeasible given the large dimensions of the covariance matrix that needs to be estimated: 242 × 242. We therefore stick to a parameter-free estimate, but this time we base it on a rolling window.

Table E.1 shows that when using a rolling-window estimate, decomposition results are similar to those using an EWMA estimate. Importantly, the key observations in the main text all hold up: the position component dominates all other components for the full sample, but volatility and crowding become much more important when considering only the top 100 and top 10 exposure changes.

E.2. Alternative Sampling Frequencies

Is high-frequency analysis important for the decomposition results presented in Table 4? Note that the volume bins were chosen such that, on average, they span 15 minutes. A higher frequency is computationally feasible but economically impossible, as "microstructure noise" starts to bias return covariance estimates (Andersen et al. 2003). Lower frequency, however, is possible, and in this section, we redo the decomposition based on volume bins that, on average, span 30 minutes or a full hour.

Table E.2 presents the results but only reports full sample and top 10 decompositions to save space. The table shows that the main results are unaffected: the position component dominates in the full sample, but volatility and crowding become important in the top 10 subsample. These results, however, become attenuated when the analysis is conducted at the lower frequency. That is, the contribution of volatility and crowding drops in the top 10 subsample. This result testifies to the importance of high-frequency analysis of CCP exposure changes to diagnose the nature of trading during brief spells of volatility spikes and extreme volume.

E.3. Alternative Sequencing in Exposure Change Decomposition

The decomposition of CCP exposure change presented in Table 4 and discussed in Section 5.1 critically depends on

Table E.1. Decomposition of Exposure Change for a Rolling-Window Estimate of Return Covariance

	EWMA estimate of Cov(R)			Rolling-window estimate of Cov(R)		
	Full sample	Top 100	Top 10	Full sample	Top 100	Top 10
Panel A: CCP exposure change decomposition in euro						
RetVola	272	10,949	69,311	−414	5,156	24,953
RetCorr	113	3,555	−89	−22	553	−1,225
PrLevel	−133	3,195	−5,324	−168	4,964	−525
TrPosition	14,255	38,002	39,445	18,956	59,148	80,765
TrCrowding	443	8,186	15,571	624	11,388	20,255
ΔExpCCP	14,949	63,887	118,914	18,976	81,208	124,223
Panel B: CCP exposure change decomposition in percentage						
RetVola	1.8	17.1	58.3	−2.2	6.3	20.1
RetCorr	0.8	5.6	−0.1	−0.1	0.7	−1.0
PrLevel	−0.9	5.0	−4.5	−0.9	6.1	−0.4
TrPosition	95.4	59.5	33.2	99.9	72.8	65.0
TrCrowding	3.0	12.8	13.1	3.3	14.0	16.3
ΔExpCCP	100.0	100.0	100.0	100.0	100.0	100.0

Notes. This table repeats the exposure-change decompositions reported in Table 4 and adds decompositions based a 50-day rolling-window estimate of return covariance instead of the EWMA estimate used in the baseline decompositions.

Table E.2. Decomposition of Exposure Change for Different Frequencies

	Baseline: 34 bins per day (15-minute intervals)		17 bins per day (30-minute intervals)		8 bins per day (1-hour intervals)	
	Full sample	Top 10	Full sample	Top 10	Full sample	Top 10
Panel A: CCP exposure change decomposition in euro						
RetVola	272	69,311	984	89,104	3,881	276,342
RetCorr	113	–89	316	14,133	683	26,155
PrLevel	–133	–5,324	–351	–20,163	–1,603	–24,318
TrPosition	14,255	39,445	38,730	101,854	123,654	327,711
TrCrowding	443	15,571	1,279	22,431	5,052	77,654
Δ*ExpCCP*	14,949	118,914	40,959	207,359	131,667	683,545
Panel B: CCP exposure change decomposition in percentage						
RetVola	1.8	58.3	2.4	43.0	2.9	40.4
RetCorr	0.8	–0.1	0.8	6.8	0.5	3.8
PrLevel	–0.9	–4.5	–0.9	–9.7	–1.2	–3.6
TrPosition	95.4	33.2	94.6	49.1	93.9	47.9
TrCrowding	3.0	13.1	3.1	10.8	3.8	11.4
Δ*ExpCCP*	100.0	100.0	100.0	100.0	100.0	100.0

Notes. This table repeats the exposure-change decompositions of Table 4 and adds decompositions based on lower frequencies. The baseline result is based on having, on average, 34 volume bins per day, which corresponds to 15-minute intervals. The added frequencies are 17 and 8 and therefore correspond to 30-minute and 1-hour intervals, respectively.

