PROGRAM PROSPECTUS & CURRICULUM

Enterprise Trading & Risk Masterclass

A 500-chapter, desk-readiness program in cross-asset trading and risk management — taught end-to-end on an integrated platform, from first principles to signed-off reserve.

Scope	500 chapters · 12 modules · 1,500+ hours · 200 labs · 100 case studies · 50 projects
Audience	Investment banks · energy & commodity trading · merchants · hedge funds · risk consultancies · quants · TRM/ETRM implementation teams
Format	Every chapter: platform login → end-to-end screen flow → theory → code → incidents → case study → interview prep → capstone
Hands-on stack	SQL · ORE/QuantLib XML · FpML · REST · feeder & report definitions
Positioning	Enterprise masterclass, not an online course. Team & individual tracks.

Disclaimer

All screens shown in this program are illustrative mock-ups created for training purposes only. They do not depict any specific commercial product and are not a manual for any live system. All counterparties, trades, prices and results are fictional and internally consistent for teaching.

About this program

The Enterprise Trading & Risk Masterclass is a complete, desk-readiness curriculum in cross-asset trading and risk management. It assumes no prior background and builds every concept from first principles to the point where a learner can sit on a desk, in a middle office, or on an implementation team and be useful on day one. The program is taught on an integrated trading & risk platform used as the single reference system throughout, so learners see every concept where it actually lives in a product — not as abstract theory.

Coverage spans the full institutional value chain: deal capture and trade lifecycle, static and market data, pricing and valuation, market and counterparty risk, XVA, collateral and margin, settlement and post-trade, accounting, regulatory reporting, integration, administration, and advanced projects. Instruments span every asset class — rates, FX, credit, equity, commodity, power, inflation and digital assets — mirroring the breadth of a tier-one platform.

How every chapter is built

Each of the 500 chapters follows the same rigorous structure, designed to produce ability rather than recall:

- **Platform entry:** every chapter opens the same way a working day does — sign in, choose the workspace, and preview the end-to-end screen flow for the task.
- **Screen-by-screen walkthrough:** the actual click-path through the platform for whatever the chapter covers, with mock screens at each step.
- **Theory from first principles:** the concepts, formulae and market context, assuming no background.
- **Working code artifacts:** SQL, ORE/QuantLib XML, FpML, REST and feeder/report definitions — real, inspectable technology.
- **Day on the Desk incidents:** real trading-incident investigations drawn from the chapter's own failure modes.
- **End-to-end case study:** one of eight recurring golden trades followed across every function.
- **Regulator & model-validation perspective:** what audit, MRM and the regulator challenge.
- **Interview preparation:** role-specific question banks with strong answers and the red flags that expose weak ones.
- **Implementation capstone:** a build-and-defend project, graded on artifacts, not recall.

Commercial positioning

This is an enterprise masterclass, not a consumer online course. It is delivered in two tracks off the same curriculum: an **Enterprise / Team** track (cohort delivery, capstone review, golden-trade case studies as team exercises) for institutions onboarding trading, risk and implementation staff; and an **Individual Professional** track (self-paced with certification) for quants, risk analysts and consultants building credentials.

The program serves a broad enterprise audience: investment banks, energy trading companies, commodity merchants, hedge funds, utilities, exchanges, clearing houses and market-infrastructure providers, independent software vendors (ISVs) building on the domain, risk and management consultancies, and graduate / analyst training programs — alongside individual quants, risk analysts and TRM/ETRM implementation teams.

The program at a glance

The scale of the program, in numbers:

500	chapters across 12 modules
1,500+	hours of structured learning
200	hands-on labs
100	worked case studies
50	implementation projects
8	recurring golden trades threaded across the program

Role pathways

The same 500 chapters are navigated differently by different roles. Each pathway is a recommended route through the modules, so a buyer can map the curriculum directly to a job on the desk or the implementation team.

Trader	Market fundamentals, trade capture, pricing, market risk, and the front-office view of XVA and limits.
Quant Developer	Pricing & valuation, the Monte Carlo engine, models (Hull-White, LMM, Heston, SABR), Greeks and AAD.
Market Risk Analyst	Risk & analytics, VaR, stress testing, PFE, backtesting and risk reporting.
Credit Risk Analyst	Counterparty credit, exposure, CVA/DVA, netting, collateral and wrong-way risk.
ETRM Consultant	End-to-end lifecycle, physical commodity operations, configuration and implementation projects.
Data Engineer	Static & market data, curves, pipelines, integration, streaming and the data model.
Platform Administrator	Administration & operations, environments, scaling, monitoring, DR and security.
Solution Architect	Architecture, integration patterns, deployment, and the full end-to-end implementation.

Skill progression and certification

Progress is measured, not vague. Each 100-chapter stage corresponds to a capability level and a certification credential, giving learners and employers concrete milestones.

Stage	Capability	Certification level
100 chapters	Junior Analyst	Foundation
200 chapters	Business Analyst	Associate
300 chapters	Consultant	Professional
400 chapters	Senior Consultant	Expert
500 chapters	Solution Architect	Enterprise Architect

The five credentials — Foundation, Associate, Professional, Expert and Enterprise Architect — are progressive: each builds on the last, and the Enterprise Architect credential requires the full implementation capstone. Institutions can certify cohorts to a defined level as part of a structured onboarding program.

What you will build

This is a build-and-defend program: the capstones and implementation projects produce working artifacts, not just answers. Across the 500 chapters a learner constructs, piece by piece, the components of a complete trading-and-risk platform:

Pricing engine	A working valuation service on ORE/QuantLib, with curves, vol surfaces and Greeks.
Risk engine	Exposure simulation, VaR, PFE and the XVA exposure cube.
Trade capture APIs	Canonical trade representation, validation, lifecycle and FpML at the boundary.
Market data pipeline	Curve construction, vol surfaces, snapshots and reconciliation.
Settlement workflow	STP settlement, exception handling and reconciliation.
Collateral simulator	CSA logic, margin calls, MPOR and residual-exposure modelling.
Reporting dashboards	Risk and XVA datamart, SQL views and reconciliation reports.
Complete TRM implementation	The capstone: an end-to-end trade flow integrated across every module.

Reference deployment architecture

The program teaches implementation on a modern, cloud-native reference architecture, so learners meet the platform-engineering stack an enterprise deployment actually runs on — not just the financial logic. The reference stack is containerised, event-driven and observable end to end.

Figure 1 — The reference deployment: a React front end and versioned APIs over Kubernetes-orchestrated services (trade, pricing, risk/XVA, settlement, collateral); Spark for distributed exposure and VaR batch; Airflow for orchestration; Kafka as the event backbone; PostgreSQL for OLTP and ClickHouse for the analytical risk datamart; and full observability via Prometheus, Grafana and centralised logging, all delivered through CI/CD.

Full curriculum — 500 chapters

The complete chapter list across all 12 modules. The exemplar chapter that follows this curriculum (Chapter 193, Valuation Adjustments / XVA) demonstrates the full format applied to one of the most complex topics in the program.

MODULE 1 · CHAPTERS 1–20

Financial Market Fundamentals

1. What Are Capital Markets?
2. Primary vs Secondary Markets
3. Role of Banks, Brokers & Exchanges
4. Understanding Financial Instruments
5. Money Market vs Capital Market
6. Basics of Bonds, Equities & FX
7. What Are Derivatives?
8. Swaps, Forwards & Options Simplified
9. Margin & Collateral Basics
10. Mark-to-Market Explained
11. Introduction to Risk Management
12. Market Data & Fixings
13. Trading Lifecycle Overview
14. Role of Front, Middle & Back Office
15. Clearing and Settlement
16. Understanding Yield Curves
17. Discounting and Compounding
18. Regulatory Overview (MiFID, EMIR)
19. Role of Technology in Trading
20. Why a Platform Exists – Solving Complexity

MODULE 2 · CHAPTERS 21–45

Introduction to the Trading & Risk Platform

21. What Is an Integrated TRM Platform?
22. History and Evolution of TRM Systems
23. Global Clients & Use Cases
24. Platform Modules Overview
25. Core Architecture Explained
26. Data Model Overview
27. Environment Setup and Installation
28. Navigating the Interface
29. Front-End Workspaces & Views
30. Core Terminologies (Trade, Deal, Workflow)
31. Introduction to Static Data
32. Market Data Handling
33. Trade Capture Basics
34. Pricing Concepts
35. Simulations and Pricing Workflows
36. Workflow and Lifecycle Events
37. Reporting Overview
38. Security and Permissions
39. Role-Based Access Control
40. Performance Architecture and Caching
41. Integration Ecosystem
42. Common User Roles and Profiles
43. System Health and Monitoring Overview
44. Common Errors and Troubleshooting
45. Section Review and Mini Case Study

MODULE 3 · CHAPTERS 46–91

Static Data & Reference Data

46. Static Data Concepts
47. Managing Instruments
48. Counterparty Master Data
49. Legal Entity Management
50. Rate Curves & Curve Definitions
51. Calendars & Business Day Conventions
52. Market Data Sources
53. Reference Data Quality
54. Static Data Governance
55. Versioning Static Data
56. Securities & Identifiers (ISIN)
57. Currency & FX Definitions
58. Holiday Calendars
59. Tenor Conventions
60. Loading Static Files (CSV/XML)
69. Static Data Reconciliation
70. Static Data Performance Tuning
71. FX Rate Sources
72. Credit Entities & Ratings
73. Index & Reference Data
74. Historical Static Data Management
75. Static Data in Workflows
76. Importing FpML for Static Data
77. XML/FpML Static Data Patterns
78. Cross-Reference Tables
79. Business Metadata & Comments
80. Data Archival Strategies
81. Using Lookup Tables
82. Multi-tenant Static Data Concerns
83. Static Data Testing

61. Static Data APIs
62. Data Mapping & Transformations
63. Reference Data for Derivatives
64. Configuration Reuse Patterns
65. Static Data Security
66. Managing Complex Instruments
67. Templates & Default Values
68. Static Data Backups
84. Static Data Change Management
85. Data Lineage for Static Data
86. Static Data Automation
87. Common Pitfalls & Fixes
88. Static Data Case Study
89. Workshop: Build Reference Data
90. Quiz: Static Data
91. Section Review

MODULE 4 · CHAPTERS 92–137

Market Data & Curves

92. Market Data Overview
93. Market Data Feed Architecture
94. Curve Construction Basics
95. Bootstrapping Yield Curves
96. Curve Instruments & Inputs
97. OIS vs LIBOR Considerations
98. Multi-curve Setup
99. Curve Interpolation Methods
100. FX Spot & Forward Curves
101. Volatility Surfaces
102. Surface Building Methods
103. Implied Volatility Handling
104. Market Data Storage Patterns
105. Tick Data & Time Series
106. Market Data Caching
107. Market Data Reloads
108. Market Data Reconciliation
109. Market Data Transformation Rules
110. Historical Market Data
111. Reference Market Data Providers
112. Handling Missing Data
113. Market Data Validation
114. Curve Shift & Scenarios
115. Stress Testing Market Data
116. Calibration Engines
117. Sensitivity to Curve Inputs
118. Market Data Governance
119. Market Data Automation
120. Integrating Market Data APIs
121. Real-time vs Snapshot Feeds
122. Latency Considerations
123. Using External Data Providers
124. Synthetic Instrument Creation
125. Market Data in Pricing Engines
126. Surface Extrapolation Techniques
127. Market Data Troubleshooting
128. Workshop: Build Curve
129. Workshop: Vol Surface
130. Quiz: Market Data
131. Case Study: Calibration
132. Deployment Patterns for Market Data
133. Monitoring Market Data Pipelines
134. Market Data Security
135. Case Study: Reconciliation
136. Tips & Best Practices
137. Section Review

MODULE 5 · CHAPTERS 138–183

Trade Capture & Booking

138. Trade Capture Basics
139. Deal vs Trade Concepts
140. Trade Types
141. Front Office Entry Screens
142. Trade Validation Rules
143. Trade Enrichment
144. Capturing Complex Trades
145. Trade Templates
146. Trade Matching
147. Trade Amendments & Cancellations
148. Trade Workflows & States
149. Booking Best Practices
150. Trade Reporting
151. Audit Trails
152. Trade Serialisation (FpML)
153. Importing Trades via Files
161. Trade Reconciliation
162. Trade Retention & Archival
163. Trade Data Lineage
164. Trade Load Performance
165. Handling Corporate Actions
166. Corporate Actions Processing
167. Workbench and Dashboards
168. Bulk Trade Operations
169. Trade Lifecycle Case Study
170. Workshop: Capture & Amend
171. Testing Trade Flows
172. Trade Security Considerations
173. Trade Authorisation & Roles
174. Trade Migration Patterns
175. Trade Error Diagnostics
176. Automation for High Volume

154. Trade Lifecycle Events
155. Booking via API
156. Trade Enrichment Plugins
157. Matching and Confirmation
158. Trade Settlements
159. Trade Billing & Fees
160. Exception Handling
177. Trade Compression & Netting
178. Trade Enrichment Best Practices
179. Case Study: Trade Lifecycle
180. Quiz: Trade Capture
181. Review Exercises
182. Practical Lab
183. Section Review

MODULE 6 · CHAPTERS 184–229

Pricing & Valuation

184. Pricing Engine Overview
185. Pricing Model Types
186. Discounting Models
187. Volatility Models
188. Pricing Parameters
189. Using Market Data in Pricing
190. Calibration Processes
191. Greeks and Sensitivities
192. Mark-to-Market vs Mark-to-Model
193. Valuation Adjustments (XVA)
194. Pricing Overrides
195. Pricing Schedules & Jobs
196. Batch Valuation
197. Real-time Pricing
198. Intraday Revaluations
199. Pricing Performance Tuning
200. Pricing Logs & Audit
201. Model Risk Considerations
202. Scenario Pricing
203. Stress Pricing
204. Pricing Workflows
205. Pricing API
206. Handling Exotic Instruments
207. Greeks Calculation
208. Pricing in Multi-currency
209. Calibration Validation
210. Pricing Case Study
211. Workshop: Configure Pricing
212. Workshop: Valuation Job
213. Instrument-Specific Pricing
214. Repo & Sec Lending Pricing
215. Commodity Pricing Considerations
216. Equity Derivative Pricing
217. Interest Rate Derivative Pricing
218. Credit Derivative Pricing
219. Pricing under Collateralization
220. Pricing Reconciliation
221. Pricing for Risk Reports
222. Real-world Pricing Examples
223. Quiz: Pricing & Valuation
224. Review Exercises
225. Practical Lab
226. Troubleshooting Pricing
227. Performance Improvements
228. Best Practices
229. Section Review

MODULE 7 · CHAPTERS 230–275

Risk & Analytics

230. Risk Engine Overview
231. Market Risk Basics
232. VaR Concepts
233. Historical vs Parametric VaR
234. Stress Testing
235. Scenario Analysis
236. P&L; Attribution
237. Risk Reports
238. Credit Risk Basics
239. Counterparty Credit Risk
240. CVA & DVA Overview
241. Collateral and Margining Impact
242. Exposure Calculations
243. Netting & Compression Effects
244. Position Limits
245. Regulatory Risk Metrics
246. Risk Data Feeds
253. Collateral Optimization
254. Sensitivity Analysis
255. Greeks in Risk Context
256. Intraday Risk Monitoring
257. Risk Model Governance
258. Regulatory Reporting
259. Model Validation
260. Risk Aggregation
261. Credit Exposure Reconciliation
262. Case Study: VaR Implementation
263. Workshop: Build Risk Report
264. Troubleshooting Risk Calculations
265. Performance Tuning for Risk
266. Best Practices
267. Quiz: Risk & Analytics
268. Review Exercises
269. Practical Lab

247. Monte Carlo Simulations
248. Scenario Generation
249. Backtesting Risk Models
250. Risk Dashboards
251. Risk Alerts & Thresholds
252. Margin Requirements
270. Integration with BI Tools
271. Risk Automation Patterns
272. Real-world Examples
273. Advanced Topics
274. Summary & Checklist
275. Section Review

MODULE 8 · CHAPTERS 276–320

Post-Trade & Settlements

276. Post-Trade Overview
277. Settlement Types
278. Payment Instructions
279. Custody & Settlement Systems
280. Clearing House Interfaces
281. Settlement Failures & Repairs
282. Reconciliation Processes
283. Corporate Actions Handling
284. Payment Netting
285. FX Settlement Lifecycle
286. STP (Straight Through Processing)
287. Settlement Notifications
288. Nostro/Vostro Reconciliation
289. Account Mapping
290. Batch Settlement Jobs
291. DVP / RVP Concepts
292. Trade Settlement Reporting
293. Exception Workflow for Settlements
294. Settlement SLA Monitoring
295. Operational Controls
296. Reconciliation Tools Integration
297. Data Matching Rules
298. Settlement Case Study
299. Workshop: Settlement Flow
300. Testing Settlements
301. Settlement Security
302. Settlement Performance
303. Automation for Settlement
304. Settlement Exceptions Demo
305. Reporting & Dashboards
306. Auditing Settlements
307. Cross-Border Settlements
308. Handling Corporate Actions in Settlement
309. Settlement Reconciliation Best Practices
310. Error Remediation Patterns
311. Regulatory Requirements
312. Third-party Integrations
313. Case Study: Settlement Failure
314. Quiz: Settlements
315. Review Exercises
316. Practical Lab
317. End-to-end Settlement Example
318. Troubleshooting
319. Section Summary
320. Section Review

MODULE 9 · CHAPTERS 321–365

Collateral & Margin

321. Collateral Basics
322. Margining Concepts
323. CSA Agreements
324. Eligible Collateral Rules
325. Collateral Optimization
326. Haircuts & Valuation
327. Margin Calls Lifecycle
328. Substitution of Collateral
329. Rehypothecation Issues
330. Collateral Settlement
331. Collateral Reporting
332. Intraday Margining
333. Collateral Dispute Handling
334. Pledge vs Transfer
335. Collateral Analytics
336. Case Study: Margin Call
337. Workshop: Collateral Workflow
338. Collateral Configuration
344. Regulatory Collateral Requirements
345. SSAs & Tri-party Solutions
346. Backtesting Collateral Rules
347. Automation Examples
348. Collateral Dispute Resolution
349. Collateral Performance Tuning
350. Quiz: Collateral
351. Review Exercises
352. Practical Lab
353. Real-world Collateral Examples
354. Troubleshooting Collateral Issues
355. Best Practices
356. Summary & Checklist
357. Section Summary
358. Additional Reading
359. Case Studies
360. Section Review
361. Exercises

339. Collateral Jobs & Scheduling
340. Exposure Calculation for Collateral
341. Integration with Custodians
342. Collateral Reconciliation
343. Collateral Optimization Strategies
362. Mini Project
363. Deep Dive Topics
364. Final Quiz
365. Section Review

MODULE 10 · CHAPTERS 366–410

Integrations & Interfaces

366. Integration Patterns Overview
367. XML Messaging Basics
368. FpML & Trade Serialization
369. REST & SOAP APIs
370. Messaging Middleware (Kafka)
371. ETL for Trade Data
372. File-based Integrations
373. Real-time Streaming
374. Third-party Connectivity
375. FIX & Market Connectivity
376. Trade Confirmations
377. Trade Lifecycle Events Integration
378. Reference Data Synchronization
379. Market Data Integration
380. Settlement System Interfaces
381. Accounting & GL Feeds
382. Reporting Exports
383. Security & Encryption
384. Authentication & Authorization
385. API Rate Limiting
386. Integration Testing
387. Monitoring Integrations
388. Error Handling Strategies
389. Idempotency & Retry Patterns
390. Contract Testing
391. Performance Tuning for Integrations
392. Data Mapping Examples
393. Workshop: Build an API Integration
394. Case Study: Real-time Feed
395. Troubleshooting Integration Issues
396. Security Best Practices
397. Governance for Integrations
398. Automation & CI for Integration Code
399. Deployment Patterns
400. Integration Checklist
401. Quiz: Integrations
402. Review Exercises
403. Practical Lab
404. Advanced Integration Topics
405. Third-party Connectors
406. Message Replay & Recovery
407. Data Contracts
408. Monitoring & Alerts
409. Section Summary
410. Section Review

MODULE 11 · CHAPTERS 411–455

Administration & Operations

411. System Architecture Overview
412. Environments & Lifecycle
413. Installation Steps
414. Backup & Restore
415. High Availability
416. Scaling the Platform
417. Patching & Upgrades
418. User Management
419. Permissions & Roles
420. Job Scheduling & Orchestration
421. Monitoring & Health Checks
422. Logging Strategies
423. Alerting & Runbooks
424. Performance Tuning
425. Capacity Planning
426. Security Operations
427. Disaster Recovery Planning
428. Environment Promotion Patterns
429. Compliance & Auditing
434. Debugging System Issues
435. Common Support Scenarios
436. Incident Management
437. Change Management
438. Runbook Examples
439. Capacity & Load Testing
440. Vendor Support Models
441. Licensing & Entitlements
442. Integrations Health Monitoring
443. Maintenance Windows
444. Automation: CI/CD pipelines
445. Environment Security Hardening
446. Hands-on: Admin Tasks
447. Troubleshooting: Common Failures
448. Operational KPIs
449. Case Study: Outage Recovery
450. Quiz: Administration
451. Review Exercises
452. Practical Lab

430. Operational Playbooks
431. Automation for Admin Tasks
432. Data Archival Processes
433. Housekeeping Jobs
453. Best Practices Summary
454. Section Summary
455. Section Review

MODULE 12 · CHAPTERS 456–500

Advanced Topics, Projects

456. Advanced Configuration Patterns
457. Complex Instrument Handling
458. Model Calibration Challenges
459. Advanced Risk Analytics
460. Performance Optimization Patterns
461. Cross-Module Workflows
462. Automation & Scripting
463. API Extensions
464. Plugin Development
465. Advanced Integration Patterns
466. Custom Reporting Engines
467. Data Warehousing
468. BI Integration & Dashboards
469. Machine Learning Use-cases
470. Advanced Monitoring & APM
471. SecOps & Hardening
472. Multi-region Deployments
473. GDPR & Data Privacy
474. Governance & Compliance Deep Dive
475. Project: End-to-End Implementation
476. Project: Migration Case Study
477. Project: Custom Integration
478. Capstone: Build a Trade Flow
479. Capstone: Risk Reporting
480. Certification Preparation
481. Interview Prep for TRM Roles
482. Resume & CV Tips
483. Teaching & Mentoring Patterns
484. Running Internal Training
485. Consulting Best Practices
486. Pricing & Commercialization
487. Scaling Training For Teams
488. Continuous Learning Paths
489. Case Studies & Lessons Learned
490. Final Exam Practice
491. Final Project Submission
492. Grading & Feedback
493. Certificate Issuance
494. Post-course Support
495. Community & Forum Access
496. Continuing Education
497. Advanced Reading List
498. Final Capstone Review
499. Graduation Checklist
500. Program Wrap-up & Next Steps

EXEMPLAR CHAPTER

Chapter 193 — Valuation Adjustments (XVA)

The full chapter format, applied end-to-end to one of the program's most complex topics. What follows is a complete, production-grade chapter — the standard every chapter in the program is built to.

