

On the Margin

Portfolio Margining at a CCP

LCH The Markets'
Partner





TABLE OF CONTENTS

- 01. Executive Summary 4
- 02. Background 5
- 03. Regulatory Requirements 7
- 04. Correlations and Portfolio Risk Management 9
- 05. Dealing with Time Varying Correlations 10
- 06. Portfolio Margining in a Default 11
- 07. Robust Portfolio Margining 12
- 08. Quantifying Portfolio Model Risk 14
- 09. Model Risk of a Portfolio Margin Model 16
- 10. Conclusion 18

Portfolio margining encourages better risk management and more efficient allocation of collateral to the greatest risks.

01. Executive Summary

CCPs set margin requirements and call margin on a wide range of financial instruments to insure against the potential default of members. Typically, margins are calculated and called on a portfolio basis — i.e., they reflect the price risk across the member's portfolio of positions, rather than being set at the level of individual positions without any reference to any other positions in the member's portfolio. Portfolio margining encourages better risk management and more efficient allocation of collateral to the greatest risks.

Calculating margins on a portfolio basis requires robust capture of correlations (or the lack of correlations) across a range of market conditions. Current regulatory standards allow portfolio margining across assets that are significantly and reliably correlated. They do not however specify the meaning of significance and reliability in exact quantitative terms.

The purpose of this paper is to describe a portfolio margining standard that gives precise meaning to the type and level of reliability required in the risk modelling of interdependent assets. It is too simplistic to assess the significance and reliability of individual correlations separately. Instead, the level and reliability of portfolio margining techniques depend on the entire correlation structure embedded in the portfolio, and require a portfolio-level assessment standard. This paper lays out the conditions under which the correlations embedded in a portfolio model are jointly significant and reliable enough to yield robust portfolio margin requirements.

To achieve this, we begin by recognizing some of the key issues underlying portfolio risk management and margining, including the nature and impact of time-varying correlations on hedging and portfolio diversification. The paper then expresses the desired model robustness in terms of the probability of under-margining due to model risk, and proposes that models can be deemed reliable if this probability remains below a certain tolerance. This form of model risk is called a type-II error, and the suggested standard is that this error should be <5% to ensure model robustness.

The adoption of such an industry standard would provide a balance between the desire for specificity and consistency on the one hand, and the reality that any standard needs to be compatible with a wide range of risk management and modelling practices on the other. Indeed, the quantitative analysis underlying this approach supports the idea that both strong and weak correlations can be valuable for risk reduction. This is only true to the extent that the margin model captures these variations robustly, including the extent to which they are reliably present during times of stress. The need for correlation stability is greater for the strong correlations underlying hedging than for the weak correlations typically exploited through portfolio diversification. Thus there is a trade-off between the desired correlation strength and the allowable instability.

Given that it is not practical to set precise numerical criteria (e.g., caps or floors) for individual correlations, we propose that each CCP should rigorously establish the isoquant curves — combinations of correlation strength and stability — that represent comparable levels of model risk. The precise definition of correlation will be portfolio- and model-dependent, and it would remain the responsibility of each CCP to develop, implement and document model risk assessment procedures to ensure the risk of type-II errors remains below the 5% standard.

The intent of this paper is to contribute to the industry debate on Portfolio Margining and to encourage debate amongst CCPs and Regulators as to how best to address this difficult issue.

02. Background

Portfolios, Risk Management and Margins

Exploiting correlations to reduce the risk of losses on a portfolio of assets is both an art and a science.

Traders and asset managers are exposed to a large number of different risk factors even within one particular asset class. For example, an interest rate derivatives trader will typically manage a large number of exposures to interest rates at different maturities along the yield curve, often in multiple currencies. Meanwhile an asset manager balancing an equities portfolio will face equity risk from many single names across multiple industries and geographies. It is common practice to manage these exposures not as the sum of separate sources of risk, but as portfolios of interrelated risks that amplify or offset each other to different degrees. Decisions to trade, hedge or close-out exposures to different risk factors are therefore typically taken on a portfolio basis — i.e., with a focus on their joint portfolio impact — and trading or hedging strategies are often executed as portfolio trades.

STRESS MANAGEMENT

Risk managers are concerned not only with spotting exploitable correlations and diversification strategies, but with ensuring that the assumed risk reduction can be relied upon in times of stress. A hedge that breaks down on volatile days is not dependable for risk management purposes. Similarly, a diversification strategy across different market segments is only effective if these segments truly diversify each other even during a market-wide sell-off. Market participants have different expectations with regards to the reliability of hedges or diversification benefits, depending on their risk appetite, need for liquidity and market perspective. But none will be entirely indifferent.

