

Risk management for whales

Rama Cont

Lakshithe Wagalath

Imperial College
London

IESEG School of Management
(LEM - CNRS), France

2015

Abstract

We propose a portfolio risk model which integrates market risk with liquidation costs. The model provides a framework for computing liquidation-adjusted risk measures such as Liquidation-adjusted VaR (LVaR). Calculation of liquidation-adjusted Value-at-Risk (LVaR) for simulated and real-life examples reveals a substantial impact of liquidation costs on portfolio risk for portfolios with large concentrated positions.

1 Liquidation risk

Quantitative models commonly used in financial risk management have mainly focused on the statistical modeling of variations in the (mark-to-)market value of financial portfolios, in order to estimate a risk measure – such as Value-at-Risk or Expected shortfall – related to market losses over a given time horizon. These risk measures are then used for determining for example capital requirements or margin requirements in order to provision for losses in extreme risk scenarios. Typically, when such losses materialize, the financial institution is led to liquidate a sizable portion of its portfolio and the realized liquidation value may be quite different from the

(pre-liquidation) market value used in the model. The difference – which represents the *liquidation cost* – can be significant if the portfolio contains large, concentrated positions. Not accounting for this liquidation cost in risk calculations may result in a serious underestimation of portfolio losses in a stress scenario.

Several risk management fiascos have been associated with the underestimation of liquidation costs and liquidity risk for portfolios with large positions. A spectacular example was provided by the JP Morgan CIO losses in 2012, when the bank suffered a \$6.2 Bn loss while unwinding CDS index positions amounting to several hundred billions of dollars in gross notional [U.S. Senate, 2013, JP Morgan, 2014]. Although JP Morgan deployed a sophisticated Value-at-Risk model for monitoring the CIO portfolio, the VaR calculation was focused on market losses and did not anticipate the market moves generated by the CIO's subsequent asset liquidations.

These considerations call for a comprehensive approach for integrating liquidation risk into portfolio risk measures; this issue is particularly relevant for financial institutions managing large portfolios.

2 A model for liquidation losses

Consider a portfolio with positions in n assets classes, whose values at date $t_k = k\Delta t$ are denoted S_k^1, \dots, S_k^n . We assume that, in the absence of systematic effects from large trades, the “fundamental” return of asset i in period $[t_k, t_{k+1}]$ is given by a random variable ϵ_k^i , where the random vectors $(\epsilon_k^1, \dots, \epsilon_k^n)_{k \leq 0}$ are assumed to be independent, with mean $m\Delta t$ and covariance matrix $\Sigma\Delta t$:

$$\frac{S_{k+1}^i - S_k^i}{S_k^i} = \epsilon_{k+1}^i \quad \text{where}$$

$$E(\epsilon_k^i) = m_i\Delta t, \quad \text{Cov}(\epsilon_k^i, \epsilon_k^j) = \Sigma_{i,j}\Delta t.$$

The *fundamental covariance* matrix Σ captures structural relations between asset returns. In the examples below, Σ will be taken constant but in practice one may also incorporate some dynamics for Σ in the model.

In absence of market impact, the value of a buy-and-hold portfolio with α_i shares of asset i , $i = 1..n$, is given by:

$$V_k = \sum_{i=1}^n \alpha_i S_k^i$$

and changes according to

$$V_{k+1} - V_k = \sum_{i=1}^n \alpha_i S_k^i \epsilon_{k+1}^i$$

Institutional portfolios are subject to *constraints*, often defined as limits on balance sheet ratios: capital ratio, liquidity ratio, leverage ratio; other examples are performance constraints, such as limits on drawdowns. Following a large loss in asset values, the constraint may be breached, in which case the portfolio may need to

be *deleveraged*, i.e. some assets need to be sold over a short time period in order to comply with the constraint. This is the phenomenon of distressed asset sales or *fire sales* [Cont and Wagalath, 2013, Shleifer and Vishny, 2011].