193.0 Getting into the platform: the end-to-end flow

Every chapter in this program begins the same way you begin a real working day: you sign in to the platform, choose your workspace, and walk the actual screen flow for the task at hand. You will never be handed a concept in the abstract without seeing where it lives in the product. This section is your map for the chapter — sign in, pick the desk, and preview the twelve stops of the XVA workflow you are about to walk end to end.

Step 1 — Sign in

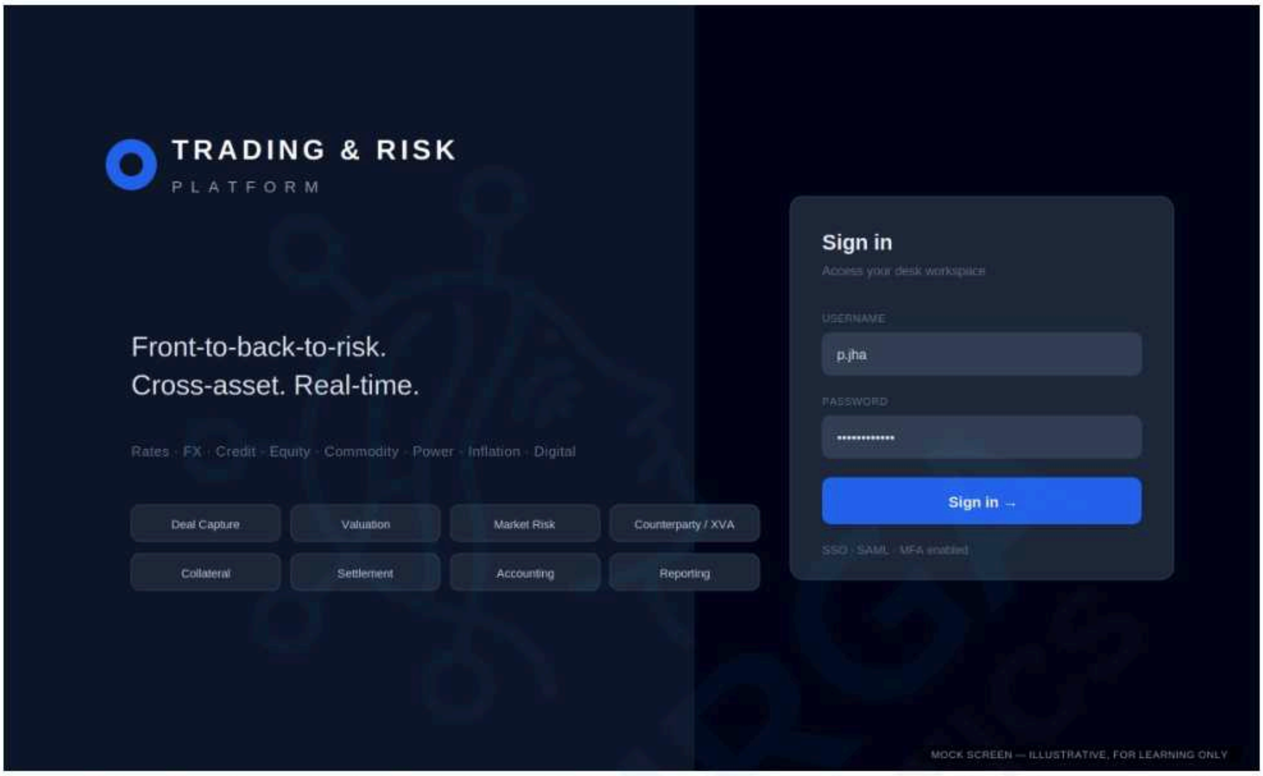

Figure 193.0a — The platform sign-in. Single sign-on with SAML and MFA; access is scoped to the desks your role entitles you to. This is the same door every chapter opens.

Step 2 — Choose your workspace

The platform is multi-tenant. After sign-in you select the tenant and desk you are working in — here, Tenant 2, the oil desk — which scopes every screen, trade and number that follows to that book. Pick the wrong workspace and you will be looking at the wrong portfolio; this is the first thing to verify when a number surprises you.

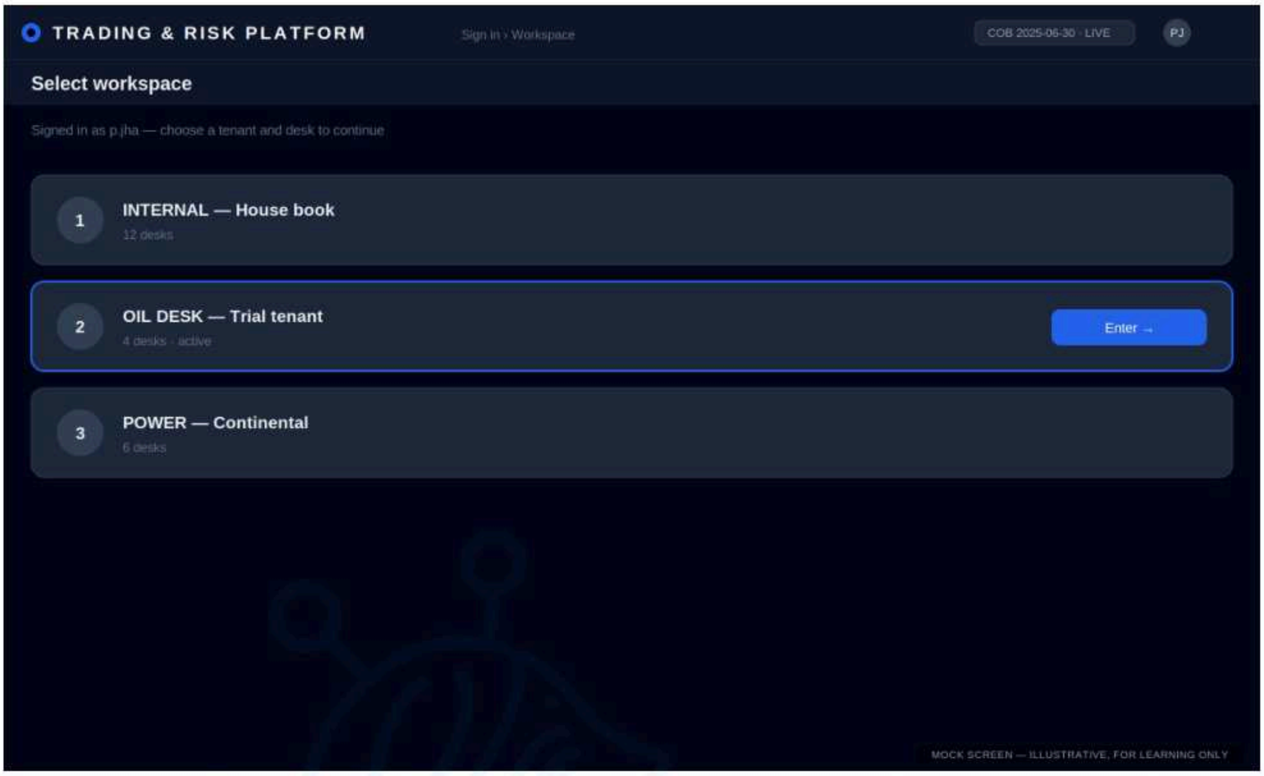

Figure 193.0b — Workspace selection. Tenant and desk scope everything downstream. For this chapter we enter Tenant 2 (oil desk), the book that holds our golden-trade Brent swap.

Step 3 — Walk the flow

Once inside, the XVA task is a defined click-path — from the trade blotter that defines scope, through netting, curves and simulation, to the cockpit, the pre-deal tool, the nightly batch and finally reserve sign-off. Here is the whole journey on one page. Every stop is a live workspace we open and read in the sections that follow; keep this map in mind so you always know where you are.

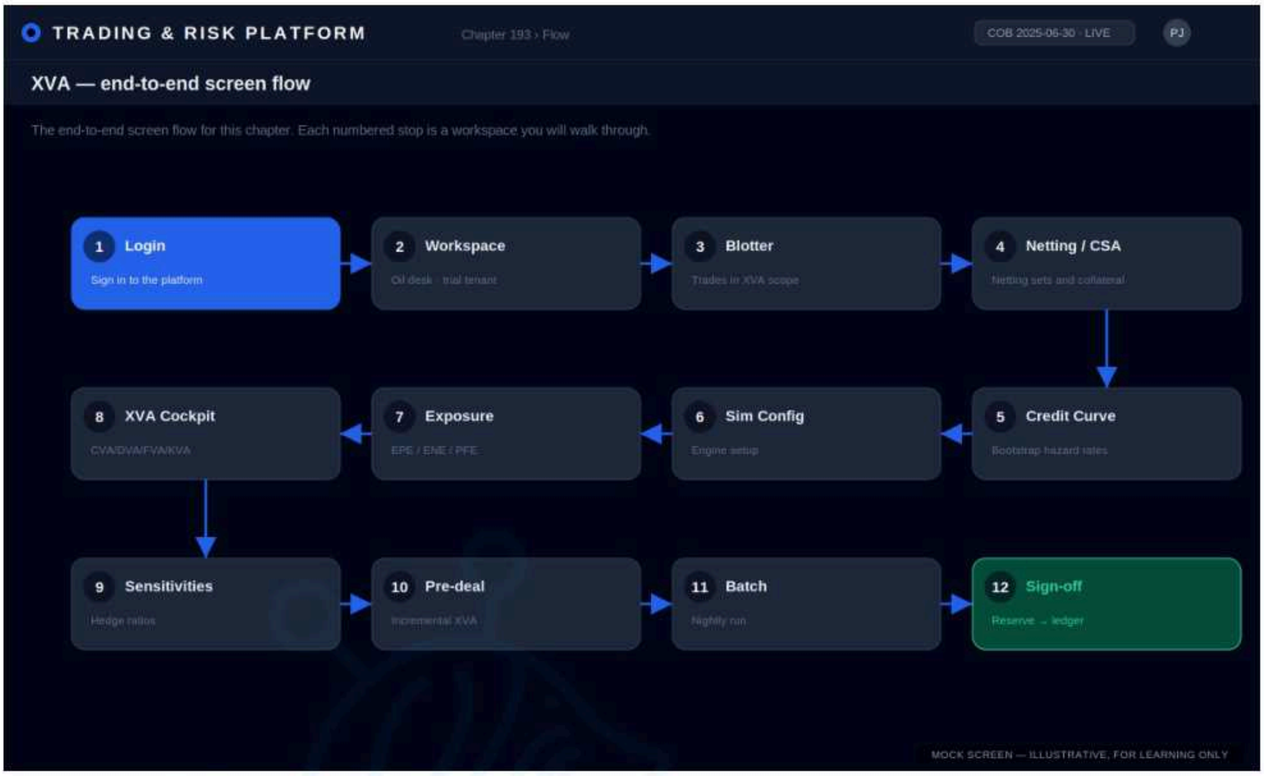

Figure 193.0c — The end-to-end XVA screen flow. Twelve stops from sign-in to signed-off reserve. Blue is where you start, teal is where you finish; each numbered node is a screen you will walk through in this chapter. This is the click-path a real analyst follows to produce and defend a CVA number.

◆ How to read every chapter in this program

Login → workspace → flow. Each chapter opens with this same three-step entry so the platform becomes second nature. By the time you finish the program you will have walked this door hundreds of times and every module's workflow will be muscle memory — which is exactly what 'desk readiness' means.

193.1 Why this chapter exists: the price is not the price

Imagine two traders on two different desks book the exact same five-year commodity swap on the same underlying, on the same day, at the same forward strike. Their pricing screens show an identical mark-to-market of, say, +USD 412,000. Yet the true economic value of those two trades to the bank can differ by tens of thousands of dollars. One is with a strongly rated counterparty under a tight collateral agreement; the other is with a weak, uncollateralised name in a region with a poor legal opinion on netting. The *textbook* price treats them as equal. Reality does not.

The gap between the idealised, risk-free valuation of a derivative and its real value once you account for the possibility that your counterparty defaults, that you must fund the position, that you post and receive collateral, and that regulators make you hold capital against the exposure — that gap is the family of **valuation adjustments**, collectively **XVA**. The 'X' is a wildcard: CVA, DVA, FVA, ColVA, MVA, KVA. Each letter is a correction to the clean price. Together they turn a model number into a number the bank can actually stand behind, reserve against, and post to its ledger.

Before 2008, most of these adjustments were either ignored or buried in a crude 'credit reserve' held centrally and rarely attributed to individual trades. The financial crisis changed that permanently. Counterparties everyone assumed were money-good — large dealers, monolines, sovereigns — either failed or came within hours of it. The market learned, expensively, that counterparty credit is not a footnote; it is a first-order component of derivative value. Accounting standards followed (fair value must reflect counterparty and own credit), then funding costs were recognised, then the regulatory capital regime made the cost of capital itself a pricing input. XVA moved from the risk department's spreadsheet to the front-office pricing screen.

◆ **Desk-readiness framing**

On a real desk, XVA is not an academic curiosity — it is a P&L line, a reserve on the balance sheet, and a charge the XVA desk levies on every new trade. If you cannot explain why the same trade costs more with counterparty A than counterparty B, you are not ready to sit on the desk. This chapter's job is to make you ready.

The one-screen mental model

Everything in this chapter rolls up into a single view. Study it before we go deep — every later section is really just explaining one tile of this cockpit and the machinery behind it.

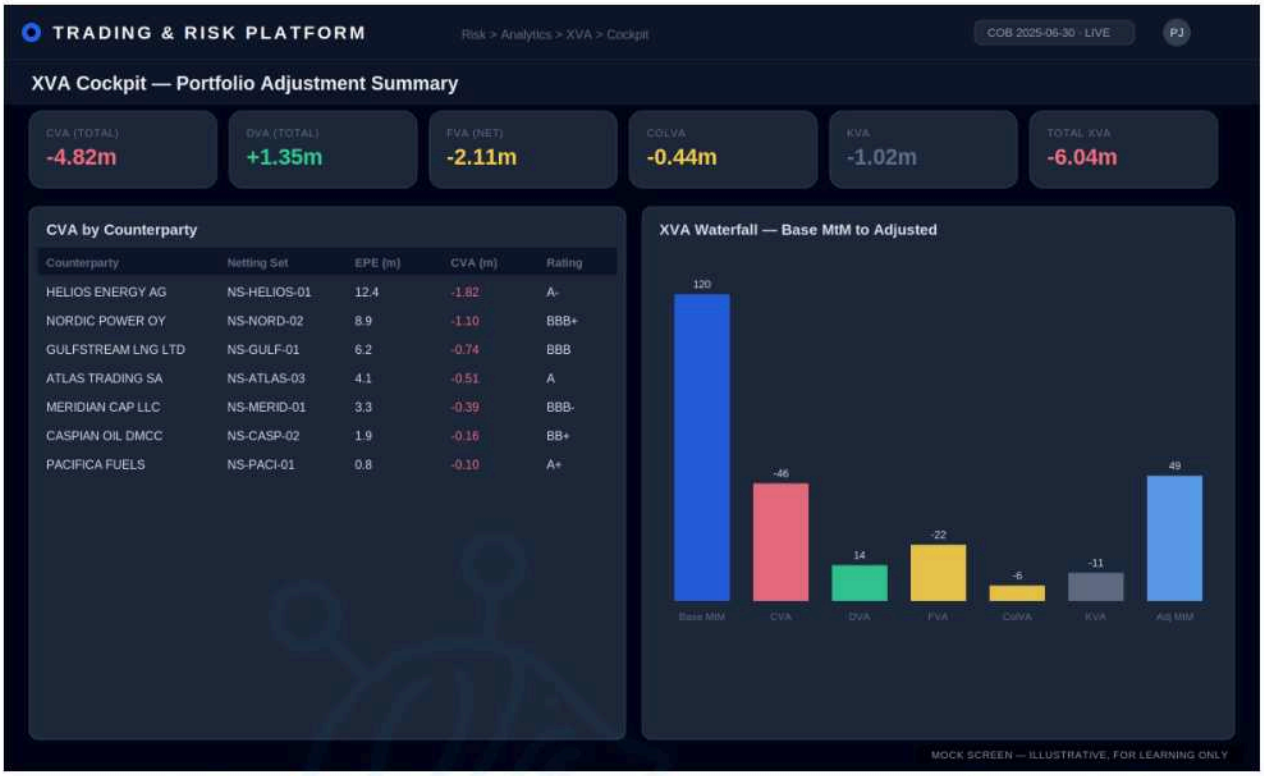

Figure 193.1 — The XVA cockpit. Clean mark-to-market enters from the left; each adjustment (CVA, DVA, FVA, ColVA, KVA) is a step in the waterfall that walks the base price down (or up) to the adjusted, reservable value. Every number on this screen is derived from the exposure simulation and curve calibration we build over the next forty pages.

Read the waterfall left to right. The blue 'Base MtM' bar is the risk-free price the pricing engine produces. CVA is a charge (we lose money if the counterparty defaults while they owe us), so it steps the value *down*. DVA is a benefit (our own default extinguishes what we owe), stepping *up*. FVA reflects the cost of funding the uncollateralised portion of the position. ColVA captures the economics of the collateral we actually post and receive. KVA is the lifetime cost of the regulatory capital the trade consumes. The rightmost bar is what the trade is genuinely worth to the institution.

What you will be able to do after this chapter

- Explain each XVA term in plain English and in formulae, and say who pays whom.
- Trace a single trade from capture, through exposure simulation, to a signed-off CVA reserve.
- Read and author the trade representations the platform uses — its canonical JSON and FpML — and drive the engine via REST.
- Query the risk datamart in SQL to reconcile a counterparty-level CVA number to the trade level.
- Configure an exposure-simulation run: models, paths, time grid, netting and collateral.
- Compute an incremental (pre-deal) XVA charge and turn it into a spread add-on for the trader.
- Stand up the nightly XVA batch, monitor it, and diagnose the three failures that break it most.

Related chapters — navigate the curriculum

- **Market Data & Curves (Ch 92–137):** the curves and vol surfaces this chapter's engine consumes.
- **Calibration Engines (Ch 116):** how the Hull-White, Gabillon and credit models are fitted to market.
- **Pricing & Valuation (Ch 184–209):** the per-path repricing that fills the exposure cube.
- **Risk & Analytics (Ch 230–275):** VaR, PFE and the counterparty-credit context around XVA.
- **Collateral & Margin (Ch 321–340):** the CSA simulation behind ColVA, MVA and residual exposure.
- **Post-Trade & Settlements (Ch 276–320):** where the golden-trade Brent swap settles.
- **Accounting & Finance (Ch 381):** the CVA reserve on the ledger under IFRS 13 / IFRS 9.

193.2 The XVA family, term by term

Let us define each adjustment precisely. For each one we give the plain-English meaning, the sign convention (does it help or hurt us), the driving inputs, and the one sentence you must be able to say out loud in an interview.

CVA — Credit Valuation Adjustment

CVA is the market value of the counterparty credit risk in your derivatives. It is the expected loss from the counterparty defaulting at some future time while the trade (or netting set) has positive value to you. Because it is a loss, CVA is a charge that reduces the value of your position. The three ingredients are: **how much they might owe you** in the future (expected positive exposure, EPE), **how likely they are to default** in each future period (the default probability implied from their credit spreads), and **how much you lose if they do** (loss given default = 1 – recovery).

Formula — discretised CVA (single netting set)

```
                 N
  CVA  =  LGD  · Σ   EPE(t_k) · DF(t_k) · [ Q(t_{k-1}) − Q(t_k) ]
                k=1

  where
    LGD        = loss given default = 1 − Recovery        (e.g. 1 − 0.40 = 0.60)
    EPE(t_k)   = discounted expected positive exposure at time t_k  (from simulation)
    DF(t_k)    = risk-free discount factor to t_k          (OIS curve)
    Q(t)       = counterparty survival probability to t    (from CDS bootstrap)
    Q(t_{k-1})−Q(t_k) = marginal default probability in the interval
```

◆ Say this in an interview

"CVA is the price of the option the counterparty holds to walk away from what they owe me. I pay for that option through a lower valuation. It rises when their credit spread widens or when my expected exposure to them grows."

DVA — Debit Valuation Adjustment

DVA is CVA seen from the mirror. It is the value of *your own* default to you: if you fail while you owe the counterparty money, you extinguish that liability, which — under fair-value accounting — is a benefit to your shareholders. DVA therefore increases the value of your position. It is symmetric to CVA but driven by your *negative* exposure (expected negative exposure, ENE) and *your* survival curve. DVA is controversial: booking a gain because your own credit deteriorated is counter-intuitive and creates a P&L that you can never actually monetise. Many desks neutralise part of it against FVA to avoid double counting.

FVA — Funding Valuation Adjustment

Derivatives are not free to hold. When a position is uncollateralised or partially collateralised, the bank must fund it — borrow cash to pay out, or hold cash it cannot reinvest at the risk-free rate. FVA is the present value of those funding costs and benefits over the life of the trade. It splits into **FCA** (funding cost adjustment, driven by expected positive exposure — the money we lend the market and must fund) and **FBA** (funding benefit adjustment, driven by expected negative exposure). Net FVA is usually a cost.

ColVA, MVA, KVA — the rest of the alphabet

- **ColVA (Collateral Valuation Adjustment):** the economics of the collateral you actually exchange under the CSA — the difference between the rate paid on posted collateral and the rate you could earn, plus the impact of collateral in the 'wrong' currency, thresholds and minimum transfer amounts.
- **MVA (Margin Valuation Adjustment):** the cost of funding *initial margin* — the extra buffer collateral posted to central counterparties or under bilateral margin rules, which is locked up and must be funded over the trade's life.
- **KVA (Capital Valuation Adjustment):** the lifetime cost of the regulatory capital the trade forces the bank to hold. Capital is not free — shareholders demand a return on it — so a trade that consumes a lot of capital for years carries a real, poundable cost.

Adjustment	Plain meaning	Sign to us	Primary driver	Key curve
CVA	Their default risk	Charge (−)	EPE × their PD	CDS / hazard
DVA	Our own default	Benefit (+)	ENE × our PD	Our CDS
FCA	Fund positive exposure	Charge (−)	EPE	Funding spread
FBA	Benefit on neg. exposure	Benefit (+)	ENE	Funding spread
ColVA	Collateral economics	Usually (−)	Coll. balance	CSA rate vs OIS
MVA	Fund initial margin	Charge (−)	IM profile	Funding spread
KVA	Cost of capital held	Charge (−)	Capital profile	Cost of equity

193.3 The unit of account: netting sets and the CSA

You do not compute CVA trade by trade. You compute it per **netting set** — the pool of trades with a single counterparty that can be legally set off against each other on default. If I have one trade worth +10 and another worth −6 with the same counterparty under an enforceable master agreement, my exposure on their default is +4, not +10. Netting is the single largest reducer of counterparty exposure in a real book, and getting the netting-set boundaries right is the difference between a defensible number and a regulatory finding.