Table E.3. Decomposition of Exposure Change for Alternative Component Sequences

	Full sample	Top 100 Δ*ExpCCP*	Top 10 Δ*ExpCCP*
Panel A: CCP exposure change decomposition in euro			
RetVola	275	11,003	69,022
	(263, 288)	(10,581, 11,427)	(65,622, 72,392)
RetCorr	115	3,612	215
	(112, 118)	(3,555, 3,669)	(–93, 534)
PrLevel	–132	3,363	–3,619
	(–136, –128)	(3,171, 3,555)	(–5,390, –1,881)
TrPosition	14,598	38,656	39,875
	(14,245, 14,951)	(37,609, 39,723)	(37,246, 42,661)
TrCrowding	93	7,253	13,421
	(–253, 439)	(6,347, 8,180)	(11,435, 15,565)
Δ*ExpCCP*	14,949	63,887	118,914
Panel B: CCP exposure change decomposition in percentage			
RetVola	1.8	17.2	58.0
	(1.8, 1.9)	(16.6, 17.9)	(55.2, 60.9)
RetCorr	0.8	5.7	0.2
	(0.8, 0.8)	(5.6, 5.7)	(–0.1, 0.4)
PrLevel	–0.9	5.3	–3.0
	(–0.9, –0.9)	(5, 5.6)	(–4.5, –1.6)
TrPosition	97.7	60.5	33.5
	(95.3, 100)	(58.9, 62.2)	(31.3, 35.9)
TrCrowding	0.6	11.4	11.3
	(–1.7, 2.9)	(9.9, 12.8)	(9.6, 13.1)
Δ*ExpCCP*	100.0	100.0	100.0

Notes. This table presents the mean and, in parentheses, the lower and upper bounds of the (relative) share of components across alternative sequences of the various components. It serves as a robustness check for Table 4, which is based on a particular economically motivated sequence. The price and trade variables are kept together as a group so the number of sequences considered is $2 \times 3! \times 2! = 24$.

the sequencing of the various components. To verify how robust the decomposition results are to alternative sequences, we redo the analysis across all possible alternatives inspired by Hasbrouck (1995). As the components belong to two groups (that are preserved in the sequencing), we end up with $2 \times 3! \times 2! = 24$ possible sequences.

The results in Table E.3 show that the decomposition results appear robust. The table reports the mean and the lower and upper bounds of each component's contribution across all 24 sequences. The distance between the lower and upper bounds seems small, as it is only a few percentage points for the relative shares reported in panel B, never exceeding 6%. The key observations in the main text all hold up: the position component dominates all other components for the full sample, but volatility and crowding become much more important when considering only the top 100 and the top 10 exposure changes.

Endnotes

[1] On February 5, 2018, VIX futures jumped 20 points, which is the largest daily increase since the 1987 stock market crash. On October 7, 2016, the British pound dropped by almost 10% in just eight minutes. On January 15, 2015, the Swiss franc rose by about 20% against the euro within five minutes after the Swiss National Bank announced that it abandoned its peg against the euro as per immediately.

[2] On September 10, 2018, the Nordic-German power spread increased by more than 17 times the average daily change, which triggered the trader's default.

[3] Wagner (2011) clarifies that these arbitrageurs could hold diversified portfolios, yet be exposed to the fire-sale channel. It is position *diversity* that is the driving force here, not the level of diversification.

[4] Menkveld (2017) extends Duffie and Zhu (2011) to focus on the tail risk in losses as opposed to mean losses.

