Cboe Clear Europe N.V.

Equity Derivatives Margin Model Description (ClearisQ)

Version 1.4.1

January 2024

Document Reference

Title	Equity Derivatives Margin Model Description (ClearisQ)
Date	January 2024
Document version	1.4.2
Document type	Methodology document
Owner	Cboe Clear Risk Management

Confidentiality

Disclaimer

This document has been prepared by Cboe Clear Europe N.V. ("**Cboe Clear**") with due care. However, discrepancies may arise between this document and the actual margin model of Cboe Clear from time to time. Cboe Clear is not liable for any errors or omissions in this document.

The Clearing Rules of Cboe Clear and the margin requirements issued to each Clearing Participant determine the obligations of each Clearing Participant at all times.

Table of Contents

1. Introduction

This document contains a description of Cboe Clear's margin model (referred to in this document as the "**ClearisQ**" model) and associated margin add-ons relevant to the clearing of equity index futures and options (collectively referred to as "equity derivatives")[1].

The margin methodologies described in this document are designed to meet the requirements set out under European Market Infrastructure Regulation ("**EMIR**"), and EMIR Regulatory Technical Standards ("**EMIR RTS**").

ClearisQ calculates core initial margin ("**IM**") requirements using a Filtered Historical Simulation ("**FHS**") methodology, based on an Expected Shortfall ("**ES**") tail measure at a confidence level of 99% and a lookback period of 700 business days. The FHS model is supplemented by a Stressed Historical Simulation ("**StressHS**") model in accordance with anti-procyclicality ("**APC**") requirements stipulated under EMIR. An APC-compliant margin is then computed as the weighted sum of the FHS and the StressHS components in line with regulatory requirements.

Where applicable, initial margin add-ons are calculated in addition to core IM to cover exposures related to wrong-way risk, liquidity risk, and large stress loss exposures.

Cboe Clear calculates margin requirements on an overnight and intraday basis for every Clearing Participant portfolio containing equity derivatives using the methodologies described in this document.

This document is intended for existing and prospective Clearing Participants, regulators, and internal users.

[1] A margin model description document relevant to cash equity products is available separately.

2. Margin Requirements

2.1 Overview

During the course of each business day, on a continuous basis, Cboe Clear calculates a Margin Requirement for each Clearing Participant portfolio containing equity derivative contracts, which is based on five components, and which must be collateralised at all times. Margin Requirements may be collateralised with eligible collateral as listed in the Cboe Clear Acceptable Collateral[2] guidelines, subject to the applicable haircut. The margin components used to calculate each Clearing Participant's Margin Requirement are:

- Initial Margin ("IM")
- Futures Variation Margin ("VM")
- Options Premium Margin ("PM")
- Unexpired Net Options Value ("NOV")
- Expired Options NOV

On the morning of each business day, Cboe Clear performs the daily cash settlement cycle in each relevant currency (CHF, EUR, and GBP), which facilitates the pass-through of cash settlement amounts pertaining to the previous clearing day, from payers to receivers. The elements comprising the daily cash settlement amounts are:

- Futures daily cash-settled obligations
- Option premium payment obligations
- Option exercise cash-settled obligations

2.2 Futures Variation Margin

Futures Variation Margin represents the intra-day accumulated gain or loss on futures contracts. On day (and time) *t,* it is calculated for a given portfolio as:

$$VM_t = \sum_{ccy} VM_{ccy,t} FX_{ccy,t} \qquad (2.1)$$

where $FX_{ccy,t}$ is the foreign exchange spot rate for currency pair *ccy/EUR* (where 1 unit of *ccy* currency buys the quoted amount of EUR currency), and $VM_{ccy,t}$ is the VM in *ccy* currency at *t* calculated as:

$$\begin{aligned} VM_{ccy,t} = & \sum_{i_{ccy}} P_{i_{ccy},t} Q_{i_{ccy},t} S_{i_{ccy}} - \sum_{i_{ccy}} P_{i_{ccy},close,t-1} Q_{i_{ccy},close,t-1} S_{i_{ccy}} \\ & - \sum_{i_{ccy}} \bar{P}_{i_{ccy},b,t}\, Q_{i_{ccy},b,t} S_{i_{ccy}} + \sum_{i_{ccy}} \bar{P}_{i_{ccy},s,t}\, Q_{i_{ccy},s,t} S_{i_{ccy}} \end{aligned} \qquad (2.2)$$

and,

i_{ccy} = futures contract i, denominated in currency ccy

$S_{i_{ccy}}$ = Contract size of options contract i_{ccy}

$P_{i_{ccy},t}$ = Market (or closing) price of futures contract i_{ccy} at *t*

$Q_{i_{ccy},t}$ = Quantity in futures contract i_{ccy} at *t*

[2] https://cdn.cboe.com/euroccp/resources/Regulation-Acceptable-Collateral.pdf

$P_{i_{ccy},close,t-1}$ = Closing price of futures contract i_{ccy} on day $t-1$

$Q_{i_{ccy},close,t-1}$ = Closing quantity in futures contract i_{ccy} on day $t-1$

$\bar{P}_{i_{ccy},b,t}$ = Weighted average price of futures contract i_{ccy} bought on day t

$Q_{i_{ccy},b,t}$ = Quantity of futures contract i_{ccy} bought on day t

$\bar{P}_{i_{ccy}s,t}$ = Weighted average price of futures contract i_{ccy} sold on day t

$Q_{i_{ccy},s,t}$ = Quantity of futures contract i_{ccy} sold on day t

Positive VM represents a gain to the Clearing Participant, whilst negative VM represents a loss. On an intra-day basis, Futures VM is aggregated across all currencies (in EUR equivalent) and collateralised throughout the clearing day.

Each morning, daily VM cash settled obligations in each currency are settled during the cash settlement cycle. On completion of the daily cash settlement cycle, the previous day's Futures VM resets to zero.

On the expiry date of the futures contract, the final daily Futures VM amount is calculated, and this daily cash settled obligation is settled within the cash settlement cycle on the following business day as usual. The final settlement amount may also be calculated using Eq. 2.2, with the final futures settlement price used in lieu of the closing price. As with the normal daily process, on completion of the daily cash settlement cycle, the previous day's Futures VM resets to zero.

2.3 Options Premium Margin

Option premiums are payable from option buyers to option sellers one business day after trade execution. As such, the settlement of the premium payment obligations occurs as part of the next day's cash settlement cycle.

Options Premium Margin represents the intra-day accumulated premiums payable/receivable on options contracts that have been traded during a clearing day. On day (and time) *t,* it is calculated for a given portfolio as:

$$PM_t = \sum_{ccy} PM_{ccy,t} FX_{ccy,t} \tag{2.3}$$

where $FX_{ccy,t}$ is the foreign exchange spot rate for currency pair *ccy/EUR* (where 1 unit of *ccy* currency buys the quoted amount of EUR currency), and $PM_{ccy,t}$ is the PM in *ccy* currency at *t* calculated as:

$$PM_{ccy,t} = \sum_{i_{ccy}} P_{j,i_{ccy},b,t}\, Q_{j,i_{ccy},b,t} S_{i_{ccy}} - \sum_{i_{ccy}} P_{j,i_{ccy},s,t}\, Q_{j,i_{ccy},s,t} S_{i_{ccy}} \tag{2.4}$$

and,

i_{ccy} = Options contract i, denominated in currency ccy

$S_{i_{ccy}}$ = Contract size of options contract i_{ccy}

$P_{j,i_{ccy},s,t}$ = Price per contract of trade *j* in option i_{ccy}, sold on day *t*

$Q_{j,i_{ccy},s,t}$ = Quantity of trade *j* in option i_m, sold on day *t*

$P_{j,i_{ccy},b,t}$ = Price per contract of trade *j* in option i_m, bought on day *t*

$Q_{j,i_{ccy},b,t}$ = Quantity of trade *j* in option i_m, bought on day *t*

On an intra-day basis, Options PM is aggregated across all currencies (in EUR equivalent) and collateralised throughout the clearing day.

Each morning, Options PM cash settled obligations in each currency are settled during the cash settlement cycle. On completion of the daily cash settlement cycle, the previous day's Options PM resets to zero.

2.4 Unexpired Net Options Value

The Unexpired Net Options Value represents the current value of all unexpired options held within a portfolio. On day (and time) *t,* it is calculated for a given portfolio as:

$$NOV_t = \sum_{ccy} NOV_{ccy,t} FX_{ccy,t} \tag{2.5}$$

where $FX_{ccy,t}$ is the foreign exchange spot rate for currency pair *ccy/EUR* (where 1 unit of *ccy* currency buys the quoted amount of EUR currency), and $NOV_{ccy,t}$ is the NOV in *ccy* currency at *t* calculated as:

$$NOV_{ccy,t} = \sum_{i_{ccy}} P_{i_{ccy},t} Q_{i_{ccy},t} S_{i_{ccy}} \tag{2.6}$$

and,

i_{ccy} = Options contract i, denominated in currency ccy

$S_{i_{ccy}}$ = Contract size of options contract i_{ccy}

$P_{i_{ccy},t}$ = Market (or closing) price of option contract i_{ccy} at t

$Q_{i_{ccy},t}$ = Quantity in unexpired option contract i_{ccy} at t (net long positon is a postive value)

The NOV represents the cash amount that would be received or paid upon the liquidation of all options within the portfolio. For this reason, any negative total NOV value must be collateralised to ensure that sufficient resources are available to close out positions in a Clearing Participant default. Positive total NOV values, on the other hand, serve to offset and reduce collateral obligations arising from the other elements within the margin requirement, subject to a floor of zero on the total margin requirement.

2.5 Exercised Net Options Value

Exercised Net Options Value reflects the accrued, but unsettled amounts arising from the exercise of options at time of expiry. On a given day (and time) *t,* it is calculated for a given portfolio as:

$$NOV_{exercised,t} = \sum_{ccy} EXO_{ccy,t} FX_{ccy,t} \tag{2.7}$$

where $FX_{ccy,t}$ is the foreign exchange spot rate for currency pair *ccy/EUR* (where 1 unit of *ccy* currency buys the quoted amount of EUR currency), and $EXO_{ccy,t}$ is the option exercise settlement obligation in *ccy* currency at *t* calculated as:

$$EXO_{ccy,t} = \sum_{c_{ccy}} \max\left(P_{c_{ref}} - K_c, 0\right) Q_c S_c + \sum_{p_{ccy}} \max\left(K_p - P_{p_{ref}}, 0\right) Q_p S_p \tag{2.8}$$

and,

c_{ccy} = Call option c denominated in currency ccy

p_{ccy} = Put option p denominated in currency ccy
$P_{c_{ref}}$ = Reference price of underlying index for call option c
K_c = Strike price of call option c
Q_c = Quantity of call option c held at time of expiry
S_c = Contract size of call option c
$P_{p_{ref}}$ = Reference price of underlying index for put option p
K_p = Strike price of put option p
Q_p = Quantity of put option p held at time of expiry
S_p = Contract size of put option p

The reference price for the underlying index is set by the exchange based on the settlement mechanism, as outlined in the contract specifications.

On an intra-day basis, Exercised NOV is aggregated across all currencies (in EUR equivalent) and collateralised throughout the clearing day.

Each morning, Exercised NOV cash settled obligations in each currency are settled during the cash settlement cycle. On completion of the daily cash settlement cycle, the previous day's Exercised NOV resets to zero.

2.6 Aggregation of Margin Requirements

On day (and time) t, the total Margin Requirement that must be collateralised for a given portfolio is calculated based on the above five components, and is floored at zero:

$$Total\ Margin_t = \max\left(IM_t - VM_t + PM_t - NOV_t - NOV_{exercised,t}, 0\right) \tag{2.9}$$

where IM_t is the Initial Margin requirement represented as a positive value.

2.7 Aggregation of Cash-Settled Obligations

On day (and time) t, the total cash amount that is due to be received (or paid) by each Clearing Participant in the next day's cash settlement cycle is calculated in each relevant currency based on the following three components:

$$\begin{aligned} Cash\ Settlement_{ccy} = &\ Futures\ daily\ cash\text{-}settled\ obligations_{ccy} \\ &+ Option\ premium\ payment\ obligations_{ccy} \\ &+ Option\ exercise\ cash\text{-}settled\ obligations_{ccy} \end{aligned} \tag{2.10}$$

where amounts due to be received by the Clearing Participant are represented as positive values, and vice versa.

The remainder of this document details the calculation of the Initial Margin requirement.

3. Initial Margin methodology overview

3.1 Initial Margin components

The core margin methodology is based on a Filtered Historical Simulation (**FHS**), Expected Shortfall (**ES**)[3] methodology. This model is then supplemented with adjustments designed to improve resilience with respect to procyclicality and netting. Finally, the Initial Margin is supplemented with three margin add-ons designed to address risks arising from wrong-way risk, illiquid positions, and positions that are large relative to the total available resources of the CCP.

3.1.1 Filtered historical simulation

The FHS methodology estimates potential losses based on historical changes (returns) in prices and implied volatilities. The margins calculated using this model represent the potential expected loss for a given portfolio based on Cboe Clear's assumed lookback period, liquidation period, and confidence level (see Section 3.2).

The key features of the FHS model are its ability to capture dynamic volatility regime changes and reproduce the joint distribution of returns, without the need to compute the covariance matrix.

The FHS methodology quantifies risk exposure at the portfolio level, i.e. allowing for risk offsets between individual instruments where appropriate. The calculation of margin at the portfolio level is subject to a portfolio margin limit rule in accordance with Article 27(4) of the EMIR RTS.

3.1.2 Anti-procyclicality measure

By construction, the Initial Margin computed using the FHS model adapts to changing market conditions. While this property is desirable from the point of view of efficient capital allocation, it can also lead to procyclical behaviour where margin requirements increase significantly over a short time period in response to market stress.

Article 28 of the EMIR RTS sets out requirements for CCPs to limit the procyclicality of margin requirements to the extent that the soundness and financial security of the CCP is not adversely affected.

To comply with regulatory requirements, Cboe Clear applies an anti-procyclicality (APC) measure in the form of a stressed historical simulation (StressHS) margin (ie. by applying the option set out under Article 28(1)(b) of the EMIR RTS).

The StressHS margin component is based on a historical simulation, expected shortfall methodology, in which simulation scenarios are obtained by combining the most recent past returns with a set of defined stressed returns. The margins calculated using this model represent the potential expected loss for a given portfolio based on stressed data observations.

The calculation of StressHS margin at the portfolio level is also subject to a portfolio margin limit rule in accordance with Article 27(4) of the EMIR RTS.

An APC-compliant margin is then computed as the weighted sum of the FHS and the StressHS components (with the stress-based component having a weight of 25%, and the FHS having a weight of 75%, in line with regulatory requirements).