Layered on top of netting is the **Credit Support Annex (CSA)** — the collateral agreement. A CSA says: when net exposure crosses a *threshold*, the out-of-the-money party posts collateral to bring exposure back down, subject to a *minimum transfer amount* and *rounding*. A perfectly collateralised, daily-margined, zero-threshold CSA drives counterparty exposure — and therefore CVA — close to zero (though never exactly, because of the margin period of risk: the gap between the last received margin and close-out). A high-threshold or uncollateralised relationship leaves exposure fully open.

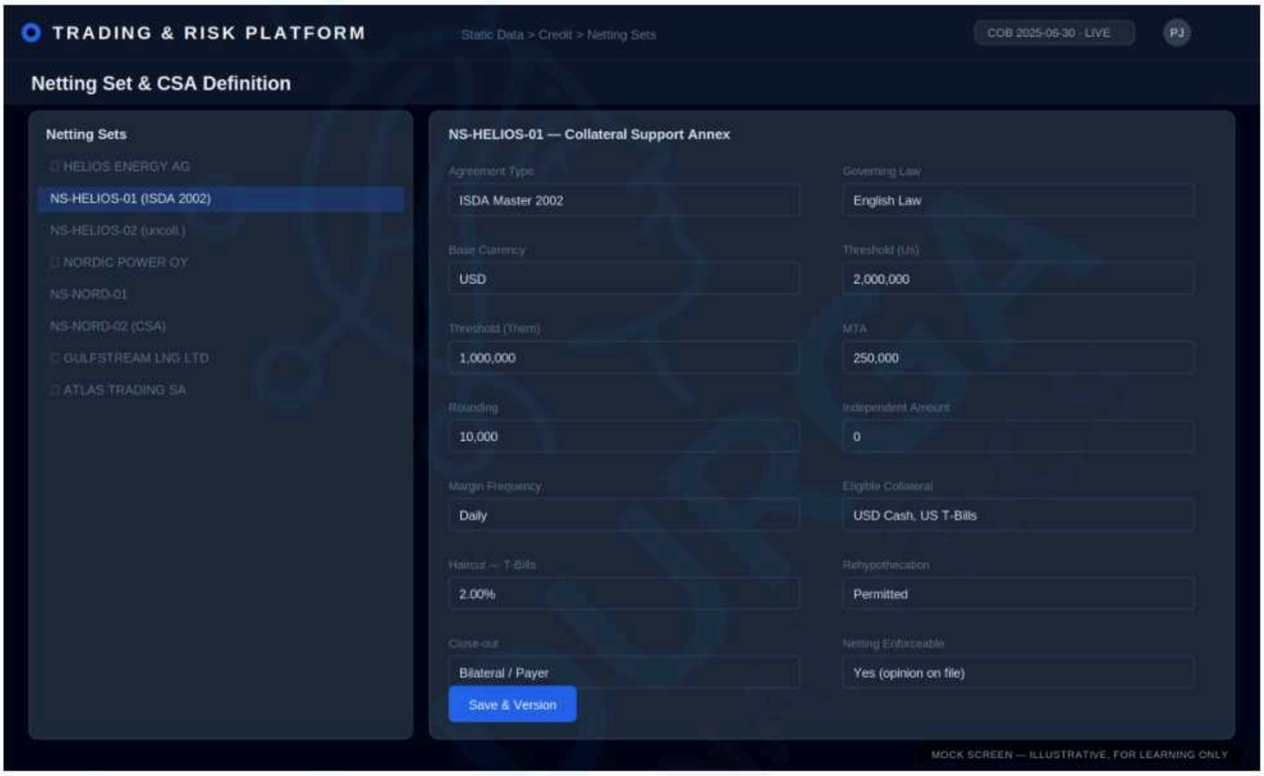

Figure 193.2 — Netting set and CSA definition. The netting hierarchy on the left groups trades under a counterparty and master agreement; the CSA panel on the right captures the collateral terms that the exposure engine reads at simulation time. Note the netting-enforceability flag — if the legal opinion does not support netting in the counterparty's jurisdiction, the engine must fall back to gross exposure.

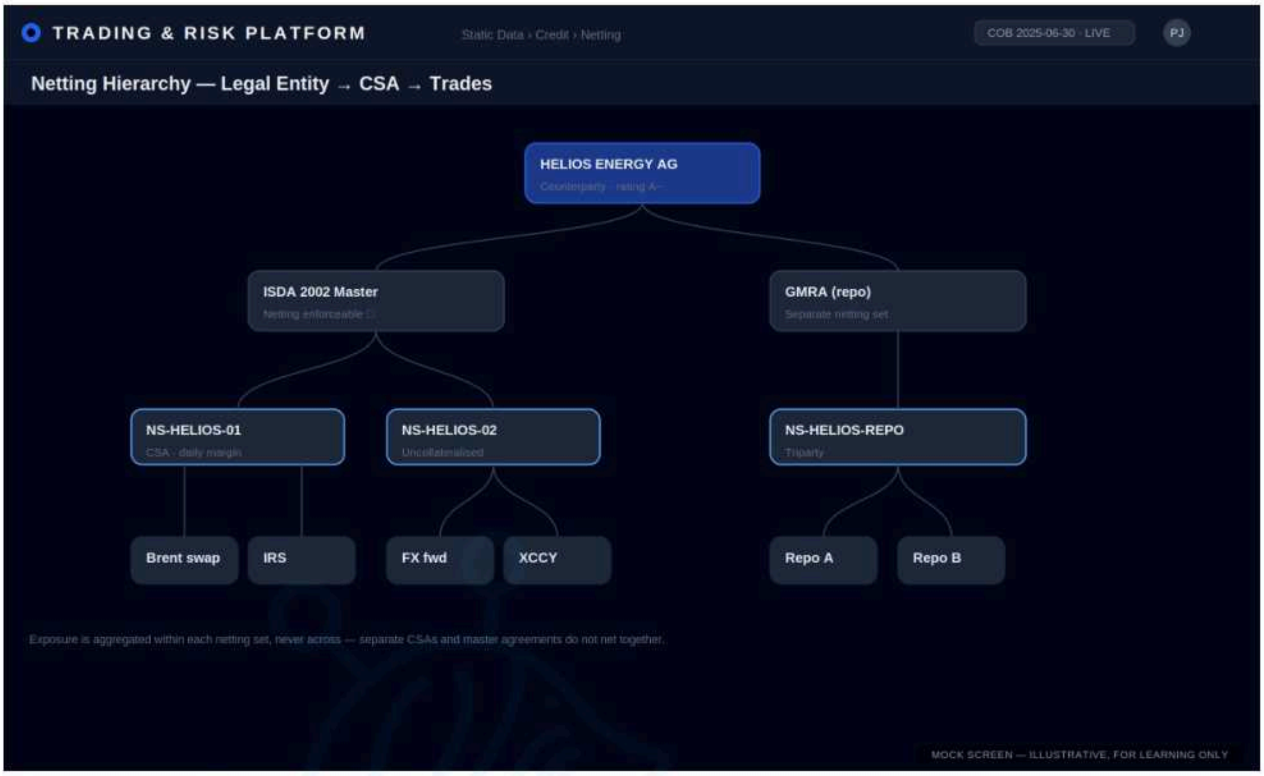

Figure 193.2b — The netting hierarchy. A single counterparty may sit above several master agreements, each with its own netting sets and CSAs; exposure nets within a set but never across separate agreements. Reading this tree correctly is the difference between a defensible netted CVA and an overstated one.

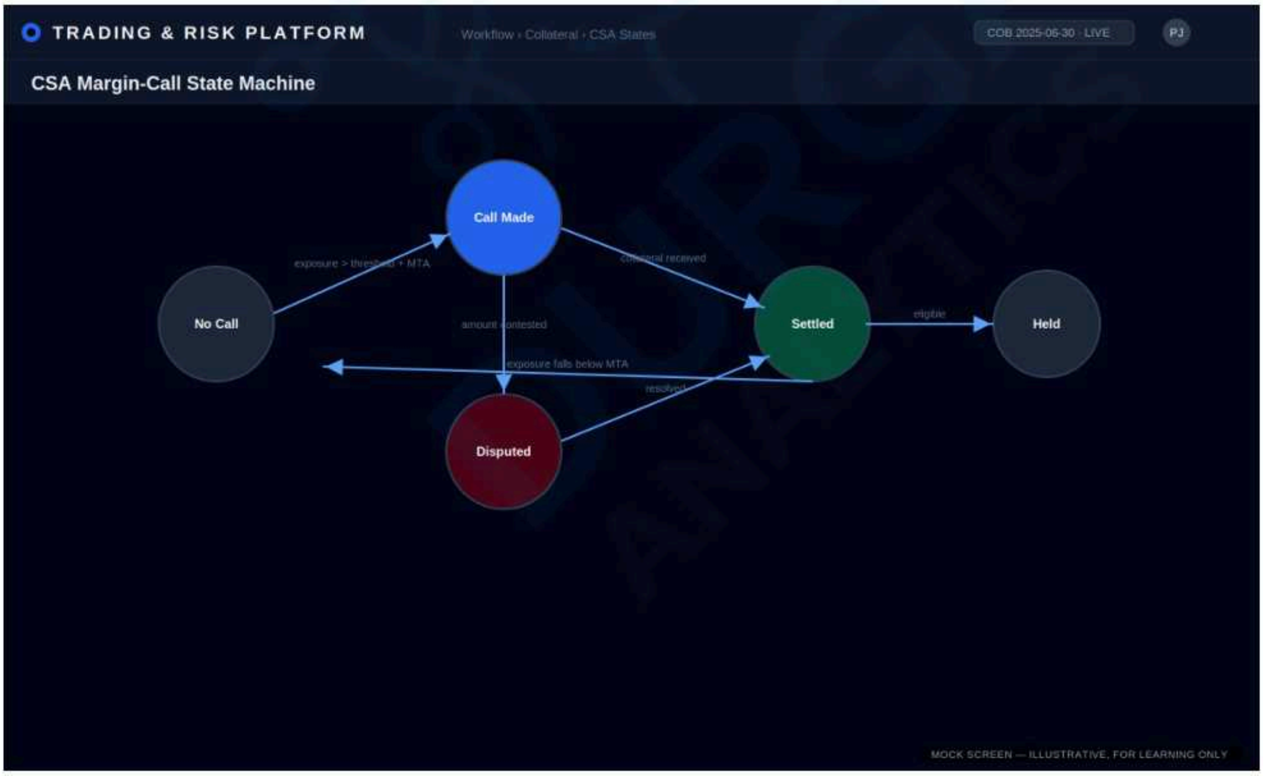

Figure 193.2c — The CSA margin-call state machine. Exposure crossing threshold plus MTA triggers a call; the call may be settled, disputed, or fall away as exposure drops. The simulation of this lifecycle — including the margin period of risk between the last received collateral and close-out — is what leaves residual exposure and therefore residual CVA even under a tight daily CSA.

The CSA fields that actually move CVA

- **Threshold:** exposure the counterparty can run before posting. Higher threshold → higher CVA.
- **Minimum Transfer Amount (MTA):** collateral only moves in blocks this size — leaves residual exposure.
- **Margin frequency:** daily vs weekly changes the margin period of risk and therefore residual exposure.
- **Eligible collateral & haircuts:** what can be posted and how it is valued; feeds ColVA.
- **Independent Amount:** over-collateralisation buffer, i.e. initial margin — feeds MVA.
- **Netting enforceability:** the on/off switch for the whole benefit. No opinion → gross exposure.

Representing the netting set and CSA in the schema

Everything on that screen is persisted in the risk store. Here is the core of a normalised schema. You will query these tables constantly when reconciling XVA numbers, so learn them cold.

SQL — netting set and CSA reference schema (risk store)

```sql
CREATE TABLE netting_set (
    netting_set_id     VARCHAR(24)  PRIMARY KEY,
    counterparty_id    VARCHAR(24)  NOT NULL REFERENCES counterparty(cpty_id),
    master_agmt_type   VARCHAR(16)  NOT NULL,    -- ISDA2002, GMRA, etc.
    governing_law      VARCHAR(24),
    base_ccy           CHAR(3)      NOT NULL,
    netting_enforce    BOOLEAN      NOT NULL DEFAULT FALSE,
    legal_opinion_ref  VARCHAR(64),
    created_ts         TIMESTAMP    NOT NULL DEFAULT now()
);

CREATE TABLE csa (
    csa_id             VARCHAR(24)  PRIMARY KEY,
    netting_set_id     VARCHAR(24)  NOT NULL REFERENCES netting_set(netting_set_id),
    threshold_us       NUMERIC(18,2) NOT NULL DEFAULT 0,
    threshold_them     NUMERIC(18,2) NOT NULL DEFAULT 0,
    mta                NUMERIC(18,2) NOT NULL DEFAULT 0,
    rounding           NUMERIC(18,2) NOT NULL DEFAULT 0,
    independent_amount NUMERIC(18,2) NOT NULL DEFAULT 0,
    margin_freq        VARCHAR(8)   NOT NULL DEFAULT 'DAILY',
    mpor_days          SMALLINT     NOT NULL DEFAULT 10,  -- margin period of risk
    collateral_ccy     CHAR(3)      NOT NULL DEFAULT 'USD',
    rehypothecation    BOOLEAN      NOT NULL DEFAULT TRUE
);

CREATE TABLE csa_eligible_collateral (
    csa_id             VARCHAR(24)  NOT NULL REFERENCES csa(csa_id),
    asset_class        VARCHAR(24)  NOT NULL,    -- CASH_USD, USTBILL, ...
    haircut_pct        NUMERIC(6,4) NOT NULL,
    PRIMARY KEY (csa_id, asset_class)
);
```

■ The netting trap that fails audits

A trade will happily book against a netting set even when the legal opinion does not support netting in that jurisdiction. If **netting_enforce** is FALSE, the exposure engine must compute exposure *gross* (sum of positive trade values) rather than *net*. A surprising number of overstated netting-benefit findings trace back to this one boolean being set optimistically. Verify it against the legal_opinion_ref before you trust any netted CVA.

193.4 The engine room: simulating future exposure

Every XVA number rests on one thing: a distribution of what each netting set will be worth at many points in the future, under many scenarios. This is the exposure simulation, and it is by far the heaviest computation in the whole system. Understand it and the rest of XVA is bookkeeping. Misunderstand it and every downstream number is quietly wrong.

From one price to a cloud of futures

The pricing engine you met earlier gives you *today's* value of a trade under *today's* market. Exposure simulation asks a different question: across thousands of possible future paths of rates, commodity prices, FX and volatility, what will this netting set be worth at each future date? The steps are always the same:

1. **Choose risk factors and models.** Interest rates (e.g. a one-factor short-rate model), commodity prices (a two-factor forward model capturing short- and long-end volatility), FX (a diffusion), and the counterparty's default intensity. Each model is calibrated to today's market.
2. **Generate scenarios.** Draw thousands of Monte Carlo paths of every risk factor forward through time, on a chosen time grid (dense in the near term, sparser far out).
3. **Reprice on every path at every grid point.** For each scenario and each future date, revalue every trade in the netting set. This is where the machine-hours go.
4. **Net and collateralise.** Sum trade values within the netting set; apply the CSA logic to simulate the collateral that would have been posted, leaving residual exposure.
5. **Aggregate into profiles.** Across scenarios, take expectations and quantiles to produce EPE, ENE and PFE curves through time.

Figure 193.3 — The output of the engine for one netting set: expected positive exposure (EPE), expected negative exposure (ENE) and the 95% potential future exposure (PFE) envelope through time. CVA integrates the EPE curve against the counterparty's default probability; DVA integrates ENE against ours; limits and regulatory capital lean on the PFE envelope. The near-term bump is the seasonality of the commodity leg; the long decay is amortisation and mean reversion.

The exposure vocabulary you must not confuse

- **EE (Expected Exposure):** average of exposure across scenarios at a future date (positive side).
- **EPE:** the time-average of EE — a single summary number, and the CVA integrand.
- **ENE:** the mirror on the negative side — drives DVA and FBA.
- **PFE (Potential Future Exposure):** a high quantile (e.g. 95%/97.5%) of exposure — drives limits.
- **Effective EPE:** a non-decreasing version of EPE used in the regulatory capital formula.
- **Peak PFE:** the maximum of the PFE curve — the number credit officers set limits against.

Configuring the run

Under the hood, the platform drives an open-source risk engine (ORE) (ORE), so the simulation is parameterised by ORE-style XML. Get the config wrong and you will burn a night of grid time producing garbage. Below is the engine-configuration screen followed by the ORE XML it generates — the artifact you will edit and check into version control.

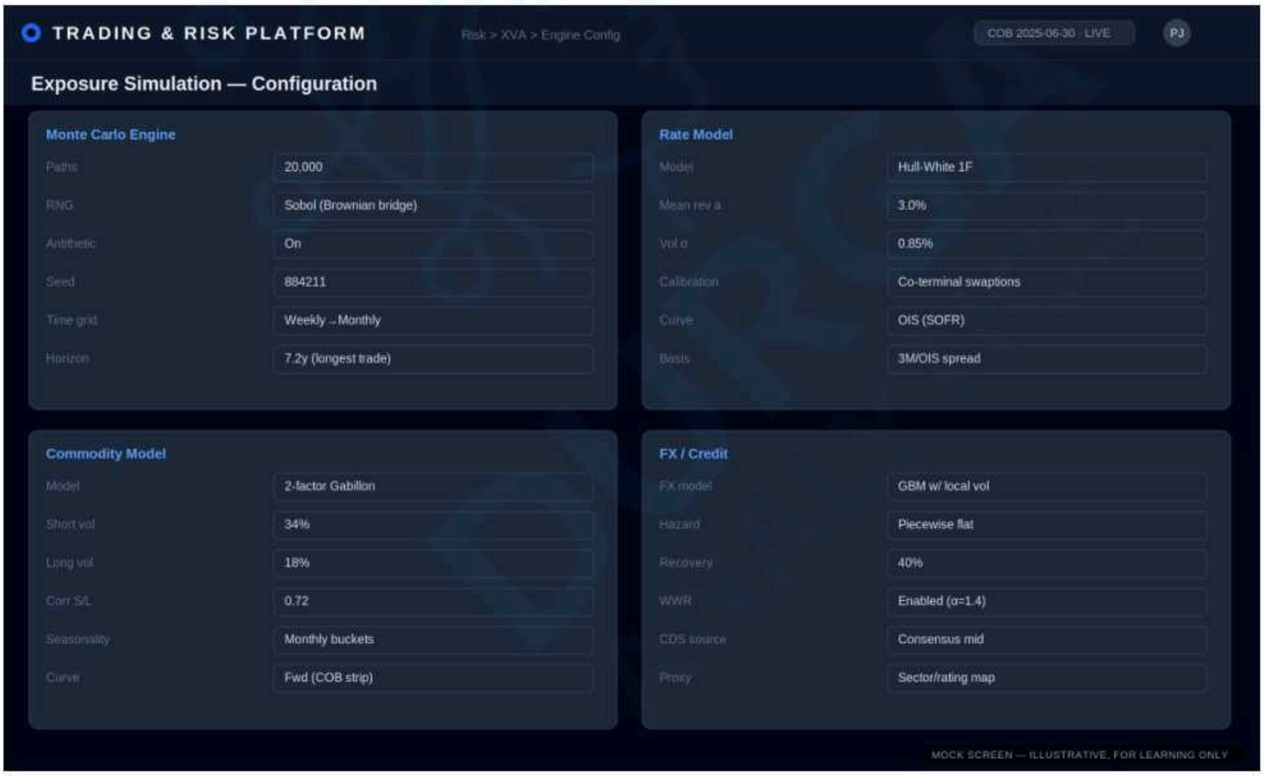

Figure 193.4 — Exposure-simulation configuration. Four model blocks (rates, commodity, FX/credit) plus the Monte Carlo controls. Every field here maps one-to-one to an element in the run configuration XML on the next page.

ORE XML — exposure/CVA run configuration (ore_xva.xml, generated by the platform)

```xml
<?xml version="1.0" encoding="UTF-8"?>
<ExposureRun id="XVA-20250630-N" cob="2025-06-30" baseCcy="USD">

  <MonteCarlo>
    <Paths>20000</Paths>
    <Generator>SOBOL</Generator>           <!-- low-discrepancy -->
    <BrownianBridge>true</BrownianBridge>
```

```
    <Antithetic>true</Antithetic>
    <Seed>884211</Seed>
  </MonteCarlo>

  <TimeGrid>
    <Rule step="1W"  until="3M"/>          <!-- weekly to 3 months -->
    <Rule step="1M"  until="2Y"/>          <!-- monthly to 2 years  -->
    <Rule step="3M"  until="MAX"/>         <!-- quarterly thereafter -->
    <Horizon source="LONGEST_TRADE"/>
  </TimeGrid>

  <Models>
    <Rates ccy="USD" model="HULL_WHITE_1F">
      <MeanReversion>0.030</MeanReversion>
      <Volatility calib="COTERMINAL_SWAPTIONS"/>
      <DiscountCurve>OIS.USD.SOFR</DiscountCurve>
    </Rates>
    <Commodity name="BRENT" model="GABILLON_2F">
      <ShortVol>0.34</ShortVol>
      <LongVol>0.18</LongVol>
      <CorrShortLong>0.72</CorrShortLong>
      <ForwardCurve>CMDTY.BRENT.ICE</ForwardCurve>
      <Seasonality bucket="MONTHLY"/>
    </Commodity>
    <FX pair="EURUSD" model="GBM_LOCALVOL"/>
    <Credit cpty="HELIOS" model="PIECEWISE_HAZARD">
      <Recovery>0.40</Recovery>
      <WrongWayRisk enabled="true" alpha="1.4"/>
    </Credit>
  </Models>

  <NettingScope>
    <Include portfolio="ALL"/>
    <NettingSet id="NS-HELIOS-01"/>
    <ApplyCollateral>true</ApplyCollateral>
    <MporDays>10</MporDays>
  </NettingScope>

  <Outputs>
    <Profile type="EPE"/><Profile type="ENE"/>
    <Profile type="PFE" quantile="0.95"/>
    <Profile type="EFFECTIVE_EPE"/>
    <Persist target="datamart" table="xva_exposure_profile"/>
  </Outputs>
</ExposureRun>
```

■ **The three settings that silently corrupt a run**

(1) Horizon shorter than the longest trade — the tail of exposure is truncated and CVA is understated. **(2) Time grid too coarse near a large cashflow or option expiry** — the exposure spike is missed entirely. **(3) ApplyCollateral=true with a stale or mis-mapped CSA** — the engine assumes collateral that legally is not there, understating exposure. None of these throw an error. All three produce a clean, wrong number. Reviewing this XML is a real desk control, not a formality.

Wrong-way risk: when exposure and default move together

The standard CVA formula quietly assumes your exposure to a counterparty is independent of that counterparty's likelihood of defaulting. Often it is not. Sell oil protection to an oil producer and your exposure to them grows exactly when oil prices — and their default risk — move against them. That positive correlation between exposure and default is **wrong-way risk (WWR)**, and it makes true CVA larger than the independent formula suggests. The simplest treatment is an *alpha* multiplier on EPE; more sophisticated engines model the correlation directly by linking the default-intensity process to the exposure drivers.

