

# Q7 AEGIS · Precision Calibration System

Two financial-modeling pillars · ten supporting techniques · one closed loop

## TWO NAMED FINANCIAL-MODELING PILLARS

### PILLAR · 01

## Bayesian Optimization

Primary search engine — finds plateaus, not peaks

- Gaussian Process posterior over the parameter landscape
- Monte-Carlo qExpectedImprovement acquisition function
- 5–10× sample-efficient vs. brute-force search
- Posterior variance reported with every shipped parameter
- Warm-start from prior PASS cycles (~26% net lift in pilot)
- Backend: BoTorch (GPU) → sklearn (fallback)

### PILLAR · 02

## Monte Carlo Simulation

Primary validation engine — does the edge survive resampling?

- Level A: path-level bootstrap of CPCV OOS distributions (always-on)
- Level B: trade-order, forward-path, sign-permutation null tests
- Tens of thousands of resampled paths per shipped configuration
- Block-bootstrap preserves serial correlation (Politis-White)
- Four MC criteria gate every SHIP decision (see below)
- Lower 95% CI bounds ship — not point estimates

## SUPPORTING TECHNIQUES

01

### CPCV

Combinatorial Purged  
K-Fold Cross-Validation,  
15 OOS paths/config (de Prado)

02

### Deflated Sharpe

DSR + PSR — adjusts for  
multiple-testing inflation;  
10000 trials → honest Sharpe

03

### Block-Bootstrap CIs

Lower 95% CI bounds on  
Sharpe / PF / MDD / trades;  
preserves serial correlation

04

### PBT Jitter

Population-Based Training:  
evolves perturbed candidates,  
selects for plateau location

05

### DoubleML

Causal feature attribution.  
Separates causal effects from  
spurious correlations

06

### Hawkes Slippage

Self-exciting point process for  
tick arrivals — bursty slippage  
matches live execution

07

### 3-Executor Parity

Pine ↔ Python ↔ Live agreement  
≥ 99.5% on entry/direction/  
exit/fill — no daylight

08

### Conformal Prediction

Distribution-free per-signal  
P(win) intervals; lower bound  
must exclude 0.50

09

### Determinism Module

Pinned RNG seeds across numpy /  
random / torch — bit-for-bit  
reproducibility across machines

10

### Online Conformal

Live-shadow drift detection with  
mathematically-valid intervals;  
halts on conformal breach

## PRECISION GATE — SHIP / WHITELIST / REJECT

### Statistical robustness

CPCV path-level lower 95% CI Sharpe  
≥ 0.5  
Conformal lower bound on per-signal P(win)  
excludes 0.50  
Jitter PnL CV across ±5% perturbation  
≤ 15%

### Monte Carlo distribution

MC drawdown upper 95% CI  
≤ 35%  
MC P(no edge)  
≤ 20%  
MC P(structural-fail tail)  
≤ 15%

### Execution reality

MC sign-permutation p-value on net P&L  
≤ 5%  
Three-executor parity (Pine ↔ Py ↔ Live)  
≥ 99.5%  
Stress-DSR / clean-DSR  
≥ 60%

All criteria clear → SHIP    Any soft-fail → WHITELIST    Structural fail → REJECT