MOLECULE

White Paper

# VALUE-AT-RISK FOR ENERGY & COMMODITY PORTFOLIOS.

A Technical and Practical Framework for Market Risk in Energy & Commodity Portfolios

# EXECUTIVE SUMMARY.

Energy and commodity portfolios exhibit complex behaviors: nonlinear exposures, volatile correlations, and physical asset linkages. Measuring market risk in such portfolios requires a VaR engine that is transparent, fast, mathematically rigorous, and aligned with operational realities.

Molecule's Value-at-Risk (VaR) engine is designed to provide real-time risk visibility to trading organizations, utilities, producers, and asset-backed market participants. The system is designed according to the following principles:

- **Transparency –** Risk teams can understand and audit the model
- **Speed –** The engine supports intraday runs
- **Accuracy –** Covariance and pricing must reflect current markets
- **Configurability –** Users can adjust horizon, confidence level, and lookback periods
- **Energy-specific modeling –** Asset-linked exposures, multi-curve structures, nonlinear derivatives

This whitepaper presents the mathematical foundations, data requirements, operational workflows, and modeling assumptions behind Molecule's VaR solution.

# 1. MARKET DATA AND REAL-TIME COVARIANCE.

Molecule VaR ingests customer market data in real time and applies a 15-minute runtime delay to ensure:

- VaR runs do not overlap
- Updated positions, not just prices, are captured
- Consistent market snapshots across all risk factors

Log returns are calculated as:

$$r_t = ln\left(\frac{P_t}{P_{t-1}}\right) \quad (1)$$

The covariance matrix Σ is then computed from these log returns:

$$\Sigma_{ij} = Cov(r_i, r_j) \quad (2)$$

These log returns are calculated with a minimum 30-day lookback window, configurable by the user. If the covariance matrix is not strictly positive definite, small or negative eigenvalues are truncated to zero (or to a small positive floor) to ensure numerical stability during factorization.

The covariance matrix is rebuilt fresh on every VaR run, making the model highly responsive to regime changes such as:

- Seasonal shifts in curves
- Rapid movements in price volatility
- Sudden changes in cross-asset correlations
- Liquidity-driven dislocations in specific tenors or instruments

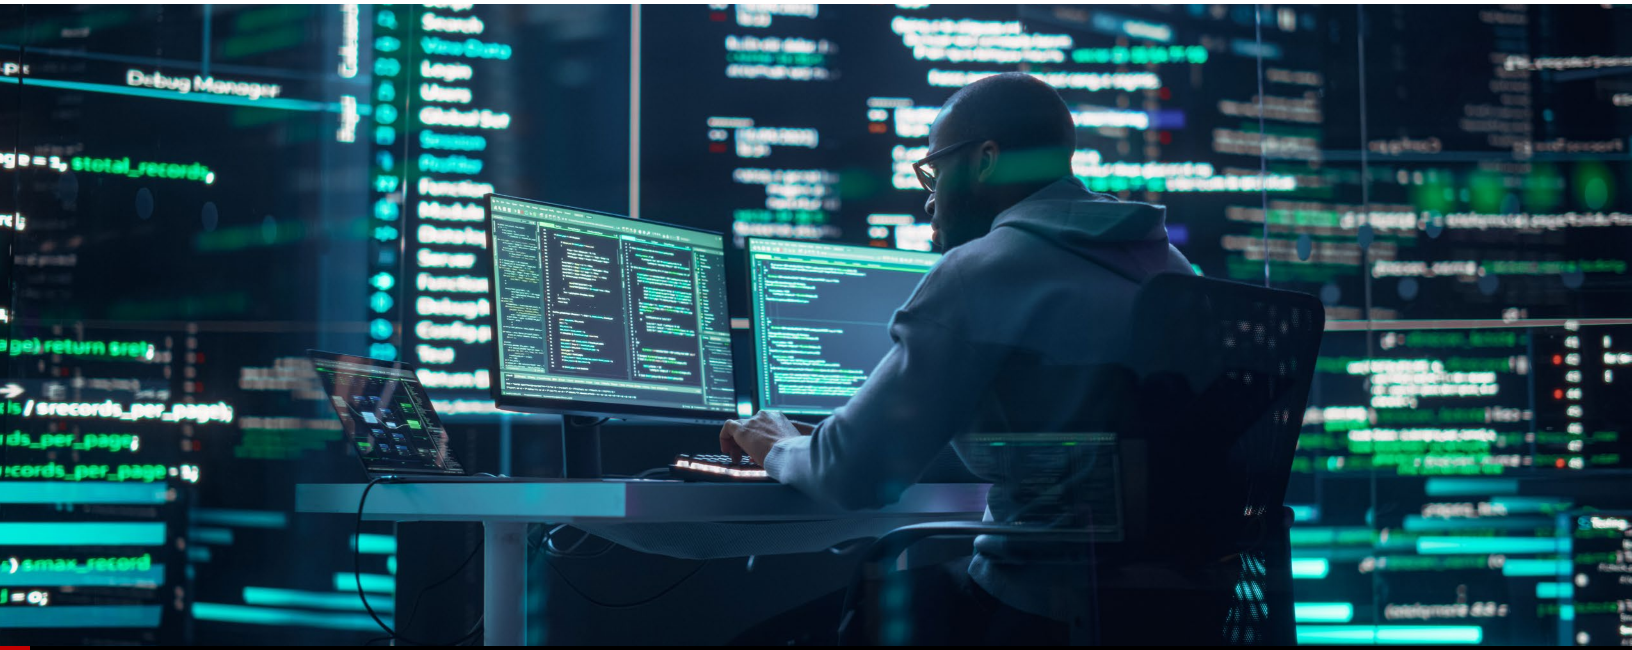

# 2. PRICING ENGINE: A HIGH-PERFORMANCE QUANTLIB-BASED FRAMEWORK.

Molecule's VaR engine uses a dedicated pricing subsystem optimized for risk calculations, distinct from the main platform's settlement engine. This design enables:

