Credit and Funding Risk from CCP trading

Leif Andersen
Bank of America Merrill Lynch

.

Joint work with A. Dickinson

April 9, 2018

Agenda

1. Introduction
2. Theory
3. Application to Client Cleared Portfolios
4. Application to House Portfolios
5. Caveats & Extensions

Introduction - (1)

- Along with the regulatory push for derivatives to be cleared on central clearing houses (=central counterparties, or *CCPs*), there are both business and regulatory interest in better quantification of the credit exposures that arise from trading on CCPs.
- CCPs have a variety of risk-based (e.g., VaR) margin policies in place, so to understand credit exposures on CCPs it is important to understand how margin affects standard credit exposure calculations.
- This is, by the way, also true for OTC deals that are traded under the new uncleared margin rules (UMR), with the ISDA-SIMM margin calculator.
- This is complicated numerically and practically: many CCPs reveal very limited (and often stale) information about their positions, margin policies, default fund methodologies, etc.

Introduction - (2)

- Rather than try to solve this complex problem with an equally complex model that is hopeless to parametrize, we will here try to cut through the clutter by making some "clever" minimal assumptions about margin policy, CCP risk concentrations, and portfolio distributions.
- This requires a bit of risk measure theory, which we will breeze through first.
- We will also briefly touch upon questions of the *optimal* setup of a CCP.
- *Caution*: ambitious to cover the topic of CCP credit and funding risk in a single talk! The talk is simplified; see Andersen and Dickinson (2018) for a more complete/general treatment[1].

---

[1] Andersen, L. and A. Dickinson, "Funding and Credit Risk with Locally Elliptical Portfolio Processes: an Application to CCPs," Working Paper, www.ssrn.com

Theory: Risk Measures - (1)

- Given a process $X(t)$ and a "lag" δ, define process increments (measure $\mathbb{P}$, filtration $\mathcal{F}_t$)

$$\Delta X(t) = X(t+\delta) - X(t).$$

- Define a time-invariant Risk Measure $\mathcal{R}$ on the distribution of increments, such that $\mathcal{R}\left(\Delta X(t)\right) = \mathcal{R}\left(\Delta X(t); \mathbb{P}, \mathcal{F}_t\right)$ is t-measurable.
- It is natural to assume that risk measures scale linearly with the size (= notional) of risk positions, so we introduce:

Definition

A risk measure $\mathcal{R}$ is *homogenous*, if for any scalar $\gamma > 0$

$$\begin{aligned} \mathcal{R}\left(\Delta X(t)\right) &\geq 0, \\ \mathcal{R}\left(\gamma \cdot \Delta X(t)\right) &= \gamma \cdot \mathcal{R}\left(\Delta X(t)\right). \end{aligned}$$

Theory: Risk Measures - (2)

- The following are examples of homogenous risk measures:
  - $\mathcal{R}(\Delta X(t)) = \mathbb{E}_t(\Delta X(t)^+)$;
  - $\mathcal{R}(\Delta X(t)) = \mathrm{stdev}_t(\Delta X(t))$ (standard deviation);
  - $\mathcal{R}(\Delta X(t)) = \mathrm{VaR}_t^p(\Delta X(t))$ (value-at-risk at conf. level p);
  - $\mathcal{R}(\Delta X(t)) = \mathrm{ES}_t^p(\Delta X(t))$ (expected shortfall at level p).
- The homogenous risk measure class is larger than the *coherent* risk measure class (of Arzner *et al.* fame).
- Given a homogenous risk measure $\mathcal{R}$, we can construct a new risk measure in many ways. For instance, if $\eta > 0$ is a constant, the following is also a homogenous risk measure:

$$\mathbb{E}_t\left((\Delta X(t) - \eta\mathcal{R}(\Delta X(t)))^+\right). \tag{1}$$

- The measure (1) is particularly relevant for CCP exposure work, as we shall see.

Theory: Locally Elliptical Processes - (1)

- Make $X(t)$ N-dimensional, $X(t) = (X_1(t), \ldots, X_N(t))^\top$.

