Climate Alpha Platform: Unified, Explainable

Architecture for Climate-Aware Portfolio Construction

Author: Avaneendra Trivedi (placeholder—verify affiliation before submission)

Affiliation: Thapar Institute of Engineering & Technology (placeholder)

Specification Edition — Document Control

This document is a specification-grade enhancement of the original Climate Alpha Platform manuscript. It preserves the full original structure and content while strengthening definitions, adding normative requirements, formal derivations, worked numerical examples, acceptance tests, and audit-ready governance artifacts. The intended use is (i) implementable system specification, (ii) reproducible research appendix, and (iii) auditable model-risk documentation.

Document Title	Climate Alpha Platform: A Unified, Explainable Architecture for Climate-Aware Portfolio Construction
Document Type	Full Technical Specification (Architecture + Math + Governance + Verification)
Version	v12 (exhaustive, spec-hardened)
Date	January 2026
Audience	Quant researchers, portfolio engineers, model validators, risk governance stakeholders
Normative Language	MUST/SHALL/SHOULD/MAY per RFC 2119 convention
Integrity	Includes traceability + acceptance tests + audit hooks
Distribution	Preprint / Public (no confidential data; theory-first)

Change Log

Version	Date	Author	Summary
Source	Nov 2025 (from original PDF)	Avaneendra Trivedi	Original manuscript and architecture
v12	Jan 2026	Regenerated by ChatGPT	Spec hardening: TOC, traceability, proofs, worked calculations, loophole controls, extracted figures
Next	TBD	—	Affiliation/author verification, code repository link, dataset registry, and validation pack

Normative Keywords

The key words “MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”, “SHOULD”, “SHOULD NOT”, “RECOMMENDED”, “MAY”, and “OPTIONAL” in this document are to be interpreted as normative requirements. The specification is written so that implementation and validation teams can directly derive acceptance tests from MUST/SHALL statements.

Glossary

Term	Definition
Climate Alpha	Residual expected return orthogonal to conventional risk factors, attributable to spatial climate signal.
Footprint	Issuer-level exposure vector to climate hazards/risks derived from facility-level mapping and aggregation.
Data Spine	Governed, versioned, point-in-time data layer joining market, issuer, facility, and hazard objects.
QUBO	Quadratic Unconstrained Binary Optimization formulation used for discrete portfolio selection.
Model Risk	Risk of adverse consequences from incorrect or misused models; requires governance, validation, and monitoring.

Table of Contents

Specification Edition — Document Control 2

Change Log 2

Normative Keywords 3

Glossary 3

Architecture Overview 8

Requirements Traceability Matrix 9

Abstract 13

1. Introduction 13

1.1. Climate risk as a priced but poorly engineered dimension 13

1.2. Gaps in current practice 14

1.3. The Climate Alpha Platform at a high level 15

1.4. Formal scope and non-scope 15

1.5. Contributions in detail 16

2. Problem Context and Design Requirements 17

2.1. Market under-reaction to physical and transition risk 17

2.2. Tooling fragmentation and operational friction 17

2.3. Governance deficit and regulatory pressure 18

2.4. Functional requirements for a climate-alpha stack 19

2.5. Non-functional requirements and design principles 20

3. Institutional and Regulatory Context 20

3.1. Prudential model-risk guidance (SR 11-7–type regimes) 20

3.1.1. Model definition and inventory requirements 20

3.1.2. Conceptual soundness and developmental evidence 21

3.1.3. Ongoing monitoring and performance thresholds 21

3.1.4. Independent validation and effective challenge 22

3.2. Climate-related disclosures and sustainable-finance regulation 22

3.2.1. TCFD and IFRS S2 22

3.2.2. EU SFDR, ESRS, and taxonomy alignment 23

3.2.3. Client and stakeholder expectations 23

3.3. AI-specific frameworks and explainability obligations 24

3.3.1. Classification as high-risk AI system 24

3.3.2. Explainability at three abstraction levels 25

3.4. Data lineage, reproducibility, and cryptographic guarantees 25

3.4.1. Lineage semantics 25
3.4.2. Replay and independent reconstruction 26
3.4.3. Cryptographic attestation 26

3.5. Consolidated design constraints 27

4. Data Spine and Lineage 28

4.1. Data domains and objects 28

4.2. Point-in-time schema and temporal semantics 29

4.3. Facility registry and hazard rasters 30
4.3.1. Facility registry structure 30
4.3.2. Hazard raster catalogue 30
4.3.3. Facility–hazard intersection semantics 31