Consider for example the case of a portfolio with equity/capital E and subject to a leverage constraint L_{\max} , representing the maximum allowed leverage ratio. Prior to a stress scenario, the leverage is given by:

$$\frac{V}{E} \leq L_{\max}.$$

In the event of a loss of l (%), the leverage ratio increases from $\frac{V}{E}$ to

$$\frac{V(1-l)}{E-lV}$$

It is then straightforward to show that if the loss exceeds a threshold

$$l^* = \frac{1}{V} \frac{EL_{\max} - V}{L_{\max} - 1} < l < \frac{E}{V}$$

then the leverage constraint L_{\max} is breached and the portfolio needs to be deleveraged: a portion $1 - \frac{L_{\max}(E-lV)}{V(1-l)}$ of holdings needs to be liquidated. The volume of assets sold thus depends on the magnitude of the loss: Figure 1 shows the volume of liquidation for a portfolio with initial leverage $\frac{V}{E} = 25$ and leverage limit $L_{\max} = 33$, as a function of the loss l (blue line). In practice, the deleveraging policy may deviate from this simple linear example, but its qualitative features remain valid: the volume of asset sales is zero for losses below a threshold, increases with the loss size and saturates beyond a certain loss level (which represents total liquidation). We represent this through a *response function* f , which represents the proportion of the portfolio which

is deleveraged, as a function of the portfolio loss (red curve in Figure 1). The fraction of the portfolio liquidated at period k in response to price moves is then $f\left(\frac{V_k}{V_0}\right) - f\left(\frac{V_k}{V_0} + \sum_{i=1}^n \frac{\alpha_i S_k^i}{V_0} \epsilon_{k+1}^i\right)$. We assume that assets are liquidated proportionally to the initial holdings (this assumption may be relaxed, see below).

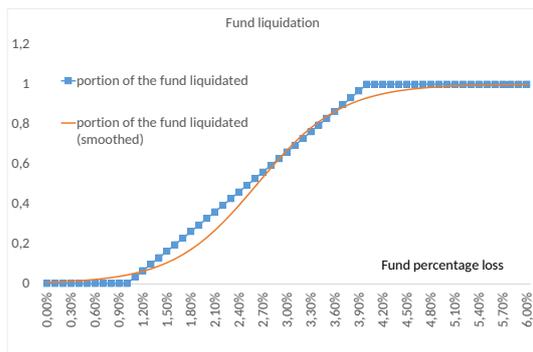


Figure 1: Volume of liquidation as a function of portfolio loss.

Liquidation of large volumes of assets has an impact on the market price [Almgren and Chriss, 2000, Kyle, 1985]: assuming this impact to be linear [Cont et al., 2014, Kyle, 1985, Obizhaeva, 2012], this leads to the following price dynamics, where the new terms correspond to the price impact of deleveraging:

$$\underbrace{-\frac{\alpha_i}{D_i} \left(f\left(\frac{V_k}{V_0}\right) - f\left(\frac{V_k}{V_0} + \sum_{j=1}^n \frac{\alpha_j S_k^j}{V_0} \epsilon_{k+1}^j\right) \right)}_{\text{Feedback from liquidation}} = \underbrace{\epsilon_{k+1}^i}_{\text{Fundamental return}}$$

where D_i represents a liquidity parameter, or market depth, for asset class i , estimated using the methodology proposed in [Obizhaeva, 2012].

Equation (1) gives a decomposition of the asset returns into a “fundamental” component and an endogenous – or self-induced – component which is generated by the fund’s own deleveraging and depends on asset liquidity. This endogenous component is zero in “normal” scenarios, but when the portfolio experiences large losses, leading to partial liquidation, this term may become non-zero, generating larger-than-expected portfolio losses and an increase in observed correlations, as described below.

As shown in [Cont and Wagalath, 2013], this model has a continuous-time limit described by a multi-asset diffusion (local volatility) model

$$\frac{dP_t^i}{P_t^i} = \mu_t^i dt + (\sigma_t dW_t)_i \quad 1 \leq i \leq n$$

where the drift μ_t^i and the instantaneous covariance $c_t = \sigma_t \sigma_t'$ are given by

$$\begin{aligned} \mu_t^i &= m_i + \frac{\Lambda_i}{2} f''\left(\frac{V_t}{V_0}\right) \frac{\langle \pi_t, \Sigma \pi_t \rangle}{V_0^2} \\ c_t &= \sigma_t \sigma_t' = \Sigma + \frac{1}{V_0} f'\left(\frac{V_s}{V_0}\right) [\Lambda \pi_s' \Sigma + \Sigma \pi_s \Lambda'] \\ &\quad + \frac{1}{V_0^2} (f')^2\left(\frac{V_s}{V_0}\right) (\pi_s' \Sigma \pi_s) \Lambda \Lambda' \end{aligned} \quad (2)$$

where $\pi_t = (\alpha_1 P_t^1, \dots, \alpha_n P_t^n)'$ is the dollar allocation of the portfolio and $\Lambda = {}^{(1)}\left(\frac{\alpha_1}{D_1}, \dots, \frac{\alpha_n}{D_n}\right)'$ represents the positions expressed as a fraction of market depth. Here we have denoted by v' the transpose of a vector (or matrix) v .¹