[5] One could argue that the tails are not riskier to a CCP because higher exposures against clearing members are insured by the latter posting higher margins with the CCP. Although this is true, it is also true that if there are losses that exceed the margin, they exceed it by a larger amount in the tail (i.e., loss given default is likely to be larger). A deeper analysis of risk *net* of margin and other forms of collateralization (e.g., the default fund) is beyond the scope of this study, as intraday margin and default-fund data are unavailable to us.

[6] We do not know how our findings compare with CCP exposures in the market for credit default swaps or for interest rate swaps for lack of evidence. We do, however, believe that our approach could be implemented for trading in these markets in spite of their different market structure (which features mostly over-the-counter trading as opposed to trading via a central limit order book).

[7] Bignon and Vuillemey (2020, Figs. 3 and A1) conduct a forensic analysis on the Paris commodity futures CCP that failed in 1974. They, for example, find that there was elevated activity (in terms of transactions) in the half year before failure, but open positions declined (measured in 1,000 tons of sugar).

[8] More precisely, Menkveld (2017) shows that there is a positive time series correlation between crowding and CCP exposure at a daily level. This paper, however, studies exposures *intradaily* and *decomposes* changes into all variables that enter the calculation of *ExpCCP*.

[9] A related set of papers that do not focus on concentration and systemic risk but rather on incentives and economic efficiency includes Koeppl et al. (2012), Acharya and Bisin (2014), Fontaine et al. (2014), Biais et al. (2016), and Huang (2019).

[10] Duffie and Zhu (2011) study the *mean* loss, which is invariant to the level of crowding—the VaR loss is not.

[11] The normality assumption yields analytic results for CCP exposure along with a natural decomposition. To stay close to normality in the data, the sample clock will run in volume time (for details, see Section 4.2).

[12] We include explicit formulas for all five components in Appendix B for completeness.

[13] Note that this is a narrower definition of crowding than the one that underlies the *CrowdIx* indicator in the Menkveld (2017) work. An increase in the correlations of security returns would lead to a higher level of *CrowdIx* because crowding in the Menkveld (2017) work is more broadly defined in terms of risk factors. We use the narrower definition here to distinguish between a change in security-return correlations and crowding due to position changes.

[14] σ_j is the square root of the jth diagonal element of the portfolio return covariance matrix Σ.

[15] Time subscripts are suppressed here for the sake of brevity.

[16] This equation corresponds to Menkveld (2017, Equation (27)). Note that there is a typo in (27), as $\sqrt{1/(2\pi)}$ should have been multiplied by σ_j instead of σ_j^2. This typo has been corrected in (12) here.

[17] ω_k is the square root of the kth diagonal element of the security-return covariance matrix Ω.

[18] The postcrisis European Market Infrastructure Regulation (EMIR) regulation in Europe requires a CCP to segregate trades on house accounts from those on client accounts as of 2013. Our data sample precedes this date, but EMCF had already implemented such segregation.

[19] More specifically, the conversion of portfolio-return correlations to portfolio-loss correlations is done with the M function in (15), which relies on assuming normality.

[20] In case of any residual trades due to imperfect grouping at the end of each day, they are included in the last bin.

[21] Easley et al. (2012) sample E-mini future returns based on volume clock and find similar evidence of partial recovery of normality.

[22] Given that overnight return variance is about four times the variance of an intraday 15-minute period, we update the covariance matrix after an overnight return R_{t-1} by $\Omega_t = (1 - 0.9984^4)\tilde{R}_{t-1}\tilde{R}'_{t-1} + 0.9984^4\Omega_{t-1}$, where $\tilde{R}_{t-1} = R_{t-1}/\sqrt{4}$.

[23] Note that aggregate loss is not normal because it is the sum of *truncated* normals.

[24] A review of the main events in this month is as follows. On May 5, mass protests erupted in Greece against the imposed austerity measures, with three deaths reported. This social unrest led to concerns that it could jeopardize the rescue package proposed by the European Union and the International Monetary Fund on May 2. To fund this intervention and future ones, the European Commission created the European Financial Stabilisation Mechanism on May 9 (European Commission 2010). On May 10, the European Central Bank announced the Securities Markets Program to address "dysfunctional" securities markets (European Central Bank 2010).