4.4. Feature engineering on the spine 31

4.5. Lineage model and /audit artefacts 31
4.5.1. Lineage metadata per object 31
4.5.2. Run-level lineage bundles 32
4.5.3. /audit endpoint behaviour 32

4.6. Cryptographic integrity and PQC signing 32
4.6.1. Hashing and signing process 32
4.6.2. Verification semantics 33
4.6.3. Integration into KPIs and controls 33

5. Methodology 33

5.1. Notation and universes 33

5.2. Factor residualization and definition of climate alpha 33
5.2.1. Cross-sectional factor model 34
5.2.2. Residual alpha and its estimation 34

5.3. Climate-aware signal estimation 34
5.3.2. Bayesian temporal smoothing 34
5.3.3. Multiple-testing and signal-quality control 34

5.4. Geospatial risk mapping (method level) 34
5.4.1. Facility-level hazard computation 34
5.4.2. Issuer-level hazard aggregation 34
5.4.3. Portfolio-level footprints 34

5.5. Portfolio construction: objectives and constraints 34
5.5.1. Mean–variance with factor and climate constraints 34
5.5.2. ES/CVaR formulation 35
5.5.3. Tracking-error-constrained climate alpha 35

5.6. Discrete selection as QUBO and quantum-inspired optimisation 35
5.6.1. Binary selection variables and energy function 35
5.6.2. Simulated bifurcation / annealing-class solvers 35
5.6.3. Classical MILP/QP parity 35

5.7. Explainability layer 35
5.7.1. Kernel-SHAP for climate-alpha signals 36
5.7.2. Constraint duals and binding-set diagnostics 36
5.7.3. Hazard-to-exposure-to-portfolio path 36
5.8. Consistency, robustness, and ablations 36
6. Governance, Security, and Compliance 36
6.1. SR 11-7 model-risk framework 36
6.1.1. Model inventory and classification 36
6.1.2. Conceptual-soundness documentation 37
6.1.3. Developmental evidence and validation packs 37
6.1.4. Monitoring, thresholds, and change control 38
6.2. EU AI Act and AI-governance alignment 39
6.2.1. AI-system boundary and classification 39
6.2.2. Technical documentation (Annex IV–style) 39
6.2.3. Logging, traceability, and risk management 39
6.2.4. Human oversight mechanisms 40
6.3. Security architecture 40
6.3.1. Data in transit and at rest 40
6.3.2. Access control and segregation of duties 41
6.3.3. Logging and tamper-evident event trails 41
6.4. Post-quantum signatures and trust guarantees 41
6.4.1. PQC scheme selection and key management 41
6.4.2. Signing of lineage bundles and reports 42
6.4.3. Verification and external assurance 42
6.5. Governance roles, RACI, and operating cycle 42
6.6. Documentation and disclosure mapping 43
7. Proof-of-Value Protocol 44
7.1. Objectives and evaluation philosophy 44
7.2. Universe construction and data windows 44
7.3. Strategy ladder (ablation design) 45
7.4. Costs, liquidity, and realistic constraints 46
7.5. Metrics and definitions 46
7.6. Statistical tests and multiple-testing controls 46
7.7. Decision gates 47
7.8. “No-cheat” constraints 47
7.9. Artefacts and reproducibility 48
8. Implementation Blueprint 48

9. Risks and Mitigations 52
9.1. Model risk 52
9.2. Data risk 53
9.3. Operational and governance risk 54
9.4. Regulatory and legal risk 54
10. Limitations and Scope 55
11. Conclusion 56
12. References 57
Appendix A. Governance Anchors and Publication-Grade Credibility Controls 59
Appendix B. Formal Derivations and Correctness Proofs 60
B.1 Residualisation as an orthogonal projection 60
B.2 QUBO mapping for cardinality-constrained selection 60
Appendix C. Worked Numerical Examples (Step-by-Step Calculations) 61
C.1 Cross-sectional residualisation (OLS) — worked example 61
C.2 CVaR (Expected Shortfall) — worked example 62
Appendix D. Loophole Analysis and Specification Hardening 63
Factor leakage 63
Look-ahead / leakage 63
Coverage drift 63
Constraint gaming 63
Scenario misuse 63
Appendix E. Figures Extracted from the Original Manuscript 64