For a CCP, a core focus of risk management is limiting the downside risk to the Clearing House and its membership in the event of a member default. Under normal operating conditions, a CCP is not directly exposed to market risk (because it acts, for every trade, as an intermediary between two clearing members). However, this situation changes dramatically when a member defaults.

In a default, the CCP is directly exposed to the market risk in the defaulting member's portfolio. While a CCP will act quickly to liquidate the defaulted portfolio, this can take a few days. In order to protect itself against portfolio losses during the liquidation (or close-out) period, the CCP requires each member to post a suitable amount of collateral sufficient to cover all but the most extreme portfolio losses.

CALCULATING MARGINS

Margins are calculated on the basis of a statistical calculation of the risk of large losses. It is best practice, as well as a regulatory requirement, to call margins at a proven high confidence level (under EMIR, this is a minimum of 99.5% for OTC derivatives and 99% for other instruments).

If a clearing member holds positions in one contract only, margins can be calculated by analysing the price risk on that one contract. More typically, members hold a portfolio of positions in many different contracts, and are therefore exposed to many different price moves or risk factors. Portfolio margining is the setting of margin requirements for the entire portfolio jointly.

Portfolio margin requirements are typically derived from a portfolio margin model, which calculates the joint risk of correlated price moves across the different contracts in a member's portfolio. Portfolio margins are therefore impacted by the co-movement of contract prices, i.e. correlations. The alternative to portfolio margining consists in setting margin requirements for each contract or risk factor separately, without recognition of risk reducing portfolio effects.

Portfolio margining incentivizes better risk management by reducing margin collateral requirements for portfolios of lower risk.

Portfolio margining is another form of portfolio risk management, this time seen through the eyes of a CCP managing the default risk of its clearing members (and their clients). If a member's portfolio contains reliable hedges or cross-asset diversification, these will likely be reflected in a lower portfolio margin requirement. Even though portfolio margining tends to reduce the financial assets available to the clearing house, it is broadly accepted that portfolio margining benefits are beneficial to all market participants including the clearing house. Portfolio margining incentivizes better risk management by reducing margin collateral requirements for portfolios of lower risk. At a market level, portfolio margining improves the allocation of resources — specifically, it helps allocate a finite amount of market-wide collateral towards those positions representing the greatest risk. Consistent minimum portfolio margining standards across CCPs prevents a race to the bottom, where traders seek out the CCP with the lowest portfolio margin requirement for a given portfolio.

However, portfolio margining is only beneficial if it recognizes hedges or diversification strategies only to the extent that they succeed in reducing risk to the clearing house. Like other risk managers, a margin risk manager will also want to know the extent to which the assumed hedging or diversification effects in a particular portfolio reflected in reduced margin requirements can be relied upon in times of stress.

03. Regulatory Requirements

Both the benefits and potential pitfalls of portfolio margining are recognised in financial markets regulation. Portfolio margining is explicitly allowed under EMIR in the European Union and under CFTC regulation in the United States. Regulation rightly places constraints on the allowable margin benefits, demanding that portfolio margined assets are significantly and reliably correlated, and that there is a sound basis for these correlations other than (potentially spurious) statistics. EMIR also requires that margin benefits covering multiple instruments passed on to members (and their clients) are capped at 80% of the calculated margin benefit, but allows 100% under certain conditions.

“ REGULATORY REQUIREMENTS FOR PORTFOLIO MARGINING

EMIR Art 41.4

A CCP may calculate margins with respect to a portfolio of financial instruments provided that the methodology used is prudent and robust.

EMIR RTS Art 27.1

A CCP may allow offsets or reductions in the required margin across the financial instruments that it clears if the price risk of one financial instrument or a set of financial instruments is **significantly and reliably correlated**, or based on equivalent statistical parameter of dependence, with the price risk of other financial instruments.

EMIR RTS Art 27.2

The CCP shall document its approach on portfolio margining and it shall at least provide that the correlation, or an equivalent statistical parameter of dependence, between two or more financial instruments cleared is shown to be **reliable** over the look back period calculated in accordance with Article 25 and demonstrates **resilience during stressed historical or hypothetical scenarios**. The CCP shall demonstrate the existence of an **economic rationale for the price relation**.