Definition

Let $C = (C_1, \ldots, C_N)^\top$ be spherically distributed with zero mean. $\Delta X(t)$ is *locally elliptical with zero mean (LEZM)*, if for some t-measurable stochastic scale matrix $A(t) \in \mathbb{R}^{N \times N}$

$$\Delta X(t) \stackrel{d}{=} A(t) C.$$

Lemma

If $\Delta X(t)$ is LEZM, then the $\mathcal{F}_t$-conditional char. function φ_t is

$$\varphi_t(k) \triangleq \mathbb{E}_t\left(e^{ik^\top \Delta X(t)}\right) = g\left(k^\top \Sigma(t) k\right), \quad k \in \mathbb{R}^N,$$

where $g : \mathbb{R} \to \mathbb{R}$ is a time-homogenous generator function and

$$\Sigma(t) = A(t) A(t)^\top.$$

Theory: Locally Elliptical Processes - (2)

- Examples: Gaussian, Student t, Cauchy, Laplace, stable, logistic, hyperbolic, variance mixtures, etc. – all with stochastic volatility.
- The Gaussian LEZM can model any Ito diffusion, since an N-dim Ito process (σ stochastic $N \times N$ matrix)

$$dX(t) = \sigma(t)dW(t)$$

will be approximately

$$\Delta X(t) \overset{d}{\simeq} \sqrt{\delta} \cdot \sigma(t) Z$$

where Z is a vector of N independent standard Gaussian r.v's (and therefore spherical).

Theory: Locally Elliptical Processes - (3)

- Using non-Gaussian processes is akin to using a Levy jump process as the driver. This allows us to model fat tails, for instance.
- A good choice in practice

$$\Delta X(t) \stackrel{d}{=} A(t)T$$

  where T is a vector of N independent Student t variables with $\nu > 2$ degrees of freedom (and $A(t)$ again can be stochastic, but t-measurable).
- The parameter ν allows us to crank tail fatness up relative to the Gaussian case (by lowering ν, of course). $\nu = \infty$ is the Gaussian limit.

Theory: LEZMs and Homogenous Risk Measures - (1)

- There are a number of interesting properties of homogenous risk measures applied to LEZM processes. See Andersen-Dickinson (2018) for complete description.
- One important feature is this:

Lemma

Let $\Delta X(t)$ be an LEZM, and let $\mathcal{R}, \tilde{\mathcal{R}}$ be homogenous risk measures. Then, for arbitrary $t, t' \geq 0$ and arbitrary $n, m \in \{1, \ldots, N\}$,

$$\frac{\mathcal{R}\left(\Delta X_n(t)\right)}{\tilde{\mathcal{R}}\left(\Delta X_n(t)\right)} = \frac{\mathcal{R}\left(\Delta X_n(t')\right)}{\tilde{\mathcal{R}}\left(\Delta X_n(t')\right)},$$
$$\frac{\mathcal{R}\left(\Delta X_m(t)\right)}{\tilde{\mathcal{R}}\left(\Delta X_m(t)\right)} = \frac{\mathcal{R}\left(\Delta X_n(t)\right)}{\tilde{\mathcal{R}}\left(\Delta X_n(t)\right)}.$$

- According to the Lemma, *ratios* of homogenous risk measures are invariant across time and across components of X. This is handy.

Theory: LEZMs and Homogenous Risk Measures - (2)

- A related handy result concerns the risk measure (1), which we restate below:

$$\begin{aligned} \mathcal{R}_{max}\left(\Delta X_n(t)\right) &= \mathcal{R}_{max}\left(\Delta X_n(t); \eta, \tilde{\mathcal{R}}\right) \\ &= \mathbb{E}_t\left(\left(\Delta X_n(t) - \eta \cdot \tilde{\mathcal{R}}(\Delta X_n(t))\right)^+\right). \end{aligned} \tag{2}$$