Architecture Overview

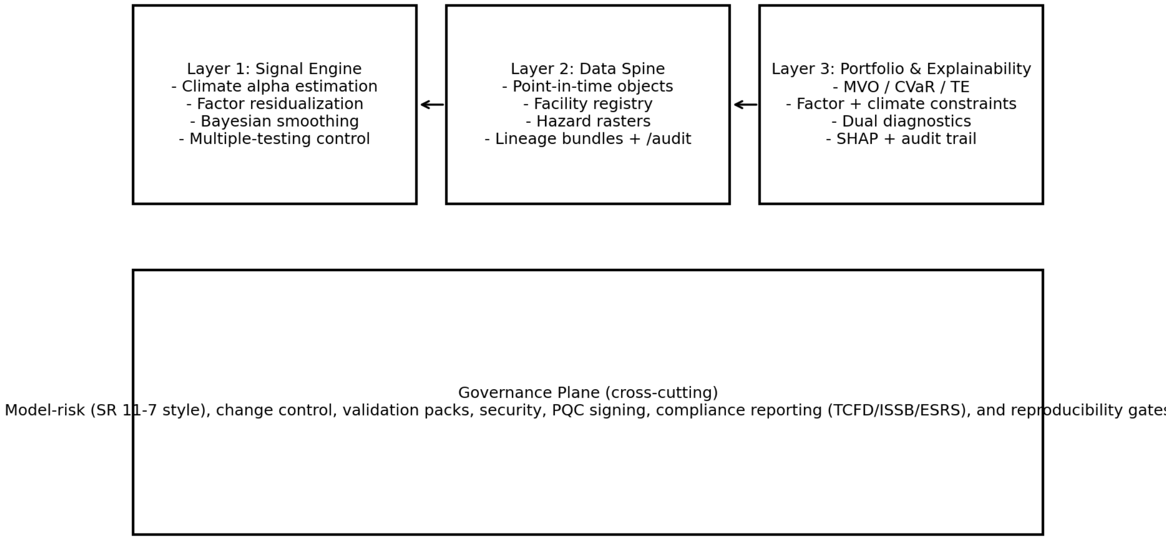

Figure 1: Climate Alpha Platform — architecture overview (data spine + alpha stack + governance stack).

Requirements Traceability Matrix

This matrix is generated from the Functional Requirements section (if present) and is intended to be completed during implementation. It provides end-to-end traceability from requirement to module/component and verification evidence.

Req ID	Requirement	Primary Module	Verification Method	Evidence Artifact
FR-001	From the above, the platform’s core functional requirements can be stated explicitly. A satisfactory climate-alpha stack must:		Unit / Integration / Backtest / Review	
FR-002	Generate factor-neutral, interpretable signals.		Unit / Integration / Backtest / Review	
FR-003	Start from a recognised multi-factor baseline.		Unit / Integration / Backtest / Review	
FR-004	Project out factors and operate on residual returns.		Unit / Integration / Backtest / Review	
FR-005	Attribute predicted residual alpha to specific climate, geospatial, and news features with		Unit / Integration / Backtest / Review	
FR-006	confidence bands.		Unit / Integration / Backtest / Review	
FR-007	Aggregate climate risk from facility to portfolio.		Unit / Integration / Backtest / Review	
FR-008	Maintain a facility registry		Unit / Integration /	

	with polygons/points and metadata.		Backtest / Review	
FR-009	Intersect hazards rasters with facility geometries under multiple scenarios.		Unit / Integration / Backtest / Review	
FR-010	Aggregate exposures to issuer and portfolio levels with documented weights.		Unit / Integration / Backtest / Review	
FR-011	Construct portfolios under joint financial and climate constraints.		Unit / Integration / Backtest / Review	
FR-012	Optimise for expected residual alpha subject to variance, ES/CVaR, liquidity, and transaction-cost		Unit / Integration / Backtest / Review	
FR-013	constraints.		Unit / Integration / Backtest / Review	
FR-014	Enforce factor neutrality and climate-footprint caps simultaneously.		Unit / Integration / Backtest / Review	
FR-015	Provide both discrete (name selection) and continuous (weights) control with deterministic		Unit / Integration / Backtest / Review	
FR-016	fallbacks.		Unit /	