Figure 193.5 — Wrong-way risk diagnostic. Each point is one simulated scenario at one time: default intensity on the horizontal axis, exposure on the vertical. An upward cloud (shown) means exposure tends to be high exactly when the counterparty is most likely to default — CVA must be scaled up. A flat cloud means independence holds and the plain formula is adequate.

193.5 Turning credit spreads into default probabilities

The exposure engine gives you the EPE curve. To turn it into CVA you need the other half of the integrand: the probability that the counterparty defaults in each future interval. That comes from their credit curve. For a name with liquid single-name CDS, you bootstrap a term-structure of default (hazard) rates directly from the CDS quotes. For a name without liquid CDS — which is most corporates and all mid-market counterparties — you build a **proxy curve** from a sector/rating/region mapping onto liquid index constituents.

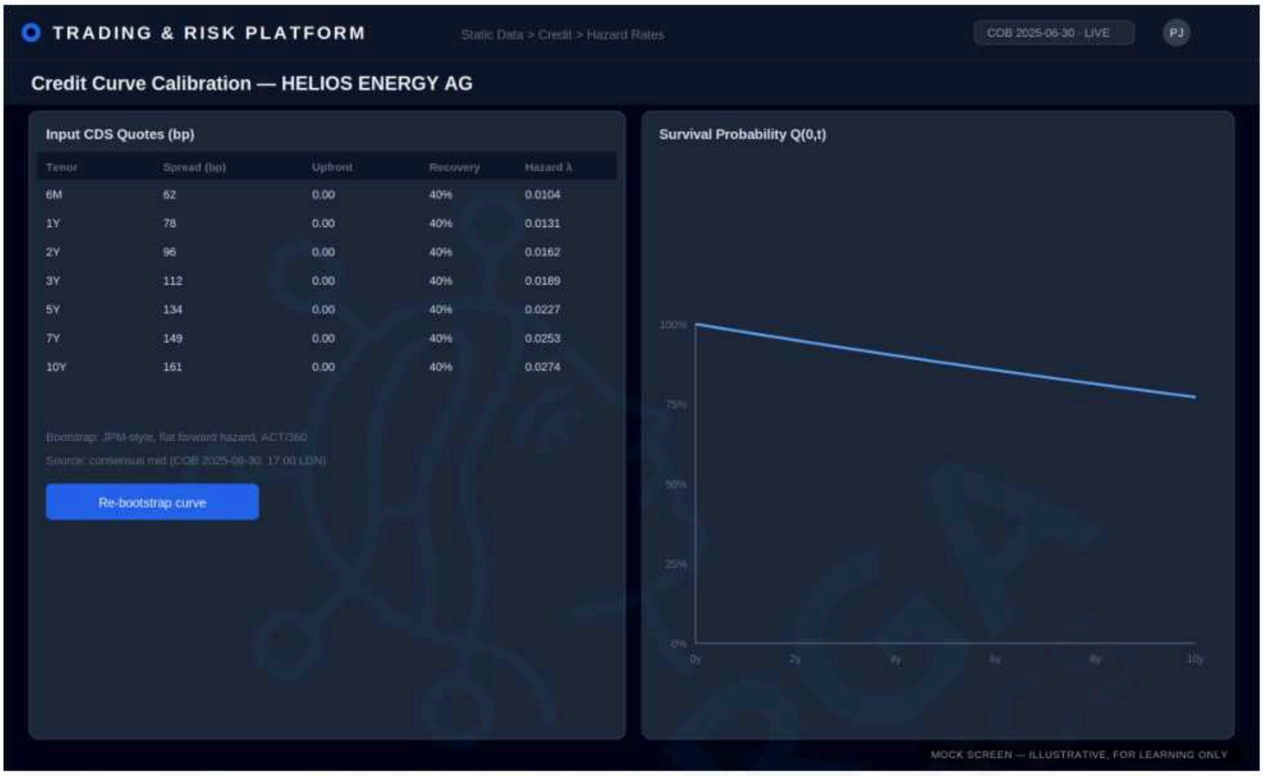

Figure 193.6 — Credit-curve calibration. Market CDS spreads by tenor (left) are bootstrapped into a piecewise-flat hazard rate term structure, which integrates to the survival probability curve Q(0,t) on the right. The marginal default probability in each simulation interval — the term the CVA formula needs — is simply Q at the start of the interval minus Q at the end.

Figure 193.6b — The bootstrapped hazard rate as a piecewise-flat term structure, and the survival curve it integrates to. Each step is the forward default intensity implied by the CDS spread for that tenor; the survival probability is the exponential of the negative cumulative hazard. This is the engine's view of the counterparty's default timing — the other half of the CVA integrand.

Figure 193.6c — The bootstrap in action. Starting from the shortest tenor, each CDS quote is solved for the hazard rate that reprices it to par, holding earlier hazards fixed, then the procedure steps outward. The highlighted band is the tenor being solved; everything to its left is already pinned down.

The bootstrap, in words then in numbers

A CDS is a contract where you pay a periodic spread in exchange for protection against a name's default. Its fair spread encodes the market's view of that name's default probability. Bootstrapping inverts this:

starting from the shortest tenor, you solve for the hazard rate that makes the CDS fairly priced, then move to the next tenor holding the earlier hazard rates fixed, and so on. The output is a survival curve $Q(0,t) = \exp(-\int\lambda\, ds)$.

Tenor	CDS mid (bp)	Bootstrapped hazard λ	Survival Q(0,t)	Marginal PD in bucket
6M	62	0.0104	0.9948	0.52%
1Y	78	0.0131	0.9869	0.79%
2Y	96	0.0162	0.9560	3.09%
3Y	112	0.0189	0.9209	3.51%
5Y	134	0.0227	0.8395	8.14%
7Y	149	0.0253	0.7599	7.96%
10Y	161	0.0274	0.6524	10.75%

Read the last column against Figure 193.1's CVA number: the CVA charge is heaviest where the product of marginal default probability and expected exposure is largest — typically the 2–5 year region, where both the exposure profile and the cumulative default probability are meaningful. That is why a five-year uncollateralised swap with a BBB name is the classic 'expensive' trade for the XVA desk.

SQL — persisting the bootstrapped curve and reading marginal PDs

```
CREATE TABLE credit_curve_point (
    cpty_id      VARCHAR(24) NOT NULL,
    cob_date     DATE        NOT NULL,
    tenor_years  NUMERIC(6,3) NOT NULL,
    hazard_rate  NUMERIC(10,6) NOT NULL,
    survival_q   NUMERIC(10,8) NOT NULL,
    recovery     NUMERIC(5,4) NOT NULL,
    curve_source VARCHAR(12) NOT NULL,   -- CDS_DIRECT | PROXY
    PRIMARY KEY (cpty_id, cob_date, tenor_years)
);

-- Marginal default probability per interval, for the CVA integral
SELECT  a.tenor_years            AS t_end,
        LAG(a.survival_q) OVER w AS q_start,
        a.survival_q             AS q_end,
        LAG(a.survival_q) OVER w - a.survival_q AS marginal_pd
FROM    credit_curve_point a
WHERE   a.cpty_id  = 'HELIOS'
  AND   a.cob_date = DATE '2025-06-30'
WINDOW  w AS (ORDER BY a.tenor_years)
ORDER BY a.tenor_years;
```

◆ Proxy curves are where reviewers dig

When curve_source = PROXY, you are asserting that an unlisted counterparty behaves like a blend of liquid names sharing its sector, rating and region. That mapping is a modelling choice with real P&L consequences, and it is the first thing model validation and audit will challenge. Always be able to state, for any proxied counterparty, exactly which index and adjustment produced its curve.

193.6 Feeding the engine: trade population and representation

The exposure engine cannot simulate what it cannot see. Every trade in scope for XVA is extracted from the platform trade store, enriched with the static data the models need, and handed to the engine in a canonical form. The platform keeps its own **canonical trade representation** — a versioned JSON object persisted in the `etrm` schema — and can emit or ingest **FpML** (the open ISDA standard) at the boundary when it needs to interoperate with clearing, confirmation or a counterparty's system. Learning to read both is a core desk skill: when a CVA number looks wrong, the trade object is the first place you look.

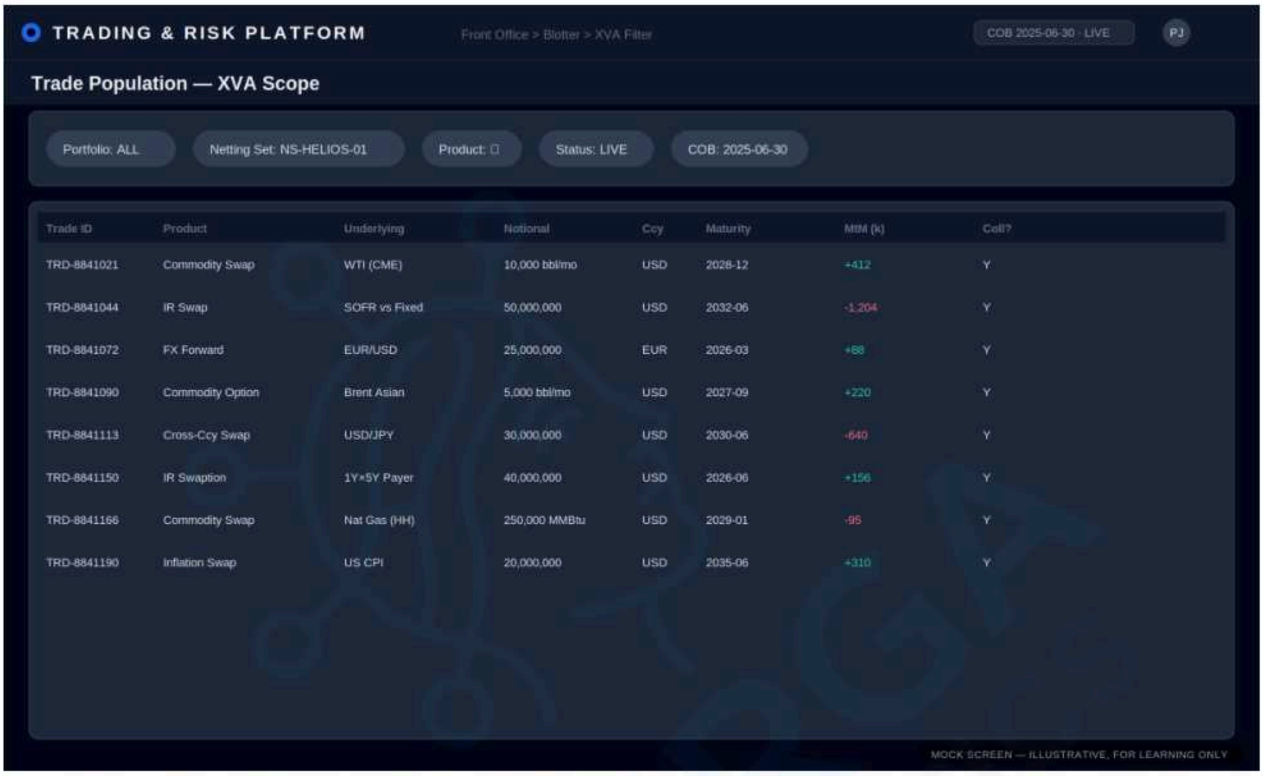

Figure 193.7 — The platform trade blotter, filtered to the XVA scope of one netting set. A cross-asset book: commodity swaps and options, interest-rate swaps and swaptions, FX and cross-currency, inflation. Each row resolves to a canonical trade object before it reaches the exposure engine. The blotter's filters define scope; a trade silently excluded here is silently missing from CVA.

The canonical trade object

Below is the first row of the blotter — a commodity swap — as the platform holds it. A header carries identity and lifecycle status; a parties block links to the netting set and CSA; a product block describes the economics. This is the object the valuation service reads; if a field is wrong here, the simulated exposure is wrong, and so is CVA.

canonical trade object (GET /api/v1/trades/TRD-8841021)


```
{
  "trade_id": "TRD-8841021",
  "tenant_id": 2,
  "version": 4,
  "status": "LIVE",
  "trade_date": "2023-11-14",
  "book": "CMDTY-OIL-01",
  "trader": "desk.oil.02",
  "parties": {
    "ourselves": "INTERNAL-DESK",
    "counterparty": "HELIOS",
```

```
    "netting_set_id": "NS-HELIOS-01",
    "csa_id": "CSA-HELIOS-01"
  },
  "product": {
    "type": "COMMODITY_SWAP",
    "direction": "PAY_FIXED",
    "underlying": { "commodity": "BRENT", "market": "ICE" },
    "notional_quantity": { "value": 10000, "unit": "BBL", "freq": "MONTHLY" },
    "fixed_price": { "value": 78.40, "ccy": "USD" },
    "float_reference": "BRENT.FUTURES.FRONT",
    "effective_date": "2023-12-01",
    "termination_date": "2028-12-31",
    "payment_calendar": ["USNY", "LOND"],
    "day_count": "ACT/360"
  }
}
```


The same trade as FpML, at the interop boundary

When the platform must speak to the outside world — a clearing house, a confirmation platform, a counterparty reconciliation — it maps the canonical object to **FpML**, the open standard. The economics are identical; only the envelope changes. You will meet FpML repeatedly across this program, so learn to recognise its shape.

FpML 5.x — same commodity swap, emitted for interoperability


```
<commoditySwap>
  <productType>CommoditySwap</productType>
  <fixedLeg>
    <payerPartyReference href="INTERNAL-DESK"/>
    <receiverPartyReference href="HELIOS"/>
    <commodity><instrumentId>BRENT</instrumentId>
      <specifiedPrice>Settlement</specifiedPrice></commodity>
    <notionalQuantity>10000</notionalQuantity>
    <quantityUnit>BBL</quantityUnit>
    <calculationPeriodsSchedule><periodMultiplier>1</periodMultiplier>
      <period>M</period></calculationPeriodsSchedule>
    <fixedPrice><price>78.40</price><priceCurrency>USD</priceCurrency></fixedPrice>
  </fixedLeg>
  <floatingLeg>
    <commodity><instrumentId>BRENT</instrumentId>
      <priceSource>ICE_FUTURES</priceSource></commodity>
    <effectiveDate>2023-12-01</effectiveDate>
    <terminationDate>2028-12-31</terminationDate>
  </floatingLeg>
</commoditySwap>
```

Fields the exposure engine cannot live without

- **netting_set_id** / **csa_id:** without these, the trade is treated standalone and uncollateralised — CVA balloons.
- **termination_date:** sets the simulation horizon; a wrong date truncates or over-extends exposure.
- **direction & fixed_price:** the sign of exposure; flip them and EPE becomes ENE.
- **underlying** / **float_reference:** binds the trade to the right simulated risk factor and curve.
- **payment_calendar** / **day_count:** place cashflows on the correct dates; drives the exposure profile shape.

◆ Why the platform keeps its own object and speaks FpML at the edge

Incumbent platforms bind you to a proprietary internal format that only their tools understand. It takes the open route: a clean canonical object you can query and inspect in PostgreSQL, plus FpML export/import so nothing is locked in. Same coverage as the incumbents, none of the black box.

The result coming back from the engine

After the ORE-based valuation service runs, results return to the platform as a structured result object per netting set: the computed adjustments plus the exposure summary statistics that feed limits and capital. This is what lands in the datamart.

Platform XVA result object (POST /api/v1/xva/results)


```
{
  "cob_date": "2025-06-30",
  "run_id": "XVA-20250630-N",
  "netting_set_id": "NS-HELIOS-01",
  "counterparty": "HELIOS",
  "base_ccy": "USD",
  "adjustments": {
    "cva":  -1820344.11,
    "dva":    512088.40,
    "fca":   -742110.02,
    "fba":    311204.55,
    "colva": -144098.10,
    "kva":   -402551.79
  },
  "exposure": { "epe_peak": 12418003, "pfe_95": 28104221, "eff_epe": 9820115 },
  "diagnostics": { "paths": 20000, "horizon_years": 7.2, "wwr_alpha": 1.4,
                   "engine": "ORE", "engine_code": "ORE_PYTHON" }
}
```

193.7 Funding: the cost of holding the position

CVA and DVA are about default. FVA is about the simpler, ever-present fact that money has a cost. When you are in-the-money on an uncollateralised trade, you are effectively lending to the market and must fund that lending at your own funding rate — above the risk-free rate. When you are out-of-the-money, you are effectively borrowing and may earn a benefit. FVA is the present value of that funding cost/benefit strip over the trade's life, computed off the same exposure profiles you already built for CVA, but weighted by your **funding spread** rather than the counterparty's default probability.

Figure 193.8 — FVA construction. The funding-spread curve (left) is your incremental cost of funds over the risk-free rate by tenor. Applied to expected positive exposure it gives the funding cost (FCA); applied to expected negative exposure it gives the funding benefit (FBA). The decomposition (right) also shows MVA — the cost of funding initial margin — which sits alongside FVA in the funding family.

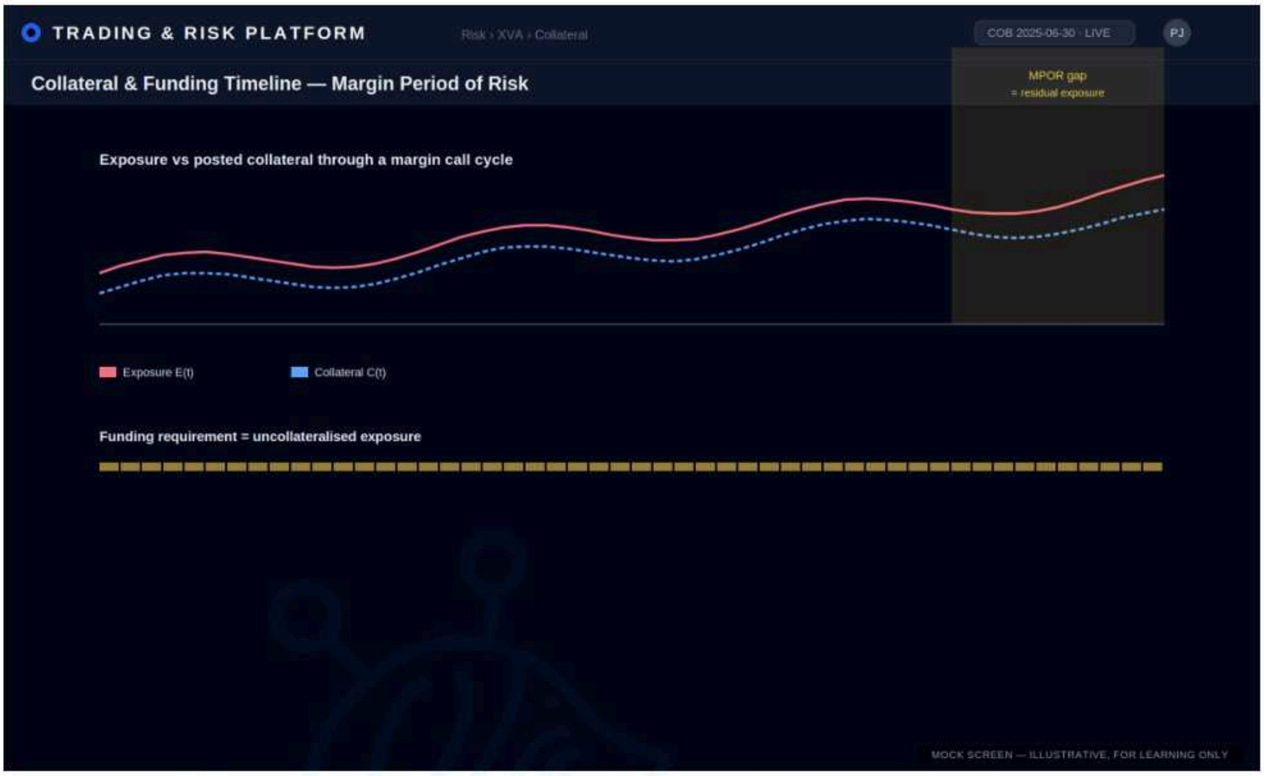

Figure 193.8b — The collateral and funding timeline. Posted collateral (dashed) tracks exposure (solid) but lags it, and during the margin period of risk a gap opens that is both residual counterparty exposure and a funding requirement. The shaded strip below shows the uncollateralised exposure that FVA prices — funding cost is charged on exactly the amount collateral does not cover.

Figure 193.8c — The collateral waterfall. Gross exposure is reduced by the threshold and MTA to give the call amount; a haircut on posted assets and the margin-period-of-risk gap leave a residual that is exactly the exposure CVA and ColVA still price after collateral.

Formula — FCA and FBA

N

```
FCA = − Σ   s_fund(t_k) · EPE(t_k) · DF(t_k) · Δt_k     (cost, on positive exposure)
        k=1
         N
FBA = + Σ   s_fund(t_k) · ENE(t_k) · DF(t_k) · Δt_k     (benefit, on negative exposure)
        k=1

FVA(net) = FCA + FBA
s_fund(t) = funding spread over risk-free at tenor t   (from Figure 193.8 left panel)
```

■ The FVA / DVA overlap — a real double-counting trap

FBA (a funding benefit on your negative exposure) and DVA (a credit benefit on the same negative exposure) are economically overlapping: both reward you for being out-of-the-money to a counterparty you might default on. Booking both in full double-counts the benefit. Desks resolve this by choosing a funding framework that nets the overlap — commonly recognising FVA and treating DVA as a memo item, or explicitly subtracting the overlap. Whatever the house choice, you must know which convention your book uses, because it changes the sign and size of the total adjustment materially.

Configuring the funding curve as a feeder

The funding spread curve is not a market observable you can buy off a screen; it is constructed by treasury from the bank's own issuance and blended into the risk store by a feeder job. Here is the feeder definition that publishes it each morning.

XML — funding curve feeder definition (feeder.funding.xml)

```xml
<Feeder id="FUNDING_SPREAD_DAILY" schedule="0 30 6 * * MON-FRI" tz="Europe/London">
  <Source type="TREASURY_FILE" path="/feeds/treasury/funding_${COB}.csv"/>
  <Mapping>
    <Column name="tenor"       target="tenor_years" type="tenor"/>
    <Column name="spread_bp"   target="spread"      type="bp_to_decimal"/>
  </Mapping>
  <Transform>
    <Blend base="OIS.USD.SOFR" add="spread" method="ADDITIVE"/>
    <Interpolation>LOG_LINEAR_DF</Interpolation>
  </Transform>
  <Publish curveId="FUNDING.USD.HOUSE" store="risk_datamart"/>
  <Validation>
    <Rule check="MONOTONIC_NONNEG" onFail="QUARANTINE"/>
    <Rule check="MAX_DAILY_MOVE_BP" value="25" onFail="ALERT"/>
  </Validation>
</Feeder>
```

193.8 The Monte Carlo engine: from Brownian motion to the exposure cube

Everything in this chapter rests on one machine: a Monte Carlo simulation of the future that reprices the book on thousands of scenarios. So far we have treated it as a black box that emits EPE and PFE. This section opens the box. It is the most mathematically demanding part of the chapter, written for the quant developer and the model validator: we derive the dynamics, discretise them, give implementation pseudocode, work a numerical example, and discuss convergence. If you have read the earlier sections, you already know *why* we simulate; here you learn *how*, precisely enough to build or validate an engine. (Cross-reference: the curve inputs come from Ch 92–137, Market Data & Curves; the pricing that runs on each path is Ch 184–192, Pricing.)