- High-speed computation
- Consistent scenario revaluation
- Configurable model assumptions

The VaR pricing stack incorporates:

- **QuantLib** for forward, swap, and option valuation (Black-76, Black-Scholes)
- **NumPy/SciPy** for linear algebra and matrix decomposition
- **Pandas** for time-series alignment

This engine supports:

- Forwards and futures (priced directly from marks)
- Swaps
- Vanilla options
- Structured instruments that reduce to analytically priced forms
- As-gen asset-linked exposures (long or short)
- The architecture enables efficient Monte Carlo revaluation.

The architecture enables efficient Monte Carlo revaluation.

# 3. HANDLING MODEL DOMAIN CONSTRAINTS (NEGATIVE PRICES AND SPREADS).

Because Molecule's VaR uses log returns, negative prices cannot be processed directly:

- Log returns are undefined for values $\leq 0$
- Multiplicative shocks distort near-zero nodes
- Correlations become unstable

When encountered:

- The affected values are excluded safely
- VaR continues without interruption
- Warnings are surfaced in the UI and diagnostics

This is a controlled, transparent behavior aligning with industry standards. Reasons for excluding negative values from calculations are surfaced in the diagnostics reports generated with every VaR run and accessible via UI.

# 4. PARAMETRIC (VARIANCE-COVARIANCE) VAR.

Parametric VaR is Molecule's fastest VaR method. It computes the first-order sensitivity of the portfolio with respect to each price factor:

$$\Delta_i = \frac{\partial V}{\partial P_i} \quad (3)$$

Given the covariance matrix $\Sigma$, the portfolio P&L variance is then:

$$\sigma^2{}_{portfolio} = \Delta^T \Sigma \Delta \quad (4)$$

The VaR at confidence level α is computed as:

$$VaR_\alpha = z_\alpha \times \sigma_{portfolio} \quad (5)$$

where $z_\alpha$ is the $\alpha - \text{quantile}$ of the standard normal distribution.

**Primary applications:**

- Linear books (forwards, swaps)
- Hedge monitoring
- Intraday risk updates

# 5. MONTE CARLO VAR: CORRELATED SHOCKS + FULL REVALUATION.

Molecule's Monte Carlo VaR accounts for nonlinearities and optionality. Each Monte Carlo scenario samples a vector of correlated log-returns over the VaR horizon, employing Cholesky decomposition of the covariance matrix.

**Monte Carlo VaR Algorithm**

1. Compute covariance matrix $\Sigma$ from historical log returns
2. Perform Cholesky decomposition: $\Sigma = LL^T$
3. For each simulation $i = 1, \ldots, N$:
   - Draw independent standard normals: $z \sim N(0, I)$
   - Generate correlated returns: $r = Lz$
   - Shock prices: $P' = P \times exp(r)$
   - Revalue portfolio at shocked prices: $V'_i$
   - Calculate P&L: $\Delta V_i = V'_i - V_0$
4. Sort simulated P&Ls: $\Delta V_{(1)} \leq \ldots \leq \Delta V_{(N)}$
5. Extract VaR at confidence level $\alpha$

Unless otherwise specified, Monte Carlo shocks are generated from daily returns and produce a one-day VaR; longer horizons apply standard variance scaling.

# 6. AS-GENERATED TRADES: INTEGRATION OF PHYSICAL ASSET RISK.

Molecule automatically incorporates as-generated trades, representing physical production exposures. These may be:

- Long exposures (generation)
- Short exposures (certain operational or contractual cases)

**Examples:**

- Power plant generation positions
- Renewable asset production profiles
- Gas supply/delivery flows

VaR evaluates these exposures alongside financial trades, yielding a holistic view of enterprise risk.

**Key modeling choice:**

VaR for as-generated trades captures market price risk, not volume uncertainty. Physical volumes are used to weight monthly market prices, producing an equivalent price exposure that is evaluated alongside financial trades.

# 7. VAR OUTPUTS AND DIVERSIFICATION ANALYTICS.

Molecule VaR produces multiple levels of aggregation:

- **Book-level VaR**

  Shows trading desk or strategy-level exposure.

- **Counterparty-level VaR**

  Reveals concentrated positions tied to specific counterparties.

- **Tenor-level VaR**

  Surfaces risk concentrated in specific future months/quarters.

**Diversification effects** are computed as:

Diversification Benefit = $\sum_i VaR_i - VaR_{portfolio}$ (6)

## 8. VAR BACKTESTING.


Molecule provides a built-in visual backtest, displaying:

- VaR forecasts
- Daily MTM of positions that existed at forecast time
- Breaches (exceptions)
- Model behavior over time

This facilitates:

- Model validation
- Regulatory reporting
- CRO oversight
- Audit defensibility

## 9. DIAGNOSTIC WORKBOOK.

Each VaR run produces an Excel workbook summarizing:

- VaR settings used
- Portfolio-level VaR
- Book, counterparty, and tenor breakdowns
- Diversification effects

The diagnostics do not include:

- Individual trades
- Simulated scenarios
- Raw historical returns

Excluding individual trades, simulated scenarios, and raw returns preserves clarity without requiring the management of large data sets from many intraday runs.