3.1.3 Margin add-ons

In addition to underlying price, and implied volatility market risk, Cboe Clear is also exposed to other specific risks which require separate treatments in the form of margin add-ons:

[3] Expected Shortfall is also commonly referred to as Expected Tail Loss (ETL) or Conditional Value-at-Risk (CVaR).

- **Wrong-Way Risk ("WWR") add-on** – reflects the risk due to a Clearing Participant holding long exposure to their own stock, including cases where this exposure arises via equity index futures and/or options on equity index futures. This risk originates from the correlation between the default risk and the credit exposure of the Clearing Participant.
- **Liquidity Risk ("LR") add-on** – reflects the risk due to concentrated and/or illiquid positions in a Clearing Participant portfolio. This add-on captures the potential price impact of liquidating such positions in the event of a default.
- **Large Position ("LP") add-on** – reflects the potential excess stress loss in a Clearing Participant portfolio over and above available financial resources. This add-on captures the excess loss exposure taking into account available mutualised Clearing Fund resources under a Cover-2 standard and the Clearing Participant's individual margin collateral posted.

3.1.4 Total IM requirement

The total IM requirement is the aggregate of the APC-compliant margin plus all the margin add-ons detailed in Section 9.

3.2 Key model assumptions

3.2.1 Confidence level

Article 24 of the EMIR RTS requires that IM shall be calculated for financial instruments that are not OTC derivatives using a minimum confidence level of 99%.

To ensure compliance with the EMIR requirement, Cboe Clear calculates margin requirements using an ES risk measure at a confidence level of 99%. For a given confidence level, the ES measure is a more conservative approach than Value-at-Risk ("**VaR**")[4] – therefore, the ES margin result is never less than VaR for a given confidence level.

3.2.2 Lookback window

Article 25 of the EMIR RTS requires that IM shall be calculated based on data covering at least the latest 12 months, and that data used shall capture a full range of market conditions including periods of stress.

Cboe Clear satisfies these two requirements by:

- using a lookback window (*N*) of 700 business days; and
- including 50 stress scenarios for computing the StressHS margin component.

3.2.3 Liquidation period (Margin period of risk)

Article 26 of the EMIR RTS requires that the defined time horizon for the liquidation of open positions shall be at least two business days for financial instruments that are not OTC derivatives.

To allow for additional conservativeness, Cboe Clear assumes a liquidation period, also referred to as margin period of risk or "**MPOR**" (*h*), of three business days when calculating margin requirements. The computation of the risk over the liquidation period is performed by taking into account possible trending or mean-reverting behaviour of consecutive returns observed over the lookback window.

3.2.4 Portfolio margin limit rule

Article 27(4) of the EMIR RTS requires that, where portfolio margining covers multiple instruments, the amount of margin shall be no greater than 80% of the difference between the sum of the margins for each product

[4] Reference to a VaR-based confidence level is implicit in EMIR.

calculated on an individual basis (ie. gross portfolio margin) and the margin calculated based on a combined estimation of the exposure for the combined portfolio (ie. net portfolio margin).

To ensure compliance with the EMIR requirement, Cboe Clear applies a portfolio margin limit rule at all times for the calculation of portfolio margin which is thus computed as the weighted sum of the gross and net portfolio margin (with the net portfolio margin having a weight of 80% in line with regulatory requirements).

3.3 Key model limitations

3.3.1 Procyclicality

Notwithstanding the mitigating impact of the APC adjustment, it is important to recognise that the Initial Margin methodology, as with any risk-based margin calculation, is sensitive and responsive to the prevailing market environment. As such, the methodology will generate higher margin requirements in times of market stress, at precisely the time that market participants may be subjected to increased market risk losses, increased credit risk or even, in some cases, operational risk arising from new, and unanticipated market conditions. While this, in itself, is a limitation, it must be weighed carefully against the benefits of a risk-sensitive measure, namely:

- coverage: the methodology ensures that margin levels cover the expected loss in observed market stress conditions;
- responsiveness: margin levels adapt automatically in an appropriate way to changing market conditions;
- efficiency: margin levels reduce automatically in times of lower volatility, subject to the dampening effect of the APC measure; and
- transparency: margin levels are determined by the data directly, and independently of any meaningful intervention on a day-to-day basis. Losses for individual scenarios may be investigated and explained.

3.3.2 Reliance on historical market data

As with any historical Value-at-Risk based methodology, the initial margin model uses, as its inputs, an extensive set of historical data points for a substantial range of risk factors. As such, it is limited by both the quality and quantity of this data and is sensitive to both errors and omissions. Cboe Clear addresses this shortcoming by subjecting historical input data to a rigorous quality assessment, and by performing validation of new market data prior to being used in the margin calculation, with a view to capturing and excluding any erroneous data points. Furthermore, the methodology includes a proxying logic (see Section 10) for addressing cases where data is unavailable, along with an accompanying proxy suppression methodology aimed at safeguarding the robustness of the calculation in the event that a subset of the required inputs are missing for a given day.

3.3.3 Sensitivity to model parameters

The methodology relies on the setting of parameters (see Appendix 1) which affect multiple aspects of the calculation. Although these levels have been determined based on extensive analysis and in line with all applicable regulations and best practice they are also, in a certain sense, arbitrary and could easily be set to different levels, which would produce different initial margin numbers.

3.3.4 Sensitivity to distributional assumptions

Although the Initial Margin methodology generates scenarios using actual observed risk factor returns, and hence avoids making explicit parametric distributional assumptions, it does rely on other non-parametric distributional assumptions e.g.:

- Futures prices and FX rates can be modelled using log-returns, with adjustments for prevailing volatility (Eq. 7.1 and 7.7)

- Implied volatilities may be modelled using absolute returns, without scaling for the volatility-of-volatility (Eq. 5.20)

Assumptions of this kind are unavoidable for any methodology in which future returns are modelled based on observed historical returns; however, they do represent a limitation to the model. If the distribution of future returns for a given risk factor were to deviate dramatically from its historical distribution over a short period of time, portfolios could experience losses in excess of those predicted. This is mitigated by the responsive nature of the model, which will pick up, and incorporate, the most severe of such scenarios and retain them indefinitely through the historical stress scenarios, such that future margin calculations will provide protection against a re-occurrence. Furthermore, the hypothetical stress-testing which drives the sizing of the default fund and, by extension, the large position add-on, does incorporate stresses in excess of those observed historically so that these stresses may also be covered.

3.3.5 Reliance on interpolation and extrapolation methodologies

The model relies on linear interpolation and extrapolation when calculating the futures returns at defined pivot points and when calculating instrument P&Ls from these pivot points. It is also used when generating implied volatilities at specified pivot points, and interpolating from these points to the points required to generate P&L for Clearing Participant portfolios. This latter step also employs spline interpolation (see Section 5.4.2). While this methodology is theoretically robust and is in line with best practice, it is also true to say that each time a form of interpolation or extrapolation is used, it introduces a small element of imprecision to the calculation which can prove difficult to accurately quantify. In addition, the choice of interpolation methodology is somewhat arbitrary, and another method might be equally defensible. The methodology aims to address this by choosing, in all cases, to use the least complex method of interpolation and extrapolation, subject to the requirements of the specific circumstances.

3.3.6 Reliance on calibration of implied volatility surfaces

As outlined in Section 5.2, the methodology uses the extended surface stochastic volatility inspired (eSSVI) methodology to construct arbitrage-free implied volatility surfaces. This is one of many methods that could have been used and relies on a number of additional assumptions regarding the behaviour of implied volatilities. Although this is a commonly-used methodology with well-documented advantages, it also has a number of limitations. The choice of calibration model is important because, if the model does not provide a sufficiently close fit to the observed volatility surface by, for example, understating the slope, curvature or time-dependence, then it is possible that the model will understate the variability of the implied volatility pivot point that are used within the construction of the options P&L scenarios. Cboe Clear mitigates this risk by conducting back-testing, which compares actual realised returns in options prices vs the returns distribution. In this way, any potential shortcoming of the model can be identified.

3.3.7 Reliance on calibration of implied repo rates and dividend amounts

As outlined in Section 6.3, the fitting of volatility surfaces for single stock options contracts necessitates the calibration of implied repo rates, and implied expected future discrete dividend payment amounts for each equity underlier. As detailed in this Section, these values are calibrated using the assumption that the implied volatility for a call and a put option on a single equity underlier with the same moneyness and maturity are the same. If this assumption does not hold true in practice, then the chosen calibration process could result in repo rates or dividends that are not aligned with the true expected values. This, is turn, could give rise to volatility surfaces that are biased, which could lead to imprecise estimates of Initial Margin. Cboe Clear mitigates this risk by monitoring the calibrated dividend and repo rates, and investigating any outliers.

4. Futures scenarios

This section describes the computation of P&L scenarios for futures to be fed into the margin calculation.

In order to capture the dynamics of futures prices, it is necessary to model the term structure of the underlying index. The main drivers of the term structure are interest rates and expected dividends. Instead of explicitly modelling these two quantities, the approach adopted is to model the futures returns directly at a set of static times-to-maturities ("pivot points"). This approach requires the following steps:

- derivation of the futures price from available option prices through put-call parity;
- derivation of the futures returns for the available maturity dates; and
- interpolation of the futures returns at the required set of maturity points.

As outlined in Section 3.2, in order to compute Initial Margins for a given product, the FHS methodology described in Section 6 requires *N* scenarios for this product over a horizon *h* (the holding period or *MPOR*). These scenarios are built from past data within a time period immediately preceding the calculation date t_0 and referred to as the lookback window.

Cboe Clear adopts a lookback window of *N* = 700 business days. More specifically, the methodology uses daily returns within the lookback period and $N^* = N + h - 1$ denotes the total number of such daily returns needed to compute *N* scenarios with MPOR *h* (see Section 7.3).

4.1 Synthetic futures prices

4.1.1 Synthetic futures motivation

Synthetic futures $F_{t,T}^i$ at time *t* for maturity date *T* and underlying *i* can be generated by combining a call and a put of maturity *T* and of same strike *K*, according to put-call parity. This property is used to compute the futures' prices with the following benefits:

- accuracy: the options market is often more liquid than the futures one (especially for equity underliers).
- consistency: some quantities of interest (e.g. margins for the options) depend both on futures and options prices (e.g. through implied volatilities) so it is desirable to have data coming from the least possible number of sources in order to avoid discrepancies.

For the current purpose, only the options where both calls and puts are available for a given maturity/strike couple (*T*,*K*) are considered. Because in-the-money ("**ITM**") options are usually less liquid, it is desirable to select pairs of call/put options with the strike close to the futures price (which is, however, yet to be determined). As a fallback heuristic, the median synthetic future price derived from the 3 couples that have the strike closer to the underlying spot price S^i is used (implicitly assuming a low value for the discount and dividend term). As a fallback, if no pairs of call/put for the same strike and maturity are available, the quoted futures price will be directly used.

4.1.2 Synthetic futures price derivation

Given the maturity-dependent, risk-free interest rate *r*, the price of a synthetic future of maturity *T* may be deduced for European options using put-call parity:

$$F_{t,T}^i = K + \left(C_{t,T,K}^i - P_{t,T,K}^i\right)e^{r(T-t)} \tag{4.1}$$

with $C_{t,T,K}^i$ and $P_{t,T,K}^i$ call and put option prices at time $t \in [1, N^* + 1]$ with maturity *T* and strike *K*.

4.2 Synthetic futures returns

4.2.1 Synthetic futures returns derivation

The futures returns are calculated from the (synthetic) future prices, as logarithmic returns. Compared to arithmetic ones, logarithmic returns have the advantage of being additive, and this property is used afterwards for both the MPOR scaling, and the FX adjustment.

Given the prices $F_{t,T}$ for futures contracts of maturity T on date t (where the underlying dependency has been dropped for ease of notation), the daily logarithmic returns are calculated as:

$$p_{t,\tau} = ln\left(\frac{F_{t,T}}{F_{t-1,T}}\right) \quad \forall t \in [1, N^*] \qquad (4.2)$$

where $\tau = T - t$ is the time-to-maturity, which is used to index the returns in the term-structure dimension. Note that returns are computed over $F_{t,T}$ and $F_{t-1,T}$ which are the prices for the same contract maturity on two consecutive days. For every time t, $S(t) = \{\tau_1, , , \tau_s\}$ represents the ensemble of all available interpolated time-to-maturity points, assumed to be in ascending order without loss of generality. This ensemble depends on the available option maturities for a given date.

4.3 Futures return pivot points

4.3.1 Motivation for pivot points

As stated above, the set of observable time-to-maturity points, collected in *S(t)*, varies from one date to another. For the purposes of this methodology, it is necessary to use these points to interpolate a set of returns at a fixed set of defined points, for two reasons:

- to have complete time series with constant statistical properties over the lookback window $t = 1,\dots,N^*$ for every risk factor; and
- to allow for efficient reconstruction of futures returns at any required maturity.

To ensure that the scenarios are free from any discontinuities that could arise from corporate actions, such as discrete dividends, interpolation is performed in the futures returns space rather than the futures price space.

4.3.2 Definition of pivot points

A set of pivot time-to-maturities $P = \tau_1, \dots \tau_{N_\tau}$ is defined to model the term structure of each index. As best practice, the selection of pivot points:

- includes the spot (zero-time-to-maturity) so that $\tau_1 = 0$;
- have more pivot times-to-maturity at shorter maturities, to reflect the greater number of contracts, and greater liquidity at shorter maturities and
- includes the farthest possible maturity, so that τ_{N_τ} = 24 months.

With these in mind, the following points were selected: {0, 1m, 2m, 3m, 6m, 9m, 12m, 18m, and 24m}.