EMIR RTS Art 27.4

Where portfolio margining covers multiple instruments, the amount of margin reductions shall be no greater than 80% of the difference between the sum of the margins for each product calculated on an individual basis and the margin calculated based on a combined estimation of the exposure for the combined portfolio. Where the CCP is not exposed to any potential risk from the margin reduction, it may apply a reduction of up to 100% of that difference. ”

04. Correlations and Portfolio Risk Management

CFTC Code of Federal Regulations Title 17, § 39.13(g)(4)(i)

A derivatives clearing organization may allow reductions in initial margin requirements for related positions if the price risks with respect to such positions are **significantly and reliably correlated**. The price risks of different positions will only be considered to be reliably correlated if there is a **theoretical basis for the correlation** in addition to an exhibited statistical correlation. That theoretical basis may include, but is not limited to, the following:

- A. The products on which the positions are based are complements of, or substitutes for, each other;
- B. One product is a significant input into the other product(s);
- C. The products share a significant common input; or
- D. The prices of the products are influenced by common external factors.

While it is tempting to demand reliability of the correlation and diversification effects underlying portfolio margining, it is not straightforward to define exactly how these should be demonstrated or measured. Are correlations reliable when they persist over days, months, years? Is it alright for correlations to vary by 5% over the course of the business cycle or by 10%, by 30%? And what if a margin model does not even measure correlations, but instead calculates margins on the basis of historical or Monte Carlo simulations of portfolio gains and losses?

A previous LCH paper (“Stress This House”) emphasized the importance of transparency, consistency in CCP risk management, margin and default fund requirements. These principles are no less important in the context of portfolio margining. Consequently it is critical to establish sound and consistent assessment criteria for portfolio margining, not only to ensure consistent interpretation of existing regulatory requirements, but to improve transparency around CCP risk management and generate a level playing field across CCPs.

It is too simplistic to translate margin model reliability into numerical limits on correlations or their stability. In part, this is because there is no single, universally applicable way to measure correlations and embed them in a margin model. More importantly, the key issue around portfolio margining is not a yes/no decision about whether to include or exclude certain assets from a portfolio margin model, established through in/out requirements on correlations. Instead, what is required is a rigorous but model-independent framework for assessing the degree to which portfolio effects are recognized in margin reductions.

Despite the technical complexities of correlations and margin modelling, a set of criteria is defined that can be applied to any CCP, irrespective of geography, asset market, or risk factor characteristics. Establishing a standardised margin model assessment framework can significantly improve market consistency and transparency, would be beneficial to CCP members, their clients and investors, and would facilitate the task of regulators and auditors and other model reviewers.

It is too simplistic to translate margin model reliability into numerical limits on correlations or their stability.

Correlation intuitively means the co-movement of prices or risk factors. When two contracts are positively correlated, you would expect a price increase in one contract to be accompanied by a price increase in the other contract. Conversely, if two contracts are not correlated, one expects prices to move independently. A price increase in one contract is then neither more nor less likely to be accompanied by a price increase in the other contract. This intuitive concept can be made more precise through statistical definitions of dependence, covariance and correlation. Both correlation and the absence of correlation have an impact on the joint price risk of a portfolio, and therefore on portfolio risk management.

A **positive correlation** can be exploited to partially reduce the risk in one contract by taking an opposite position in the other contract. The increased probability that a loss in the original (long) position is accompanied by a gain in the offsetting (short) position means that the likelihood of a net loss on the combined position is reduced. This concept is the basis of hedging, and its precise quantification is a standard technique of modern derivatives risk measurement and management.

The **absence of correlation** between two contract prices can also be exploited — this time not to hedge a position, but to diversify risk. If an investment or trading position is split between two uncorrelated contracts, risk goes down. This is because a loss on one contract is now not certain — and in the case of low correlations not more likely than a gain on the second contract. In other words, a joint position is somewhat less risky than either of the two contracts separately. This form of risk reduction — the diversification of idiosyncratic risk — is also a staple of risk management.

Both high and low correlations can be exploited for risk reduction. A correlation is significant in the context of portfolio modelling, not when it is strong (i.e., numerically high), but when it is **statistically significant** — i.e. not the result of a spurious coincidence (or lack of coincidence) of price moves, or due to temporary effects which break down during periods of stress. Low as well as high correlations can be statistically significant for portfolio risk management. The principal interest to a risk manager is not whether correlations are high or low, but whether they can be captured robustly in a portfolio model, and the extent to which they are enduring, particularly at times of stress.

05. Dealing with Time Varying Correlations

The degree of correlation or co-movement between asset prices is rarely constant throughout the business cycle. As has been observed in many different markets, correlations can suddenly increase (e.g. in a flight to quality, when investors pull out of a range of high-risk assets that were previously uncorrelated). Conversely, assets that are strongly correlated can suddenly de-correlate (e.g. when a currency peg is broken). In the first case, the risk of a correlated stress undermines the efficacy of portfolio diversification — the diversification effect of uncorrelated investments disappears precisely when it is most needed. In the second case, the risk of decorrelation means that hedged positions become unhedged at times of stress. In either case, the possibility of sudden changes in correlation levels means that care must be taken to ensure the model does not overstate any assumed risk management benefits during times of stress.