193.8.0 Stochastic differential equations and Ito intuition

Before the models, the language they are written in. A stochastic differential equation (SDE) describes how a quantity evolves when part of its change is random. The generic form is a drift term (the predictable pull) plus a diffusion term (the random shock scaled by volatility):

$$dX(t) = \mu(X,t)\,dt + \sigma(X,t)\,dW(t)$$

A generic SDE: deterministic drift μ plus random diffusion σ driven by Brownian motion W.

The one piece of stochastic calculus a practitioner truly needs is **Ito's lemma**, and the intuition behind it is simpler than the formalism. In ordinary calculus, if you change variables to f(X), the change df is f′ dX. For a random X driven by Brownian motion that is incomplete, because Brownian motion is so jagged that its squared increment does not vanish: $(dW)^2$ behaves like dt, not like zero. Ito's lemma adds exactly that second-order correction:

$$df = f'(X)\,dX + \tfrac{1}{2} f''(X)\,(dX)^2, \quad \text{with } (dW)^2 \to dt,$$

$$\text{so } df = \left[\, \mu f' + \tfrac{1}{2} \sigma^2 f'' \,\right] dt + \sigma f'\,dW.$$

Ito's lemma: the $\tfrac{1}{2}\sigma^2 f''$ term is the correction Brownian roughness forces on you.

That extra half-sigma-squared-f-double-prime term is why a lognormal asset has a $-\tfrac{1}{2}\sigma^2$ drift adjustment, why variance accumulates linearly in time, and why the exact Hull-White discretisation of Section 193.8.3 has the variance it does. You do not need to derive Ito's lemma to build an engine, but you must recognise the correction term when you see it in a model's dynamics — its presence or absence is a common validation check.

> ◆ **Practical takeaway, not a proof**
> The one thing to remember: randomness of the Brownian kind makes $(dW)^2 = dt$, and every consequence in derivatives pricing — drift adjustments, variance scaling, the shape of exact discretisations — flows from that single fact. When a model's simulated variance does not match its stated σ, an Ito term has usually been dropped.

193.8.5b Variance reduction, treated properly

Section 193.8.5 listed the variance-reduction levers; here is why each works, because a validator will ask you to justify them, not just name them. Monte Carlo error falls as $1/\sqrt{N}$, so anything that lowers the variance constant buys accuracy far more cheaply than more paths.

Technique	Mechanism	Typical gain	Caveat
Antithetic variates	Pair each Z with –Z; average cancels odd error	~2× effective paths	Only helps near-symmetric payoffs
Control variates	Subtract a correlated quantity with known mean	Large if correlation high	Needs an analytic control
Sobol (QMC)	Low-discrepancy fill of sample space	Up to ~1/N in low dim	Degrades in high effective dim
Brownian bridge	Load big Sobol dims onto big time steps	Better QMC convergence	Construction overhead
Stratified sampling	Force even coverage of key dimension	Cuts tail noise	Bookkeeping complexity
Common random numbers	Reuse paths across bumps	Stable Greeks	Must fix seed & grid

Control variates deserve a word because they are the most powerful and least used. If you can find a related instrument whose price you know analytically and whose simulated value is highly correlated with the target, you simulate the difference (which has far lower variance) and add back the known mean. For CVA, a common control is the exposure of a simplified, analytically tractable version of the netting set. The variance reduction is proportional to the square of the correlation, so a control correlated at 0.95 with the target removes about 90% of the variance.

193.8.11 Greeks within XVA

XVA is not computed once and forgotten; the desk hedges its daily P&L, which means it needs the sensitivities of the adjustment to every market factor — the XVA Greeks. These are harder than vanilla Greeks because the adjustment is itself an expectation over a simulation, so a sensitivity is the derivative of a Monte Carlo estimate. Three approaches, in increasing sophistication:

- **Bump-and-revalue:** shift one input by a small amount, re-run the whole simulation on the same paths (common random numbers), and difference. Simple and exact to first order, but each Greek costs a full re-simulation — for a large risk ladder this is the dominant compute cost of the whole XVA function.
- **Pathwise derivatives:** differentiate the payoff along each path analytically where the payoff is smooth. Cheaper, but breaks at kinks (the max(·,0) in exposure) and needs care at discontinuities.
- **Adjoint algorithmic differentiation (AAD):** propagate derivatives backward through the entire computation in one pass, producing *all* sensitivities at a cost of a small multiple of a single valuation. This is the production standard for large XVA books; it is what makes a full CVA risk ladder tractable overnight.

$$\text{Bump-and-revalue:}\quad \partial CVA/\partial x \approx [\, CVA(x+h) - CVA(x-h)\,] / (2h),$$

same random paths for x+h and x–h so simulation noise cancels in the difference.

Central-difference CVA Greek with common random numbers — the noise-cancelling trick.

◆ **Why AAD is a career-relevant skill**

The single biggest computational problem in XVA is producing a full, stable risk ladder fast enough to hedge on. AAD is the answer the whole industry converged on, and quant developers who can implement or validate an adjoint pass are in demand precisely because it is the bottleneck. Cross-reference: the hedging use of these Greeks is Section 193.8F (Sensitivities and hedging).

193.8.1 The random driver: Brownian motion

Every diffusion model in the engine is built from one primitive: standard Brownian motion (a Wiener process) $W(t)$. Its defining properties are that it starts at zero, has independent increments, and each increment over a horizon is normally distributed with variance equal to the length of the horizon:

$$W(0) = 0, \quad W(t) - W(s) \sim N(0, t - s) \quad \text{for } 0 \le s < t,$$

and increments over disjoint intervals are independent.

The Wiener process — the engine's only source of randomness.

In discrete time on a grid $0 = t_0 < t_1 < \dots < t_M$, we generate the process by drawing independent standard normals $Z_k \sim N(0,1)$ and setting the increment $\Delta W_k = W(t_k) - W(t_{k-1}) = \sqrt{\Delta t_k} \cdot Z_k$, where $\Delta t_k = t_k - t_{k-1}$. That single line — scale a standard normal by the square root of the step — is the seed of the entire simulation. Everything else is a transformation of these increments.

193.8.2 From one factor to many: correlated risk factors

A real book depends on many risk factors at once — rates, several commodity curves, FX, credit — and they move together. We therefore need a vector of correlated Brownian motions. Let Σ be the correlation matrix of the d risk-factor drivers. We produce correlated normals by Cholesky decomposition: factor $\Sigma = L L^T$ with L lower-triangular, draw a vector of independent standard normals Z, and set the correlated draw as L Z.

$$\Sigma = L L^T \ \text{(Cholesky)}, \quad Z \sim N(0, I_d), \quad X = L Z \sim N(0, \Sigma).$$

Correlated shocks from independent normals via the Cholesky factor L.

The correlated increment vector at step k is then $\Delta W_k = \sqrt{\Delta t_k} \cdot L\ Z_k$. The correlation matters enormously for XVA: an oil producer counterparty whose credit spread is positively correlated with the oil price you are exposed to creates wrong-way risk (Section 193.8.8), and that correlation lives in Σ.

◆ **Validator checkpoint — the correlation matrix**

Σ must be symmetric and positive semi-definite or the Cholesky factorisation fails. In practice estimated correlation matrices are often *not* PSD (noise, mismatched histories), so engines apply a repair (nearest-correlation-matrix projection). A validator will ask: how is Σ estimated, over what window, and what repair is applied? An unrepaired or silently-clipped matrix is a common finding.

193.8.3 The interest-rate model: Hull-White, derived and discretised

For rates the workhorse is the one-factor Hull-White (extended Vasicek) model. The short rate $r(t)$ mean-reverts to a time-dependent level with constant volatility, driven by a Brownian motion. The continuous-time stochastic differential equation is:

$$dr(t) = [\ \theta(t) - a \cdot r(t)\ ]\ dt + \sigma\ dW(t),$$

Hull-White 1-factor: a = mean-reversion speed, σ = volatility, θ(t) fits today's curve.

The mean-reversion term $-a \cdot r$ pulls the rate back toward $\theta(t)/a$; the $\theta(t)$ function is chosen so the model reprices today's discount curve exactly (a no-arbitrage fit). To simulate it we use the exact conditional distribution rather than a crude Euler step, because Hull-White has a closed-form transition: given $r(t_{k-1})$, the next value $r(t_k)$ is normal with a known mean and variance:

$$r(t_k) = r(t_{k-1})\ e^{-a\Delta t} + \alpha_k + \sigma \cdot \sqrt{[(1-e^{-2a\Delta t})/(2a)]} \cdot Z_k,$$

$$\text{where } \alpha_k \text{ absorbs } \theta \text{ and the curve fit; } Z_k \sim N(0,1).$$

Exact discretisation of Hull-White — no Euler discretisation error.

Using the exact transition matters: it removes the time-stepping bias that an Euler-Maruyama scheme would introduce, so the only error left is Monte Carlo sampling error, which we control with path count. This is the kind of choice a model validator checks first — exact scheme versus Euler, and the resulting bias.

193.8.4 Beyond one factor: LMM, Heston, SABR

Hull-White is adequate for many exposure calculations, but richer products and richer risk demand richer models. Three you must be able to name and place:

- **LMM (LIBOR / forward Market Model):** instead of a single short rate, it evolves the whole set of forward rates directly, each with its own volatility and pairwise correlations. It captures the full shape and de-correlation of the curve — essential for exposure on products sensitive to many points on the curve (Bermudan swaptions, callable structures). The cost is high dimensionality: dozens of correlated Brownian factors.
- **Heston:** a stochastic-volatility model for equity/FX where variance itself follows a mean-reverting square-root (CIR) process, $dv = \kappa(\theta - v)dt + \xi\sqrt{v}\ dW$. It produces a volatility smile and, crucially for exposure, fat-tailed price distributions that a constant-vol model understates.
- **SABR:** the market-standard parameterisation of the volatility smile (α, β, ρ, ν) used to interpolate and extrapolate implied vols by strike. In an exposure engine SABR usually supplies the *calibrated volatility surface* that a simulation model is then fitted to, rather than being the diffusion itself.

Model	Simulates	Best for	Cost / caveat
Hull-White 1F	Short rate	General rates exposure	Single factor; limited curve shape
LMM	Forward rates	Callable / multi-point rate risk	High dimension; heavy calibration
Heston	Price + stochastic var	Equity/FX smile & tails	Extra params; Feller condition
SABR	Implied-vol smile	Vol surface interpolation	A surface, not a path model
Gabillon 2F	Commodity fwd curve	Oil/gas short- & long-vol	Seasonality handled separately

◆ **Cross-reference**
Calibration of these models to market instruments is developed in Ch 116 (Calibration Engines) and Ch 184–209 (Pricing & Valuation). Here we assume a calibrated model and focus on simulation.

193.8.5 Path generation, step by step

We can now assemble the full path-generation loop. For each of N paths we walk the time grid, drawing correlated shocks and evolving every risk factor with its exact (or best-available) scheme. The pseudocode below is the heart of the engine.

Algorithm 1 — correlated multi-factor path generation

```
INPUT:  time grid {t_0..t_M}, models with calibrated params,
        correlation matrix Sigma (d x d), path count N, RNG
PRECOMPUTE: L = cholesky(Sigma)         # d x d lower triangular

for n in 1..N:                           # loop over paths
    state = initial_risk_factor_values()  # r0, F0, fx0, ...
    for k in 1..M:                       # loop over time steps
        dt = t[k] - t[k-1]
        Z  = draw_standard_normals(d)     # d independent N(0,1)
        dW = sqrt(dt) * (L @ Z)           # correlated increments
        for each risk factor j:
            state[j] = evolve_exact(model[j], state[j], dt, dW[j])
        store F(n, k) = state             # cube slice
return F                                  # N x M x (risk factors)
```

Variance-reduction levers every production engine uses

- **Low-discrepancy sequences (Sobol):** fill the sample space more evenly than pseudo-random draws, cutting error faster than $1/\sqrt{N}$ in low effective dimension.
- **Brownian bridge construction:** assigns the most significant Sobol dimensions to the largest time steps, improving convergence for path-dependent exposure.
- **Antithetic variates:** for every draw Z also use –Z; the averaged pair has lower variance at no extra model evaluations.
- **Fixed seed & common random numbers:** reuse the same paths across sensitivity bumps so noise cancels in the difference (essential for stable Greeks).

Figure 193.8a — A handful of the 20,000 simulated Brent forward-price paths from Algorithm 1. Each faint line is one scenario ω; the dashed line is the Monte Carlo estimate of the mean E[F(t)]. The spread of the fan at each date is what future exposure is computed from — wide fan, large potential exposure.

193.8.6 Repricing on each path: the exposure cube

Path generation gives us the state of the world on every scenario and date. Now we reprice. For each path ω, each time t_k, and each trade, we compute the mark-to-market conditional on that simulated state. The result is a three-dimensional object — the **exposure cube** — indexed by path, time and trade:

$$V[\,\omega, t_k, \text{trade}\,] \;=\; \text{PV of trade at time } t_k \text{, on scenario } \omega.$$

The exposure cube: the central data structure of counterparty risk.

Figure 193.8b — The exposure cube and its three collapses. Sum over the trade axis (within a netting set) to get net portfolio value per scenario and date; take the positive part and average over paths to get expected exposure EE(t_k); time-average EE to get EPE. Every XVA number is a projection of this cube.

The collapses are precise. First, net within the netting set on each scenario and date. Then apply collateral (simulate the CSA logic to get residual exposure). Then take expectations and quantiles across scenarios:

$$\text{Netted value: } V_{ns}(\omega,t_k) = \Sigma_{\{trade \in ns\}} V[\omega,t_k,trade]$$

$$\text{Exposure: } E(\omega,t_k) = \max(V_{ns}(\omega,t_k) - C(\omega,t_k), 0)$$

$$\text{Expected exposure: } EE(t_k) = (1/N)\, \Sigma_{\{\omega\}} E(\omega,t_k)\cdot DF(t_k)$$

$$\text{PFE: } PFE_q(t_k) = \text{q-th quantile of } \{ E(\omega,t_k) : \omega \}$$

From cube to profiles: net, collateralise, then take moments across scenarios.

With EE(t_k) in hand, the discretised CVA of Section 193.2 is a sum over the time grid against the counterparty's marginal default probabilities — the two halves of the integrand finally meet. (Cross-reference: the collateral simulation C(ω,t_k) is developed in Ch 321–340, Collateral & Margin.)

193.8.7 Regression methods: Longstaff–Schwartz and American exposure

A subtlety breaks the simple picture: many trades embed **early-exercise** options — Bermudan swaptions, callable notes, American options. To reprice such a trade at a future date on a path, you need its *continuation value*: the expected discounted payoff of holding rather than exercising. That is a conditional expectation inside the simulation, and computing it with a nested inner simulation on every node is prohibitively expensive.

The **Longstaff–Schwartz** least-squares Monte Carlo (LSM) method solves this elegantly. It approximates the continuation value by *regressing* realised discounted future payoffs on a set of basis functions of the current state, across all paths at once. The fitted regression gives, for every path, an estimate of the continuation value — compare it to the immediate exercise value to decide the exercise policy, then value the trade under that policy.

Algorithm 2 — Longstaff–Schwartz continuation value (backward pass)

```
for k = M down to 1:                        # backward in time
    in_money = paths where exercise_value(k) > 0
    X = basis_functions(state[in_money, k]) # e.g. 1, S, S^2, ...
    Y = discounted_future_cashflow[in_money] # realised continuation
    beta = least_squares_fit(X, Y)          # regress Y on X
    cont_value = X @ beta                    # fitted continuation
    exercise = exercise_value(k) > cont_value
    update cashflows where exercise is True
return value = mean(discounted_cashflows)
```

◆ **Why this matters for exposure**

For a callable trade, the exposure profile depends on the simulated exercise policy: once a path exercises, the trade is gone and contributes no further exposure. Getting the LSM regression right — enough basis functions, stable fit — directly shapes the EPE of any callable book. A validator will probe the choice of basis functions and check for regression instability in the tails.

193.8.8 Wrong-way risk, formally

Section 193.4 introduced wrong-way risk (WWR) qualitatively. Now we place it in the cube. The independent-CVA formula assumes exposure E(t) and the counterparty default time τ are independent, so $E[EE(t)\cdot dPD(t)]$ factorises. WWR is precisely the failure of that factorisation: exposure and default intensity are correlated.

$$\text{Independent: } CVA = LGD \cdot \Sigma\_k\ EE(t\_k)\cdot DF(t\_k)\cdot \Delta PD\_k$$

$$\text{With WWR: } CVA = LGD \cdot \Sigma\_k\ E[\ E(t\_k)\mid \tau = t\_k\ ]\cdot DF(t\_k)\cdot \Delta PD\_k$$

WWR replaces unconditional expected exposure with exposure conditional on default at that time.

Two engine approaches. The pragmatic one is an **alpha multiplier**: scale EPE by $\alpha > 1$ to approximate the conditional-on-default uplift (the regulatory floor for α is 1.4 under the internal model method). The rigorous one is **modelled WWR**: link the counterparty's default-intensity process to the exposure drivers directly — for example, make the hazard rate a function of the simulated commodity price — so that in scenarios where your exposure is high, the simulated default probability is also high. The cube then carries the correlation natively and no external multiplier is needed.

Figure 193.8c — Why commodity WWR is acute. Exposure profiles differ sharply by underlying: an oil swap carries a seasonality bump and amortising tail; a gas portfolio swings with winter/summer demand; power options are front-loaded and decay fast; emissions rise monotonically with policy tightening. When the counterparty's business is the same commodity, exposure and default move together — wrong-way risk is structural, not incidental.

Figure 193.8e — Wrong-way versus right-way risk. Left: exposure rises with the counterparty's default intensity (positive correlation) — CVA must be scaled up. Right: exposure falls as default intensity rises (right-way) — CVA is mitigated. The correlation lives in the matrix Σ and is what the alpha multiplier or a modelled hazard link captures.

193.8.9 A worked numerical example

Numbers make the machinery concrete. Take a single-trade netting set, one year, a flat exposure of EE = USD 10m at each of four quarterly points, a flat discount factor of 1.0 for simplicity, LGD = 0.60, and the HELIOS marginal default probabilities from Section 193.5. Then CVA is a four-term sum:

Quarter t_k	EE(t_k) (USD m)	ΔPD_k	LGD	Contribution (USD)
0.25	10.0	0.0026	0.60	15,600
0.50	10.0	0.0026	0.60	15,600
0.75	10.0	0.0033	0.60	19,800
1.00	10.0	0.0033	0.60	19,800
Total CVA	—	—	—	70,800

Each row is LGD × EE × ΔPD; the CVA is their sum, USD 70,800. Now add wrong-way risk with $\alpha = 1.4$: every EE becomes 14m, and CVA scales to USD 99,120 — a 40% increase from correlation alone, no change in trade or spread. This single arithmetic is the whole chapter in miniature: exposure × default × loss, summed over time, adjusted for correlation.

193.8.10 Convergence: how many paths is enough?

Monte Carlo estimates carry sampling error that shrinks as the number of paths N grows. For a mean estimated from N independent paths, the standard error falls as $1/\sqrt{N}$: to halve the error you must quadruple the paths. That square-root law governs the whole cost of the engine.

> SE(CVA_hat) ≈ s / √N,   where s = sample std-dev of per-path CVA.
>
> Halving SE ⇒ 4× the paths. This is why variance reduction (§193.8.5) pays.

Monte Carlo error and the square-root-of-N law.

Figure 193.8d — CVA estimate with Monte Carlo standard-error bars as path count rises from 1,000 to 40,000. The estimate settles onto the converged value and the error bars shrink as 1/√N. The desk runs 20,000 paths as the production trade-off between accuracy and overnight compute budget; a reduced-path fallback exists for grid-exhaustion nights (Section 193.10).

◆ Validator checkpoint — convergence evidence

"20,000 paths" is not an answer; it is a claim that must be evidenced. Model validation expects a convergence study showing the estimate is stable and the standard error acceptable at the chosen N, especially in the tail (PFE quantiles converge more slowly than the mean EPE). Keep this study as audit evidence; it is the first thing asked for in a model-approval review.

193.8B The regulatory dimension: accounting CVA vs capital

Everything so far produces the **accounting CVA** — a fair-value reserve under IFRS 13, the price of counterparty credit embedded in fair value. Regulators require a separate and different number: **capital** held against the risk that CVA itself moves. Confusing the two is the classic tell of someone who has learned the mechanics without the regime. This section maps the regulatory landscape a quant or validator must navigate.

The accounting anchor: IFRS 13

IFRS 13 defines fair value as an exit price and requires it to reflect counterparty (and own) credit risk. That is the accounting mandate for booking CVA and DVA at all: a derivative's fair value is not its risk-free value but its value net of the expected loss from counterparty default. The accounting CVA computed in this chapter is what satisfies IFRS 13; it flows to the P&L and balance sheet through the sign-off of Section 193.12.

The capital regimes: Basel III, Basel IV and the CVA charge

Basel III introduced a specific capital charge for CVA risk after the 2008 crisis, when the majority of counterparty-credit losses came not from actual defaults but from CVA mark-to-market moves as spreads widened. Basel IV (the finalised framework) revised the approaches. A bank computes its CVA capital under one of these, in decreasing order of sophistication:

Approach	Full name	Who uses it	Character
SA-CVA	Standardised approach to CVA	Larger books, with approval	Sensitivities-based; rewards hedging
BA-CVA	Basic approach to CVA	Smaller / less complex books	Formula on exposure; limited hedge credit
—	(IMM-CVA withdrawn)	—	Internal-model CVA removed in Basel IV

Note the direction of travel: Basel IV **removed** the internal-model approach for CVA capital, leaving the standardised SA-CVA and basic BA-CVA. SA-CVA is sensitivities-based — it aggregates CVA sensitivities to credit, rates, FX and commodity risk factors under prescribed risk weights and correlations, and it gives capital relief for eligible CVA hedges, which is why desks that actively hedge CVA (Section 193.8 sensitivities) prefer it. BA-CVA is a simpler formula on exposure with limited hedge recognition.

The exposure side: SA-CCR and IMM

CVA capital needs an exposure input — exposure at default (EAD). That comes from one of two counterparty-exposure regimes: the **Standardised Approach for Counterparty Credit Risk (SA-CCR)**, a prescribed formula turning netting-set and collateral terms into EAD; or the **Internal Model Method (IMM)**, where a bank with approval uses its own simulated exposure (the very cube of Section 193.8) to compute effective EPE and hence EAD. The consistency the regulator prizes is that the exposure driving accounting CVA, IMM capital and limits comes from one coherent engine — not three different numbers.