			Integration / Backtest / Review	
FR-017	Expose explainability and diagnostics.		Unit / Integration / Backtest / Review	
FR-018	SHAP-style per-asset attributions for signals.		Unit / Integration / Backtest / Review	
FR-019	Constraint duals and binding sets for portfolios.		Unit / Integration / Backtest / Review	
FR-020	Hazard-to-exposure paths for climate footprints.		Unit / Integration / Backtest / Review	
FR-021	Produce governance-grade artefacts.		Unit / Integration / Backtest / Review	
FR-022	Model inventory entries for each component.		Unit / Integration / Backtest / Review	
FR-023	Development and validation documents suitable for SR 11-7.		Unit / Integration / Backtest / Review	
FR-024	EU AI Act technical annexes and transparency notes.		Unit / Integration / Backtest / Review	
FR-025	Lineage bundles with cryptographic signing.		Unit / Integration / Backtest / Review	
FR-026	Without all five classes of functionality, a system may be able to		Unit / Integration / Backtest / Review	

	demonstrate attractive backtests but will fail either regulatory scrutiny or operational robustness.			

Abstract

This paper introduces a governed, explainable architecture for climate-aware portfolio construction that integrates three tightly coupled layers. First, a signal engine estimates cross-sectional expected-return differentials ($\Delta\mu$) that are residual to standard multi-factor models, with climate, geospatial, and news features entering only in the residual space. Second, a geospatial risk layer maps facility-level hazard exposures (flood, wildfire, heat, wind, drought, coastal and fluvial/pluvial risk) into issuer- and portfolio-level footprint vectors under alternative climate scenarios. Third, a constructor performs quantum-inspired discrete selection on a cardinality-constrained QUBO and then refines weights via convex mean–variance and ES/CVaR optimisation under full classical MILP/QP parity, enforcing factor, liquidity, turnover, and climate-footprint constraints exactly. The platform exposes per-asset feature attributions (kernel-SHAP), per-constraint duals and binding sets, and cryptographically signed data lineage compliant with SR 11-7 model-risk guidance and emerging EU AI transparency obligations. I formalize the residualization operator, regularized and Bayesian signal estimation, the hazard-to-exposure mapping, the QUBO energy function, and the convex refinement problems. I then specify a proof-of-value protocol with train/walk-forward design, ablation ladder, and target performance gates (Information Ratio uplift of 0.15–0.30 vs a factor-neutral baseline and Max Drawdown reduction of 10–20%). The paper is theory-first: no code or empirical results are released, but the architecture, mathematical formulations, and evaluation protocol are sufficient for immediate implementation and pilot deployment by practitioners.

Keywords: Climate finance; Portfolio optimization; Explainable AI; Geospatial risk; QUBO; Conditional Value-at-Risk; Model risk management; Data lineage JEL Codes: G11, G12, G14, C61, C63, Q54

1. Introduction

1.1. Climate risk as a priced but poorly engineered dimension

Physical and transition risks from climate change are now explicitly recognized as financially material. Facility-level hazards such as riverine and coastal flooding, wildfire, heat waves, windstorms, drought and low-flow events, and storm surge alter firms' operating costs, capital expenditure, insurance premia, and supply-chain reliability. Policy and technology shocks—carbon pricing, technology substitution, litigation, and consumer preference shifts—add a second transition channel. These shocks propagate into cash flows, discount rates, and expected defaults over multi-decade horizons. Capital markets partially reflect these risks in prices, but the transmission is noisy and slow. Climate exposures are heterogeneous across firms, facilities, and supply chains; disclosures remain incomplete; and vendor products often compress rich physical and transition information into a single opaque "ESG" or "climate score". As a result, a large fraction of climate risk is either conflated with standard style factors (value, quality, sector, size) or treated as a marketing overlay rather than a systematic input into portfolio optimisation.