Equation (2) shows that the dependence structure of asset returns is (temporarily) modified during liquidation: the realized covariance matrix is equal to the

¹For detailed mathematical proofs we refer to [Cont and Wagalath, 2013].

fundamental covariance matrix Σ plus a (liquidity-dependent) excess covariance term which depends on the volume of assets liquidated in each asset class relative to market depth. This generates larger-than-expected losses and amplifies the realized volatility of the portfolio, precisely in bad scenarios where it is compelled to engage in fire sales.

3 Liquidity-adjusted VaR

As shown above, if a portfolio has large positions (relative to market depth), one cannot ignore the impact of possible liquidations on market dynamics when assessing the portfolio's risk. This impact is size-dependent and, unlike usual risk calculations based on VaR or Expected Shortfall, leads to a nonlinear scaling of portfolio risk with the size of its positions. Our model provides a systematic approach for taking into account liquidation risk when assessing the risk of a portfolio. As the following example shows, the resulting adjustments to portfolio risk can be quite substantial.

Consider for instance a portfolio, with leverage constraint $L_{\max} = 33$, initial leverage 25 and positions in three asset classes with independent returns, assumed to be Gaussian with respective annualized volatilities 10%, 20% and 30%. The market depths for these asset classes are taken to be \$1,000, \$100 and \$10 Bn respectively. As a benchmark, the estimated market depth for the SPY, the main ETF tracking the S&P500, is close to \$1,000 Bn.

The loss distribution for this portfolio (Figure 2), based on 10,000 scenarios simulated from model (1), exhibits heavy tails associated with liquidation scenarios. We define the (99%) **Liquidity-adjusted**

VaR as the 99% quantile of this loss distribution. Figure 3 displays the one-day 99% LVaR for the portfolio as a function of portfolio size, when all notional positions are increased proportionally, compared to a 99% VaR based solely on market risk.

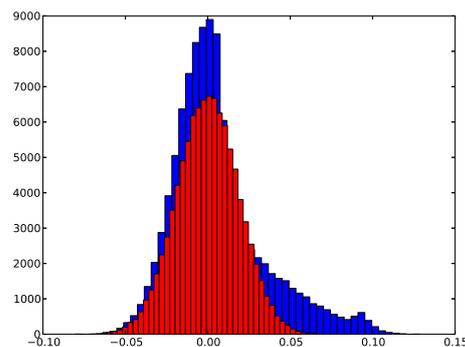


Figure 2: Distribution of portfolio losses with (blue) and without (red) feedback from liquidations.

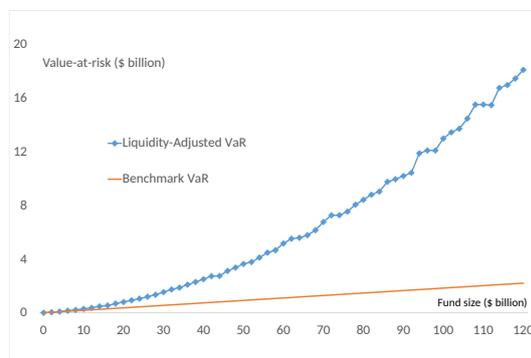


Figure 3: Liquidity-adjusted Value-at-Risk for a sample portfolio with 3 asset classes.

Whereas the traditional benchmark VaR is, as expected, linear in portfolio size, the liquidation-adjusted VaR computed using our model is not: it is convex as a function of portfolio size and is much larger than a linear VaR for large portfolio. The differ-

ence between the two numbers reflects the liquidity risk of the portfolio. For a portfolio with small positions relative to market depth, liquidation-adjusted VaR is close to a traditional VaR measure. However, for a leveraged portfolio with large, concentrated positions comparable to or larger than market depth, liquidation-adjusted VaR can be significantly (in our example, up to 10 times) larger than the usual VaR.