- We shall be interested in relating values of the risk measure when $\eta = 0$ to values of the risk measure for arbitrary η.
- For this, define first the time 0 observable ratio (note that $\eta = 0$ in the denominator:

$$q = \frac{\tilde{\mathcal{R}}\left(\Delta X_n(0)\right)}{\mathcal{R}_{max}\left(\Delta X_n(0); 0, \tilde{\mathcal{R}}\right)}. \tag{3}$$

Theory: LEZMs and Homogenous Risk Measures - (3)

Lemma

Consider the risk measure (2) and let q be as in (3) with $\eta > 0$. If $\mathcal{R}_{max}$,$\tilde{\mathcal{R}}$ are homogenous and $\Delta X_n(t)$ is LEZM, then

$$\mathcal{R}_{max}\left(\Delta X_n(t); \eta, \tilde{\mathcal{R}}\right) = \mathcal{R}_{max}\left(\Delta X_n(t); 0, \tilde{\mathcal{R}}\right) \cdot \Omega(q \cdot \eta), \qquad (4)$$

where Ω is a function independent of $\tilde{\mathcal{R}}$ that only depends on the parameters in the generating function g of the LEZM.

- So: if we know the function Ω and the "easy" profile $\mathbb{E}_t\left((\Delta X_n(t))^+\right)$, $\forall t$, then all we need is the time 0 value of the (arbitrary) risk measure $\tilde{\mathcal{R}}\left(\Delta X_n(0)\right)$ in order to compute the entire profile $\mathbb{E}_t\left(\left(\Delta X_n(t) - \eta \cdot \tilde{\mathcal{R}}(\Delta X_n(t))\right)^+\right)$, $\forall t$.
- There is a generic result for Ω. Here we show only the Student t (and Gaussian) LEZM:

Theory: LEZMs and Homogenous Risk Measures - (4)

Lemma

Assume that $\mathcal{R}^0$ is the expected positive part and that the risk driver distribution μ_δ is Student t with $\nu > 2$ degrees of freedom. Then,

$$\Omega(q) = \frac{\mathcal{J}_\nu\left(q\mathcal{J}_\nu\left(0\right)\right)}{\mathcal{J}_\nu\left(0\right)} = \frac{\left(\nu-1\right)\mathcal{J}_\nu\left(\frac{q\nu f_\nu(0)}{\nu-1}\right)}{\nu f_\nu\left(0\right)}, \qquad (5)$$

where

$$\mathcal{J}_\nu\left(x\right) = \frac{\nu + x^2}{\nu - 1} f_\nu\left(x\right) - xF_\nu\left(-x\right) \qquad (6)$$

and f_ν, F_ν are the PDF and CDF (resp.) of the standard Student t distribution with ν degrees of freedom. The Gaussian limit is:

$$\lim_{\nu\to\infty}\Omega(q) = -q\Phi\left(-\frac{q}{\sqrt{2\pi}}\right) + \sqrt{2\pi}\phi\left(\frac{q}{\sqrt{2\pi}}\right).$$

Theory: LEZMs and Homogenous Risk Measures - (5)

Figure: Exposure Scale Ω vs. Degrees of Freedom ν

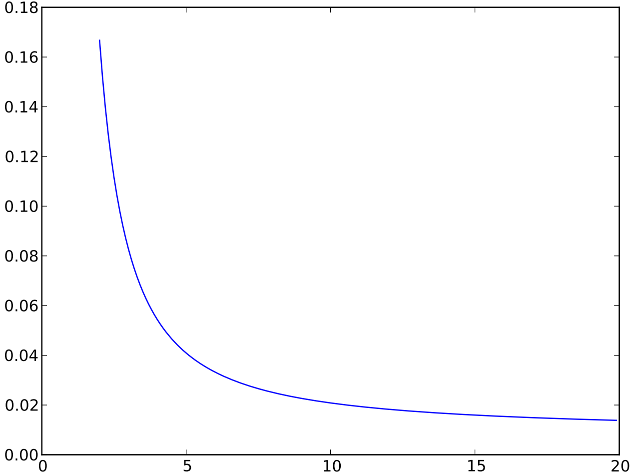

- Parameters: $\eta = 1$ and $q = 0.1715$.