The differential impact of correlation and de-correlation stresses on hedges and diversification effects is summarised in the table below.

Correlation type	Portfolio Impact	Downside Risk
Strong, positive correlations	Risk can be reduced through offsetting (long/short) positions	Decorrelation stress increases risk by reducing effectiveness of hedges
Weak or no correlations	Risk can be reduced by spreading position across uncorrelated (or weakly correlated) assets	Correlation stress increases risk by reducing the effects of portfolio diversification

For most markets, the observation that strong correlations sometimes are diminished, and that weakly correlated assets sometimes become correlated, is simply a fact of life. It does not imply that hedging and diversification tools become entirely useless — it simply means their effectiveness at reducing risk is diminished. A prudent risk manager will take the less-than-perfect effectiveness of hedging and diversification strategies into account, for example by subjecting portfolios to an adequate stress testing regime, and reducing or re-balancing exposures where appropriate. In addition, a proactive risk manager will keep a keen eye out for any signs that markets that normally move in tandem are diverging, or conversely that normally unrelated markets are experiencing a correlated stress.

06. Portfolio Margining in a Default

While the effectiveness of hedges and diversification strategies is important to all risk managers, it is particularly relevant to a CCP setting margin requirements. Margins provide a financial cushion to protect the clearing house and its members, in the unlikely event that one or more members default, and these defaults are most likely to occur during periods of severe market stress. The size of the financial cushion required to achieve a desired level of protection (typically against all but the 0.5% worst case losses in the defaulting member's portfolio) therefore depends critically on the co-movement of portfolio assets during stress — i.e. stressed correlations or decorrelations.

Besides being more focused on tail risk than the typical risk manager, CCPs have additional concerns. In the event of a member's default, the CCP's primary responsibility is to minimize losses through a rapid but orderly closeout of a defaulting member's positions. Sitting tight until the storm blows over is simply not an option. This further concentrates the CCP's risk exposure to times of peak stress — in the two to five days following a member's default.

The need for an orderly de-risking process incentivizes the CCP and its members to conduct its default management processes on a portfolio basis — i.e. not to break any hedges or diversification inherent in the portfolio. Conversely, it also implies that margin benefits can only be recognized for those hedges and diversification effects that can be preserved during the portfolio unwind.

This last requirement limits portfolio margin management to those portfolios that can be priced, managed and auctioned or closed out as a whole. In the cases where, for reasons of liquidity or to ensure optimal bids, a portfolio is segmented into sub-portfolios for auctioning, extra care would need to be taken to ensure that sub-portfolios remain hedged or diversified, and that there is no material survivor bias as the portfolio is unwound (i.e., the lowest risk sub-portfolios are successively auctioned, and the clearing house and its surviving members are left with the highest risk positions).

A key principle in portfolio margining is therefore that margin calculations are consistent with the default management process, and that those positions are margined together that can be liquidated together. In practice, this means that portfolio margining is restricted to portfolios of similar assets, e.g. rates or equities but not both. As discussed in "Stress This House", LCH allows portfolio margin offsets within, but not across, these broad asset classes (rates, equities, credit).

Finally, a CCP needs to ensure that its margin requirements and closeout procedures do not amplify pro-cyclicality. This is both a regulatory requirement, as well as in the direct interest of the clearing house and its membership, in order to ensure that positions can continue to be cleared and margined without exacerbating the squeeze on collateral during times of stress. There is a well-known trade-off between this need to limit pro-cyclicality, and the desire to be maximally responsive to temporary fluctuations in risk. In the context of portfolio margining, this means that time-varying portfolio benefits can be recognized in self-correcting portfolio margin models, but only if they do not fluctuate too rapidly.

07. Robust Portfolio Margining

The Need for Model Risk Assessment

To recap, this paper argues that it is reasonable for the benefits of portfolio risk management, both those due to hedging (high correlations) and diversification (low correlations), to be reflected in reduced portfolio margin requirements. However, the level of margin reduction allowed should be limited so as to not overstate the extent to which these portfolio effects are reliable under stressed conditions, actionable during the default management process, and in line with the need to avoid pro-cyclicality.

The next step is to determine how to use these intuitive requirements to develop a rigorous portfolio model assessment criteria?

It is worth re-emphasizing at this point, that we are not looking for a binary (yes/no) decision on the scope of where portfolio margining is permitted, but rather a framework to assess which portfolio modelling assumptions are prudent.