This mismatch—material risk, weak engineering—is the starting point for the Climate Alpha Platform. The working hypothesis is that there exists a persistent climate-risk mispricing wedge: a difference between the compensation investors demand for bearing physical and transition risk, and the compensation embedded in current prices once standard factors are controlled for. Capturing this wedge requires infrastructure that is simultaneously:

- Factor-neutral by construction, so that any excess performance is not merely repackaged exposure to traditional risk premia.
- Granular and spatial, operating at facility or asset-level rather than only at issuer averages.
- Explainable and auditable, so that risk, compliance, and regulators can interrogate every step from raw data to final portfolio.
- Implementable under constraints, producing portfolios that meet liquidity, turnover, and climate-footprint caps with deterministic fallbacks.

The Climate Alpha Platform is designed as such an infrastructure layer rather than as a single trading strategy.

1.2. Gaps in current practice

Three gaps motivate the architecture.

1. Market under-reaction and conflation with style factors. Empirical climate-finance studies show non-zero climate risk premia, but signals are contaminated by industry, region, and standard style tilts. Without strict residualization, "climate alpha" is often indistinguishable from a value, quality, or sector bet. Existing products rarely demonstrate factor-neutral out performance under robust tests; marketing materials highlight raw returns, not residual alpha net of multi-factor models.
2. Tooling fragmentation and black boxes. The production stack is fractured across vendors and internal teams. Climate data is purchased as one set of feeds, alpha models are maintained separately, and portfolio construction is another silo. Integration typically consists of bolting a climate screen or tilt onto an existing optimizer. Vendor scores are often non-reproducible black boxes with limited documentation and incomplete coverage. This fragmentation makes it difficult to embed climate awareness deep in the alpha engine and optimisation layer, and almost impossible to provide line-by-line audit trails.
3. Governance and documentation deficits. Regulators expect SR 11-7-grade documentation, TCFD/IFRS-aligned climate disclosures, and, under the EU AI Act, traceable and explainable AI behavior. Most ESG-labelled offerings do not ship with a comprehensive model inventory, lineage documentation, validation packs, or effective-challenge artefacts. Claims of "ESG alpha" often collapse once standard factors and transaction costs are accounted for, and there is no deterministic, constraint-clean fallback when new methods fail. These gaps define the design problem: construct a governed climate-alpha stack that is

integrated, transparent, and implementation-ready, not a loose federation of data feeds and heuristics.

1.3. The Climate Alpha Platform at a high level

The Climate Alpha Platform is a tri-stack architecture with a shared data spine and governance envelope:

4. Signal Engine (XAI). A cross-sectional expected-return engine that produces factor-neutral return differentials ($\Delta\mu$) by applying a residualizer to standard multi-factor models and then regressing those residual returns on climate, geospatial, and news features. It outputs both point estimates and uncertainty, together with kernel-SHAP style attributions per asset.
5. Geospatial Risk Layer. A spatial-finance module that maps facility polygons and points to global hazard rasters (flood, drought, wildfire, heat, storm) and aggregates exposures from facility to issuer to portfolio under scenario shocks. It yields exposure vectors and Jacobians that feed both the signal engine and optimisation constraints.
6. Constructor. A discrete-plus-continuous portfolio construction engine. It performs quantum-inspired discrete selection over a QUBO energy function encoding alpha, risk, factor neutrality, and climate-footprint penalties, then refines weights via classical mean–variance or ES/CVaR optimisation. A MILP/QP parity path guarantees zero constraint violations with identical constraint sets.

Each layer writes to a provenance-rich data model and is wrapped by governance, audit, and API services. The platform is not a monolithic application: it is a specification for how to build a climate-aware signal-to-portfolio pipeline that can survive scrutiny from quantitative researchers, risk controllers, and supervisors.

1.4. Formal scope and non-scope

The paper is deliberately theory and architecture-first. It provides:

- Mathematical formulations for residualization, climate-aware signal estimation, hazard-to-exposure mapping, QUBO energy, and convex refinement.
- Logical design of APIs, lineage, and governance artefacts sufficient for a regulated institution to implement.
- A standardized proof-of-value protocol guiding empirical pilots.

The paper does not provide:

- Production code, parameter estimates, or empirical backtest results.
- Hardware-specific implementations (quantum annealers, GPU clusters).
- A full survey of every climate metric or vendor.

Omission of code and experiments is intentional: the target is a working-paper-grade conceptual and architectural reference that can be implemented independently by multiple institutions, each with its own data vendors, governance processes, and risk appetites.