4.3.3 Interpolation from available data points to defined pivot points

A linear interpolation is used:

$$p_{t,\tau_j} = \frac{\tau_j - \tau_a}{\tau_b - \tau_a} p_{t,\tau_a} + \frac{\tau_b - \tau_j}{\tau_b - \tau_a} p_{t,\tau_b} \quad \forall j \in [1, \dots, N_\tau] \qquad (4.3)$$

where $\tau_a < \tau_j$ and $\tau_j < \tau_b$, representing the two time-to-maturities that are available in *S(t)* and closest to τ_j. Because the spot is included, it is always possible to find a lower bound to interpolate from. For longer-dated maturities, however, it may be necessary to extrapolate from the furthest available data point. In the absence of a suitable upper pivot, flat extrapolation is used which is supported by the observation that long-dated futures typically exhibit very strong correlations:

$$p_{t,\tau_j} = p_{t,\tau_S} \quad \text{if} \quad \tau_j > \tau_s = \max\left(S(t)\right) \tag{4.4}$$

Utilising both interpolation and, where necessary, extrapolation guarantees that these "futures return pivot points" can always be defined. A collection of N_τ futures returns is then stored at each date needed to build the lookback window, corresponding to $N^* \times N_\tau$ returns. For computational convenience, the returns over the pivot times-to-maturity can be collected into a matrix where row *k* contains the returns on day $t - k + 1$ and column *j* contains the returns for the *j*th pivot time-to-maturity:

$$p = \begin{pmatrix} p_{t,\tau_1} & p_{t,\tau_2} & \cdots & p_{t,\tau_{N_\tau}} \\ p_{t-1,\tau_1} & p_{t,\tau_2} & \cdots & p_{t,\tau_{N_\tau}} \\ \cdots & \cdots & \ddots & \cdots \\ p_{t-N^*+1,\tau_1} & p_{t-N^*+1,\tau_2} & \cdots & p_{t,\tau_{N_\tau}} \end{pmatrix} \tag{4.5}$$

Using these 1-day, unscaled returns (p) it is possible to construct a matrix of *MPOR*-scaled hypothetical returns ($\hat{r}$) representing the modelled distribution of scaled future returns at each pivot point over the 3-day *MPOR* horizon. Sections 7.2 and 7.3 provide further details on this transformation.

$$p \to \hat{r} = \begin{pmatrix} \hat{r}_{t,\tau_1} & \hat{r}_{t,\tau_2} & \cdots & \hat{r}_{t,\tau_{N_\tau}} \\ \hat{r}_{t-1,\tau_1} & \hat{r}_{t,\tau_2} & \cdots & \hat{r}_{t,\tau_{N_\tau}} \\ \cdots & \cdots & \ddots & \cdots \\ \hat{r}_{t-N^*+1,\tau_1} & \hat{r}_{t-N^*+1,\tau_2} & \cdots & \hat{r}_{t,\tau_{N_\tau}} \end{pmatrix} \tag{4.6}$$

4.4 Interpolation from pivot points to required maturities

For every underlying index, a collection of *MPOR*-scaled returns over the lookback window has been computed at pivot maturities $\tau_1,\dots,\tau_{N_\tau}$ as detailed above. It is therefore possible to re-interpolate the returns at the required time-to-maturity τ. The same linear interpolation as in Section 4.3 is employed:

$$\hat{r}_{t,\tau} = \frac{\tau - \tau_a}{\tau_{j_b} - \tau_{j_a}} \hat{r}_{t,\tau_{j_a}} + \frac{\tau_b - \tau}{\tau_{j_b} - \tau_{j_a}} \hat{r}_{t,\tau_{j_b}} \quad \forall t \in [1, \dots, N^*] \tag{4.7}$$

where $\tau_{j_a} < \tau < \tau_{j_b}$ is the closest bracketing using time-to-maturities available in the pivots *P(t)*. Extrapolation is not required for this step, because the longest-dated pivot point (24m) represents the longest possible maturity for a traded market futures contract.

4.5 Futures P&L scenarios

Given $F_{t_0}(\tau)$, the price on today's date t_0 of a futures contract at time-to-maturity τ, a set of price scenarios and a set of P&L scenarios is computed. The set of scenarios over horizon δ_h (spanning *h* business days) for the futures prices (levels), denoted by $\widehat{\cdot}^{\,t}$, reads:

$$\hat{F}^t_{t_{0+h}}(\tau - \delta_h) = F_{t_0}(\tau)\, e^{\hat{r}_{t_0,\tau}} \quad \forall t \in [1, N] \tag{4.8}$$

The set of P&L scenarios (variations), denoted by $\widehat{\Delta^h\ \cdot}^{\,t}$, reads:

$$\widehat{\Delta^h F}^t_{t_0}(\tau) = \hat{F}^t_{t_0+h}(\tau - \delta_h) - F_{t_0}(\tau) \quad \forall t \in [1, N] \tag{4.9}$$

In case of mismatch between the currency of the index and the Clearing Participant margining currency, Eq. 4.9 is replaced with the variation below, which incorporates the corresponding FX scenario:

$$\widehat{\Delta^h} F_{t_0}^t(\tau) = \widehat{FX}_{t_0+h}^t \left(\hat{F}_{t_0+h}^t(\tau - \delta_h) - F_{t_0}(\tau)\right) \qquad \forall t \in [1, N] \tag{4.10}$$

As with the futures return scenarios, FX scenarios are generated using the FHS approach which is detailed in Section 7.4. The calculation of returns in Eq. 4.10 reflects the fact that the FX exposure only affects VM, and not the full notional value of the futures contract.

Once these variations have been calculated, the position-level P&L scenarios may be calculated by multiplying by the position size and quantity:

$$X_{i_\tau}^t = \widehat{\Delta^h} F_{t_0}^t(\tau) Q_{\mathrm{i},} S_{\mathrm{i}} \qquad \forall t \in [1, N] \tag{4.11}$$

where:

- $X_{i_\tau}^t$ is the position-level P&L scenario (t) for futures contract *i* with time-to-maturity τ
- $\widehat{\Delta^h} F_{t_0}^t(\tau)$ is the single-contract-level P&L scenario (t) for futures contract *i* with time-to-maturity τ
- Q_i is the quantity of future *i*, in a given portfolio
- S_i is the contract size of option *i,*

The position P&Ls can then be aggregated across all futures positions within the portfolio, along with those of the options positions, (see Section 5.4.3) to calculate the overall portfolio-level P&L scenarios:

$$X_t = \sum_i X_i^t \qquad \forall t \in [1, N] \tag{4.12}$$

It is these portfolio-level P&L scenarios that are carried forward to the next stage of the FHS calculation (see Section 7.6).

5. Equity Index Options scenarios

As is the case for equity index futures, the FHS calculation for options requires N=700 portfolio P&L scenarios. More specifically, the calculation requires $N^* = N + h - 1$ daily P&L returns per instrument, where h denotes the *MPOR*. This section describes the computation of these scenarios, which will be used within the margin calculation described in Section 7.6.

Futures prices and implied volatilities are key risk-factors required in the options pricing formula. Futures price scenarios have been covered in the previous section. This section addresses implied volatility and option price scenarios specifically, focusing on vanilla European options.

5.1 Inversion of implied volatility

Modelling non-linear products requires computing the implied volatilities out of available options. The process of deriving implied volatilities from listed option market data requires the following inputs:

- risk-free interest rate curve: inferred from vanilla interest rate products (sourced directly from market data vendors);
- futures term structure: derived from listed option prices using the put-call parity formula (see Section 4.1); and
- vanilla option prices: used to invert the option pricing formula to derive the implied volatilities.

5.1.1 Pricing formula inversion

The derivation of implied volatility is based on the inversion of an option pricing formula. As the pricing formula takes as input the underlying volatility, one looks for the volatility that matches the observed option price, resulting in the implicit definition of the so-called implied volatility:

$$\sigma_{impl} \quad \text{solves} \quad \left|P\left(\sigma_{impl}\right) - P^{Market}\right| = 0 \tag{5.1}$$

Even in the simplest case of European options, there is no analytical solution for this inversion problem and the solution is computed numerically. As the option price is monotone in the volatility parameter, the unique solution is usually reached in a few iterations. If the optimizer does not converge to a solution, the data point is disregarded; this can happen because the option price violates some arbitrage constraints (e.g. price lower than the intrinsic value).

5.1.2 European options pricing formula

In the case of European option, the implied volatility can be derived by inverting Black's formula, which gives the price of such option as a function of volatility and futures price. It reads, in its general form for both calls and puts:

$$P\left(F_{t,\tau}, \sigma\right) = e^{r\tau}\left(\phi F_{t,\tau} \mathrm{N}(\phi d_1\right) - \phi K \mathrm{N}(\phi d_2)) \tag{5.2}$$

$$d_1 = \frac{ln\left(\frac{F_{t,\tau}}{K}\right) + \frac{1}{2}\sigma^2\tau}{\sigma\sqrt{\tau}} \tag{5.3}$$

$$d_2 = d_1 - \sigma\sqrt{\tau} \tag{5.4}$$

where ϕ = 1 for a call option and ϕ = 1 for a put option, $F_{t,\tau}$ is the futures price at time t for time to maturity τ, r is the τ-dependent risk-free rate of the underlying currency, K is the strike and σ is the implied volatility.

5.2 Fitting a volatility surface: extended Surface Stochastic Volatility Inspired model

5.2.1 Arbitrage-free implied volatility surfaces

The subsequent computations require implied volatilities at pairs (τ,m) of time-to-maturity and moneyness[5] points, that do not correspond to directly observable points, (i.e. quoted, options) and hence require interpolation. Rather than performing pointwise interpolations, it is deemed more accurate and efficient to calibrate a full volatility surface in an integrated manner. One of the main motivations is to ensure that the implied volatility surface is arbitrage-free, i.e. that the prices associated to any pairs of points (τ_1,m_1) and (τ_2,m_2) of the volatility surface do not yield riskless profit.

For the calibration of implied volatility surfaces, the mixture of lognormals model[6] is adopted as it guarantees both flexibility and arbitrage-free surfaces.

This section briefly summarizes the eSSVI model of the implied volatility surface. In this model, the implied volatility of the underlying stock observed at each option expiry date is parametrized in terms of total variance *w* as:

$$w(k|\theta,\rho,\psi) = \frac{\theta}{2}\left(1 + \rho\varphi k + \sqrt{(\varphi k + \rho)^2 + (1 - \rho^2)}\right) \tag{5.5}$$

where:

k is the log-forward option moneyness,
θ is the at-the-money forward total variance parameter,
ρ is the correlation between the price of the underlier and its implied volatility, which controls the skew of the implied volatility slice,
φ controls the curvature of the slice.

Each slice is therefore described by three parameters, and it is necessary to impose conditions to ensure that the model does not permit any form of butterfly or calendar arbitrages.

Following the methodology outlined in Cohort, Corbetta, Martini and Laachir (2018), for a given option expiry date, the log-moneyness (k^*) and the implied total variance (θ^*) of the observed option closest to the at-the-money are calculated. This leaves only two parameters are left to be calibrated for each expiry date (ρ and ψ) because the at-the-money forward total variance parameter can be approximated as $\theta \cong \theta^* - \rho\psi k^*$,

The following conditions are required to prevent butterfly arbitrages:

$$-1 < \rho < 1 \tag{5.6}$$

$$0 < \psi < min\left(\frac{4}{1+|\rho|}, \frac{-2\rho k^*}{1+|\rho|} + \sqrt{\frac{4\rho^2(k^*)^2}{(1+|\rho|)^2} + \frac{4\theta^*}{1+|\rho|}}\right) \tag{5.7}$$

Also, in order to ensure that the at-the-money total variance parameter (θ) is always positive, the following additional constraint needs to be imposed:

$$\psi < \frac{\theta^*}{\rho k^*} \quad if \quad \rho k^* > 0 \tag{5.8}$$

Given the fitted model parameters for the preceding expiry date, (denoted with an underscore), the following conditions are needed in order to prevent calendar spread arbitrage:

---

[5] Defined as the natural logarithm of the ratio of the option strike and the underlying futures.

[6] D. Brigo, F. Mercurio, and G. Sartorelli. Alternative asset-price dynamics and volatility smile. *Quantitative Finance*, 3(3):173-183, 2003.

$$\theta \geq \underline{\theta} \tag{5.9}$$

$$\psi \geq \underline{\psi} \tag{5.10}$$

$$\left|\frac{\rho\psi - \rho\underline{\psi}}{\psi - \underline{\psi}}\right| \leq 1 \tag{5.11}$$

These last two can be combined as:

$$\psi \geq \max\left(\underline{\psi}, \frac{\underline{\psi} - \rho\underline{\psi}}{1-\rho}, \frac{\underline{\psi} + \rho\underline{\psi}}{1-\rho}\right) \tag{5.12}$$

Finally, the condition in equation 5.12 can be expressed as:

$$\psi < \frac{\theta^* - \underline{\theta}}{\rho k^*} \quad if \quad \rho k^* > 0 \tag{5.13}$$

$$\psi > \frac{\theta^* - \underline{\theta}}{\rho k^*} \quad if \quad \rho k^* < 0 \tag{5.14}$$

5.2.2 Fitting algorithm

The calibration procedure, as presented in Cohort, Corbetta, Martini and Laachir (2018), is performed as follows. The fitting is performed on the earliest expiry date first and then on each subsequent expiry date in order of occurrence imposing, at each date, the time-arbitrage constraints based on the parameters calibrated for the previous expiry date. Within each maturity, the optimization is performed using the following steps:

1. calculate k^* and θ^* from observed option prices,
2. perform an outer grid search on ρ,
3. using the calibrated value for ρ, perform an inner one-dimensional optimization over ψ with the bounds given in the previous section.

The objective function of the optimization consists of the L^1 norm of the differences between the model price and the market price of the options:

$$\min_{\rho}\min_{\psi|\rho}\left\|\underline{P}^{model}(\rho,\psi) - \underline{P}\right\| \tag{5.15}$$

The following adjustments were applied to the algorithm above when calibrating the volatility surfaces using options data prior to go-live in order to improve the robustness of the results in cases where historical data was scarce.

1. introducing a regularization term consisting of the L^1 distance from the implied volatility slice given by the previous day fitted parameters (at slice if no previous parameters are available),
2. replacing the time-arbitrage and buttery constraints, with soft constraints in the objective function,
3. performing a grid search for values around θ^*

Overall the calibration procedure is simple, robust and helps avoiding local minima.

5.2.3 Monitoring of model fit

The methodology relies on the algorithm outlined above to calibrate a volatility surface that is a good fit for the observed market data. For this reason, it is important to implement a check to ensure that the fit is good i.e. that the implied volatilities inferred from the calibrated surface are closely aligned with those calculated using observed market prices. To perform this check, the square root of the average squared error of the implied volatility is calculated for each index, on each business day.

$$\Lambda = \sqrt{\frac{1}{N_o}\sum_{o=1,\ldots,N_o}(\sigma_o^{market} - \sigma_o^{model})^2} \quad (5.16)$$

where,

$\sigma_o^{market(model)}$: the market (model) implied volatility for option *o*

N_o: the number of options quoted

This metric is compared to a defined threshold value (ϖ). If this threshold is exceeded, the goodness of fit is deemed to be inadequate, and the implied volatilities at each pivot point are calculated by simple linear interpolation using the observable points nearest to each pivot point. The threshold value is defined based on analysis of historical surface calibration outcomes to identify instances where the calibration has been affected by poor or outlying data.

5.3 Implied volatility pivot points

Consider an option of price *P(τ,m)* on today's date t_0, with time-to-maturity τ and moneyness *m*. In order to compute a risk measure, the margin model requires a distribution of possible outcomes from a set of P&L scenarios. For this option, the latter are computed from scenarios of:

- the price of the futures contract with option time-to-maturity τ; and
- the implied volatility at option time-to-maturity τ and moneyness m.

In order to derive implied volatilities at any combination of (τ;m), a set of pivot points is defined on a grid of time-to-maturities $(\tau_j)_{1\le j\le N_\tau}$ and moneyness $(m_k)_{1\le k\le N_m}$. These pivot points are responsible for driving the dynamics of the entire implied volatility surface and interpolation is used to derive an implied volatility level at any arbitrary point of the surface.