There is a continuum from the significant margin reduction of strong correlations on long/short portfolios, to the more modest impact of portfolio diversification across weakly correlated risk factors to the negligible risk reduction in cases where correlations are known to break down with high probability at times of stress. The last example is a limiting case where the interdependence of some risk factors may be so uncertain that they are excluded from the portfolio model altogether, and margined on a standalone basis. But this limiting case is not sufficiently general to serve as sound basis for a broadly applicable portfolio model assessment framework. There is no minimum or maximum level of correlation that makes a portfolio model suitable or unsuitable.

It is tempting to look for criteria involving numerical estimates of correlations and correlation stability. For example, why not require that assets included in a portfolio model should be modelled through stressed rather than average correlations (“add some conservatism to the correlations”)? There are several complications, both theoretical and practical, that make this avenue unworkable. First, as seen in the examples above, a correlation assumption that appears conservative for one portfolio may be generous for another portfolio. Members relying on hedges to reduce the risk in a balanced portfolio will desire strong correlations, whereas members relying on diversification across a range of directionally similar positions (long/long/long) will desire low correlations. What makes a correlation assumption conservative is therefore portfolio dependent.

More fundamentally, while correlation is an intuitive concept, there is no unique way to define and compare correlations across different portfolio models, especially if correlations are time-varying. The simplest mathematical definition of correlation (the pairwise, linear correlation between two assets) adequately captures correlations only when prices are assumed to have a simple and unvarying interdependence (so-called multivariate normal, stationary distributions with very limited tail risk). As soon as tail risk and time-varying behaviour is allowed, no single technical definition of correlation will be applicable to all portfolios or to all varieties of portfolio modelling. For this reason, sophisticated portfolio models tend to be based on simulation techniques, rather than calculation formulas. Specifically, they do not have tuneable correlations that can be set to more or less conservative values, and they generally do not even calculate correlations explicitly.

How does one assess whether a model-based portfolio margin benefit is a prudent reflection of a real, enduring risk reduction in the portfolio?

Finally, the asset interdependence that gives rise to hedging and diversification effects in typical clearing portfolios is an emergent property of the entire portfolio, rather than being due to the co-movement between pairs of assets. Large portfolios will often contain dozens or even hundreds of assets (or, more accurately, risk factors), and it is the time-varying, stressed interdependence across all of these that determines the residual portfolio risk. Here again, this modelling challenge is generally tackled with simulation techniques rather than parametric models, and there is no practical way to disentangle these simulations to assign the emergent portfolio margin benefit to some pairwise asset correlations, rather than others.

The focus therefore should be on a more productive question: how does one assess whether a model-based portfolio margin benefit is a prudent reflection of a real, enduring risk reduction in the portfolio? How does one assess, not whether individual asset correlations are reliable or unreliable, but whether the **overall portfolio model** is sufficiently reliable to underpin a given level of margin benefit at the desired confidence level?

The key to assessing adequacy of portfolio margining is to delve into the model deeply enough to understand exactly how the co-movement of risk factors affects the modelled tail risk, and to quantify the model risk inherent in any portfolio risk measurement procedure.

08. Quantifying Portfolio Model Risk

Setting a tolerance on type II errors

There are a variety of techniques available to assess the robustness of any risk model. One approach is sensitivity analysis: understanding and assessing the dependence of model outcomes on data inputs, model assumptions and any calibration or other parameters. Another technique is backtesting: testing the performance of a model under a variety of historical and hypothetical scenarios. Backtesting is particularly suitable for a risk model designed to measure the potential for losses at a target confidence level. For example, if a margin model is calibrated to estimate daily portfolio losses at a 99.5% confidence level, one can run the model under historical conditions (using, say, historical data for the last 500 trading days) and test whether the calculated margins are adequate to cover losses 199 out of every 200 trading days. The number of margin exceptions in a historical backtest thus serves as a test statistic to assess model performance.

For purposes of model risk assessment, historical backtesting may be insufficient — for example, if inadequate historical data is available, or if the data does not capture the types of stresses under which the model is designed to perform. It is therefore prudent, particularly during the model design and calibration phase, to supplement history with hypothetical scenarios to capture the full range of operating conditions under which the model will be used. Care must then be taken to ensure that the statistics of these hypothetical scenarios are representative of the true, forward-looking risk environment. In other words, a representative set of hypothetical scenarios needs to be constructed to ensure that extreme scenarios are neither under- nor over-represented in the backtesting scenario set.