CCP Exposures: Overview - (1)

- In CCP trading, the CCP is basically the counterparty to all trading agreements:

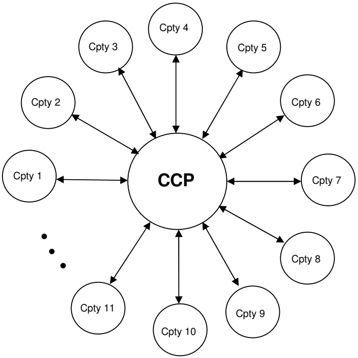

CCP Exposures: Overview - (2)

- The CCP has members (so-called clearing members, or *CMs*), who are associated with subsets of the total portfolio of trades. This association comes in two forms:
  - A member's own trades with the CCP (the *house trades*)
  - A list of client portfolios that the member has cleared with the CCP, and is servicing on behalf of the CCP (*client cleared trades*).
- Credit exposures arise from both of these position types. We start with the latter.

Client Clearing - (1)

- Assume that a client of a CM has a portfolio of trades with values $V_c(t)$, as seen from the CCP's perspective.
- The CM is responsible for covering any losses to the CCP should the client default.
- Two separate collateral mechanisms are used to limit the exposure: **variation margin** and **initial margin**. Let the time t value of each of these be $M_c(t)$ and $I_c(t)$, respectively. Variation margin always tracks the value of the portfolio, so $M_c(t) = V_c(t)$.
- If a client default takes place at time t, it will take a period of δ to settle the claim; on $[t, t + \delta]$ margins are *not* updated. For most CCPs: δ is around 2-10 bdays.
- The exposure associated with a default at time t is therefore:

$$E_c(t) = (V_c(t + \delta) - (M_c(t) + I_c(t)))^+ = (\Delta V_c(t) - I_c(t))^+ .$$

Client Clearing - (2)

- For CVA calculations we need the expectation profile

$$\mathrm{EE}_c(t) \triangleq \mathbb{E}\left(E_c(t)\right) = \mathbb{E}\left(\left(\Delta V_c(t) - I_c(t)\right)^+\right).$$

- The initial margin $I_c(t)$ updates dynamically to cover settlement risk on $[t, t+\delta]$, and is often based on VaR or ES – but there is rarely much transparency into the CCPs policies. This makes computation of $\mathrm{EE}_c(t)$ very challenging.
- But if I_c is computed from a risk-measure $\mathcal{R}_I$, then

$$\mathbb{E}_t\left(E_c(t)\right) = \mathcal{R}_{max}\left(\Delta C_c(t); 1, \mathcal{R}_I\right).$$

- So if V_c is driven by an LEZM and the CCP uses a homogenous risk measure to compute I_c, then from (4)

$$\begin{aligned}\mathrm{EE}_c(t) &= \mathbb{E}\left(\mathbb{E}_t\left(E_c(t)\right)\right) = \Omega(q)\mathbb{E}\left(\mathcal{R}_{max}\left(\Delta C_c(t); 0, \mathcal{R}_I\right)\right) \\ &= \Omega(q)\mathrm{EE}_c^*(t),\end{aligned}$$

where

$$\mathrm{EE}_c^*(t) \triangleq \mathbb{E}\left(\Delta V_c(t)^+\right), \quad q = \frac{I_c(0)}{\mathrm{EE}_c^*(0)}.$$

Client Clearing - (3)

- In this expression, we always know what $I_c(0)$ is, and we can always run our standard variation-margin-only Monte Carlo simulations to establish the "usual" expected exposure profile $\mathrm{EE}_c^*(t)$.
- If we assume that the portfolio is locally Gaussian or locally Student t, we get $\Omega(q)$ explicitly from earlier results.
- If the client has a recovery rate of R_c and a forward hazard rate function of $\lambda_f^c(t)$, the exposure profile $\mathrm{EE}_c(t)$ can be turned into a (unilateral) CVA number by the usual computation (P is the discount function):

$$\mathrm{CVA}_c = (1 - R_c) \int_0^\infty P(t) e^{-\int_0^t \lambda_f^c(u)\, du} \mathrm{EE}_c(t) \times \lambda_f^c(t)\, dt.$$