1.5. Contributions in detail

Relative to both academic literature and existing products, the Climate Alpha Platform contributes:

- A governed, factor-neutral climate-alpha definition. Expected returns are defined on residual returns after projecting out standard factors; climate features enter only in this residual space. This eliminates the most common failure mode in ESG strategies—uncompensated style risk disguised as climate behavior.
- A facility-level hazard framework integrated into optimization. Hazard rasters, facility geometries, and issuer mappings are treated as first-class mathematical objects. Portfolio-level climate footprints appear both as constraints and as drivers of signal features, rather than as post-hoc labels.
- A quantum-inspired, regulator-compatible constructor. The use of QUBO and simulated bifurcation/ annealing allows exploration of high-dimensional discrete combinations under cardinality and exposure constraints. Classical parity via MILP/QP ensures every production weight vector has an exact, explainable, deterministic counterpart.
- Reg-grade governance and lineage baked into architecture. Model inventory, SR 11-7

documentation, EU AI Act technical annexes, and post-quantum signatures are not afterthoughts; they are explicit components of the design and outputs (e.g., audit packs, lineage bundles). This level of integration—signals, spatial risk, optimization, and governance—is, to the best of the author's knowledge, absent from existing open literature and vendor disclosures, which typically treat these as separate products. The rest of the paper is organised as follows. Section 2 formalises the problem context and derives functional and non-functional design requirements for a climate-alpha stack. Section 3 maps prudential, climate-disclosure, and AI-governance regimes into concrete architectural constraints. Section 4 defines the bitemporal data spine, facility registry, and hazard–exposure lineage that all components share. Section 5 develops the mathematical core of the platform: residualisation, climate-aware signal estimation, geospatial exposure integration, and the QUBO-plus-parity constructor. Section 6 describes the governance, monitoring, and cryptographic attestation layer. Section 7 specifies the proof-of-value protocol that institutions can use to test the architecture on their own data. Section 8 sets out the implementation blueprint and API surface.

Section 9 catalogues model, data, and governance risks together with mitigants. Section 10 states the limitations and scope of the current specification. Section 11 concludes, and Section 12 lists references.

2. Problem Context and Design Requirements

2.1. Market under-reaction to physical and transition risk

Physical risk operates at facility granularity. A firm with diversified facilities may experience highly uneven exposure: one plant in a floodplain, another in a wildfire-prone region, others in relatively safe locations. Portfolio-level risk depends on the joint spatial distribution of these assets across all holdings. Legacy ESG metrics compress this complexity into coarse issuer-level scores, usually based on self-reported emissions and generic location proxies. This ignores:

- Intra-issuer heterogeneity (different facilities in different hazard regimes).
- Non-linearities in hazard impact (threshold effects, correlated extremes).
- Path dependence (transition plans, adaptation investments).

Transition risk further amplifies spatial effects: carbon pricing, zoning regulations, and technology bans often apply at regional or national levels; firms with assets concentrated in affected jurisdictions face sharper shocks. Empirically, pricing signals derived from such complex, high-dimensional risk drivers are noisy. They co-move with sector bets (e.g., under-weighting traditional energy), style factors (e.g., quality vs junk), and macro factors (e.g., exposure to real rates via utilities and infrastructure). Without a clean separation between climate-specific effects and standard premia, investors cannot distinguish signal from noise. The platform addresses this by:

- Modeling climate exposure at the facility level and aggregating to issuers and portfolios with

explicit weights.

- Using residualization to remove factor contributions before fitting climate-aware alphas.
- Treating climate exposure and sensitivity as structured features with their own priors and

uncertainty, not as scalar add-ons.

2.2. Tooling fragmentation and operational friction

In a typical asset-management stack, climate data and analytics live in separate silos:

- A vendor API delivers hazard scores or ESG ratings.
- Quant teams ingest these into a data warehouse with limited tracing back to raw rasters or facility lists.
- Alpha models operate on factor, price, and fundamental data, rarely calling climate data except as occasional filters.
- Portfolio construction engines accept a limited set of “extra constraints” and cannot reason about the full structure of climate-driven exposures. Consequences:

- Latency and brittleness. Changes in vendor methodology propagate slowly and opaquely; historical time-series are often restated without clear attachment to a specific methodology version.
- Limited explainability. When performance deviates from expectations or when regulators ask for root-cause analysis, teams struggle to reconstruct which climate inputs influenced which trades.
- Duplicated governance. Every team maintains its own lineage and documentation, if any; there is no single "source of truth" for how climate information enters decisions.