The previous calculations are based on the assumption of proportional liquidations. In practice, financial institutions may choose other deleveraging strategies. For example, one may establish a pecking order in which various asset classes are liquidated; this is part of the Basel 3 requirements for financial institutions when establishing their “living will”.

Figure 4 shows that the resulting (liquidation-adjusted) portfolio risk depends on the chosen exit strategy. This suggests that the exposure of a financial institution to a stress scenario depends on its plans for dealing with this stress scenario. Though not surprising, this goes against the conventional approach which attempts to measure portfolio risk based on a static snapshot of a portfolio’s positions, rather than considering how the positions will be eventually unwinded.

4 Taming the Whale

In 2012, JP Morgan’s Chief Investment Office (CIO) had accumulated large protection-buyer positions in high yield CDS indices, such as HY.11 and HY.10, together with a large ‘offsetting’ protection-seller position in the CDX IG9 index which peaked at \$278 Bn in gross notional value [U.S. Senate, 2013, JP Morgan, 2014], and roughly one-fourth of this size in HY in-

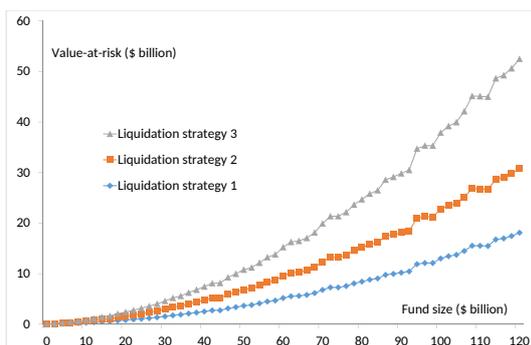


Figure 4: Liquidation-adjusted Value-at-Risk for different liquidation strategies. 1: liquidating most liquid asset first. 2: proportional liquidation. 3: liquidating least liquid asset first.

stances. While the VaR of the position was estimated to be around \$500 million, its progressive liquidation between end of March and August 2012 resulted in reported losses of \$6.2 Bn.

Reports on the CIO losses have focused on mismanagement, lack of transparency inside the organization, mismarking of positions and spreadsheet errors [U.S. Senate, 2013]. But the way risk was computed and provisioned for was not the main focus of recommendations in any of the reports. Yet, according to the report by JP Morgan’s Management Task Force [JP Morgan, 2014], the risk of these positions was evaluated using a Value-at-Risk metric which scales linearly with portfolio size and liquidation risk was not provisioned.

Figure 5 shows the one-year realized correlation between CDX IG9 and CDX IG10, two closely related indices, calculated on a rolling window. The breakdown of this correlation after March 2012, i.e. when the CIO starts liquidating its massive positions in CDX IG9, is a signature of the mar-

ket impact of CIO’s trading and suggests that the concepts of endogenous risk and liquidation-adjusted risk measure are relevant for this case.

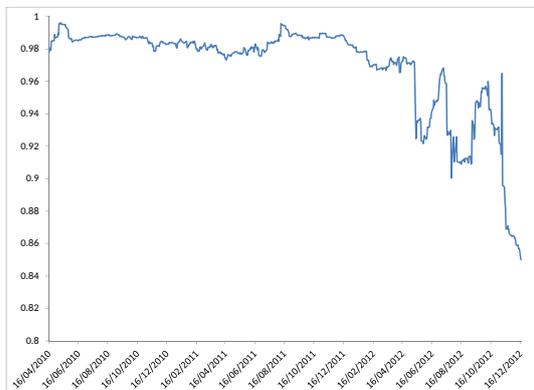


Figure 5: 1-year realized correlation between CDX IG9 and CDX IG10 returns

The liquidation of the \$278 Bn positions on CDX IG9 lasted around 5 months [U.S. Senate, 2013] which enables us to calibrate the average daily liquidation rate $\alpha f' = \frac{\$ 278 \text{ Bn}}{5 \times 20}$ in equation (2).