[25] For completeness, we also conducted these robustness analyses for the empirical results on the second and third hypotheses. Again, the results do not change qualitatively. To conserve space, we decided to only provide those robustness results upon request.

[26] Note that each element ω_{ij} of the covariance matrix Ω can be written as $\rho_{ij}\omega_i\omega_j$, where ρ_{ij} denotes the elements of the accompanying correlation matrix. This should clarify what homogeneity or a partial derivative with respect to ω_k is.

References

Acharya V, Bisin A (2014) Counterparty risk externality: Centralized vs. over-the-counter markets. *J. Econom. Theory* 149(January): 153–182.

Adrian T, Brunnermeier MK (2016) Covar. *Amer. Econom. Rev.* 106(7): 1705–1741.

Amini H, Filipović D, Minca A (2015) To fully net or not to net: Adverse effects of partial multilateral netting. *Oper. Res.* 64(5): 1135–1142.

Amini H, Filipović D, Minca A (2020) Systemic risk in networks with a central node. *SIAM J. Financial Math.* 11(1):60–98.

Andersen TG, Bollerslev T, Diebold FX, Labys P (2003) Modeling and forecasting realized volatility. *Econometrica* 71(2):579–625.

Ané T, Geman H (2000) Order flow, transaction clock, and normality of asset returns. *J. Finance* 55(5):2259–2284.

Beetsma R, Giuliodori M, De Jong F, Widijanto D (2013) Spread the news: The impact of news on the European sovereign bond markets during the crisis. *J. Internat. Money Finance* 34(April):83–101.

Benos E, Payne R, Vasios M (2020) Centralized trading, transparency and interest rate swap market liquidity: Evidence from the implementation of the Dodd-Frank Act. *J. Financial and Quant. Anal.* 55(1):159–192.

Bhanot K, Burns N, Hunter D, Williams M (2014) News spillovers from the Greek debt crisis: Impact on the eurozone financial sector. *J. Banking Finance* 38(January):51–63.

Biais B, Heider F, Hoerova M (2016) Risk-sharing or risk-taking? Counterparty risk, incentives, and margins. *J. Finance* 71(4): 1669–1698.

Bignon V, Vuillemey G (2020) The failure of a clearinghouse: Empirical evidence. *Rev. Finance* 24(1):99–128.

Bollerslev T (1990) Modelling the coherence in short-run nominal exchange rates: A multivariate generalized arch model. *Rev. Econom. Statist.* 72(3):498–505.

Brunnermeier MK, Pedersen LH (2009) Funding liquidity and market liquidity. *Rev. Financial Stud.* 22(6):2201–2238.

Candelon B, Sy MAN, Arezki MR (2011) Sovereign rating news and financial markets spillovers: Evidence from the European debt crisis. Working Paper No. 11–68, International Monetary Fund, Washington, DC.

Capponi A, Cheng WA, Rajan S, et al (2019) Firm capital dynamics in centrally cleared markets. *Math. Finance*, ePub ahead of print November 6, 2019, https://doi.org/10.1111/mafi.12236.

Clark PK (1973) A subordinated stochastic process model with finite variance for speculative prices. *Econometrica* 41(1):135–155.

Committee on Payments and Market Infrastructures (CPMI) and International Organization of Securities Commissions (IOSCO) (2012): Principles for financial market infrastructures, April. https://www.bis.org/cpmi/publ/d101a.pdf.

Committee on Payments and Market Infrastructures (CPMI) and International Organization of Securities Commissions (IOSCO) (2017) Resilience of central counterparties (CCPs): Further guidance on the PFMI. Accessed July 31, 2020, https://www.bis.org/cpmi/publ/d163.pdf.

Cox RT (2015) Central counterparties in crisis: The Hong Kong futures exchange in the crash of 1987. *J. Financial Market Infrastructures* 4(2):73–98.

Crego JA (2019) Why does public news augment information asymmetries? Working paper, Tilburg University, Tillburg, Netherlands.