House Trades - (1)

- By executing house trades with the CCP, a CM also has exposure to the CCP itself. More specifically, a CM has exposure to defaults of other CMs, since default losses from the house positions of a CM get allocated to surviving CMs.
- Assume that there are $M+1$ CMs, $m = 0, \ldots, M$; we shall generally take the perspective of the 0th CM (w.l.o.g).
- There are several protection mechanisms in play to protect surviving CMs from a clearing member default :
  - Each CM posts variation and initial margin to the CCP. Let $I_m(t)$ be the initial margin of the mth CM.
  - Each CM contributes a certain amount $D_m(t)$ to a *default fund*.
  - The CCP has a (small) amount of equity. We ignore this.

House Trades - (2)

- Define exposure originating from the mth CM as

$$E_m(t) = (\Delta V_m(t) - I_m(t) - D_m(t))^+$$

and let the mth CM have a hazard rate forward curve $\lambda_f^m(t)$.

- Any default losses get allocated to surviving CMs by the size of their default fund contributions, so the effective exposure, as seen by CM 0, to a default of a CM is (approximately)

$$\frac{D_0(t)}{D_{tot}(t)} E_m(t), \quad D_{tot}(t) \triangleq \sum_{m=0}^{M} D_m(t).$$

- Under certain technical assumptions about wrong-way risk, recovery, and about how defaulting members get replaced by new members after default, one can show that the following is a meaningful definition of CM 0's CVA:

$$\mathrm{CVA}_0 = \int_0^\infty P(t) \mathbb{E}\left(\frac{D_0(t)}{D_{tot}(t)} \sum_{m=1}^{M} E_m(t)\, \lambda_f^m(t) \right) dt.$$

House Trades - (3)

- Unfortunately, we are severely information-constrained as it pertains to this calculation, since CCPs publish very little information. Basically, all that CM 0 knows at time 0 is:
  - Its own time 0 initial margin $I_0(0)$ and its own default contribution $D_0(0)$.
  - Its own portfolio process $V_0(t)$, $t \geq 0$ (we assume zero new investments through time, as usual)
  - The total initial margin $I_{tot}(0) = \sum I_m(t)$ (albeit at a lag)
  - The total default fund $D_{tot}(0)$ (albeit at a lag)
  - The identities of the other CMs at time 0.
- It does NOT know:
  - Initial margin and default fund contributions of other CMs
  - The house portfolios (and the portfolio processes $V_m(t)$) of the other CMs.
- Seems pretty hopeless to execute the CVA_0 calculation...

House Trades - (4)

- Rather than relying on heroic assumptions about unknown CM portfolios etc., we will use the scaling relationships of our Theory section to get through the calculation.

Assumption

1. $\Delta V(t) = (\Delta V_0, \ldots \Delta V_M(t))^\top$ *is LEZM.*
2. $D_m(t) = \mathcal{R}_D (\Delta V_m(t))$ *and* $I_m(t) = \mathcal{R}_I (\Delta V_m(t)), \forall m,$ *where* $\mathcal{R}_D$,$\mathcal{R}_I$ *are homogenous risk measures.*

- It is clear that $D_m + I_m$ is also a homogenous risk measure, as is $E_m(t) = (\Delta V_m(t) - D_m - I_m)^+$. By the ratio invariance of homogenous risk measures on LEZMs, it follows that for all t,

$$\frac{D_0(t)}{D_{tot}(t)} = \frac{E_0(t)}{\sum_{m=0}^{M} E_m(t)}.$$

House Trades - (5)

- If we insert this into the CVA expression, we get

$$\mathrm{CVA}_0 = \int_0^\infty P(t)\mathbb{E}\left(E_0(t)\frac{\sum_{m=1}^M E_m(t)\,\lambda_f^m(t)}{\sum_{m=0}^M E_m(t)}\right)\,dt$$
$$\approx \int_0^\infty P(t)\mathbb{E}\left(E_0(t)\bar{\lambda}_f(t)\right)\,dt$$

where (notice numerator sum range)

$$\bar{\lambda}_f(t) = \frac{\sum_{m=1}^M E_m(t)\,\lambda_f^m(t)}{\sum_{m=1}^M E_m(t)}.$$

- $\bar{\lambda}_f(t)$ is essentially an exposure-weighted average of the CM hazard rates for $m = 1, \ldots, M$, and will likely not depend much on exposure moves.