5.3.1 Implied volatility levels at pivot points

The initial inputs for the option scenarios are quoted prices $\left((P_{t,o})_{1\le o\le N_{o_t}}\right)_{1\le t\le N*}$ for each date *t* in the lookback window and for whatever available time-to-maturity and moneyness of a generic option *o*.

On each date *t*, a volatility surface can be fitted using the algorithm outlined in Section 5.2.2. Note that, for the purposes of this fitting process, if quoted option prices are available for both the call and put option for a given maturity and moneyness point, the out-of-the-money option will be used and the in-the-money-option will be ignored. This is because, as a general rule, out-of-the money options are typically more liquid than in-the-money options. Once this volatility surface has been fitted, the implied volatilities can then be derived at each pivot time-to-maturity and moneyness point:

$$(P_{t,o})_{1\le o\le N_{o_t}} \rightarrow (\lambda_i)_{1\le i\le n}, (\bar{\mu}_{i,j}, \eta_{i,j})_{1\le i\le n, 1\le j\le m} \rightarrow (\sigma_{t,j,k})_{1\le j\le N_\tau, 1\le k\le N_m} \quad \forall t \quad (5.17)$$

As a result of this process, a time-series of implied volatilities at each implied volatility pivot point is derived.

5.3.2 Implied volatility returns at pivot points

Having fitted the implied volatility surface on all dates $t \in [1, N]$, it is then possible to compute returns at each pivot points as:

$$\Delta\sigma_{t,\tau_j,m_k} = \sigma_{t+1,\tau_j,m_k} - \sigma_{t,\tau_j,m_k} \quad \forall j,k \quad (5.18)$$

for all dates t. Note that the implied volatility returns are calculated as the arithmetic difference between each pair of volatilities, as opposed to the case of futures, described in Eq. 4.2, where logarithmic returns are used.

All implied volatility returns can be collected in matrix form as below, with the elements of each row containing returns at every combination of time to maturity (τ) and moneyness (m).

$$\Delta\sigma = \begin{pmatrix} \Delta\sigma_{t,\tau_1,m_1} & \Delta\sigma_{t,\tau_1,m_2} & \cdots & \Delta\sigma_{t,N_\tau,N_m} \\ \Delta\sigma_{t-1,\tau_1,m_1} & \Delta\sigma_{t-1,\tau_1,m_2} & \cdots & \Delta\sigma_{t-1,N_\tau,N_m} \\ \cdots & \cdots & \ddots & \cdots \\ \Delta\sigma_{t-N^*+1,\tau_1,m_1} & \Delta\sigma_{t-N^*+1,\tau_1,m_2} & \cdots & \Delta\sigma_{t-N^*+1,N_\tau,N_m} \end{pmatrix} \quad (5.19)$$

5.4 Options P&L scenarios

5.4.1 Implied volatility scenarios at pivot points

Using these historical returns, a set of *MPOR*-scaled volatility scenario returns is calculated by summing the historical returns over the t+h-1 consecutive observations.

$$\widehat{\Delta\sigma}_{t,j,k} = \sum_{\theta=t}^{t+h-1} \Delta\sigma_{\theta,j,k} \quad (5.20)$$

This is performed over all time, moneyness and maturity points to produce a comprehensive set of scenario returns at all pivot points for all *N* historical scenario dates.

$$\Delta\sigma \rightarrow \widehat{\Delta\sigma} = \begin{pmatrix} \Delta\hat{\sigma}_{t,\tau_1,m_1} & \Delta\hat{\sigma}_{t,\tau_1,m_2} & \cdots & \Delta\hat{\sigma}_{t,N_\tau,N_m} \\ \Delta\hat{\sigma}_{t-1,\tau_1,m_1} & \Delta\hat{\sigma}_{t-1,\tau_1,m_2} & \cdots & \Delta\hat{\sigma}_{t-1,N_\tau,N_m} \\ \cdots & \cdots & \ddots & \cdots \\ \Delta\hat{\sigma}_{t-N^*+1,\tau_1,m_1} & \Delta\hat{\sigma}_{t-N^*+1,\tau_1,m_2} & \cdots & \Delta\hat{\sigma}_{t-N^*+1,N_\tau,N_m} \end{pmatrix} \quad (5.21)$$

Using these *MPOR*-scaled volatility returns, a set of *N* volatility surface scenarios denoted by $\hat{\sigma}$ is computed by adding the current implied volatility to the corresponding volatility return at each maturity and moneyness point for all dates $t \in [1, N]$ in the lookback window.

$$\hat{\sigma}^t_{t_0+h,\tau_j,m_k} = \sigma_{t_0,\tau_j,m_k} + \widehat{\Delta\sigma}^t_{\theta,\tau_j,m_k} \qquad \forall j,k,t \quad (5.22)$$

The superscript *t* recalls that the scenario is built from historical data at date index *t*.

5.4.2 Implied volatility scenarios at arbitrary point

To recover the scenarios at the target time-to-maturity τ and moneyness *m*:

$$\left(\hat{\sigma}^t_{t_0+h,j,k}\right)_{j,k} \rightarrow \hat{\sigma}^t_{t_0+h}(\tau, m) \quad (5.23)$$

a two-step process is performed as followed, (dropping the time and scenario notations):

1. a pair of implied volatility scenarios $\left(\hat{\sigma}_{j_a}(m), \hat{\sigma}_{j_b}(m)\right)$ are built at the target moneyness m using spline interpolation (with linear extrapolation) for the two closest time-to-maturity slices τ_{j_a} and τ_{j_b} at which pivot points are defined,
2. the implied volatility scenario at data time-to-maturity τ is derived from these two points using a linear interpolation based on total variance $V = \sigma^2\tau$:

$$\hat{\sigma}^2(\tau, m)\tau = \frac{\tau - \tau_{j_a}}{\tau_{j_b} - \tau_{j_a}} \hat{\sigma}^2_{j_a}(m)\tau_{j_a} + \frac{\tau_{j_b} - \tau}{\tau_{j_b} - \tau_{j_a}} \hat{\sigma}^2_{j_b}(m)\tau_{j_b} \quad (5.24)$$

in order to guarantee that the scenarios produced are arbitrage-free.

This interpolation is applied for the scenarios associated with all dates $t \in [1, N]$ in the lookback window.

5.4.3 Option P&L scenarios

Each P&L scenario is created with risk-factor scenarios from the same date t in the past, according to the following formula:

$$\widehat{\Delta^h P}^t_{t_0}(\tau, m) = \widehat{FX}^t_{t_0+h} P\left(\hat{F}^t_{t_0+h}(\tau - \delta_h), \hat{\sigma}^t_{t_0+h}(\tau - \delta_h, \hat{m}) + \sigma^*(t_0, \tau, m) - \sigma(t_0, \tau, m)\right) - FX(t_0) P\big(F(t_0, \tau), \sigma^*(t_0, \tau, m)\big) \tag{5.25}$$

where,

$\hat{F}$ are the futures scenarios given in Eq. 4.8,

$\hat{m}$ are the adjusted moneyness given $\hat{F}$,

σ^* is the implied volatility derived from the quoted market price,

σ is the fitted model implied volatility,

$\widehat{FX}$ and FX denote the scenarios and initial FX rate respectively.

Time-to-maturity and moneyness are consistently adjusted over the scenarios. Contrary to futures contracts, FX is applied for the whole option value, which is in keeping with the margining process, and accurately reflects the impact of FX exposure on the potential closeout cost in the event of a Clearing Participant default. Note that the implied volatility scenarios are corrected by the term $\sigma^*(t_0, \tau, m) - \sigma(t_0, \tau, m)$ which accounts for any potential initial mismatch between fitted and observed implied volatility (and prices).

Once the P&L scenario has been defined for a single option, it is possible to scale by the position quantity and the contract size to calculate the position-level P&L scenarios for this contract:

$$X^t_{i_{\tau,m}} = \widehat{\Delta^h P}^t_{t_0}(\tau, m) Q_{\mathrm{i},} S_{\mathrm{i}} \qquad \forall t \in [1, N] \tag{5.26}$$

where:

$X^t_{i_{\tau,m}}$ is the position-level P&L scenario (t) for option *i* with time-to-maturity τ and moneyness *m*

$\widehat{\Delta^h P}^t_{t_0}(\tau, m)$ is the single-contract-level P&L scenario (t) for option contract *i* with time-to-maturity τ and moneyness *m*

Q_i is the quantity of option *i*, in a given portfolio

S_i is the contract size of option *i,*

The position-level P&L scenarios can then be aggregated across all positions within the portfolio, along with those of the futures positions, (see Section 4.5) to calculate the overall portfolio-level P&L scenarios:

$$X_t = \sum_i X^t_i \qquad \forall t \in [1, N] \tag{5.27}$$

It is these portfolio-level P&L scenarios that are carried forward to the next stage of the FHS calculation (see Section 7.6).

6. Single Stock Options Scenarios

This chapter describes the computation of P&L scenarios for single stock options for margining purposes. The main difference between single stock options and equity index options arises due to the fact that single stock options are American options and, as such, may be exercised by the buyer on any day up to, and including, the expiry date. A more complex option pricing model must therefore be used to ensure that the price reflects this feature, and fully incorporates the impact that discrete dividend payments will have on the option value over its life.

The key market risk factors that drive that drive changes in the value of single stock options are i) the spot price of the underlying and ii) the implied volatility. The spot price scenarios used to compute single stock options scenarios are generated using the ClearisQ methodology for cash equities which is covered in the Cboe Clear Margin Model Description (Cash Equities and Exchange-Traded Products). The remainder of this chapter addresses the following issues:

- pricing of American options,
- derivation of the implied discrete dividends and repo rates,
- fitting of the implied volatility surface,
- generation of option P&L scenarios.

6.1 Required Inputs

The inputs used to generate single stock option scenarios for a given underlying are historical observations of:

- the underlying price,
- the option prices,
- the risk-free rate curve,
- the expected dividend dates.

6.2 American Option Pricing Formula

Unlike European options, there is no exact closed-form solution for pricing American options. There have been published, however, numerous analytical approximations and numerical schemes to price American options with, and without, discrete dividends. For the purposes of this methodology, a tree-based approach has been selected due to its simplicity and transparency, which has led to its widespread use within the industry. In the following section, three different tree-based approaches are presented, two of which incorporate a treatment of discrete dividends. The Cox-Ross-Rubinstein tree is introduced first, as it serves to illustrate many features used in the latter two approaches.

6.2.1 Cox-Ross-Rubinstein Tree Algorithm

ClearisQ implements a tree pricing algorithm with a binomial tree structure as in Cox-Ross-Rubinstein[7]. In the latter a geometric Brownian Motion with drift μ and volatility $\sigma$is approximated by a tree in which the price can, at each node, move either up or down according to the following parameters:

$$u = 1/d = e^{\sigma\sqrt{\Delta t}} \tag{6.1}$$

$$p_u = \frac{e^{\mathrm{r}_t \Delta t} - d}{u - d} \tag{6.2}$$

where:

u and *d* denote the up and down total return movements respectively,

[7] John C Cox, Stephen A Ross, and Mark Rubinstein. Option pricing: A simplified approach. Journal of financial Economics, 7(3):229{263, 1979.

Δt is the time length of each step, derived as T/N, where T is the option time-to-maturity and N is the chosen number of steps,

p_u is the probability of an up-movement.

r_t is the risk-free interest rate applicable between times t and $(t + \Delta t)$.

This approach ensures that the distribution of the stock price at expiry matches the target distribution, while keeping the tree recombining and ensuring price convergence as $\Delta t \to 0$.

After constructing the tree for the underlying price, the option value at each pricing node at the last time-step can be easily evaluated using the final option payoff formulas:

$$Payoff_{call} = max(S - K, 0) \quad (6.3)$$
$$Payoff_{put} = max(K - S, 0) \quad (6.4)$$

For the pricing nodes at the time-step immediately prior to expiry, the option value is then derived as the discounted, probability-weighted average of the option values at the two successive nodes. For an American option, this value has to be further compared against the intrinsic value at the node and the maximum between the two is selected. Solving backwards over the whole tree, the option value at the origin node is derived. An example of an underlying tree and the associated option one is shown in Figure 6.1.

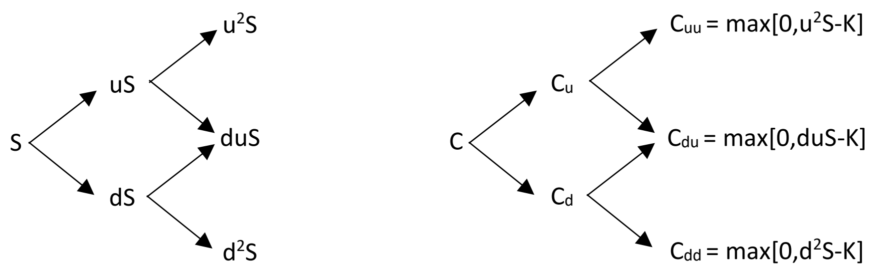

(a) Tree for the underlying price

(b) Tree for the European call option

Figure 6.1: Trees for underlying and option for a European call option with N=2. When solving backwards, the value C_u, for example, may be calculated as:

$$C_u = e^{-r\Delta t}(p_u C_{uu} + (1 - p_u)C_{du}) \quad (6.5)$$

The pricing tree methodology implemented within ClearisQ uses the technique proposed by Broadie and Detemple[8] to smooth the option payoff and improve convergence. Using this method, the recursion is started at step *N-1*, with the European option values evaluated using the Black-Scholes pricing formula and the underlying equity prices at the corresponding price nodes, instead of starting with the terminal option payoff at step *N*. This approach is designed to avoid the discontinuity of the terminal payoff and ensure that the calculated price of the option is less sensitive to the specific choice of time increment and pricing nodes. This enhancement is possible due to the fact that, over the discrete tree time grid, the option is, in practice, European at the second-to-last time step, as no exercise would be possible between these time steps.

The Cox-Ross-Rubinstein method described above would be useful for pricing options on any underlying stock that doesn't pay discrete dividends, however the inclusion of dividends presents a problem. On the ex-date of a dividend there is typically a jump downwards in the price of a stock which reflects the fact that any shares purchased on or after that date are not eligible to receive the dividend. When these discrete jumps are modelled within a pricing tree, they may lead to additional instances where it is optimal for an option holder

[8] M. Broadie, J. Detemple, American option valuation: new bounds, approximations, and a comparison of existing methods, The Review of Financial Studies, Winter 1996 Vol. 9, No. 4, pp. 1211–1250

to exercise early. This, by extension, can lead to a change in the calculated option value. For this reason, it is necessary to use a method that incorporates these expected dividends.