Once a sufficiently large set of representative (historical or hypothetical) scenarios is available, it is possible to define a new test statistic: not just the likelihood that losses exceed calculated margins, but the probability that calculated margins over- or underestimate the “true” margins required to cover the 99.5% loss tail of the full scenario set. In a sense, this test statistic serves as a “meta-statistic” for the model. It measures the risk that a model underestimates the “true” loss potential. By analogy with hypothesis testing, this is called a type-II model error.

The terminology of type I and type II errors is derived from the statistics of hypothesis testing. A type I error occurs when a true hypothesis is rejected (a false positive), whereas a type II error means that a false hypothesis is mistakenly accepted. This terminology can be applied to model backtesting, including testing under hypothetical scenarios, if we set as the null hypothesis that “the portfolio model calculates adequate margins for the given portfolio”. The table below shows the different possible testing outcomes, and highlights the risk of a type II error, i.e. the risk of under-margining.


		Null Hypothesis (margins are adequate)	
		TRUE	FALSE
Margin Model	REJECTED	Type I error	Correct Inference
	ACCEPTED	Correct Inference	Type II error

It is straightforward to quantify the model risk inherent in simple margin models. For example, imagine we build a model for a single asset at a 99.5% confidence level by setting margins equal to a suitable multiple of the standard deviation of price changes observed over the last 250 trading days (i.e. one year). If the statistics of daily price changes in this asset is sufficiently regular (viz. if daily price changes are independent of each other, and normally distributed with the same volatility every day), it can be shown that setting the margin requirement equal to 2.6 standard deviations provides a (very nearly) unbiased estimate of the true 99.5% quantile of the loss distribution. But it can also be shown that the risk of a type-II error with this simple model — the risk that the calculated margin underestimates the true loss quantile — is about 50%. The intuitive explanation for this is as follows: if one runs the margin model many times on representative (historical or hypothetical) sequences, the margin called will be correct on average, but too low about half of the time (and too high the other half of the time), due to statistical noise in the data. This level of model risk is too high for practical purposes — most risk managers will want to be more than 50% confident that they have called sufficient margin.

The source of model risk in the above example is not hard to pin-point: it arises because a 250-day sample is too short to reliably capture the full statistics of extreme price moves. What if margins were calibrated by selecting a longer time period — say, 500 or 1,000 days? This will certainly reduce the volatility of the margin estimate (by roughly a factor of two if we use 1,000 days or four years). However, it will still make it about 50% likely that the model underestimates the true margin requirement given a particular four year period. This will be true as long as the model is calibrated to yield an unbiased estimate of the true loss potential. In order to reduce model risk, it is therefore necessary to add a modest amount of conservatism to the model — for example, by adding a margin buffer equal to a set percentage of the base margin obtained from the calculation rule above.

More generally, a model developer faces the following trade-off: in order to reduce model risk below acceptable levels, the model needs to be somewhat over-calibrated on average. More conservatism reduces the risk of underestimation but increases the model over estimate on average. While excessive conservatism is not desirable, it is reasonable that model risk should be limited to a suitably low tolerance for type-II errors, say 5%. The figure below shows the level of conservatism required to bring model risk below this level, in terms of the percentage margin add-on required as a function of the number of trading days (or look back period) used in the simulation. The model risk add-on required is about 3% for a 10-year look back period (2,500 days). In mathematical terms, the line in the figure is an isoquant: it shows positions of equivalent model risk (high margin add-on for short look back periods, and vice versa). The trade-off between margin add-on and lookback period represents a real-world trade-off for assets whose price volatility is time-varying. Capturing a time-dependent volatility will require a shorter lookback period, and therefore a higher model risk add-on. Time dependent volatilities are not unacceptable per se, but require higher margins.

Figure 1. Trade-off between model conservatism and the length of the lookback period



09. Model Risk of a Portfolio Margin Model

It is straightforward conceptually, although technically more complex, to apply similar principles to the assessment of a portfolio margin model. In essence, the procedure has three steps:

01. Define a suitable testing range consisting of large numbers of (historical or hypothetical) scenarios. This testing range should cover a wide range of operating conditions, including sufficiently severe stresses at (and some distance beyond) the target confidence level of the model;
02. Run the margin model on random subsets of this testing range, and count how often the model underestimates the full loss potential at the target confidence level across the entire testing range;
03. If the likelihood of type-II error (underestimation of the true loss potential) is above acceptable tolerances (e.g., 5%), determine the margin add-on required to reduce model risk to a more comfortable level.

The most difficult of these three steps is the first one. Establishing a representative testing set of correlated price movements across a wide range of market conditions will require either very long historical time series, bootstrapping (generating new scenarios from existing historical ones) or statistical scenario generation (Monte Carlo simulation of a suitably parameterised multivariate process). However, once a suitable set of testing scenarios has been obtained, quantifying model risk and determining the level of conservatism required to keep model risk within acceptable bounds is relatively straightforward.