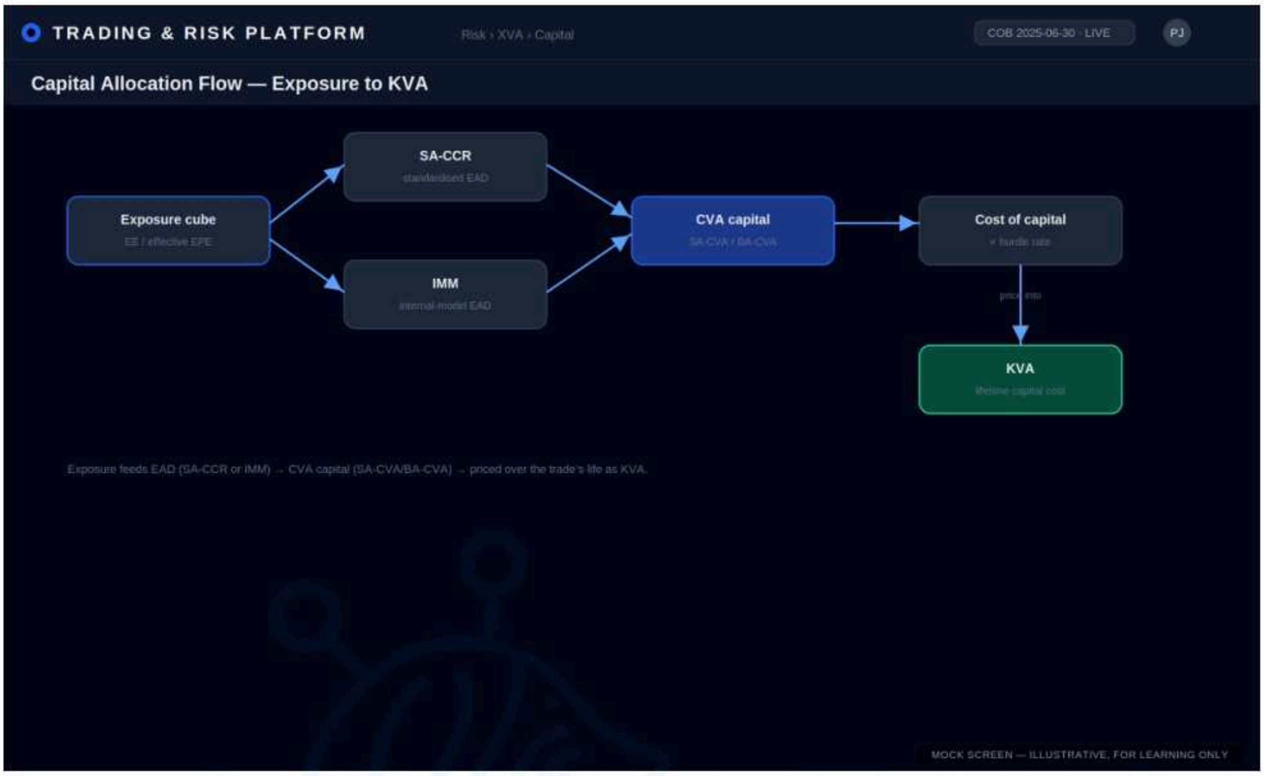

Figure 193.8Bd — Capital allocation flow. The exposure cube feeds EAD (via SA-CCR or IMM), which feeds CVA capital (SA-CVA or BA-CVA), which is priced over the trade's life as KVA. This chain is why hedging and netting, which lower CVA capital, also lower KVA.

Acronym	Governs	Key output	Relationship to this chapter
IFRS 13	Accounting fair value	CVA / DVA reserve	The number this chapter computes
SA-CVA	CVA capital (sensitivities)	CVA risk capital	Uses the sensitivities of §193 (Greeks)
BA-CVA	CVA capital (basic)	CVA risk capital	Formula on EAD; little hedge credit
SA-CCR	Counterparty EAD (standard)	Exposure at default	Prescribed EAD from netting/CSA
IMM	Counterparty EAD (internal)	Effective EPE → EAD	Uses the simulated exposure cube
FRTB-CVA	CVA under the trading-book review	Aligned CVA capital	Ties CVA capital to market-risk framework

IFRS 9 interactions

IFRS 13 mandates the fair-value CVA/DVA this chapter computes, but a second accounting standard interacts with it: **IFRS 9**, which governs impairment (expected credit losses) and hedge accounting. Two interactions matter for an XVA function. First, the same counterparty credit deterioration that widens CVA (a fair-value effect on derivatives) can also drive IFRS 9 expected-credit-loss provisions on any loans or receivables to that counterparty — the institution must avoid double-counting the same credit event across the trading book (CVA) and the banking book (ECL). Second, if the desk designates CVA hedges under IFRS 9 hedge accounting, the hedge effectiveness testing and the treatment of hedge P&L flow through a different part of the accounts than the CVA itself. The practical point for a quant or controller: know which book a credit effect lives in, and ensure the CVA reserve and any ECL provision are reconciled, not additive. (Cross-reference: Ch 381, Accounting & Finance, develops the ledger treatment.)

◆ **Capital optimisation — where quant meets the balance sheet**

Because SA-CVA rewards hedging and recognises netting, and because KVA (Section 193.2) prices the lifetime cost of that capital, the XVA desk actively **optimises capital**: choosing trades, hedges and netting arrangements that minimise CVA capital and therefore KVA. A trade that reduces CVA capital can carry a negative KVA charge — the desk should want it. This is the commercial payoff of getting the regulatory mechanics right, and it is increasingly what enterprise XVA teams are measured on.

193.8C Implementation artifacts: the data and interfaces behind XVA

This chapter emphasises inspectable technology. Here are the concrete artifacts a team building or validating an XVA capability works with daily: the schema that stores exposure and results, the views that serve reporting, the API contract, the streaming messages, and the batch orchestration. Together they are the implementation skeleton of everything above.

Exposure and XVA tables

SQL — core XVA schema (exposure + results), etrm schema

```
CREATE TABLE etrm.xva_exposure_profile (
    run_id          VARCHAR(32)  NOT NULL,
    cob_date        DATE         NOT NULL,
    netting_set_id  VARCHAR(24)  NOT NULL,
    time_years      NUMERIC(8,4) NOT NULL,   -- grid point t_k
    epe             NUMERIC(20,4),           -- expected positive exposure
    ene             NUMERIC(20,4),           -- expected negative exposure
    pfe_95          NUMERIC(20,4),
    eff_epe         NUMERIC(20,4),
    PRIMARY KEY (run_id, netting_set_id, time_years)
);

CREATE TABLE etrm.xva_result (
    run_id          VARCHAR(32)  NOT NULL,
    cob_date        DATE         NOT NULL,
    netting_set_id  VARCHAR(24)  NOT NULL,
    counterparty    VARCHAR(24)  NOT NULL,
    base_ccy        CHAR(3)      NOT NULL,
    cva  NUMERIC(20,4), dva  NUMERIC(20,4),
    fca  NUMERIC(20,4), fba  NUMERIC(20,4),
    colva NUMERIC(20,4), mva NUMERIC(20,4), kva NUMERIC(20,4),
    wwr_alpha       NUMERIC(6,3),
    engine_code     VARCHAR(16),             -- ORE_PYTHON | ORE_EXE
    signoff_status  VARCHAR(12)  DEFAULT 'COMPUTED',
    PRIMARY KEY (run_id, netting_set_id)
);
```

A reporting view: total XVA by counterparty

SQL — view rolling results up to counterparty level

```
CREATE VIEW etrm.v_xva_by_counterparty AS
SELECT  cob_date, counterparty, base_ccy,
        SUM(cva)   AS cva_total,
        SUM(dva)   AS dva_total,
        SUM(fca + fba)            AS fva_net,
        SUM(colva) AS colva_total,
        SUM(kva)   AS kva_total,
        SUM(cva + dva + fca + fba + colva + kva) AS total_xva
FROM    etrm.xva_result
WHERE   signoff_status = 'APPROVED'
GROUP BY cob_date, counterparty, base_ccy;
```

The API contract and a REST payload

The engine is driven and queried over REST. The contract below is the request to launch a run and the shape of the result resource; front-office pre-deal tools and reporting both consume it.

REST — XVA run API contract (OpenAPI-style excerpt)

```
POST /api/v1/xva/runs
  request:  { cob_date, netting_sets[], config_ref, mode: FULL|PREDEAL }
  response: 202 { run_id, status: QUEUED, links: { self, results } }

GET  /api/v1/xva/runs/{run_id}
  response: 200 { run_id, status: RUNNING|DONE|FAILED, progress_pct }

GET  /api/v1/xva/results?run_id=...&netting_set_id=...
  response: 200 {
    "netting_set_id": "NS-HELIOS-01",
    "adjustments": { "cva": -1820344.11, "dva": 512088.40,
                     "fva": -430905.47, "kva": -402551.79 },
    "exposure": { "epe_peak": 12418003, "pfe_95": 28104221 },
    "diagnostics": { "paths": 20000, "se_cva": 41220, "engine": "ORE" }
  }
```

Streaming: incremental results on the bus

As the batch completes each netting set, it publishes a result event so downstream consumers (limits, reporting, the ledger feed) react without polling. A representative message:

Kafka — topic xva.results.v1 (message value, JSON)


```
{
  "event": "XVA_RESULT_READY",
  "run_id": "XVA-20250630-N",
  "cob_date": "2025-06-30",
  "netting_set_id": "NS-HELIOS-01",
  "counterparty": "HELIOS",
  "total_xva": -6040311.00,
  "se_cva": 41220,
  "emitted_at": "2025-07-01T05:12:44Z",
  "schema": "xva.results.v1"
}
```


Batch orchestration as a DAG

The nightly run of Section 193.10 is a directed acyclic graph of stages with explicit dependencies. Expressed as a scheduler DAG:

Batch DAG — nightly XVA (dependency graph)

```
market_snapshot  >>  [trade_extract, scenario_gen]
trade_extract    >>  mtm_baseline
[mtm_baseline, scenario_gen]  >>  exposure_aggregate
exposure_aggregate  >>  collateral_apply
collateral_apply    >>  [cva_dva, fva_colva, kva]
[cva_dva, fva_colva, kva]  >>  datamart_publish
datamart_publish    >>  raise_signoff_workflow

# partitioned by netting_set with on_failure='continue_and_flag'
# so one bad set is quarantined, not fatal to the run
```

193.8C.1 Orchestration and observability

The batch DAG of Section 193.8C is expressed here in an Airflow-style Python definition — the concrete artifact an operations engineer commits — followed by the logging and monitoring signals a production XVA run emits so that failures are seen before the committee, not during it.

Airflow-style DAG definition (nightly_xva.py)

```
with DAG('nightly_xva', schedule='15 22 * * MON-FRI',
         catchup=False, default_args=args) as dag:

    snapshot = PythonOperator(task_id='market_snapshot',
                              python_callable=take_snapshot)
    extract  = PythonOperator(task_id='trade_extract',  ...)
    scen     = PythonOperator(task_id='scenario_gen',   ...)
    expo     = PythonOperator(task_id='exposure_agg',
                              pool='grid', retries=1)
    cva      = PythonOperator(task_id='cva_dva', ...)
    fva      = PythonOperator(task_id='fva_colva', ...)
    kva      = PythonOperator(task_id='kva', ...)
    publish  = PythonOperator(task_id='datamart_publish', ...)
    signoff  = TriggerDagRunOperator(task_id='raise_signoff', ...)

    snapshot >> [extract, scen]
    [extract, scen] >> expo >> [cva, fva, kva] >> publish >> signoff
```

Structured log events + monitoring signals (JSON lines)

```
{ "ts":"2025-07-01T05:12:44Z", "run":"XVA-20250630-N",
  "stage":"exposure_agg", "level":"INFO",
  "netting_sets_done":318, "netting_sets_total":407,
  "grid_nodes_busy":24, "eta_min":18 }

{ "ts":"...", "stage":"cva_dva", "level":"WARN",
  "netting_set":"NS-CASP-02", "msg":"proxy curve used",
  "curve_source":"PROXY", "action":"flagged_for_review" }

# Monitoring dashboard tiles (from these events):
#   run_progress_pct, stage_duration_p95, grid_utilisation,
#   quarantined_netting_sets, cva_se_max, eta_to_cutoff
```

◆ What a monitoring dashboard must surface

The three signals that predict a missed cut-off (Incident 6) are grid utilisation, stage duration versus its historical p95, and ETA to cut-off. A dashboard that shows these lets operations invoke the reduced-path fallback *before* the deadline rather than explaining a miss after it. Observability is not a nicety here — it is the difference between a controlled fallback and a committee with no number.

193.8D Commodity-specific XVA: why the underlying changes everything

This program serves energy and commodity trading as much as capital markets, and commodity XVA has features that rates/FX desks never meet. The exposure profile — and therefore CVA — is shaped by the physical reality of the underlying. Figure 193.8c showed the profiles; here is why each behaves as it does and what it means for the adjustment.

- **Oil swaps (Brent, WTI):** liquid forward curves with mild seasonality; exposure builds with price volatility and amortises as monthly settlements roll off. Wrong-way risk is acute when the counterparty is a producer (exposure rises as the price that stresses them rises).
- **Gas portfolios (Henry Hub, TTF):** pronounced winter/summer seasonality drives an oscillating exposure profile; storage optionality adds path-dependence. Regional basis risk multiplies the number of correlated factors in the cube.
- **LNG contracts:** long-dated, delivery-linked, often oil-indexed or hub-indexed with destination flexibility — a long build-up exposure and complex optionality that stresses the LSM machinery.
- **Power options:** non-storable, so prices are spiky and mean-reverting; exposure is front-loaded and decays fast. Load shape and peak/off-peak blocks add structure the model must respect.
- **Emissions (EUA):** policy-driven monotone drift; exposure tends to rise over the compliance horizon. Correlation to the counterparty's carbon-intensity is a wrong-way channel.
- **Freight & weather derivatives:** short-dated, mean-reverting, illiquid; exposure is spiky and hard to hedge, so residual CVA is proportionally larger.

What commodity XVA demands that rates/FX XVA does not

- **Seasonality** modelled explicitly (monthly buckets), or exposure spikes are missed.
- **Two-factor forward models** (e.g. Gabillon) to separate short-end and long-end volatility.
- **Physical delivery & storage optionality** feeding path-dependent exposure through LSM.
- **Structural wrong-way risk** — the counterparty's business often *is* the underlying.
- **Basis and regional factors** that enlarge the correlation matrix and the cube's factor axis.

193.8D.1 Worked commodity examples

Generic capital-markets XVA training stops at rates and FX. For an energy and commodity desk the instruments themselves change the exposure, so here are concrete worked cases that distinguish this program.

- **Brent / WTI swap:** a fixed-for-floating oil swap settling monthly. Exposure grows with price volatility and the remaining strip, amortising as months settle. If the counterparty is a producer, your in-the-money scenarios (price falls) coincide with their stress — wrong-way risk; if a refiner (price rises stresses them), it can be right-way.

- **Henry Hub gas swap:** the same structure on US gas, but with strong winter/summer seasonality baked into the forward curve, so the exposure profile oscillates rather than decays smoothly. Storage counterparties add optionality.
- **Spark spread:** the joint exposure to power minus the gas-and-carbon cost of generating it — a power-station economics trade. Exposure depends on the *correlation* between power and gas (and EUA) prices, so the correlation matrix Σ of Section 193.8.2 directly drives the CVA. A de-correlation event widens the spread distribution and the exposure.
- **LNG contract:** long-dated, often oil- or hub-indexed with destination flexibility; the optionality (divert cargoes to the best market) is an American-style feature that needs Longstaff–Schwartz (Section 193.8.7) to value the exposure correctly.
- **Freight derivative (FFA):** short-dated, illiquid, mean-reverting; exposure is spiky and hard to hedge, so residual CVA is proportionally large and model risk is high.
- **EUA carbon allowance:** a policy-driven asset with monotone upward drift; exposure tends to rise over the compliance horizon, and correlation to the counterparty's carbon intensity is a wrong-way channel.
- **Renewable PPA:** a long-dated power-purchase agreement with a wind or solar generator. Exposure is shaped by *volume* uncertainty (weather-driven generation) as much as price, so the model must capture a quantity risk that classic financial swaps do not have — a genuinely commodity-specific exposure driver.

◆ The spark-spread and PPA lesson

A spark spread's CVA lives in the correlation matrix; a renewable PPA's CVA lives in weather-driven volume uncertainty. Neither has an analogue in rates/FX XVA. An analyst who models a PPA as a fixed-volume power swap will misstate exposure badly. This is exactly the specialisation an energy trading firm is hiring for — and exactly what a generic course cannot teach.

◆ The commodity XVA analyst's edge

A capital-markets XVA analyst who moves to an energy desk and treats a gas book like a rates book will understate exposure badly — they will miss the seasonality, the storage optionality and the structural wrong-way risk. Commodity XVA is not rates XVA with different tickers; the physics of the underlying is in the exposure profile. This is a genuine specialisation, and a marketable one.

193.8E Model validation and governance for XVA

An XVA model that cannot be validated cannot be used. Enterprise teams increasingly hire specifically for model risk management (MRM) around XVA, and the material below is what those roles do. Validation is independent of the model-building desk and follows a defined lifecycle from challenge to approval, with audit evidence at every step. (Cross-reference: the sign-off workflow that operationalises this is Section 193.12.)

The validation toolkit

Technique	What it does	XVA application
Challenger model	Independent re-implementation	Re-derive CVA a different way; compare
Benchmarking	Compare to a reference/vendor	EE/CVA vs an alternative engine
Sensitivity analysis	Perturb inputs, observe outputs	Stability of CVA to σ, a, α, N
Backtesting	Compare model to realised	Realised exposure vs predicted PFE
Convergence study	Estimate vs path count	Evidence N is sufficient (§193.8.10)
Conceptual soundness	Review of assumptions	Model choice fits the products

Backtesting deserves emphasis: for counterparty exposure, the regulator expects the bank to compare *realised* exposure against the model's predicted quantiles over time. If realised exposure breaches the predicted PFE more often than the confidence level implies, the model is under-conservative and the IMM approval is at risk. Backtesting exceptions are counted, reported and explained — exactly as market-risk VaR is backtested (cross-reference Ch 249, Backtesting Risk Models).

The governance lifecycle

From model build to approved production use

1. **Development & documentation:** the desk documents assumptions, derivations and limitations.
2. **Independent validation:** MRM applies the toolkit above and writes findings, rated by severity.
3. **Remediation:** the desk addresses findings; unresolved ones become documented limitations with conditions.
4. **Model approval:** a model-risk committee approves the model for a defined scope, with an expiry.
5. **Ongoing monitoring:** periodic revalidation, backtesting, and performance monitoring against triggers.
6. **Audit evidence:** every step — the convergence study, backtest results, sign-offs — is retained for internal audit and the regulator.

Audit evidence an XVA function must be able to produce on request

- The convergence study justifying the production path count.
- The correlation-matrix estimation and repair methodology.
- Backtesting results for exposure, with exceptions counted and explained.
- The proxy-curve mapping and rationale for every illiquid counterparty.
- The FVA/DVA overlap treatment and the house funding convention.
- Model-approval minutes, scope, limitations and expiry.

◆ **Why this makes you employable**

Model validation and XVA governance are among the fastest-growing seats in risk. A candidate who can speak fluently about challenger models, exposure backtesting, convergence evidence and the approval lifecycle — not just the CVA formula — is interviewing for a more senior role than one who can only compute. Sections 193.8, 193.8B and 193.8E are written to put you there.

193.8F Living with XVA: sensitivities and hedging

Computing the reserve once a night is necessary but not sufficient. The XVA desk runs a P&L that moves every day as markets move, and it must hedge that P&L. To hedge, it needs **sensitivities**: how much does CVA change for a one-basis-point move in the counterparty's credit spread (credit delta), in interest rates (IR delta), in the commodity forward (commodity delta), in volatility (vega)? These are the same Greeks you met in pricing, but computed on the adjustment itself, and they are numerically expensive because each one requires re-running or bump-and-revaluing the whole exposure simulation.

Figure 193.9 — CVA sensitivities. The headline tiles give the desk its hedge ratios; the ladder breaks credit delta down by tenor so the desk knows where on the CDS curve to hedge. The 2–5 year concentration is typical: that is where the product of exposure and marginal default probability peaks, so that is where CVA is most sensitive to the counterparty's spread.

Credit delta is hedged, where possible, by buying single-name or index CDS protection: if the counterparty's spread widens, CVA rises (a loss), but the protection gains, offsetting it. IR and commodity deltas are hedged in the underlying rates and commodity markets. The practical difficulty is that XVA sensitivities are second-order, path-dependent and cross-gamma-heavy: a move in commodity prices changes not just commodity delta but also the exposure profile and hence credit delta. This cross-gamma is why XVA desks are among the most computationally demanding on the trading floor.

How XVA sensitivities are produced (and why they cost so much)

- **Bump-and-revalue:** shift one input, re-run the simulation, difference the result. Simple, exact, brutally expensive.
- **Pathwise / adjoint (AAD):** propagate derivatives through the simulation analytically — one pass gives all Greeks. Complex to implement, orders of magnitude faster.
- **Grid reuse:** keep the same random paths across bumps so noise cancels in the difference. Essential for stable Greeks.
- **Cross-gamma:** the reason XVA cannot be hedged perfectly — exposure itself moves with the market.

193.9 Pricing the next trade: incremental XVA

So far we have valued an existing book. But the highest-value use of XVA is *before* a trade is done. When a salesperson quotes a new trade, they must know what XVA charge to build into the price — otherwise the desk wins business that destroys value. The key insight is that XVA is **not additive**: because of netting, a new trade's XVA charge depends entirely on what is already in the netting set. A new trade that offsets existing exposure can carry a *negative* incremental charge — the desk should pay to do it.

Incremental XVA is computed as the XVA of the netting set *with* the proposed trade minus the XVA of the set *without* it. This is why the pre-deal tool must re-simulate the whole netting set, not just price the new trade in isolation.