Figure 6 compares the 95% 5-month Value-at-Risk with and without liquidation adjustments, for a portfolio with positions in CDX IG9 and offsetting positions in HY11 with a 1:4 ratio, as a function of gross notional in CDX IG9. For a notional value of around \$278 Bn in CDX IG9, which is the size of the London Whale’s positions in Q1 2012, we find a liquidation-adjusted Value-at-Risk around \$10.5 Bn, significantly larger than the benchmark Value-at-Risk, which is of the order of \$0.5 Bn, the difference being entirely attributable to liquidation costs. This suggests that the London Whale losses could have been anticipated in a more realistic manner if liquidation costs had been properly accounted for in the risk calculations.

The fact that the realized loss was ac-

tually smaller than the 95% Liquidity-adjusted VaR is revealing: in fact, the realized loss of \$ 6.2 Bn corresponds to the 80% quantile of the portfolio, which shows that we do *not* require an extreme market move, or ‘tail event’, to generate such a large loss! Given such a large, leveraged position, even a moderate shock can lead to a large loss. The magnitude of the loss is due not to the extreme nature of the initial shock but to the market impact resulting from the sheer size of the portfolio.

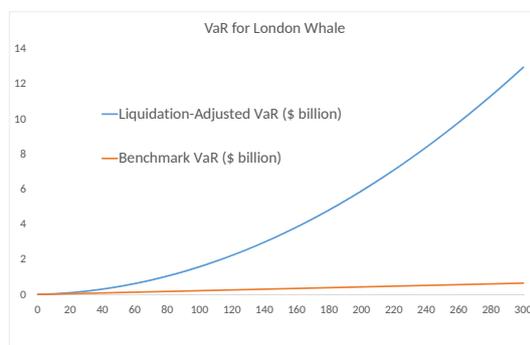


Figure 6: 95% 5-month VaR for positions in CDX IG9 (size in Bn \$).

5 Conclusion

We have proposed a tractable modeling framework for including liquidation costs in the risk analysis of financial portfolios: our approach consists in modeling the impact of liquidation by adding a price impact term in a ‘base model’ for the dynamics of asset returns. In the examples described above, the base model is a (Black-Scholes) model with constant coefficients and IID returns and the price impact is linear, but this is simply an example and both these ingredients may be modified. These extensions –nonlinear price impact, presence of multiple sources

of price impact, non-constant coefficients – are explored in a companion paper [Cont and Wagalath, 2014].

These ingredients provide the building blocks of an operational framework for calculating portfolio risk measures which correctly account for liquidation costs, exhibit nonlinear scaling with portfolio size and distinguish between liquid and illiquid positions.

References

[Almgren and Chriss, 2000] Almgren, R. and Chriss, N. (2000). Optimal execution of portfolio transactions. *J. Risk*, 3:5–39.

[Cont et al., 2014] Cont, R., Kukanov, A., and Stoikov, S. (2014). The price impact of order book events. *Journal of financial econometrics*, 12(1):47–88.

[Cont and Wagalath, 2013] Cont, R. and Wagalath, L. (2013). Running for the exit: distressed selling and endogenous correlation in financial markets. *Mathematical Finance*, 23(4):718–741.

[Cont and Wagalath, 2014] Cont, R. and Wagalath, L. (2014). Fire sales forensics: measuring endogenous risk. *Mathematical Finance*, <http://onlinelibrary.wiley.com/doi/10.1111/mafi.12071/>

[Kyle, 1985] Kyle, A. (1985). Continuous auctions and insider trading. *Econometrica*, 53:1315–1335.

[Obizhaeva, 2012] Obizhaeva, A. (2012). Liquidity estimates and selection bias. *Working Paper*.

[Shleifer and Vishny, 2011] Shleifer, A. and Vishny, R. (2011). Fire sales in

finance and macroeconomics. *Journal of Economic Perspectives*, 25(1):29–48.

[JP Morgan, 2014] JP Morgan (2014). JP Morgan Management Task Force report regarding 2012 CIO losses, Technical report.

[U.S. Senate, 2013] U.S. Senate Staff report (2013). JP Morgan Chase Whale trades: a case history of derivatives risk and abuses.

Rama Cont is Professor of Mathematics and Chair of Mathematical Finance at Imperial College London.

Lakshithe Wagalath is Associate professor of Finance at IESEG School of Management (France).

The authors would like to thank the Isaac Newton Institute for Mathematical Sciences, Cambridge, for support and hospitality during the programme *Systemic Risk: Mathematical Modelling and Interdisciplinary Approaches* where work on this paper was undertaken, supported by EPSRC grant no EP/K032208/1.