The amount and type of effort required to assess the risk of type-II errors will depend on the number of risk factors in the portfolio, and the nature of the inter-dependence between these risk factors. For some portfolios, it is useful to fit historical price changes into the particular type of dependence structure exhibited by multivariate GARCH processes. This type of process is easy to simulate and allows one to capture a range of time-varying behaviours that are qualitatively and quantitatively similar to observed price and correlation spikes.

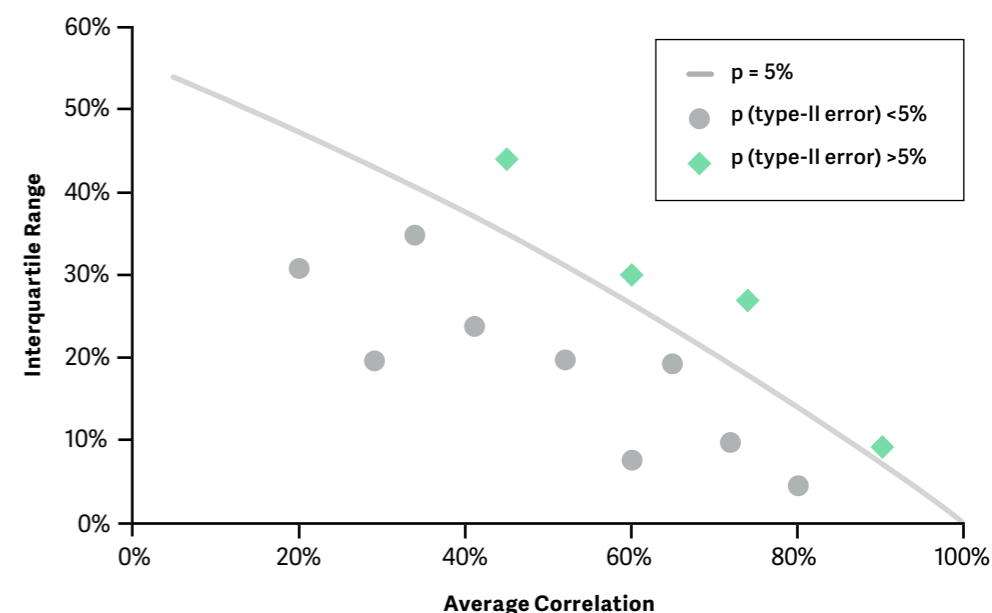
Figure 2 illustrates the application of these concepts to a particular portfolio margin model of correlated assets. The model is designed to calculate margins at a 99.5% confidence level for a portfolio of interest rate derivatives. Derivatives on interest rates of different maturities exhibit a range of time-varying correlations. For some pairs of risk factors, the correlations are very strong, in excess of 80%. Other pairs (for example, rates of very different maturities), have weak correlations, and lead to significant diversification benefits. If a correlation for a particular pair of assets varies over a wide range, it will be subject to more model risk. Conversely, correlations that are very narrowly distributed around their average will contribute in a very limited way to model risk.

The heavy orange line in figure 2 shows combinations of correlations and correlation ranges for which model risk, as measured by the risk of a type-II error, is 5%. This isoquant line was obtained by testing the margin model across a wide range of simulated interest rate paths. The volatility or range of correlations on the vertical axis is measured by means of the inter-quartile range, i.e. the difference between the 25% highest and 25% lowest correlation observed over a suitably long time period.

Also shown in Figure 2 are some risk factor pairs for which model risk falls below this level, and a few risk factor pairs that do not meet the 5% model risk standard. The model risk on portfolios containing these latter pairs would need to be mitigated by a suitable margin add-on. As the figure highlights, the model risk is not driven by either the level of correlations or the volatility of correlations separately, but by the combined effect of both. Indeed, there is seen to be a trade-off between strength and persistence of correlations. Strong correlations can result in large offsets between long and short positions, and these are only reliable if correlations stay within a very narrow range. Weak correlations lead to more modest portfolio benefits, which do not depend as critically on the exact value of the correlations, and can therefore accommodate a wider range of observed correlations for an equivalent level of model risk.

This example is illustrative only, and each CCP would need to decide how to incorporate this kind of model risk assessment in its overall testing programme, depending on the type and range of exposures included in portfolio margining, the desired confidence level of the model, and the tolerance for model risk.