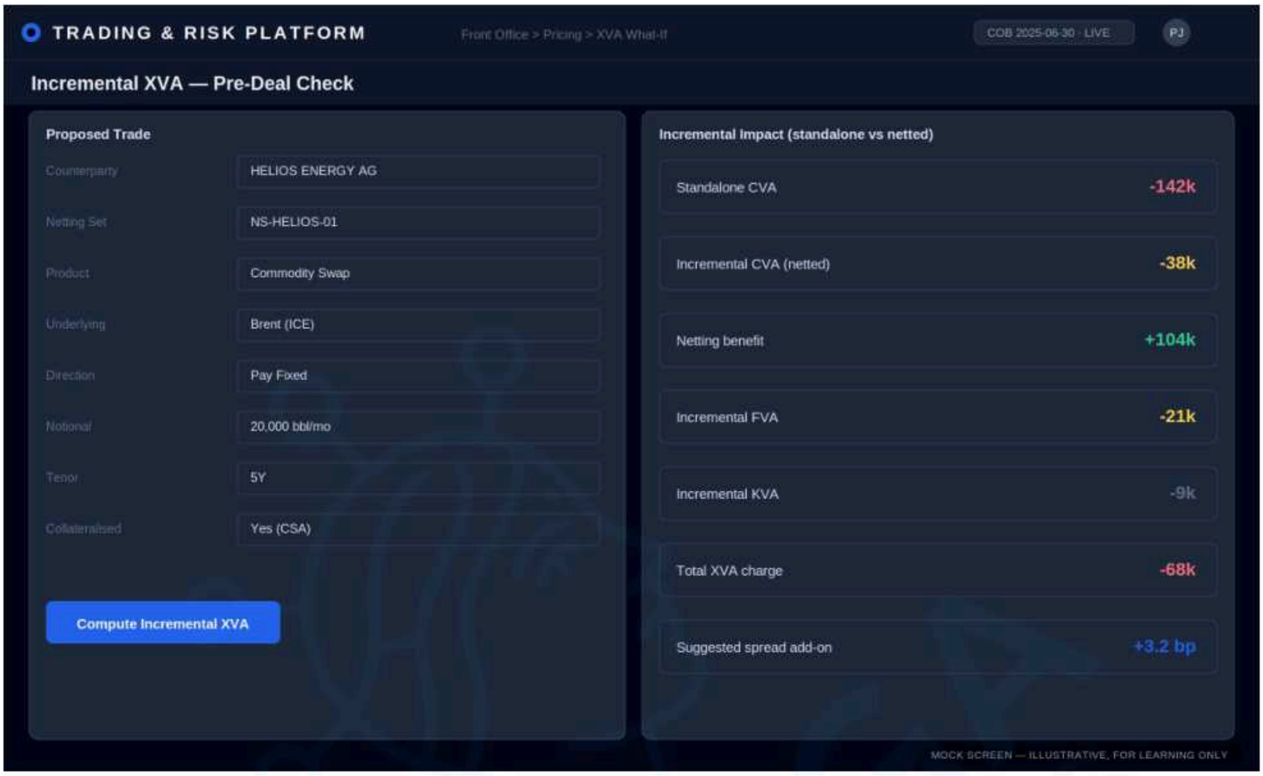

Figure 193.10 — The pre-deal (incremental) XVA tool. A proposed 5-year Brent swap is entered on the left. On the right, its standalone CVA of −142k collapses to an incremental −38k once netting against the existing book is recognised — a 104k netting benefit. The tool converts the total −68k charge into a 3.2 bp spread add-on the salesperson can quote. This is XVA doing its job at the point of sale.

Platform pre-deal request (POST /api/v1/xva/incremental)

```
{
  "mode": "PRE_DEAL",
  "cob_date": "2025-06-30",
  "netting_set_id": "NS-HELIOS-01",
  "baseline_from_cache": "XVA-20250630-N",    // reuse last night’s grid
  "proposed_trade": {
    "product": {
      "type": "COMMODITY_SWAP",
      "direction": "PAY_FIXED",
      "underlying": { "commodity": "BRENT", "market": "ICE" },
      "notional_quantity": { "value": 20000, "unit": "BBL", "freq": "MONTHLY" },
      "termination_date": "2030-06-30"
    }
  },
  "return": ["cva", "fva", "kva", "spread_addon"]
}
```

◆ **Why pre-deal reuses the nightly grid**

Re-simulating 20,000 fresh paths while a salesperson holds the phone is too slow. The pre-deal tool reuses the previous night's scenario grid (BaselineFromCache) and only reprices the incremental trade on those existing paths. This makes the quote come back in seconds instead of minutes, at the cost of using slightly stale market data — an acceptable trade-off for an indicative charge that will be refreshed on the next full run.

193.10 Running it for real: the nightly XVA batch

Everything so far comes together in the nightly batch — the orchestrated end-to-end run that, every night after close of business, produces the official XVA numbers the bank reserves against. This is where the theory meets operations, and where most of the actual pain of running an XVA function lives. The batch is a pipeline: snapshot the market, extract and revalue trades, generate scenarios, aggregate exposure, integrate the adjustments, then publish to the datamart for reporting and sign-off.

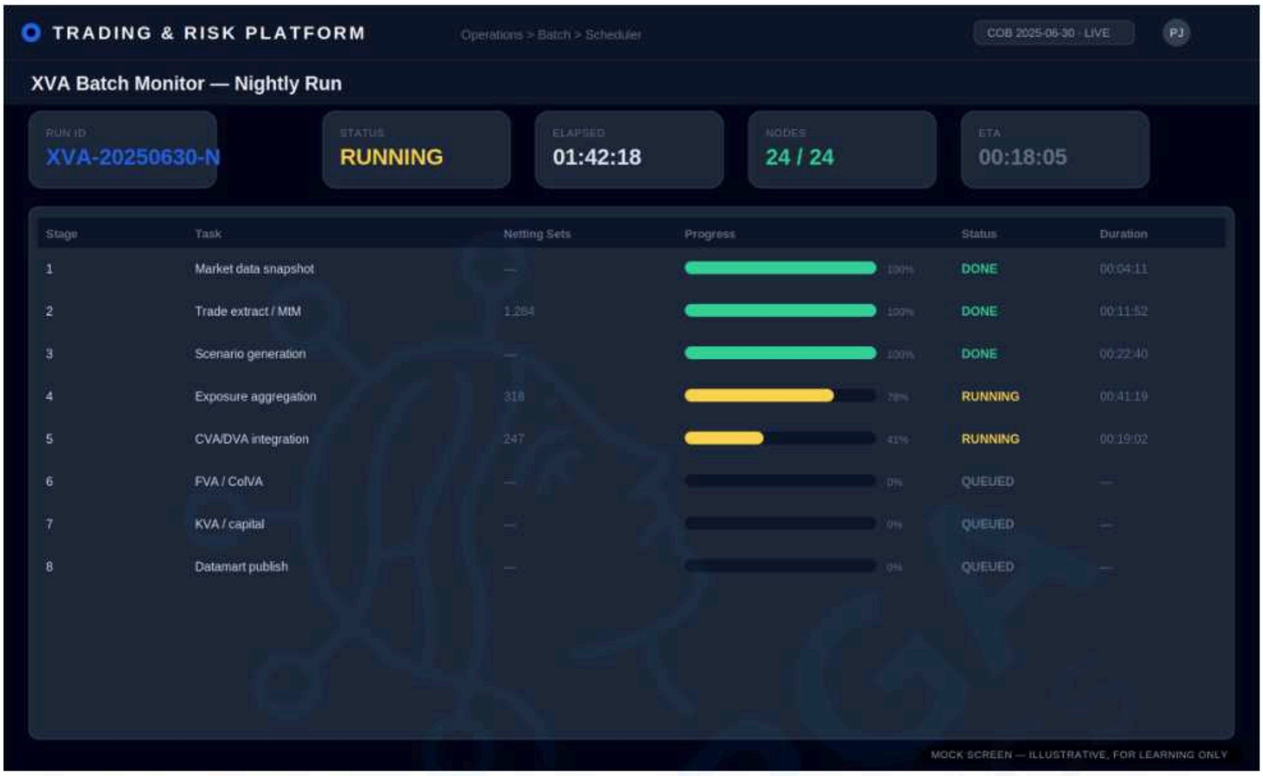

Figure 193.11 — The XVA batch monitor mid-run. Stages execute in dependency order across a 24-node grid. Scenario generation (stage 3) and exposure aggregation (stage 4) dominate the wall-clock time. The operator watches for stalled stages and node failures; a single failed netting set in stage 4 can block the whole downstream integration if not handled with partial-completion logic.

The orchestration definition

The batch itself is declared as a scheduled workflow with explicit stage dependencies, resource requests and failure policies. Below is a trimmed version of that definition — the artifact an operations engineer owns and tunes.

XML — nightly batch orchestration (batch.xva-nightly.xml)

```xml
<Batch id="XVA_NIGHTLY" schedule="0 15 22 * * MON-FRI" tz="Europe/London" grid="24">
  <Stage id="1" name="MarketSnapshot" retries="2">
    <Task ref="snapshot.curves"/>  <Task ref="snapshot.vols"/>
    <Task ref="snapshot.credit"/>
  </Stage>
  <Stage id="2" name="TradeExtract" dependsOn="1" retries="1">
    <Task ref="extract.trades" filter="LIVE" scope="XVA"/>
    <Task ref="mtm.baseline"/>
  </Stage>
  <Stage id="3" name="ScenarioGen" dependsOn="1" parallelism="24">
    <Task ref="mc.generate" config="xva-run.config.xml"/>
  </Stage>
  <Stage id="4" name="ExposureAgg" dependsOn="2,3" parallelism="24"
         partitionBy="NETTING_SET" onPartitionFail="CONTINUE_AND_FLAG">
```

```
    <Task ref="exposure.aggregate"/>
    <Task ref="collateral.apply"/>
  </Stage>
  <Stage id="5" name="CvaDva"  dependsOn="4"><Task ref="xva.credit"/></Stage>
  <Stage id="6" name="FvaColVA" dependsOn="4"><Task ref="xva.funding"/></Stage>
  <Stage id="7" name="Kva"      dependsOn="4"><Task ref="xva.capital"/></Stage>
  <Stage id="8" name="Publish"  dependsOn="5,6,7">
    <Task ref="datamart.publish" table="xva_result"/>
    <Task ref="workflow.raise" process="XVA_RESERVE_SIGNOFF"/>
  </Stage>
  <OnFailure notify="xva-ops@bank" escalateAfter="45m"/>
</Batch>
```

The three failures that break the XVA batch most often

- **Missing market data (stage 1):** a curve or CDS quote did not arrive. The snapshot proceeds with a stale point and every downstream number inherits the error. Fix: hard-fail the snapshot on any missing critical curve, never substitute silently.

- **A trade the engine cannot price (stage 4):** one exotic with a bad or unmapped payoff throws, and without partition isolation it takes the whole grid down. Fix: partitionBy=NETTING_SET with CONTINUE_AND_FLAG so one bad set is quarantined, not fatal.

- **Grid exhaustion / timeout (stage 3–4):** a spike in book size or path count blows the compute budget and the run misses its cut-off. Fix: capacity headroom, path-count governance, and a documented fallback to a reduced-path run.

193.11 Reporting and the reconciliation you will actually do

The batch publishes to the XVA datamart — the reporting layer where the numbers become reports, dashboards and downstream feeds to finance and risk. Your daily reality as a desk or middle-office analyst is querying this datamart to answer one recurring question: *why did the number change, and does it reconcile?* The most common task is reconciling a counterparty-level CVA back to the trades that make it up, and explaining the day-on-day move.

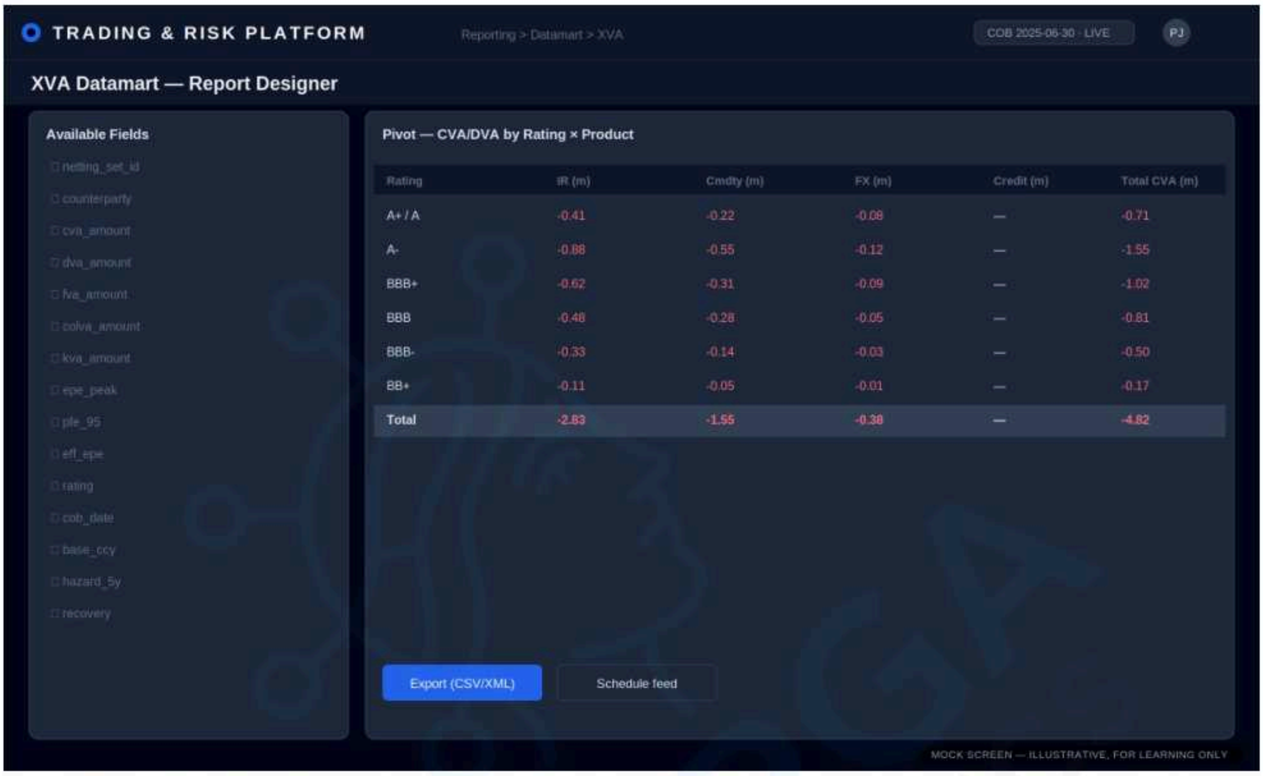

Figure 193.12 — The XVA datamart report designer. Fields from the published result set are pivoted into a CVA breakdown by rating and product. This same store is queried directly in SQL for reconciliation and P&L attribution; the report designer is just a friendly face over the tables below.

The reconciliation query

Here is the workhorse query: pull the published CVA for a counterparty, decompose it against the exposure profile, and compare to the prior day so you can attribute the move. Commit this pattern to memory — you will type versions of it every morning.

SQL — CVA reconciliation and day-on-day attribution

```sql
WITH today AS (
  SELECT netting_set_id, counterparty, cva_amount, dva_amount,
         fva_amount, colva_amount, kva_amount, epe_peak, pfe_95
  FROM   xva_result
  WHERE  cob_date = DATE '2025-06-30'
), prior AS (
  SELECT netting_set_id, cva_amount AS cva_prev, epe_peak AS epe_prev
  FROM   xva_result
  WHERE  cob_date = DATE '2025-06-27'   -- previous business day
)
SELECT  t.counterparty,
        t.netting_set_id,
        t.cva_amount,
        p.cva_prev,
        t.cva_amount - p.cva_prev              AS cva_move,
```

```
        t.epe_peak,
        p.epe_prev,
        round(100.0 * (t.epe_peak - p.epe_prev)
              / nullif(p.epe_prev,0), 1)       AS epe_move_pct,
        (t.cva_amount + t.dva_amount + t.fva_amount
         + t.colva_amount + t.kva_amount)      AS total_xva
FROM    today t
JOIN    prior p USING (netting_set_id)
ORDER BY abs(t.cva_amount - p.cva_prev) DESC;   -- biggest movers first
```

Counterparty	CVA today	CVA prior	Move	Likely driver
HELIOS ENERGY	–1,820,344	–1,655,900	–164,444	Spread widened 11 bp
NORDIC POWER	–1,102,004	–1,180,220	+78,216	Exposure amortised
GULFSTREAM LNG	–740,110	–690,050	–50,060	New swap added
ATLAS TRADING	–510,880	–512,010	+1,130	Immaterial

◆ The reconciliation discipline that catches real errors

A CVA move always decomposes into a market piece (their spread moved), an exposure piece (the profile changed), and a population piece (trades were added, matured or amended). If a move cannot be attributed to one of those three, it is a red flag — usually a data problem, a mis-mapped netting set, or a pricing failure quarantined in the batch. 'Unexplained' is not an acceptable attribution on a desk.

Publishing onward: the report feed definition

Finance needs the signed-off reserve for the ledger; regulatory reporting needs the exposure numbers for capital. Those consumers are fed by scheduled report definitions off the datamart. Here is the finance feed.

XML — report/feed definition (report.xva-ledger.xml)

```
<Report id="XVA_LEDGER_FEED" runAfter="XVA_NIGHTLY.Publish">
  <Query>
    SELECT cob_date, counterparty, netting_set_id,
           cva_amount, dva_amount, fva_amount, colva_amount, kva_amount
    FROM   xva_result WHERE cob_date = :COB AND signoff_status = 'APPROVED'
  </Query>
  <Guard>
    <Rule check="ALL_ROWS_APPROVED" onFail="HOLD_FEED"/>
  </Guard>
  <Output format="ISO20022" schema="ledger-posting"
          target="/feeds/finance/xva_reserve_${COB}.xml"/>
  <Reconcile totalTo="gl.account.XVA_RESERVE" tolerance="0.01"/>
</Report>
```

Note the guard: the feed to finance is **held** until every row carries an APPROVED sign-off status. XVA numbers do not reach the ledger unreviewed — which brings us to the control that closes the loop.

193.12 Governance: from a number to a reserve

A computed XVA is not yet a reserve. Because it feeds the P&L and the balance sheet, it passes through a controlled sign-off workflow with independent price verification and senior risk approval before it can be posted. This is not bureaucracy for its own sake — it is the control that stands between a model output and a number the auditors and regulators will accept.

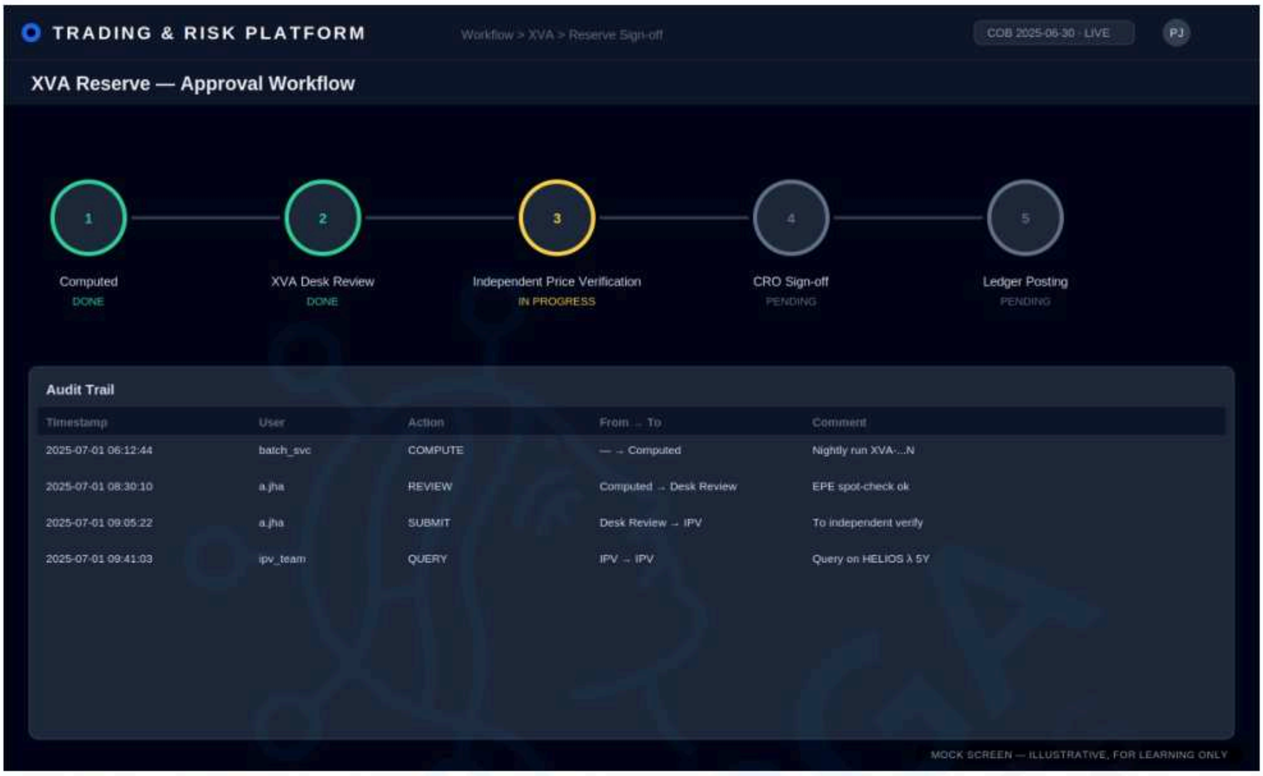

Figure 193.13 — The XVA reserve sign-off workflow with its audit trail. The number moves through desk review, independent price verification (IPV), and CRO sign-off before ledger posting. Every transition is logged with user, timestamp and comment — the evidence an auditor will ask to see. A number stuck in IPV (as shown) cannot post; the feed from the previous section stays held.

What each gate in the sign-off actually checks

- **Desk review:** does the number make sense? Spot-check the biggest movers against market moves.
- **Independent Price Verification (IPV):** an independent function re-derives key inputs (curves, spreads) from independent sources and confirms the result is within tolerance.
- **CRO sign-off:** senior risk accepts the reserve and any model limitations flagged during the run.
- **Ledger posting:** only now does the reserve hit the balance sheet, via the guarded feed.

193.13 The pitfalls that separate juniors from desk-ready analysts

Symptom	Root cause	How you catch it
CVA far too high on a name	Netting flag off; exposure computed gross	Check netting_enforce and legal opinion ref
CVA ignores a live trade	Trade outside blotter/extract filter	Reconcile trade count vs front-office book
Exposure spike missing	Time grid too coarse near a cashflow	Inspect the grid rules around key dates
DVA + FVA double-counts	Overlap not neutralised	Confirm the house funding convention
Collateral over-assumed	Stale/mis-mapped CSA in engine	Diff engine CSA vs signed CSA terms
Unexplained day move	Data/pricing error in batch	Attribute to market/exposure/population
Proxy curve challenged	Sector/rating mapping undocumented	Keep the mapping and rationale on file
Batch misses cut-off	Path count / book size spike	Path governance + reduced-path fallback

◆ The habit that makes you trusted on the desk

For every XVA number you present, be able to state its three drivers — exposure, credit/funding, and population — and reconcile it to the trade level in SQL. Analysts who can do this are handed responsibility quickly. Analysts who quote the datamart number without being able to defend it are not. This chapter has given you every tool to be the former.

193.14 Hands-on lab: price the incremental XVA on a new swap

You will now do end-to-end what the desk does before every trade. Work through the steps; the deliverable is a spread add-on you could quote to a salesperson, plus the SQL that proves your netting benefit.