6.2.2 Hull Method

Hull's approach[9] builds upon the methodology of Cox, Ross and Rubinstein and aims to incorporate expected discrete dividends by assuming that the underlying price S_t can be represented as:

$$S_t = S_t^* + \sum_{i;t_i>t} df(t,t_i)D_i \quad (6.6)$$

where:

S_t^* is the stochastic portion of the stock price, which follows a geometric Brownian Motion with a drift equal to the risk-free rate minus the repo rate,

$df(u,v)$ is the risk-free discount factor between times *u* and *v*,

t_i is the time of the *i*-th dividend,

D_i is the size of the *i*-th dividend,

This model divides the stock price into two components: One component reflects the present value of the expected future dividends and is assumed to be fixed. The other, remaining component is assumed to be stochastic, with a value that is described with a geometric Brownian Motion with a drift equal to the risk-free rate.

This model admits a closed-formed solution for the European price, similar to the standard Black-Scholes formula with the strike adjusted for the dividends. American options can be priced using a standard tree for S_t^*, constructed as described in Section 6.2.1, and the option payoff is calculated at each node evaluating S_t as defined in Equation 6.3.

Although analytically tractable, the model has a major drawback: the calculated option prices, and consequently the implied volatilities, are assumed to be dependent on all future dividends, even those beyond the option's expiry date. This dependence on all future dividends makes it impossible to use a bootstrapping process to estimate future dividends using the observed option prices. Furthermore, the addition of, or removal of an anticipated future dividend from the calculation can induce a spurious jump in the implied volatility time series. Additions and removals of this kind would naturally occur over time as some options expire and new options are listed.

6.2.3 Vellekoop Method

Vellekoop and Nieuwenhuis[10] approach models the price of the underlying stock as a geometric Brownian Motion with deterministic[11] jumps observed at dividend times, i.e.

$$A_t = -\sum_{i;t_i \leq t} D_i \quad (6.7)$$

[9] Section 21.3 of John C Hull. Options futures and other derivatives. Pearson Education India, 9th edition, 2015.,

[10] Michel H Vellekoop and Johannes W Nieuwenhuis. Efficient pricing of derivatives on assets with discrete dividends. Applied Mathematical Finance, 13(3):265{284, 2006.

[11] In cases where the price immediately prior to the expected dividend is less than the expected dividend, the dividend is assumed to be capped at the equity price, and hence the equity price will jump to zero upon payment of the dividend. This assumption is easily embedded in the numerical algorithm described below.

$$dS_t = r_t S_{t-}dt + \sigma S_{t-}dW_t + dA_t \tag{6.8}$$

where:

D_i is the amount of the dividend paid at time *i*

$0 < t_i \leq t$ is the set of times until the dividends.

This model is a natural extension of the Black-Scholes dynamics to the discrete dividend case. In contrast to Hull's methodology, each option price depends only on future dividends prior to the option's maturity, which makes it possible to adopt a bootstrapping procedure to infer future dividends from quoted option prices. This approach, however, does not allow for the calculation of option prices using a closed form solution and hence it is necessary to employ a more sophisticated tree algorithm, which includes an interpolation step at each dividend date.

The example below illustrates firstly how this model can be used to price an option in a case where future dividends, repo rates and implied volatility are all known with certainty. It then illustrates how this process can be inverted to infer these inputs using a set of observed option prices. In both cases, the example uses the observed equity spot price, as well as the risk-free interest rate.

To demonstrate this approach, the example below describes a case where there is only one dividend before an option expiry date, defined by the time-to-payment (t_1) and value (D_1). The forward price of the underlying stock at time *t* for maturity *T* is defined as:

$$F_t(T) = S_t e^{\int_t^T (r_s - q_s)ds} - \sum_{i:t<t_i<T} e^{\int_{t_i}^T r_s ds} D_i \tag{6.9}$$

Where q_s represents the difference between the risk-free rate and the interest rate that is derived from options prices for a specific underlying stock, and option expiry date. This parameter is hereafter referred to as the "repo rate".

On the pricing tree the dividend is mapped to closest time node on the discrete time grid which is greater than the time-to-payment $(\tilde{t}_1)$ as $\left(D_1 e^{r(\tilde{t}_1 - t_1)}\right)$. Firstly, a standard tree is built with initial price $e^{-rT}F_0(T)$, drift *r* and volatility σ^2. The tree is solved backwards until $\tilde{t}_1$ to get a vector of equity prices, $\underline{s}$, and option values, $\underline{o}$, at the time-node immediately after the dividend. A new tree is then constructed up to $\tilde{t}_1$ with initial price $e^{-r\tilde{t}_1}F_0(\tilde{t}_1)(= S_t)$ resulting in an underlying price grid, $\underline{s}^-$, at the time-node immediately prior to the dividend. The associated vector of option values at the time-node immediately prior to the dividend, $\left(\underline{o}^-\right)$, are then obtained as:

$$\underline{o}^- = max\{f\left(max\left(\underline{s}^- - D_t, 0\right)|\underline{s}, \underline{o}\right), g\left(\underline{s}^-\right)\} \tag{6.10}$$

where

$f\left(x_1|\underline{x_0}, \underline{y_0}\right)$ is an interpolant which, given input vectors of paired observations (x_0 and y_0), returns a vector of expected values of *y* for a given vector x_1

$g(.)$ is a function which returns the payoff from exercise.

The chosen interpolation method is linear, as tested in Vellekoop and Nieuwenhuis' original article[12][13]. Given $\underline{s}^-$ and $\underline{o}^-$, a new pricing tree is built in order to derive the current option price at time zero. This method may be extended to include multiple dividends, with option prices at each associated ex-date being similarly

[12] Michel H Vellekoop and Johannes W Nieuwenhuis. Efficient pricing of derivatives on assets with discrete dividends. Applied Mathematical Finance, 13(3):265{284, 2006.

[13] The input data points $\underline{s}$ and $\underline{o}$ are enriched with additional points based on the asymptotic values of the options in order to avoid extrapolation.

interpolated using the method outlined above. This interpolation step ensures that the pricing grid at these points is regularised with respect to the price of the underlying, such that the pricing tree recombines. Without this step, the number of pricing nodes would increase dramatically such that the approach would be impractical from a computational perspective.

6.2.4 Derivation of implied volatility

As with European options implied volatility levels are interpolated by minimizing the difference between the model price, calculated using the chosen option pricing model, and the observed market price.

$$\begin{gathered} \sigma_{impl} \quad solves \quad \left|P\left(\sigma_{impl}\right) - P^{Market}\right| = 0 \\ for \quad P(\sigma_{impl} = 500\%) \geq P^{Market} > \begin{cases} S_0 - K & for\ call\ options \\ K - S_0 & for\ put\ options \end{cases} \end{gathered} \tag{6.11}$$

$$\sigma_{impl} \quad solves \quad \left|P\left(\sigma_{impl}\right) - P^{Market}\right| = 0 \quad for \quad 0\% < \sigma_{impl} \leq 500\% \tag{6.12}$$

A standard Brent algorithm, bounded within [0%; 500%], is adopted to find the option implied volatility. For this step, the smoothing technique, introduced in Section 6.2, is crucial to avoid numerical issues, due to the non-monotonicity of the tree pricing formula, and to ensure fast convergence.

Note that the derived implied volatility values depend on the pricing formula used for the derivation and its underlying assumptions. For an American option in the exercise region, it is not possible to derive a single implied volatility because of the fact that multiple values of the implied volatility will all result in model prices that match the observed market price. For calibration purposes, options in the exercise region, identified as options with prices less than, or equal to, their intrinsic values are explicitly excluded from the surface-fitting algorithm. In addition, when generating the Greeks and scenarios for these options, the implied volatility is evaluated based on the fitted volatility surface rather than being calculated for each option independently.

6.3 Derivation of Discrete Dividends

Unlike European options, the expected ex-dates and the expected size of discrete dividends have to be considered when pricing American options in order to properly incorporate the possibility of early exercise. The expected ex-dates of the dividends can be inferred from historical dividend dates, or public announcements. Given these dates, it is possible to use observed option prices to infer the expected value of the dividends using additional numerical optimization routines and model assumptions.

Inferring the dividend values in this way, rather than relying on external sources (e.g. projections by equity analysts), ensures that the model will produce consist estimates of the dividend values given the same set of inputs. As with implied volatilities, dividends inferred in this way should be interpreted as a model artefact, rather than a realistic forecast.

The chosen calibration method relies on achieving consistency between the implied volatility of the American call and put option pairs with the same strike and expiry date. The implied dividends and repo rates are jointly derived by an algorithm that minimizes the sum of absolute differences[14] between the implied volatility of the pairs of call and put options with the same strike and maturity[15]. i.e.

$$SAD_{\tau_i} = \sum_j \left|\sigma_{call,\tau_i,K_{ij}} - \sigma_{put,\tau_i,K_{ij}}\right| \tag{6.13}$$

[14] The algorithm uses a smooth approximation to the sum of the absolute differences to ensure convergence.

[15] It should be noted that the size of dividends inferred in this way will vary based on the model assumptions.

where:

SAD_{τ_i} is the sum of absolute differences for expiry date τ_i

$\sigma_{call,\tau_i,K_{ij}}$ is the implied volatility of the call option with expiry date τ_i and strike K_{ij}

$\sigma_{put,\tau_i,K_{ij}}$ is the implied volatility of the call option with expiry date τ_i and strike K_{ij}

The calibration algorithm below uses the following notation. The American options are defined on a time-to-maturity grid $\mathrm{T} = \{\tau_i, i = 1, \dots, n\}$ and, for each maturity, on a strike grid $K_i = \{K_{ij}, j = 1, \dots, m_i\}$. The times to the dividend ex-dates are defined as $\mathrm{T}^D = \{\tau_k^D, k = 1, \dots, d\}$. The algorithm bootstraps the dividend values and the equity repo rates starting with the earliest expiry date, and then processes each successive expiry date in turn. The result of the algorithm is a set of dividend values associated with the input dividend dates and a set of repo rates associated with each expiry date. The repo rate, as it is defined here can reflect any other market effect, other than the dividend, which is not captured within the model and which would affect the forward price of the underlying. The algorithm incorporates a regularization term (*p*) over the repo term structure when they are calibrated jointly with the dividend values, to handle the fact that the repo rate and the dividend have a similar effect on the implied volatility of the options.

Algorithm 6.1: Algorithm to infer dividends $\underline{d}^*$ and repo rates $\underline{y}^*$ from the prices of American options

> $\tau^B \leftarrow 0, \overline{y} \leftarrow 0$
> for $\tau^M \epsilon \mathrm{T}$ do:
>
>   select N put-call pairs on this maturity closest in strike to the underlying price
>   if no put-call pairs available then
>   continue
>   end if
>
> $\mathrm{T}^* = \{\tau^D \epsilon\ \mathrm{T}^D \mid \tau^B < \tau^D \leq \tau\}$
>   if T^* is not empty then:
> $d^*, y^* \leftarrow \arg min \sum |\sigma_{put}(d, y) - \sigma_{call}(d, y)| + p|y - \overline{y}|$
>   Collect d^* for each $\tau^D \epsilon\ \mathrm{T}^*$
>   Collect y^* for τ^M
>   else:
> $y^* \leftarrow \arg min \sum \left|\sigma_{put}(y) - \sigma_{call}(y)\right|$
>   Collect y^* for τ^M
>   end if
>
> $\tau^B \tau^M, \overline{y} \leftarrow y^*$
> end for

6.4 Fitting of the Implied Volatility Surface

6.4.1 De-Americanization of Option Prices

Once dividends and repo rates have been inferred, it is possible to use the observed prices of American option to infer the prices of equivalent European options on the same underlier ("de-Americanisation")[16]. This approach consists of calculating the implied volatility of an American option and using this to price the equivalent European option with the same contract specifications and market inputs.

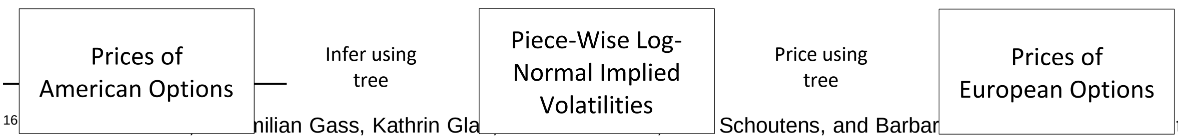

[16] … nilian Gass, Kathrin Gla … Schoutens, and Barbar … to american options: numerical investigation of the de-americanization method. Quantitative Finance, 18(7):1091-1113, 2018.

Note that the procedure is performed for each option separately assuming a piece-wise log-normal dynamics and a different volatility for each maturity and strike. This is an approximation since a different underlying dynamic is assumed for each combination of maturity and strike. Fitting a consistent model for the underlying, or equivalently the implied volatility surface, directly on American option prices is not feasible due to the significant computational burden required to optimise the model parameters in the absence of a closed-form pricing formula. De-Americanisation allows for the separation of i) the step of deriving European option prices from the corresponding American option prices and ii) the fitting of a model on the latter.

6.4.2 Fitting of European Prices using the Extended Surface SVI model

Given the European prices derived in the previous section, a model for the whole volatility surface is calibrated in a similar manner to the procedure for Equity Index Options (see Section 5.2).

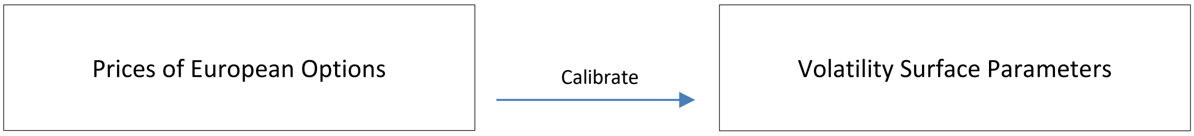

6.4.3 Pivot-Points Derivation

Given the model fitted in the previous section, the implied volatility at the selected pivot points can be evaluated, similarly to Index Options.

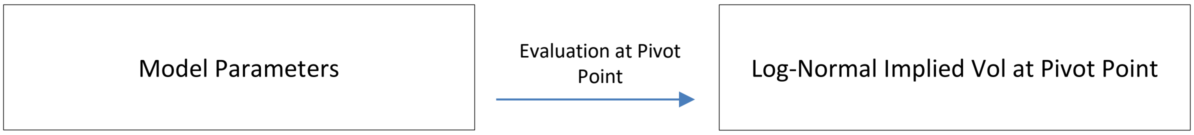

These implied volatilities at the pivot points are transformed back into their American counterparts. This step consists in the inverse of the procedure described in Section 3.4.1. The piece-wise log-normal implied volatilities are therefore the risk-factors to be modelled. Note that working directly with the piece-wise log-normal implied volatilities instead of the standard Black-Scholes implied volatilities allows to avoid performing an inversion over each scenario, while in both cases any guarantee of absence of arbitrage is lost while generating scenarios.