Figure 2.
Trade-off between average correlation and correlation range



10. Conclusion

A standardised assessment framework for portfolio margining

Robust margining of a portfolio of related assets requires a margin model that captures the effects of correlations (or their absence) reliably throughout the business cycle and especially during times of stress. This paper has reviewed the key issues and considerations that need to be addressed when assessing the reliability of a margin model. It is shown that model reliability is less about the level or even stability of individual correlations, and more about whether and how the model captures the varying levels of interdependence of asset prices for different portfolios at different times. Margin reliability can be assessed and constrained quantitatively by estimating model risk along the set principles, specifically the risk of under margining due to a type-II error.

The list below outlines the key elements of LCH's portfolio margining assessment framework. In addition to a quantitative model risk assessment, this framework also includes broad requirements about the availability of reliable prices and consistency between margining and default management procedures. These further requirements are intuitive and relatively self-explanatory. While the framework does not explicitly echo regulatory requirements around an economic rationale or theoretical basis underlying jointly margined assets, such requirements are implicit in several of the framework principles. In the absence of such fundamental price relationships, it is unlikely that portfolios can be jointly priced, default managed and reliably modelled.

The framework sketched here is consistent with existing regulatory requirements, in that it allows portfolio margining within broad asset classes, but recognizes margin benefits only to the extent that they are reliably present at times of stress. In fact, the framework goes beyond current regulatory standards, by setting well-defined and quantitative criteria for the significance and reliability of correlations in the context of margin modelling. Specifically, a correlation offset or diversification benefit is allowed within this framework, as long as it can be modelled with a type-II error below 5%.

An error analysis of the type suggested here can also help to meet EMIR requirements that full portfolio margin benefit is passed on to members only if a CCP is able to show that this does not pose unacceptable risks (article 27.4 of the EMIR Regulatory Technical Standards). The type-II error analysis would provide a suitable basis to show that any remaining risk is within acceptable tolerances. For CCPs unable to show this, for example because the margin model is a bottom-up, risk factor by risk factor model, rather than a full portfolio simulation, reducing margin benefits by 20% (i.e., passing through only 80% of the calculated benefit) may serve as a suitable fallback option.

We suggest that the adoption of a standardised portfolio margining assessment framework along these lines by regulators and CCPs would significantly enhance transparency and consistency across the clearing industry. It gives precise meaning to the type and extent of statistical testing required for portfolio margin models to be deemed reliable, without being prescriptive with respect to the exact mathematical implementation of the testing regime. It would remain the responsibility of each CCP to construct the scenarios required for robust model testing, and to demonstrate model reliability across a suitable operating range to its own satisfaction as well as that of its auditors and regulators.

There will always remain an element of art, or at least business judgment, in the development, implementation and assessment of portfolio risk models. Establishing a common understanding of the more quantitative aspects of portfolio margining would clear the air to allow firm and frank discussion of the judgment behind the numbers.

MODEL RISK FRAMEWORK: 5 REQUIREMENTS

01. Reliable and representative price data on all contracts in the portfolio

A robust source of actual prices needs to be available for all contracts in the portfolio. In some instances, this may require techniques for interpolating or inferring implied prices for less liquid contracts. If so, these techniques need to be consistently applied across the entire portfolio.

02. Ability to price entire portfolio across a wide range of historical and hypothetical scenarios

If some contract prices are interpolated or inferred, the techniques used need to enable portfolio pricing during hypothetical scenarios. All segments of the portfolio need to remain liquid at all times.

03. Portfolio margining aligned with default management procedures

In particular, only those positions are margined jointly that can be exited jointly in the event of a member default. In practice, this limits portfolio margining to portfolios within, but not across, the major asset markets (rates, equities, credit).

04. Ability to quantify model risk/ type II errors

Model testing procedures based on historical backtesting, bootstrapping or hypothetical scenarios in order to estimate the probability of a type-II error (underestimation of margins due to statistical noise in the model inputs).

05. Margin add-ons as required to keep model risk below 5%

If the risk of a type-II error is above 5%, a suitable margin add-on is required to cover this risk.

DISCLAIMER

This paper On the Margin (the "white paper") has been prepared by the LCH Group Holdings Limited and may not be copied or reproduced, whether in whole or in part, without the prior written consent of LCH Group Holdings Limited. Copyright and any other intellectual property rights in and to the white paper are vested in LCH Group Holdings Limited and its affiliates as appropriate. The contents of the white paper are not intended, and should not be construed, as investment, tax or legal advice. Although all reasonable care has been taken in the preparation of the white paper, no representation, warranty or condition, express or implied, is made or given by or on behalf of LCH Group Holdings Limited, its affiliates, directors or any other person as to the accuracy, completeness or fairness of the information or opinions contained in the white paper and no responsibility or liability is accepted for any such information or opinions.

Copyright © LCH Limited 2021. All rights reserved.