Daniel C (1973) One-at-a-time plans. *J. Amer. Statist. Assoc.* 68(342): 353–360.

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

Duffie D, Scheicher M, Vuillemey G (2015) Central clearing and collateral demand. *J. Financial Econom.* 116(2):237–256.

Easley D, López de Prado MM, O'Hara M (2012) Flow toxicity and liquidity in a high-frequency world. *Rev. Financial Stud.* 25(5): 1457–1493.

Engle R (2002) Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *J. Bus. Econom. Statist.* 20(3):339–350.

European Central Bank (2010) ECB decides on measures to address severe tensions in financial markets. Press release, European Central Bank, Frankfurt, Germany.

European Commission (2010) The European stabilization mechanism. Memo /10/173, European Commission, Brussels, Belgium.

Fontaine JS, Perez-Saiz H, Slive J (2014) How should central counterparty clearing reduce risk? Collateral requirements and entry restrictions. Report, Bank of Canada, Ottawa, ON, Canada.

Glasserman P, Moallemi CC, Yuan K (2016) Hidden illiquidity with multiple central counterparties. *Oper. Res.* 64(5):1053–1176.

Gromb D, Vayanos D (2002) Equilibrium and welfare in markets with financially constrained arbitrageurs. *J. Financial Econom.* 66(2–3): 361–407.

Hansen PR, Lunde A (2006) Realized variance and market microstructure noise. *J. Bus. Econom. Statist.* 24(2):127–161.

Hasbrouck J (1995) One security, many markets: Determining the contributions to price discovery. *J. Finance* 50(4):1175–1199.

Huang W (2019) Central counterparty capitalization and misaligned incentives. Working paper, Bank for International Settlements, Basel, Switzerland.

Jones RA, Perignon C (2013) Derivatives clearing, default risk and insurance. *J. Risk Insurance* 80(2):373–400.

Jones C, Kaul G, Lipson M (1994) Information, trading, and volatility. *J. Financial Econom.* 36(1):127–154.

Khandani AE, Lo AW (2007) What happened to the quants in August 2007? *J. Investment Management* 5(4):5–54.

Khandani AE, Lo AW (2011) What happened to the quants in August 2007? Evidence from factors and transactions data. *J. Financial Markets* 14(1):1–46.

Kim O, Verrecchia RE (1994) Market liquidity and volume around earnings announcements. *J. Accounting Econom.* 17(1–2):41–67.

Koeppl T, Monnet C, Temzelides T (2012) Optimal clearing arrangements for financial trades. *J. Financial Econom.* 103(1): 189–203.

Loon YC, Zhong ZK (2014) The impact of central clearing on counterparty risk, liquidity, and trading: Evidence from the credit default swap market. *J. Financial Econom.* 112(1): 91–115.

Loon YC, Zhong ZK (2016) Does dodd-frank affect otc transaction costs and liquidity? Evidence from real-time cds trade reports. *J. Financial Econom.* 119(3):645–672.

Lopez JAC, Harris JH, Hurlin C, Pérignon C (2017) Comargin. *J. Financial Quantitative Anal.* 52(5):2183–2215.

Menkveld AJ (2016) Systemic risk in central clearing: Should crowded trades be avoided? Working paper, Vrije Universiteit, Amsterdam.

Menkveld AJ (2017) Crowded positions: An overlooked systemic risk for central clearing parties. *Rev. Asset Pricing Stud.* 7(2): 209–242.

Menkveld AJ, Pagnotta E, Zoican MA (2015) Does central clearing affect price stability? Evidence from nordic equity markets. Working paper, Imperial CollegeLondon, London.

Mink M, De Haan J (2013) Contagion during the Greek sovereign debt crisis. *J. Internat. Money Finance* 34(April):102–113.

Shleifer A, Vishny RW (1997) The limits of arbitrage. *J. Finance* 52(1): 35–55.

Stein JC (2009) Presidential address: Sophisticated investors and market efficiency. *J. Finance* 64(4):1517–1548.

Wagner W (2011) Systemic liquidation risk and the diversity–diversification trade-off. *J. Finance* 66(4):1141–1175.