House Trades - (6)

- $\bar{\lambda}_f(t)$ can be approximated deterministically as something like

$$\bar{\lambda}_f(t) \approx \hat{\lambda}_f(t) \triangleq \frac{\sum_{m=1}^{M} c_m \, \lambda_f^m(t)}{\sum_{m=1}^{M} c_m}. \qquad (7)$$

  where the c_m are constants that express the guessed average relative magnitude of the clearing members house positions.
- These constants can be found pragmatically from league tables, market caps, or similar.
- The precise method is likely non-critical, since the λ_f^m curves rarely disperse too much.
- In the CVA expression, all that is missing is now $\mathbb{E}\left(E_0(t)\right)$.

House Trades - (7)

- Risk measure ratio invariance for LEZMs shows that

$$\frac{I_0(t)}{D_0(t)} = \frac{I_0(0)}{D_0(0)} \triangleq \alpha.$$

- Defining the variation-margin-only profile $\mathrm{EE}^*(t) = \mathbb{E}\left(\Delta V_0(0)^+\right)$, we get (as for client-cleared trades)

$$\begin{aligned}\mathbb{E}\left(E_0(t)\right) &= \mathbb{E}\left(\left(\Delta V_0(t) - (1+\alpha)I_0(t)\right)^+\right) \\ &= \mathrm{EE}^*(t)\Omega\left(q(1+\alpha)\right),\end{aligned}$$

where

$$q = \frac{I_0(0)}{\mathrm{EE}^*(0)}.$$

House Trades - (8)

- So **to summarize**, our expression for CM 0's CVA to the other CMs of the CCP is simply:

$$\mathrm{CVA}_0 = \Omega\left(q(1+\alpha)\right) \int_0^{\infty} P(t)\mathrm{EE}^*(t)\hat{\lambda}_f(t)\, dt,$$

where $\mathrm{EE}^*(t)$ is the variation-margin-only exposure profile, $\hat{\lambda}_f(t)$ is given in (7), and

$$\alpha = \frac{I_0(0)}{D_0(0)}, \quad q = \frac{\mathrm{EE}^*(0)}{I_0(0)}.$$

- As long as we can compute the variation-margin-only exposure profile $\mathrm{EE}^*(t), \forall t$ – which we can – and know $I_0(0)$ and $D_0(0)$ – which we do – we are done!

Caveats & Extensions - (1)

- In practice, CCP margin calculations are only "partially" dynamic, as they are often (by design) slow to update their volatility estimates.
- This has ramifications for stress testing and can be incorporated into the model (see Andersen/Dickinson).
- In practice, default fund calculations will consider both house and client positions of each CM (not just house positions).
- This can *also* be incorporated into the model, without too much difficulty (Andersen/Dickinson)
- We can also incorporate contagion effects, as well as take into account more accurately how defaulting members do (or do NOT) get replaced on default.
- We can also more accurately model distribution of member size/risk inside a CCP; in particular, we can incorporate the so-called *Herfindal-Hirschmann* index into the CVA calculation.

Caveats & Extensions - (2)

- From the CM perspective, one might think that the CMs want initial margin as high as possible – this will limit their exposure to both other CMs and to clients. However, there is a cost – the so-called *MVA* (margin valuation adjustment) – to the CMs from having themselves post more margin.
- As it turns out, this cost – which hits the CM shareholders directly – is much (much) bigger than the CVA credit, for a typical CCP.
- This trade-off can be computed explicitly in our framework, and a cost-optimal margin policy can be identified.
- One can also contemplate the trade-offs of margin-vs.-default-fund protection.
- These policy topics require a bit more work, and is left for another day! Andersen-Dickinson has some results for those curious.