1. **Set the scene.** Load netting set NS-HELIOS-01 and confirm its CSA terms and netting enforceability. Record the baseline total XVA from the last nightly run.
2. **Author the trade.** Build the canonical trade object for a new 5-year Pay-Fixed Brent swap, 20,000 bbl/month, against NS-HELIOS-01. Validate every field the engine needs (Section 193.6).
3. **Compute standalone XVA.** Submit the trade as its own netting set. Record standalone CVA, FVA, KVA.
4. **Compute incremental XVA.** Submit the incremental request reusing the nightly grid. Record the incremental figures and the netting benefit (standalone minus incremental).
5. **Convert to a spread.** Turn the total incremental charge into a running spread add-on over the trade's notional and tenor. This is your quote.
6. **Prove it in SQL.** Write a query against xva_result showing the netting-set total before and after, and confirm the incremental charge equals the difference.
7. **Sensitivity check.** Bump the counterparty's 5Y CDS spread by 10 bp and re-run. Confirm the direction and rough magnitude of the CVA change against the credit delta from Section 193.8.

Lab success criteria

- Your incremental CVA is materially smaller than standalone — you can state the netting benefit in dollars.
- Your SQL reconciles the incremental charge to the difference of two netting-set totals exactly.
- Your 10 bp bump moves CVA in the right direction and is consistent with the reported credit delta.
- You can defend your spread add-on to a sceptical salesperson in two sentences.

193.15 Section review — check yourself

1. In one sentence each, define CVA, DVA, FVA, ColVA, MVA and KVA, and give the sign of each to your institution.
2. Why is CVA computed per netting set rather than per trade? What single flag can silently destroy the netting benefit?
3. Walk the five steps of exposure simulation from risk-factor models to EPE/ENE/PFE profiles.
4. Given a survival curve and an EPE profile, write the discretised CVA sum and identify each term.
5. Explain wrong-way risk with a commodity example and state one way an engine accounts for it.
6. Name the three settings in the run configuration that can silently corrupt an exposure run.
7. Describe the FVA/DVA overlap and why booking both in full double-counts a benefit.
8. Why does the pre-deal tool reuse the nightly scenario grid, and what is the trade-off?
9. List the three failures that most often break the nightly batch and the config that mitigates each.
10. Describe the sign-off gates between a computed CVA and a posted reserve, and what each verifies.

◆ **Where this chapter sits in the program**

You have now connected pricing (Module 6) to counterparty and market risk (Module 7) through the adjustments that live between them. The next chapters take each XVA term deeper — the capital and regulatory machinery behind KVA, the standardised approach to CVA capital, and the collateral optimisation that minimises ColVA and MVA across a whole book.

193.16 Day on the Desk: seven XVA incidents to investigate

Knowing the theory is not the job. The job is what happens at 08:40 when a number looks wrong and the desk head wants an answer before the CVA committee at 09:30. These three incidents are drawn from the failure modes taught in this chapter. Work each one as an investigation: read the evidence, form a hypothesis, and decide what you would check next *before* reading the resolution. This is exactly how a desk-ready analyst thinks.

Incident 1 — “HELIOS CVA doubled overnight and nobody traded”

CASE FILE #XVA-2291

Reported: 08:41, by the XVA desk head

Symptom: NS-HELIOS-01 CVA moved from –1.82m to –3.61m in one COB

Population: Trade count unchanged (1,284 both days); no amendments logged

Market: HELIOS 5Y CDS moved +3bp only — nowhere near enough to double CVA

Batch: Completed clean, no quarantined netting sets, all stages DONE

Your investigation path

1. Rule out population: trade count and amendment log are flat, so it is not a new or changed trade.
2. Rule out credit: a 3bp spread move cannot double CVA; the credit curve is not the driver.
3. That leaves exposure. Pull the EPE profile for NS-HELIOS-01 today vs the prior COB and overlay them.
4. Observe: today’s EPE is roughly double across the whole curve — a level shift, not a shape change.
5. A uniform doubling of exposure with no market move points at collateral: the CSA is no longer being applied.
6. Check the run config: is ApplyCollateral true? Is the CSA still mapped to the netting set?

■ Resolution

The CSA reference on NS-HELIOS-01 had been re-versioned by a static-data change the previous afternoon, and the new version failed to link (a dangling csa_id). The exposure engine, finding no CSA, computed **uncollateralised** exposure — roughly double. The fix was to repair the CSA link and re-run; the control lesson is that static-data changes to CSAs must trigger an XVA re-validation, not wait for the next nightly surprise. This is the stale/mis-mapped-CSA trap from Section 193.4, seen from the desk.

Incident 2 — “The new swap made CVA go down”

CASE FILE #XVA-2307

Reported: By a sales-trader querying a pre-deal quote

Symptom: Pre-deal tool returned a NEGATIVE incremental CVA on a new Brent swap — a credit

Reaction: Trader assumed a bug: "how can adding risk reduce the charge?"

Trade: New 5Y Pay-Fixed Brent swap into NS-HELIOS-01, which already holds a Receive-Fixed swap

Your investigation path

1. First question: is this actually wrong, or is it correct and counter-intuitive?
2. Look at the existing netting set: it is net long exposure via a Receive-Fixed Brent position.
3. The new trade is Pay-Fixed — the opposite direction — so it offsets existing exposure.
4. Because CVA is computed on NET exposure, reducing net exposure reduces CVA: the credit is real.
5. Confirm with the incremental math: standalone CVA of the new trade is negative-signed, but the netting benefit exceeds it.
6. Explain to the trader: this is netting doing exactly what it should. The desk should welcome the trade.

◆ **Resolution**

Not a bug — a teaching moment. XVA is non-additive because of netting (Section 193.9). A trade that offsets existing exposure carries a negative incremental charge and the desk should be willing to price it keenly. The analyst who can explain this calmly, with the net-exposure argument, earns the trader's trust. The one who 'fixes' the non-bug destroys value.

Incident 3 — "Our CVA and the counterparty's don't match"

CASE FILE #XVA-2318

Reported: During a collateral/valuation dispute with the counterparty

Symptom: Our CVA on the shared netting set is materially larger than theirs

Context: Both sides claim to use the same CDS quotes and the same trades

Stakes: A live dispute; the difference feeds a contested collateral call

Your investigation path

1. Same trades, same CDS — so the difference is in an assumption, not the inputs. Enumerate the assumptions.
2. Recovery rate: confirm both sides use 40%. A different recovery changes LGD and scales CVA directly.
3. Wrong-way risk: are we applying an alpha multiplier they are not? That inflates our EPE relative to theirs.
4. Close-out assumption and margin period of risk: a longer MPOR raises residual exposure and CVA.
5. Compare the two exposure profiles if they will share them; the divergence point localises the cause.
6. Document which assumption differs — that becomes the negotiating position in the dispute.

◆ Resolution

The counterparty was not modelling wrong-way risk; we applied an alpha of 1.4 because HELIOS is an oil producer and our exposure to them rises with the oil price that also drives their default. Our number was arguably *more* correct, but the point of the exercise is that CVA differences between counterparties are almost always assumption differences — recovery, WWR, MPOR, close-out — not input errors. Being able to localise them fast is what resolves disputes.

Incident 4 — "The CDS curve didn't arrive and CVA looks fine"

CASE FILE #XVA-2334

Reported: By the market-data controller during morning checks

Symptom: A counterparty's single-name CDS feed failed overnight, yet its CVA published normally

Concern: "If the curve is missing, why is there a number at all?"

Batch: Completed clean — no missing-data alert raised

Your investigation path

1. Confirm the feed genuinely failed: check the market-data snapshot log for the CDS series.
2. A published CVA with a missing curve means the engine substituted something — find what.
3. Inspect curve_source on the credit_curve_point row: it flipped from CDS_DIRECT to PROXY or STALE.
4. If STALE, the engine silently reused yesterday's curve — the number is real but not today's.
5. If PROXY, it fell back to a sector/rating map without flagging — defensible only if documented.
6. Escalate: the danger is not a wrong number today but a silent substitution nobody signed off.

■ **Resolution**

The snapshot had a soft-fallback rule that reused the prior curve on a missing feed, with the alert mis-configured to WARN rather than QUARANTINE. The CVA was computed on a one-day-stale curve. The fix was procedural, not numerical: missing *critical* curves must hard-fail the snapshot (Section 193.10), never substitute silently. A number you cannot trace to today's market is worse than no number, because it looks fine.

Incident 5 — "The simulation didn't converge"

CASE FILE #XVA-2351

Reported: By the quant on shift, from the run diagnostics

Symptom: CVA for a large callable book swung 18% between two consecutive nightly runs

Market: Curves and spreads barely moved; population unchanged

Diagnostics: Standard error on CVA reported unusually high for this netting set

Your investigation path

1. A large swing with a flat market and flat population points at the estimator, not the inputs.
2. Check the reported Monte Carlo standard error — here it is large relative to the CVA itself.
3. This book is callable, so the Longstaff–Schwartz regression drives exposure (Section 193.8.7).
4. Inspect the regression: too few basis functions, or instability in the deep out-of-the-money tail.
5. Confirm the path count is adequate for the tail — PFE-driven books need more paths than EPE-driven ones.
6. Re-run with more paths and a richer basis; confirm the standard error falls and the swing disappears.

◆ **Resolution**

The callable book's exposure was dominated by a regression that was under-specified in the tail, so the exercise policy flipped between runs on noise. More paths and additional basis functions stabilised the continuation-value estimate and collapsed the swing. The lesson is the convergence discipline of Section 193.8.10: a number without a standard error is a number you cannot defend, and callable books are where convergence bites hardest.

Incident 6 — "The overnight batch missed its cut-off"

CASE FILE #XVA-2360

Reported: By XVA operations at 06:05, ahead of the CVA committee

Symptom: The nightly run was still in exposure aggregation at cut-off; results not published

Trigger: Book size spiked after a large new-business day; path count unchanged

Impact: No signed-off reserve available for the morning committee

Your investigation path

1. Locate the stalled stage on the batch monitor — exposure aggregation, running long.
2. Correlate with the trigger: a jump in netting sets or trades raises the compute per stage.
3. Check grid utilisation — were all nodes engaged, or did a partition failure serialise the work?
4. Decide the immediate action: invoke the documented reduced-path fallback to produce an indicative number.
5. Flag the reduced-path run clearly so the committee knows its accuracy is provisional.
6. Root-cause for tomorrow: capacity headroom, path governance, or partition-isolation config (Section 193.10).

◆ **Resolution**

The new-business spike pushed exposure aggregation past the compute budget. Operations invoked the reduced-path fallback to give the committee an indicative reserve on time, clearly flagged as provisional, and the full-path number followed later that morning. The durable fixes are the three from Section 193.10: capacity headroom, path-count governance, and a documented fallback. Missing a cut-off is an operational failure, not a modelling one — and it is prevented operationally.

Incident 7 — "The reserve sign-off was rejected"

CASE FILE #XVA-2377

Reported: By independent price verification (IPV), blocking ledger posting

Symptom: IPV rejected the CVA reserve for one counterparty at the verification gate

Reason: IPV's independently sourced 5Y CDS spread differed from the desk's by 9bp

Effect: The finance feed is held; the reserve cannot post (Section 193.11)

Your investigation path

1. This is a control working as designed — the number is blocked until the discrepancy is resolved.
2. Compare the two spreads: desk uses consensus mid; IPV uses an independent source — which is right?
3. Check timing: were the two quotes struck at the same snapshot time? A stale mid explains small gaps.
4. Check the curve point: is 9bp at 5Y material to this counterparty's CVA? Quantify the CVA impact.
5. If the desk mid is stale or off-source, re-snapshot and re-run; if IPV's is the outlier, evidence it.
6. Document the resolution in the audit trail — that record is what the regulator and audit will read.

◆ Resolution

The desk's consensus mid was struck fifteen minutes before the official close and had drifted; IPV's close-time source was correct. The desk re-snapshotted, CVA moved within tolerance, and the reserve posted. The point of the incident is that a rejected sign-off is not a failure of the desk — it is the governance of Section 193.12 doing exactly its job. The professional response is to resolve and document, not to argue the gate away.

193.17 End-to-end case study: one trade, every function

Throughout this program we follow a small set of **golden trades** across every discipline, so you see the same deal accumulate meaning as it moves through the institution. This chapter follows **GT-6**, the Brent commodity swap you have already met (TRD-8841021). Here we trace how XVA touches it — and how XVA connects to the functions taught in the chapters before and after this one. When you meet GT-6 again in the settlement and reporting chapters, this is the trade you will recognise.

GOLDEN TRADE GT-6 — Brent Commodity Swap (TRD-8841021)

Economics: Pay-Fixed 78.40 USD/bbl vs Brent front future, 10,000 bbl/month, 5Y

Counterparty: HELIOS ENERGY AG (A–), under NS-HELIOS-01 / ISDA 2002 + CSA

Base MtM at COB: +412,000 USD (clean, risk-free valuation)

Why it matters for XVA: Oil producer counterparty → wrong-way risk; long-dated → material exposure

Function (chapter)	What happens to GT-6	XVA linkage
Capture (Ch 33)	Booked as canonical trade object; netting set + CSA linked	Wrong netting link would balloon CVA
Valuation (Ch 184)	Clean MtM of +412k from the pricing engine	This is the base the XVA waterfall walks from
Market risk (Ch 230)	Contributes commodity delta and vega to the book	Same risk factors drive exposure simulation
Exposure (Ch 193)	EPE peak 12.4m; PFE95 28.1m over 7.2y	The CVA integrand
XVA (Ch 193)	CVA –1.82m share; WWR alpha 1.4 applied	The charge levied on the trade
Collateral (Ch 321)	Daily margined under the CSA; residual = MPOR	ColVA and residual CVA
Settlement (Ch 293)	Monthly cash settlement vs float fixing	Settlement risk in the MPOR window
Accounting (Ch 381)	CVA posted as a fair-value reserve	Hits the P&L and balance sheet

Read the table as a story, not a grid. The same 412k clean value that the pricing engine produces in Module 6 is the leftmost bar of the XVA waterfall in this chapter; the commodity delta the market-risk desk hedges in Module 7 is driven by the same simulated Brent paths that generate this trade's exposure; the residual exposure left after daily margining in Module 9 is what keeps this trade's CVA above zero despite the CSA. XVA is not a silo — it is the connective tissue between valuation, risk, collateral and finance, and GT-6 is how you feel that connection.

◆ The golden-trade habit

When you are learning a new function, ask: *what does it do to GT-6?* Anchoring every abstract concept to a concrete trade you already understand is the fastest way to build the mental model a desk needs — and it is how senior practitioners actually reason: from the deal outward, never from the formula inward.

193.18 Through the validator's eyes: regulator and model validation

Every number in this chapter will be challenged by people whose job is to doubt it: independent model validation (MRM), internal audit, and ultimately the regulator. Learning XVA without learning how it is *challenged* leaves you half-ready. This section puts you in the validator's chair. For each area, know what they ask and what a weak answer looks like.

What model validation challenges

Area	The challenge they raise	What a strong response shows
Proxy curves	"Justify the sector/rating map for illiquid names"	Documented mapping + backtest vs realised
Wrong-way risk	"Is alpha=1.4 conservative or convenient?"	Empirical basis or explicit conservatism
MPOR	"Why 10 days? Show it fits the margin flow"	Consistency with CSA and dispute history
Convergence	"Are 20,000 paths enough for the tail?"	MC error bars; PFE stability analysis
Model choice	"Why Hull-White 1F for rates here?"	Fit to co-terminal swaptions; limitations noted
FVA/DVA overlap	"Prove you are not double-counting"	Explicit framework and overlap treatment

What the regulator cares about

The prudential regulator's lens is capital and consistency, not elegance. For counterparty credit and CVA, the regime distinguishes the **accounting CVA** you reserve (fair value, taught in this chapter) from the **regulatory CVA capital charge** — computed under a standardised approach (SA-CVA) or basic approach — which sizes the capital you must hold against CVA volatility. Likewise, exposure at default for the counterparty-risk charge is set under a standardised approach (SA-CCR) unless an internal model is approved. The regulator's recurring demand is *consistency*: the exposure that drives your accounting CVA, your capital, and your limits should come from one coherent framework, not three different numbers.

◆ **The distinction you must not blur**

Accounting CVA (a fair-value reserve on the balance sheet, the subject of this chapter) and the regulatory CVA capital charge (capital held against CVA risk) are **different numbers computed for different purposes**. Confusing them in an interview or a validation meeting is an immediate tell that someone has learned the mechanics without the context. Know which one you are talking about, always.

193.19 Interview preparation: XVA questions by role

These are the questions actually asked for XVA-adjacent roles, grouped by the seat you are interviewing for. Each comes with the answer a strong candidate gives and the *red flag* that exposes a weak one. Rehearse them out loud — XVA interviews test whether you can explain, not just compute.

For a quant / XVA modelling seat

Q. Why is CVA computed on a netting set rather than trade by trade?
A. Because on default, trades under an enforceable master agreement net against each other, so exposure is the net positive value, not the sum of positive values. Computing per trade ignores the single largest exposure reducer in a real book.
Red flag: *Answering 'to save compute' — it is about legal economics, not performance.*

Q. How does wrong-way risk change the CVA calculation?
A. It breaks the independence assumption between exposure and default. When exposure rises exactly as the counterparty's default probability rises, true CVA is higher than the independent formula gives. You capture it with an alpha multiplier on EPE or by correlating the default intensity to the exposure drivers directly.
Red flag: *Not being able to give a concrete example (e.g. selling oil protection to an oil producer).*

For a risk / counterparty-credit seat

Q. What is the difference between EPE, PFE and effective EPE?
A. EPE is the time-average of expected positive exposure and drives CVA. PFE is a high quantile of exposure and drives credit limits. Effective EPE is a non-decreasing version of EPE used in the regulatory capital formula.
Red flag: *Using the three terms interchangeably — they answer different questions for different consumers.*

Q. A counterparty's CVA jumped but their spread barely moved. Where do you look?
A. Rule out population (new or amended trades), then rule out credit (small spread move can't explain a large jump), which leaves exposure — most often a collateral or CSA mapping problem that made the engine compute uncollateralised exposure.
Red flag: *Jumping straight to 'it must be the spread' without a structured elimination.*

For a middle-office / product-control seat

Q. How do you reconcile a counterparty-level CVA back to the trades?
A. Query the datamart for the netting-set CVA, decompose the day-on-day move into market, exposure and population components, and confirm each trade's contribution ties out. An unexplained move is a red flag for a data or pricing error, not something to sign off.
Red flag: *Treating 'unexplained' as acceptable — on a desk it never is.*

For an ETRM/CTRM implementation-consultant seat

Q. How would you design the trade representation feeding an XVA engine?

A. A clean canonical object holding the economics, netting-set and CSA links, persisted and queryable, with FpML export at the interop boundary — so nothing is locked to a proprietary format and the engine reads one consistent shape. On the platform that is exactly the pattern: canonical JSON in the etrm schema, ORE XML to the engine, FpML at the edge.

Red flag: *Proposing a bespoke undocumented format that only one tool can read.*

193.20 Capstone project: stand up an XVA reserve, end to end

This capstone requires you to **build and defend**, not recall. You will produce working artifacts and a sign-off pack for a small netting set — the same deliverables a real XVA onboarding would demand. Recall-based quizzes prove you read the chapter; this proves you can do the job. Budget four to six hours.

Deliverable 1 — Configuration (buildable artifact)

1. Author the ORE-style run configuration for a netting set of three golden trades (GT-1 rates, GT-3 XCCY, GT-6 Brent).
2. Choose and justify the time grid, path count, and models for each risk factor. Enable collateral and set an MPOR.
3. Enable wrong-way risk for GT-6 only, and state your alpha with a one-line justification.

Deliverable 2 — Computation and reconciliation (must tie out)

1. Run (or hand-compute on a supplied grid) the CVA, DVA and net FVA for the netting set.
2. Write the SQL that reconciles the netting-set CVA to the sum of trade contributions.
3. Produce the day-on-day attribution splitting the move into market, exposure and population.

Deliverable 3 — The sign-off pack (defend it)

1. Assemble a one-page reserve pack: the numbers, the key assumptions (recovery, WWR, MPOR), and known limitations.
2. Write the validator responses to the six challenges in Section 193.18 for YOUR configuration.
3. Prepare a two-sentence defence of your total XVA to a sceptical CRO.

Capstone grading criteria (implementation, not recall)

- Your configuration is internally consistent and every non-default choice is justified.
- Your SQL reconciliation ties the netting-set total to trade contributions exactly — no rounding hand-waving.
- Your attribution explains 100% of the day-on-day move; nothing is left ‘unexplained’.
- Your validator responses survive a follow-up challenge, not just the first question.
- You can defend the WWR and MPOR choices specifically, because those are what a validator attacks first.

◆ What passing this capstone means

If you can configure the run, tie out the reconciliation, and defend the assumptions to a validator, you can sit in an XVA middle-office or desk-support seat and be useful on day one. That is the bar this program holds you to — and it is why an institution pays for it.

193.21 Glossary of abbreviations

The XVA vocabulary is dense with acronyms. Keep this reference to hand; every term below is used in this chapter and recurs across the program.

Term	Expansion	One-line meaning
EE	Expected Exposure	Average exposure across scenarios at a future date
EPE	Expected Positive Exposure	Time-average of EE; the CVA integrand
ENE	Expected Negative Exposure	Mirror of EPE on the negative side; drives DVA/FBA
PFE	Potential Future Exposure	High quantile of exposure; drives credit limits
Eff. EPE	Effective EPE	Non-decreasing EPE used in regulatory capital
MPOR	Margin Period of Risk	Gap between last collateral and close-out
CSA	Credit Support Annex	The collateral agreement under a master agreement
MTA	Minimum Transfer Amount	Smallest collateral block that moves
OIS	Overnight Index Swap	Basis for risk-free discounting
CDS	Credit Default Swap	Instrument whose spread implies default probability
LGD	Loss Given Default	1 – recovery; the loss fraction on default
CVA	Credit Valuation Adjustment	Charge for counterparty default risk
DVA	Debit Valuation Adjustment	Benefit from one's own default risk
FVA	Funding Valuation Adjustment	Cost/benefit of funding the position
FCA / FBA	Funding Cost / Benefit Adj.	FVA on positive / negative exposure
ColVA	Collateral Valuation Adjustment	Economics of collateral actually exchanged
MVA	Margin Valuation Adjustment	Cost of funding initial margin
KVA	Capital Valuation Adjustment	Lifetime cost of regulatory capital held
WWR	Wrong-Way Risk	Exposure and default positively correlated
EAD	Exposure at Default	Regulatory exposure measure for capital
SA-CCR	Std Approach Counterparty Credit Risk	Prescribed EAD from netting/CSA
IMM	Internal Model Method	Own-model exposure for capital (with approval)
SA-CVA / BA-CVA	Std / Basic Approach to CVA	The two CVA capital approaches
AAD	Adjoint Algorithmic Differentiation	All Greeks in one backward pass
SDE	Stochastic Differential Equation	Drift + diffusion dynamics of a factor
LSM	Least-Squares Monte Carlo	Longstaff–Schwartz regression for early exercise

End of Chapter 193 — Valuation Adjustments (XVA)