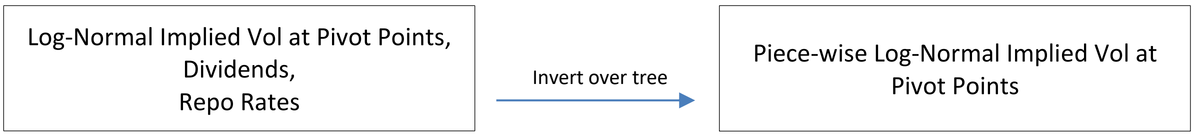

Finally, implied volatility returns are derived by applying first differences to the reconstructed pivot point time-series.

6.5 Single Stock Option P&L scenarios

The methodology for generating single stock option P&L scenarios is identical to that used to generate scenarios for equity index options, described in Section 5.4, apart from the fact that, in subsection 5.4.3 the Black formula used for pricing the options is replaced with the piece-wise lognormal tree-pricing formula described above.

6.6 Corporate Actions

6.6.1 Introduction

If a corporate action occurs with respect to a cash equity that is the underlier of a single stock option, then it may be necessary to apply changes to the contract specifications of that option position to ensure that both parties to the option trade are not unduly disadvantaged. The principal underlying any such changes is that the value of the position should not be affected. It is necessary to apply adjustments to the data used to construct the single stock options scenarios to ensure that these scenarios are not unduly affected.

6.6.2 Adjustments to single stock options due to corporate actions

In the case of a special dividend, rights issue or stock-split the affected options will be adjusted in several ways:

- An adjustment factor will be calculated representing the ratio between the value of the stock after the corporate action relative to the value of the stock prior to the corporate action.
- The strike price of each affected option will be multiplied by the adjustment factor.
- In the case of a stock-split or reverse stock-split, the lot size of each affected option will be divided by the adjustment factor.

In the case of a demerger or spin-off, a position in an option on a single stock may be replaced with a position in an option on a basket of two or more underlying stocks. In these cases, it is not necessary to adjust either the strike or the lot size of the option contract.

In each of the cases mentioned above, these changes would be affected on the ex-date of the corporate action and would be applied to all positions open as of end-of-day on the prior business date. In practice, this change would be affected by the following series of steps:

1. The creation of a new option for each existing option, which would have i) a new, incremented, version number that allows for them to be distinguished from the options listed prior to the corporate action, and ii) new contract specifications that reflect the adjustments detailed above.
2. The replacement of positions in the old option series with positions in the new option series.
3. The delisting of the old options series.

In the case of a special dividend, rights issue or stock-split, the adjustments detailed above would not lead to a change in the implied volatility of the option position. i.e. the implied volatility on the option position created as a result of the corporate action should, all other things being equal, be identical to the implied volatility of the option prior to the corporate action. For this reason, it is not necessary to apply any adjustment to the implied volatility pivot-points when generating the options scenarios. Corporate actions of this sort do, however, lead to a jump in the price of the underlying stock. If ClearisQ was to use a historical time-series of stock prices that had not been adjusted to correct for this jump, it would be interpreted as a market-driven price change which could lead to a spurious increase in the Initial Margin. For this reason ClearisQ will use a historical time-series of underlying equity prices that has been adjusted to correct for the impact of corporate actions when generating the single stock options scenarios[17]. For clarity, the historical observed values of the implied volatility will not re-calculated using these adjusted stock prices, as to do so would introduce a distortion to the historical implied volatility and, by extension, the option P&L scenarios.

In the case of a spin-off, the price of the basket of stocks on the date of the corporate action will be identical to the price of the underlying stock immediately prior to the corporate action. Following the corporate action, ClearisQ will need to fit a volatility surface for, and construct option P&L scenarios for the basket option. In

[17] ClearisQ also uses historical equity prices that have been adjusted for corporate actions when calculating scenarios and, by extension, Initial Margin for cash equity positions.

order to do this it is necessary to calculate the price of the basket of equities for each date. This is done by taking the weighted sum of the constituents of the basket.

$$P_{basket} = \sum_i P_i W_i \quad (6.14)$$

where:

P_{basket} is the price of the equity basket,

P_i is the price of constituent *i* of the basket,

W_i is the weight of constituent *i* of the basket.

In order to calculate option P&L scenarios ClearisQ will use:

- a combined series of historical prices comprised of i) the single stock prices observed prior to the corporate action, and ii) the basket prices observed after the corporate action
- a combined series of historical implied volatilities comprised of i) the single stock implied volatility levels observed prior to the corporate action and ii) the basket implied volatility levels observed after the corporate action.

7. Filtered Historical Simulation

The FHS methodology is based on the simulation of the distribution of future returns (scenarios) by sampling the realized past returns, scaled by the ratio of past and most recent volatility estimate. Compared to a historical simulation approach (HS), which directly identifies the scenarios with the past returns and thus assumes that returns are independent and identically distributed (i.i.d.), the FHS directly models the volatility dynamics, without making any parametric distributional assumptions. This point is a strong advantage as it allows to capture non-trivial statistical properties of real returns distributions, such as asymmetry (skewness), fat-tails (kurtosis), volatility clustering, and tail dependence. Finally, the FHS methodology is not limited to equity products and/or linear instruments and it will be here applied to futures and options.

The computation of initial margins within the chosen FHS-CVaR model extends the methodology applicable to cash equity products cleared by Cboe Clear. In this section, we describe in detail the computation of IM within the FHS model.

7.1 Data inputs

The inputs for the scenario generation steps are the time-series of the relevant risk-factors. In the case of equity index derivatives, the relevant risk factors are:

- Futures relative logarithmic returns at pivot time-to-maturity points (see Section 4);
- Implied volatility absolute returns at pivot time-to-maturity and moneyness points (see Section 5); and
- FX spot rate relative logarithmic returns to account for potential mismatches between the risk factor time-series currency and the Clearing Participant margining currency.

7.2 Volatility and residuals estimation

One of the primary aims of the IM methodology is to capture the changes in volatility, and to use these to produce future return scenarios that are consistent with the present level of volatility[18]. To achieve this, the FHS incorporates a dynamic measure of market volatility.

The starting assumption of the FHS model is that the returns are well modelled as:

$$p_t^i = \hat{\sigma}_t^i \varepsilon_t^i \tag{7.1}$$

where ε_t^i are normalised residuals (unit variance) and $\boldsymbol{\hat{\sigma}_t^i}$ is the forecasted instrument volatility on day t.

The volatilities are estimated from the historical returns of futures pivot points (Eq. 4.5) using an Exponentially Weighted Moving Average model (EWMA):

$$\left(\hat{\sigma}_t^i\right)^2 = \lambda\left(\hat{\sigma}_{t-1}^i\right)^2 + (1-\lambda)\left(p_{t-1}^i\right)^2 \quad , \qquad t \geq 2 \tag{7.2}$$

The normalised residuals are then estimated as:

$$\hat{\varepsilon}_t^i = \frac{p_t^i}{\hat{\sigma}_{t+1}^i} \tag{7.3}$$

Repeating the calculations (7.2) and (7.3) for all rows of the futures returns matrix (Eq. 4.5) gives the following matrix of normalised residuals:

[18] Only applicable to futures price returns.

$$\hat{\varepsilon} = \begin{pmatrix} \hat{\varepsilon}_t^1 & \hat{\varepsilon}_t^2 & \cdots & \hat{\varepsilon}_t^M \\ \hat{\varepsilon}_{t-1}^1 & \hat{\varepsilon}_{t-1}^2 & \cdots & \hat{\varepsilon}_{t-1}^M \\ \vdots & \vdots & \ddots & \vdots \\ \hat{\varepsilon}_{t-N^*+1}^1 & \hat{\varepsilon}_{t-N^*+1}^2 & \cdots & \hat{\varepsilon}_{t-N^*+1}^M \end{pmatrix} \tag{7.4}$$

The estimation of the volatilities according to Eq. (6.2) requires the specification of a seed, the initial value $\hat{\sigma}_1^i$. Given the fully available return time series p_j with j = 1, …, t, the seed value σ_1^2 is computed as the average over the first N_0 squared returns and the values are propagated using recursion from there on. The selected value for N_0 is 200 days, which corresponds approximately to twice the characteristic period of the EWMA for λ = 0.99, and therefore reflects the timescale of the volatility process.

7.3 Filtered scenarios generation

Once determined, the normalised residuals can be used to generate a set of hypothetical *MPOR*-scaled (ie. *h*-day) return scenarios for a defined number of scenarios N (where h = 3 and N = 700).

To obtain *h*-day return scenarios, we sum 1-day residuals over *h* consecutive days in order to capture possible auto-correlations, ie. trending or mean-reverting behaviours. The sums of residuals are then multiplied by the last available volatility, which under the EWMA volatility process assumption, is the one-day ahead forecast volatility. Note that the summation across residuals requires the use of logarithmic returns.

The *MPOR*-scaled hypothetical return scenario for instrument *i* and scenario *k* is calculated as:

$$\hat{r}_k^i = \hat{\sigma}_{T+1}^i \sum_{j=0}^{h-1} \hat{\varepsilon}_{T-k+1-j}^i \tag{7.5}$$

In matrix notation, the *MPOR*-scaled hypothetical returns are represented as:

$$\hat{r} = \begin{pmatrix} \hat{r}_1^1 & \hat{r}_1^i & \cdots & \hat{r}_1^I \\ \hat{r}_k^1 & \hat{r}_k^i & \cdots & \hat{r}_k^I \\ \vdots & \vdots & \ddots & \vdots \\ \hat{r}_N^1 & \hat{r}_N^i & \cdots & \hat{r}_N^I \end{pmatrix} \tag{7.6}$$

Each column of (6.6) represents a set 700 scenarios of hypothetical 3-day returns for instrument *i*.

7.4 FX treatment

For products denominated in a currency different from the portfolio base currency (ie. EUR), simulations of the relevant currency pairs are also required. The latter can be obtained by modelling the foreign exchange (FX) spot rate return p_t^{fx} for currency pair *fx* as:

$$p_t^{fx} = \sigma_t^{fx} \varepsilon_t^{fx} \tag{7.7}$$

The *MPOR*-scaled hypothetical return scenario for currency pair *fx* and scenario *k* is then calculated in the same manner as described in Sections 6.2 and 6.3, with instrument *i* being replaced with currency pair *fx*. These fx return scenarios $(\hat{r}_t^{fx})$ are then applied to the corresponding current FX spot rate $FX(t_0)$ to generate a set of *N* FX scenarios:

$$\widehat{FX}_{t_0+h}^t = FX(t_0)\left(exp\left(\hat{r}_t^{fx}\right) - 1\right) \tag{7.8}$$

The FX scenarios are then used, along with the corresponding futures and options return scenarios to calculate a set of P&L scenarios in the portfolio base currency. Eq. 4.10 and 5.25 illustrate how this is done for futures and options respectively.

7.5 P&L scenarios

7.5.1 Futures P&L scenarios

The process to generate contract-level P&L scenarios $\widehat{\Delta^h F^t_{t_0}}(\tau)$ for a future of time-to-maturity τ for all dates $t \in [1, N]$ of the lookback window is described in detail in Section 4.5.

The returns scenarios used to generate these contract-level P&L scenarios are calculated using the FHS process described in Sections 6.2 and 6.3.

7.5.2 Option P&L scenarios

The process to generate contract-level P&L scenarios $\widehat{\Delta^h P^t_{t_0}}(\tau, m)$ for an option of time-to-maturity τ and moneyness m for all dates $t \in [1, N]$ of the lookback window is described in detail in Section 5.4.

The inputs are return scenarios for the risk factors, namely futures returns at pivot time-to-maturity points and implied volatilities returns at pivot time-to-maturity and moneyness points. As described in Section 5.3.2, implied volatility scenarios are calculated based on absolute returns, rather than logarithmic, and are left unscaled in line with a pure HS methodology. The futures returns are calculated using FHS, as described in Sections 6.2 and 6.3, and are fully consistent with the calculation for futures P&L scenarios.

7.6 Margin calculation

7.6.1 Expected Shortfall risk measure

Margin requirements are obtained by calculating an ES risk measure for the portfolio based on the matrix of simulated P&L scenarios.

Given a random vector of portfolio P&L or return scenarios $\vec{x} = (x_1, \dots, x_N)$, the ES risk measure with confidence level α is estimated as follows[19]:

1. Sort $\vec{x}$ in ascending order, where $x_i \leq x_{i+1}$:

$$\vec{x} = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_N \end{pmatrix} \tag{7.9}$$

2. Compute $N_\alpha = \lfloor N(1-\alpha) \rfloor$ where $\lfloor y \rfloor$ is the closest integer smaller or equal to y. For example, if $\alpha = 99\%$ then $N_\alpha = 7$ (for $N = 700$).

3. The ES (or conditional value-at-risk) is then estimated as the average of all scenarios up to, and including scenario N_α

$$ES_\alpha = \frac{1}{N_\alpha} \sum_{n=1}^{N_\alpha} x_n \tag{7.10}$$

[19] Note that flooring to the closest integer is more conservative than applying interpolation to estimate the empirical quantile.

7.6.2 Portfolio margin limit rule

The initial margin for a given portfolio is subject to a portfolio margin limit rule in accordance with Article 27(4) of the EMIR RTS.

Accordingly, the portfolio IM is calculated as a weighted sum of the gross and net portfolio margins:

$$IM^{FHS} = (1-c)IM^{FHS}_{gross} + c.IM^{FHS}_{net} \tag{7.11}$$

where $c = 0.8$, and IM is expressed as a negative value. Note that the additive property of the ES tail measure guarantees that $IM^{FHS}_{net} > IM^{FHS}_{gross}$.

The gross portfolio margin IM^{FHS}_{gross} is the sum of the individual instrument margins:

$$IM^{FHS}_{gross} = \sum_{i=1}^{I} f\,(w_i) \tag{7.12}$$

and the net portfolio margin IM^{FHS}_{net} is the margin of the netter portfolio including the effect of diversification:

$$IM^{FHS}_{gross} = f\left(\sum_{i=1}^{I} w_i\right) \tag{7.13}$$

where f denotes the ES function described in Section 6.3.

7.6.3 Derivatives netting rule

As an important particular case of netting, derivatives with the same underlying are considered as products for which full netting is allowed because their correlation stems by construction from a sufficiently fundamental and stable reason. As a result, all components of the portfolio that have the same underlying can be aggregated at scenario level as if they were one instrument and the aggregation described in the previous section applies to the netted scenarios across underliers.

8. Anti-Procyclicality Measure

Article 28 of the EMIR RTS sets out requirements for CCPs to limit the procyclicality of margin requirements to the extent that the soundness and financial security of the CCP is not adversely affected.

To comply with these regulatory requirements, Cboe Clear applies an anti-procyclicality (APC) measure in the form of a stressed historical simulation (StressHS) margin which aims to reduce the procyclicality of the stand-alone FHS IM (ie. by applying the option set out under Article 28(1)(b) of the EMIR RTS).

The APC-compliant margin is then computed as the weighted sum of the FHS and the StressHS margins.

This section describes in detail the computation of the StressHS margin component and resultant APC-compliant margin.

8.1 StressHS margin

The StressHS margin component is based on a historical simulation, expected shortfall methodology. Simulation scenarios are obtained by combining a set of 650 of the recent past returns with a set of 50 returns that correspond to defined historical dates of exceptional market stress. The margins calculated using this model represent the potential expected loss for a given portfolio in stressed market conditions.

The incorporation of a less-variable margin component, derived from a relatively stable set of historical periods of stress provides for a conservative margin estimate that is less sensitive to the current dynamics of market volatility. As a result, margin requirements are more conservative during low volatility periods and less variable during high volatility periods.

8.1.1 Data inputs, scenario generation, and P&L scenarios

The StressHS computations are based on a combination of a set of historical stress scenarios with a set of the most recent scenarios, which together generate the set of simulated scenarios.

The 50 historical stress scenarios correspond to the Cboe Clear stress-testing scenario dates, which are defined by Cboe Clear as part of its overall stress testing framework. The selection and review of these stress scenarios is performed on a periodic basis and is subject to internal risk management and regulatory governance. Currently, Cboe Clear's stress-testing scenario suite contains 50 historical stress scenarios covering all relevant markets that Cboe Clear clears.

The 50 historical stress scenarios are complemented by the $(N - 50)$ most recent 3-day unscaled scenarios in the lookback period to form a total of N unscaled scenarios. The total length of the simulation window is $N = 700$ scenarios[20], and therefore the most recent 650 unscaled scenarios are included[21].

In accordance with ESMA regulatory guidelines, the StressHS margin component is based on unscaled (unfiltered) historical scenarios, and therefore the EWMA normalisation process described in Section 6 is not used for the generation of StressHS scenarios.

The processes described in Sections 4 and 5 can be used to directly generate a set of unscaled *h*-day P&L scenarios for the required number of recent lookback scenarios $N - 50$ (ie. 700 - 50 = 650) and for each stress period.

---

[20] The inclusion of 50 stress scenarios in the total simulation window of 700 scenarios is therefore sufficient to cover the tail of the distribution at the chosen confidence level (99%).

[21] Without duplication of scenarios, ie. where a scenario in the $(N - 50)$ lookback period corresponds with a stress scenario, it is treated as an invalid scenario for the purposes of scenario generation.

8.1.2 Margin calculation

Finally, the StressHS margin, IM^{SHS}, is computed similarly to Section 6.6, where the gross and net portfolio margin estimates are obtained by calculating ES risk measures for the portfolio based on the simulated P&L scenarios in matrix w^{SHS}, and the portfolio IM is calculated as a weighted sum of the gross and net portfolio margins in accordance with the portfolio margin limit rule.

Accordingly, the portfolio IM is calculated as a weighted sum of the gross and net portfolio margins:

$$IM^{SHS} = (1-c)IM^{SHS}_{gross} + c.IM^{SHS}_{net} \tag{8.1}$$

where,

$$IM^{SHS}_{gross} = \sum_{i=1}^{I} f\,(w_i^{SHS}) \tag{8.2}$$

$$IM^{SHS}_{gross} = f\left(\sum_{i=1}^{I} w_i^{SHS}\right) \tag{8.3}$$

and where again c = 0.8, f denotes the ES function described in Section 6.6.1, the derivatives netting rule described in Section 6.6.3 applies, and IM is expressed as a negative value.

8.2 APC-compliant margin

Once the StressHS margin component has been calculated as per Eq. 7.1, it is combined with the FHS margin calculated as per Eq. 6.11 to produce the unfloored APC-compliant margin in accordance with Article 28(1) of the EMIR RTS:

$$IM^{APC}_{unfloored} = (1-\eta)IM^{FHS} + \eta.IM^{SHS} \tag{8.4}$$

where η = 0.25.

The final, floored APC margin is then calculated as the max of the floored APC and the IM^{FHS} margin component. This step is undertaken to ensure that, during times of heightened market volatility, the APC margin will never drop below the level of the unadjusted IM^{FHS}.

$$= \max\left(IM^{APC}_{unfloored}, IM^{FHS}\right) \tag{8.5}$$

9. Treatment of equity collateral (Collateralised Underlying Exposure)

As part of the single stock options clearing service, Cboe Clear offers Clearing Participants the ability to provide equity securities for the purposes of generating risk and margin offsets against cleared single-name equity option positions. Margin offsets of this kind will only be applicable in cases where the single stock options positions reference the equity provided, and any offset will be limited by the extent to which simulated price-shifts in these equity securities serve to offset the simulated price-shifts in the single-stock options.

9.1 Equity collateral P&L scenarios

When equity collateral is provided by a Clearing Participant, Cboe Clear will generate FHS and StressHS scenarios for each equity, in line with the ClearisQ methodology for cash equities which is covered in the Cboe Clear Margin Model Description (Cash Equities and Exchange-Traded Products). Note that, per Section 6, this methodology is also employed to generate the equity spot price scenarios that are used to generate the single stock options price scenarios.

9.2 Calculation of Derivatives Initial Margin Offset

Once equity collateral FHS and StressHS P&L scenarios have been calculated, these are combined with the corresponding single stock option FHS and StressHS P&L scenarios to determine whether an offset exists, and what the value of the offset should be. To do this, it is necessary to sum the P&L of all options on a given underlying. The FHS P&L scenarios for single stock options instruments referencing a given underlying can be expressed as:

$$\hat{x} = \begin{pmatrix} \hat{x}_1^1 & \hat{x}_1^i & \dots & \hat{x}_1^I \\ \hat{x}_k^1 & \hat{x}_k^i & \dots & \hat{x}_k^I \\ \vdots & \vdots & \ddots & \vdots \\ \hat{x}_N^1 & \hat{x}_N^i & \cdots & \hat{x}_N^I \end{pmatrix} \tag{9.1}$$

Each column of $\hat{x}$ represents a set 700 FHS P&L scenarios for single stock option instrument *i* referencing a given underlying equity. These scenarios are summed by scenario, resulting in a vector containing a single netted P&L observation for each of the 700 FHS and StressHS scenarios. When performing this summation, the proxy scaling factor described in Section 10 is applied (i.e. any P&L scenario which is flagged as proxied, and which generates a risk offset will be multiplied by the proxy scaling factor ϕ).

$$\hat{x} = \begin{pmatrix} \sum_i \hat{x}_1^i \\ \sum_i x_k^i \\ \vdots \\ \sum_i \hat{x}_N^i \end{pmatrix} \tag{9.2}$$

Similarly, the P&L scenarios generated for the corresponding equity collateral can be expressed as:

$$\widehat{m} = \begin{pmatrix} \widehat{m}_1 \\ \widehat{m}_k \\ \vdots \\ \widehat{m}_N \end{pmatrix} \tag{9.3}$$

For each scenario, the adjusted scenario P&L is then calculated based on the corresponding single stock option and equity collateral P&L scenarios using the formula below.

$$\hat{d}_k = \begin{cases} \min\left(\sum_i \hat{x}_k^i + \widehat{m}_k, 0\right) & if \qquad \sum_i \hat{x}_k^i < 0 \; and \; \widehat{m}_k > 0 \\ \sum_i \hat{r}_k^i & \qquad otherwise \end{cases} \tag{9.4}$$

Applying the offset in this way ensures that the equity collateral is only used within the model in cases where it serves to offset the risk of the corresponding single stock option positions. The inclusion of the equity collateral in this way will never lead to an increase in Initial Margin, nor will it serve to offset the risk on any other instruments within the portfolio. With this adjustment applied to all scenarios, the vector of adjusted P&L scenarios for the given equity underlying can be expressed as:

$$\hat{d} = \begin{pmatrix} \hat{d}_1 \\ \hat{d}_k \\ \vdots \\ \hat{d}_N \end{pmatrix} \tag{9.5}$$

This vector of adjusted scenario P&Ls is then carried forward to the subsequent stages of the Expected shortfall calculation detailed in Section 7.6.1. This same adjustment process is carried out for the FHS and StressHS scenarios and, in each case, the resulting vector of adjusted P&L scenarios is used in both the net and gross Expected Shortfall calculation.

10. Treatment of missing data

In this chapter the treatment of missing data is presented. Two cases are considered:

- Forward-filling: missing values after the first available date for each individual risk factor time-series are forward filled. As a new observation is collected, the latter will results in practice in a return from the last available observation. From a scenarios perspective, a risk factor with multiple missing values will results in many zeros and few large returns, with the latter largely driving the margin value[22]. This approach is therefore considered conservative.
- Proxying: missing data before the first available date for each individual risk factor time-series are proxied, i.e. filled, in some cases with a scaling factor, with the returns of a similar available time-series. The proxying methodology differs across risk factors.

In the rest of this chapter, the proxying methodology for the latter case, is described for each risk factor type.

10.1 Proxying for equity underlying

For equity underlying, a proxy instrument is chosen based on its domicile and sector[23]. A set of proxy indexes, each mapped to a country and/or a sector, is provided by the user. For each underlying, the proxy is selected among the available ones trying to match in the following priority order:

- country and sector,
- sector,
- country.

If no matches are found, a fallback proxy index is used. The choice of using domicile and sector indexes balances out the needs of having a tractable set of proxies, an extended coverage and a clear mapping between instruments and proxies.

Our approach takes into account the fact that indexes are by construction less volatile than single equities and therefore a rescaling coefficient is applied. Explicitly for an instrument *X* with associated proxy index *I*, a missing return at time *t* is proxied according to:

$$\rho_t^X = \beta \phi \rho_t^I \tag{10.1}$$

Here $\beta \epsilon \{-1,1\}$ is the sign of the correlation between *X* and *I*, i.e. β = 1 if $\text{Corr}(\rho_t^X, \rho_t^I) > 0$ and β = -1 if $\text{Corr}(\rho_t^X, \rho_t^I) < 0$ where Corr(.) denotes the Pearson correlation. In order to limit the estimation errors a minimum of n_β = 20 valid returns is required to compute β, otherwise a default value of β = +1 is used.

The proxy scaling factor ϕ in Eq. (6.1), which can be interpreted as the ratio of the instrument and index volatilities, determines the initial margin level and is set conservatively equal to 3, based on sensitivity analysis conducted on IPO stocks. As more underlying returns become available, the impact of ϕ decays exponentially, as per the EWMA volatility process.

---

[22] Note that, in the event that a zero risk factor return is observed, the EWMA formulation adopted in CleaRisQ uses the previously available volatility value which ensures that multiple missing observations do not result in a decrease in the estimated volatility and, by extension, the associated FHS margin.

[23] Defined according to the GICS Sector classification.

10.2 Proxying for index future curves

Proxying of the future curve returns is required for a given underlying index when on a given date:

- no underlying return is available, and
- no future returns are available.

If at least one return is observed, the standard interpolation/extrapolation rule is used to derive the risk factors at the pivot-point constant time-to-maturities. Note that extrapolated returns are marked as “proxied". When no returns over the whole term structure are observed, proxying is applied. For index futures, this only occurs over the stress date intervals and not over the standard lookback windows used both for daily margining and historical backtesting. A fallback index is defined for each stress date interval. If no returns are observed for given index and stress interval, the pivot-point returns from the fallback index are used over the whole term structure of that proxied index and they are all marked as “proxied". No scaling is applied, since the volatilities across indexes are comparable.

10.3 Proxying for implied volatilities

Proxying of the implied volatility surface is required for a given underlying when on a given date:

- no options data are available at all, or
- the criteria to perform the fitting are not met,)e.g. minimum number of options per maturity, minimum number of maturities, etc.), or
- the fitting procedure failed.

If the fitting process is successful, all pivot points are computed according to the fitted model. If it was not possible to fit a model on a given date, two different approaches apply for indexes and equities.

10.3.1 Proxying for index implied volatilities

If the underlying is an index, proxying will be required only over the stress date intervals and not over the standard lookback windows used both for daily margining and historical backtesting. A fallback index is defined for each stress date interval. If implied volatility returns cannot be derived for given index and stress interval, the pivot-point returns from the fallback index are used over the whole surface of that proxied index and they are all marked as “proxied". No scaling is applied, since the volatilities across indexes are comparable.

10.3.2 Proxying for equity implied volatilities

If the underlying is an equity, the proxy index is selected based on the following priority order:

1. index containing the underlying as one of its constituents (if multiple indexes, select the one with least constituents),
2. index matching the domicile of the underlying,
3. fallback index.

Proxying is performed for all implied volatility pivot points returns using the ratio of the EWMAs of the underlyings as scaling factor:

$$\rho_t^{X,\tau,m} = \frac{\sigma_t^X}{\sigma_t^I} \rho_t^{I,\tau,m} \qquad \forall \tau, m \tag{10.2}$$

where:

X denotes the proxied underlying,

I is the index underlying,

$\rho_t^{j,\tau,m}$ is the return at time t of the implied volatility pivot point defined at time-to-maturity and moneyness m for the generic underlying j,

σ_t^j is the EWMA volatility at time t of the generic underlying j.

11. Margin add-ons

In addition to underlying price volatility risk, Cboe Clear is also exposed to other specific risks which require separate treatment in the form of margin add-ons. The following margin add-ons are applied by Cboe Clear and included in the total initial margin requirement:

- Wrong-Way Risk (WWR) add-on
- Liquidity Risk (LR) add-on
- Large Position (LP) add-on

This section describes the methodology for calculating the above margin add-ons.

11.1 Wrong-way risk (WWR) add-on

The WWR add-on addresses the specific wrong-way risk arising when a Clearing Participant has a long exposure to its own stock or those of an institution belonging to the same financial group. In this case, the position attracts a 100% fixed margin, reflecting the fact that in case of default the position could be worthless.

In the case of index-based products, it is possible that a Clearing Participant's stock enters the composition of the index. In that case, their weight in the index α is used to evaluate the loss assuming that, in a default, the Clearing Participant's stock prices go to zero, leading to the following add-on:

$$WWR_i = \max\left(\Delta(S,t_0)\alpha v_i - \frac{1}{2}\Gamma(S,t_0)(\alpha v_i)^2, 0\right) \qquad (11.1)$$

where αv_i is the P&L of the underlying assuming the above, and where Δ and Γ are the aggregated deltas and gammas of the derivatives written on the underlying involving wrong-way risk. In theory, all the positions used to compute the initial margin should be reduced by α within the IM methodology but, for computational efficiency, this correction is neglected, making the overall approach slightly more conservative.

In the case of single stock options written on underlyings linked to a Clearing Participant, a scenario is constructed wherein the price of the underlying is assumed to drop to zero. In this scenario, the prices of all put options for those underlyings will be equal to their strikes, and the prices of all call options for those underlyings will be equal to zero. If, under this scenario, the P&L of all options written on a given underlying is negative, then this P&L will be multiplied by minus one and added to the WWR add-on. In the case that the P&L is positive, no change is made to the WWR add-on.

For a given Clearing Participant, the portfolio WWR add-on is the sum of all position WWR add-ons:

$$WWR_{portfolio} = \sum_{i=1}^{I} WWR_i \qquad (11.2)$$

11.2 Liquidity Risk (LR) add-on

The LR add-on addresses the specific risk due to concentrated and/or illiquid positions in a Clearing Participant portfolio. This add-on captures the potential price impact of liquidating such positions in the event of a Clearing Participant default.

In such a case, liquidation of large and/or illiquid positions may incur additional costs beyond those captured in the standard margin computations, which address market risk only. The Liquidity Risk add-on described below is designed to capture the key risk factors on those portfolios, and to reflect the expected hedging and liquidation strategy that would be applied by Cboe Clear in the event of a Clearing Participant default. Liquidation costs arising from offsetting positions across separate maturity buckets are accounted for within the vega liquidity risk component on a gross basis; in particular, positions of contracts expiring within MPOR are accounted for as a separate and stand-alone maturity bucket to prevent any offset benefit against other longer-dated maturities being allowed. The total liquidity risk add-on for each underlier for each portfolio is calculated as the sum of the directional delta liquidity risk add-on, the delta calendar spread liquidity risk add-on, the vega bid-mid liquidity risk add-on and the vega concentration risk add-on.

11.2.1 Directional delta liquidity risk add-on

The delta liquidity risk add-on is calculated as the sum of the directional delta risk add-on and the calendar spread liquidity risk add-on. In order to compute the directional delta LR add-on, the following additional time-series is needed for each underlying *i*:

- the percentage bid-ask spread $Sp_{i,t}$, defined as:

$$Sp_{i,t} = \frac{Ask_{i,t} - Bid_{i,t}}{0.5(Ask_{i,t} + Bid_{i,t})} \tag{11.3}$$

- the daily traded volume[24] denoted as $V_{i,t}$.

In the case of equity index futures, these values are calculated based on the front month future contract. In the case of single stock underliers, they are calculated based on the underlying equity. The directional delta LR add-on computed on day *t* for underlying instrument *i*, as percentage of the instrument price, is defined as:

$$LR\%_{delta,i,t} = 0.5\overline{Sp}_{i,t} + \mu_i \sigma_{i,t+1} \sqrt{\frac{|\Delta_{i,t}|}{\bar{V}_{i,t}}} \tag{11.4}$$

where:

$\overline{Sp}_{i,t}$ is the simple average bid-ask spread computed over the last 250 observations of instrument *i*,
$\Delta_{i,t}$ is the total net delta for all options and futures in underlying *i* to be liquidated[25],
μ_i is a coefficient measuring market impact, with default value equal to 1,
$\bar{V}_{i,t}$ is the average traded volume computed over the last N_{liq} observations of V_t, and
$\sigma_{i,t+1}$ is the EWMA volatility computed on day *t*, as for instrument *i* (Eq. 6.2).

In absolute terms, the delta LR add-on for a given underlier instrument *i* is given by:

$$LR_{,i,t} = v_{i,t}\Delta_{i,t}C_i(LR\%_{futures,i,t}) \tag{11.5}$$

Where:
$v_{i,t}$ is the last valid value of the position on instrument *i*.
$\Delta_{i,t}$ is the net delta of the position in underlying instrument *i* to be liquidated on day *t* (expressed as the number of contracts)
C_i is the contract multiplier or lot size of underlying instrument *i.*

[24] Where required, the primary listing of the instrument is used to avoid underestimating the instrument liquidity.

[25] For each Clearing Participant and for each index, position size is determined aggregating the deltas of all positions on the index across all accounts, so long and short positions are allowed to cancel out.

The first part of the LR add-on accounts for the *exogenous* liquidity component, which does not depend on the position size and it covers half of the bid-ask spread to reflect the use of mid-prices in the margin computations.

The second part of the LR add-on accounts for the endogenous liquidity component, which considers the market impact of liquidating the position given the position size[26]. As the position to be liquidated gets larger, it will incur additional costs.

In cases where data inputs are missing:

- If an observation $SP_{i,t}$ (or $V_{i,t}$) is not available, the missing value is set equal to the last available rolling average $SP_{i,t-1}$ (or $V_{i,t-1}$).
- Where observations of $SP_{i,t}$ are not available at the start of the observation window, the spread time series is pre-filled with conservative default values, Sp_i^*, set to 5% in order to start the rolling average calculation (typical bid-ask quotes for large caps are in the range of few percentage basis points).
- Similarly, where observations of $V_{i,t}$ are not available at the start of the observation window, the volume time series is pre-filled with conservative default values Q_i/π, where π = 0.2. This implies that on the first trading day, $Q_i/V_i = \pi$. In other words, it is assumed that any position in an IPO stock constitutes a fraction π of the unknown tradeable volume[27]. This assumption fades out progressively as actual volume data are observed. For new IPOs, the volumes over the first week of trading are discarded as these are not representative of the standard market trading activity.

For a given Clearing Participant, the total directional delta LR add-on on day *t* is calculated as:

$$LR_{portfolio,delta,t} = \sum_{i=1}^{I} LR_{delta,i,t}\, FX_{i,t} \tag{11.6}$$

Where:
$FX_{i,t}$ is the FX rate from the currency of underlying instrument *i* to the base currency (EUR).

11.2.2 Delta calendar spread liquidity risk add-on

The directional delta liquidity add-on is designed to address the liquidity cost associated with hedging the directional delta risk on the portfolio. The delta calendar spread liquidity add-on is designed to address the residual liquidity risk associated with liquidating the calendar spread risk that remains after the directional delta has been hedged. In order to estimate the potential liquidation cost of the residual calendar spread risk, it is first necessary to define a set of maturity buckets.

Table 11.1: Delta calendar spread liquidity add-on maturity buckets

Bucket	Time to maturity start (years)	Time to maturity end (years)
1	0	0.25
2	0.25	0.5
3	0.5	0.75
4	0.75	1
5	1	1.5
6	1.5	Infinity

[26] The shape of the market impact function has been widely investigated in the literature and the square-root form with γ_i = 1 adopted in Eq. 5.4 has been shown to perform well with respect to observable trading data. Commercial applications using the square-root form of the market impact include for instance Barra and Bloomberg.

[27] In practice, trades larger than 10% of the daily volume are rarely observed.

The residual calendar spread delta is then calculated using the formulas below. This represents half the gross delta remaining on the portfolio after the net delta has been hedged using futures within the first maturity bucket, or using the underlying itself (in the case of single stock options). This is equivalent to the delta on one leg of the calendar spread which would be required to hedge the residual delta at each maturity bucket.

$$H_i = \sum_m \Delta_{m,i} \tag{11.7}$$

$$\Delta_{m,i,h} = \begin{cases} \Delta_{m,i} - H_i & m = 1 \\ \Delta_{m,i} & m \neq 1 \end{cases} \tag{11.8}$$

$$\Delta_{calendar,i} = \frac{\sum_{m,h} \left| \Delta_{m,i,h} \right|}{2} \tag{11.9}$$

Where:
$\Delta_{m,i}$ is the aggregated net delta for the *m*-th maturity bucket.
H_i represents the total directional hedge trade which is executed in the underlying itself (in the case of single stock options), or in the front month future (in the case of equity index futures).
In order to calculate the calendar spread liquidity add-on, the calendar spread delta must be multiplied by the bid-ask spread of a calendar position. The bid-ask spread of a calendar position is not directly observable, but can be estimated as a fixed percentage of the bid-ask of either i) the front month future or ii) (in the case of single stock options), the underlying equity.

$$\overline{Sp}_{calendar,i,t} = \overline{Sp}_{i,t}.K \tag{11.10}$$

The parameter K represents the ratio between the bid-ask spread of a calendar position to the bid-ask spread of an outright position in the underlier. This parameter is calibrated periodically based on the relative volatility of calendar spread positions vs outright positions in the underlier. The default value is 0.25, which was determined based on the 10 most liquid European equity index futures contracts.

11.2.3 Vega bid-mid risk add-on

The vega risk on a portfolio may be hedged by executing at-the-money straddles at each maturity bucket. The vega bid-mid add-on addresses the bid-to-mid cost associated with executing these hedging trades. A predefined constant $\overline{\sigma'_{m,\iota}}$ is a conservative estimate of both:

i) the bid-to-mid cost of hedging one unit of vega for the 1st maturity bucket in underlying instrument *i*, and

ii) the bid-to-mid cost of hedging one unit of vega-neutral at-the-money straddle calendar spreads between the 1st maturity bucket, and the *m*-th maturity bucket. (i.e. at at-the-money straddle calendar spread with an exposure of one unit of vega long (short) in the 1st maturity bucket and 1 unit of vega short (long) in the *m*-th maturity bucket.

The vega bid-mid add-on for the *m*-th maturity bucket is calculated as:

$$LR_{vega\ bid-mid,t,m,i} = v_{m,i}\overline{\sigma'_\iota} \tag{11.11}$$

Where:
$v_{m,i}$ is the aggregated vega for the *m*-th vega maturity bucket for all option positions on underlying instrument *i*.

$$\overline{\sigma'_\iota} = \max_m\left(\overline{\sigma'_{\iota,m}}\right) \tag{11.12}$$

$\overline{\sigma'_{m,\iota}}$ is calibrated periodically for each instrument and maturity bucket based on the average of the following metric over the preceding 250 trading days. This metric is designed to capture the bid-to-mid on the at-the-money straddle, after the bid-to-mid cost of any residual delta risk has been removed.

$$\sigma_{m,i}^{l} = 0.5.\Big(\big(C_{atm,ask,i,m} + P_{atm,ask,i,m}\big) - \big(C_{atm,bid,i,m} + P_{atm,bid,i,m}\big) - \big|\Delta_{c,atm,i,m} + \Delta_{p,atm,i,m}\big|.\left(i_{ask} - i_{bid}\right)\Big) \tag{11.13}$$

Where:
$C_{atm,ask,i,m}$ is the ask price of the at-the-money call option on underlying instrument *i* at a maturity within bucket *m*
$P_{atm,ask,i,m}$ is the ask price of the at-the-money put option on underlying instrument *I* at a maturity within bucket *m*
$\Delta_{c,atm,i,m}$ is the delta on the at-the-money call option on underlying instrument *I* at a maturity within bucket *m*.
$\Delta_{p,atm,i,m}$ is the delta on the at-the-money put option on underlying instrument *I* at a maturity within bucket *m*.
i_{ask} is the ask price of underlying instrument *i.*
i_{bid} is the bid price of underlying instrument *i.*

Although the implied volatilities observed at different maturities are generally highly correlated, a conservative estimation of the total vega hedging costs is calculated by taking the maximum of:

1. the total hedging cost for all long vega maturities, and
2. the total hedging cost for all short vega maturities.

Very short dated vega (<1 week to maturity) is treated separately and, as a result, the add-on for these positions will be calculated independently of other maturities.

The overall the vega bid-mid liquidity component is calculated as:

$$LR_{vega\ bid-mid,t} = \max\left(\sum_{m=1}^{M} \max\left(LR_{vega\ bid-mid,t,m}, 0\right), \sum_{m=1}^{M} \max\left(-LR_{vega\ bid-mid,t,m}, 0\right)\right) \tag{11.14}$$

11.2.4 Vega concentration risk add-on

In addition to the bid-to-mid cost, it is also necessary to apply an add-on which reflects the additional cost that could be incurred when hedging a vega exposure which is large relative to the average days traded volume. In order to assess this, the vega concentration add-on is calculated as:

$$LR_{vega\ concentration,t,m,i} = \left|v_{i,t}\right| \mu_i \varrho_{i,t+1} \sqrt{\frac{\left|v_{i,t}\right|}{\overline{W}_{i,t}}} \tag{11.15}$$

Where:
$v_{i,t}$ is the aggregate vega for all option positions on underlying instrument *i*
μ_i is the market impact coefficient with default value of 1

$\varrho_{i,t+1}$ is the projected daily volatility of volatility for underlier *i* on day *t*+1. This is calculated using the EWMA volatility of the 6-month ATM implied volatility pivot point, using the same seed vol period and lambda parameter as is used for the price of underlying instruments.

$\overline{W}_{i,t}$ is the average traded vega for a given date and underlying computed over the last N_{liq} observations of V_t,. The daily traded vega ($W_{i,t}$) is calculated as:

$$W_{i,t} = \sum_j \alpha_{j,t}\beta_{j,t} \tag{11.16}$$

Where:
$\alpha_{j,t}$ is the vega of option *j* on day *t*
$\beta_{j,t}$ is the traded volume of option *j* on day *t*

11.2.5 Total Liquidity Risk Add-On

The overall liquidity cost incurred is finally calculated as the sum of the liquidity cost of the directional delta, the calendar spread delta and the vega contribution:

$$LR_{portfolio,t} = LR_{directional\ delta,t} + LR_{calendar\ delta,t} + LR_{vega\ bid-mid,t} + LR_{vega\ concentration,t} \tag{11.17}$$

The liquidity risk add-on will be calculated as the sum of the liquidity risk ass-on for each underlying independently, with no hedging benefit recognised for offsetting short and long positions with different underliers.

11.3 Large Position (LP) add-on

The LP add-on reflects the potential excess stress loss in a Clearing Participant portfolio over and above available financial resources. This add-on captures the excess loss exposure taking into account available resources in the equity derivatives segment of the mutualised Clearing Fund (under a Cover-2 standard) and the Clearing Participant's individual margin collateral posted.

For a given Clearing Participant, the portfolio LP add-on is calculated as follows:

$$LP_{portfolio} = min[0, MaxSTL - TM + \varsigma_i . DF] \tag{11.18}$$

where,

$MaxSTL$ is the maximum equity derivatives-related stress loss observed for the clearing member (expressed as a negative value).

TM is the total equity derivatives-related margin requirement including WWR for the Clearing Participant (expressed as a negative value).

DF is size of the equity derivatives segment of the Cboe Clear Clearing Fund.

ς_i is defined by Clearing Participant, with a maximum (default) value of 0.45.

The LR add-on has been removed from the total margin requirement to ensure that this figure is comparable with the stress scenarios, which deal exclusively with market risk.

12. Treatment of instruments expiring during the MPOR

As stated in Section3.2.3, the methodology assumes an MPOR of 3 business days. For this reason, it is important that the methodology considers the increase in risk that could arise due to the expiry of instruments during the margin period of risk. For this reason, the Initial Margin is calculated as the maximum of:

i) The margin calculated across all positions in the portfolio and

ii) The margin calculated on only the positions that are not within the MPOR.

Appendix A – Margin model parameters

Table A.1 – FHS and StressHS margin model parameters

Parameter	Symbol	Value
Lookback window	N	700 days
EWMA decay factor	λ	0.99
EWMA seed period	-	200 days
Portfolio margin limit coefficient	c	0.8
StressHS margin coefficient	H	0.25
Proxy scaling factor	ϕ	0.8
Volatility surface fit threshold	ϖ	0.6%