The Climate Alpha Platform's data spine, lineage model, and API surface are designed explicitly to dissolve this fragmentation. Signals, hazards, exposures, and portfolios are all built atop the same schema and provenance framework, allowing consistent reasoning and audit.

2.3. Governance deficit and regulatory pressure

Regulators are not only asking whether climate risk is considered; they are starting to inspect how. Under SR 11-7, a climate-aware portfolio model is subject to the same expectations as any other risk or pricing model:

- Conceptual soundness with clear documentation.
- Developmental evidence showing robustness and performance.
- Ongoing monitoring and definable performance thresholds.
- Independent validation and challenger models.

Under IFRS S2, banks and asset managers must disclose climate-related risks, metrics, and targets with sufficient granularity and reproducibility. The EU AI Act adds requirements on transparency, logging, risk management, and human oversight for AI in high-risk applications. Existing ESG offerings fall short:

- They provide summary reports without underlying code, data lineage, or replication packs.
- They rarely expose constraint duals, binding sets, or optimisation logs.
- They are not shipped with living documentation that evolves with regulations and models.

This governance deficit is not an accessory concern; it is a structural barrier to deploying climate-aware strategies at scale in systemically important institutions. The architecture presented here treats governance, lineage, and compliance as first-class design objects, not post-hoc paperwork.

2.4. Functional requirements for a climate-alpha stack

From the above, the platform's core functional requirements can be stated explicitly. A satisfactory climate-alpha stack must:

7. Generate factor-neutral, interpretable signals.

- Start from a recognised multi-factor baseline.
- Project out factors and operate on residual returns.
- Attribute predicted residual alpha to specific climate, geospatial, and news features with confidence bands.

8. Aggregate climate risk from facility to portfolio.

- Maintain a facility registry with polygons/points and metadata.
- Intersect hazards rasters with facility geometries under multiple scenarios.
- Aggregate exposures to issuer and portfolio levels with documented weights.

9. Construct portfolios under joint financial and climate constraints.

- Optimise for expected residual alpha subject to variance, ES/CVaR, liquidity, and transaction-cost constraints.
- Enforce factor neutrality and climate-footprint caps simultaneously.
- Provide both discrete (name selection) and continuous (weights) control with deterministic fallbacks.

10. Expose explainability and diagnostics.

- SHAP-style per-asset attributions for signals.
- Constraint duals and binding sets for portfolios.
- Hazard-to-exposure paths for climate footprints.

11. Produce governance-grade artefacts.

- Model inventory entries for each component.
- Development and validation documents suitable for SR 11-7.
- EU AI Act technical annexes and transparency notes.
- Lineage bundles with cryptographic signing.

Without all five classes of functionality, a system may be able to demonstrate attractive backtests but will fail either regulatory scrutiny or operational robustness.

2.5. Non-functional requirements and design principles

Beyond functionality, the platform must satisfy non-functional requirements:

- Determinism with controlled stochasticity. Monte-Carlo or annealing components must be seeded and logged so that any run is reproducible to within numerical tolerances. Parity MILP/QP solutions provide deterministic references.
- Latency bounds. Portfolio construction and diagnostics must complete within operational windows (seconds for moderate-sized universes, tens of seconds for larger MILP backstops), enabling intraday or daily rebalancing without bottlenecks.
- Scalability and coverage. The data spine must accommodate global universes (≥98% of target assets) and dense hazard rasters without sacrificing tractability.
- Vendor-agnosticism. The architecture assumes no privileged vendor; all external data is wrapped in lineage and schema-versioning, allowing swap-in/swapt-out of vendors without architectural changes.
- Audibility and human-overridable decisions. Every automated decision must be interrogable and, if necessary, overridden by human risk committees. APIs and consoles are designed accordingly: they surface enough internal state to support override decisions, not just final weights. These requirements translate directly into the design choices in subsequent sections: the residualizer structure, the modular geospatial pipeline, the QUBO + parity--MILP constructor, and the governance and API surfaces.

3. Institutional and Regulatory Context

This section defines the external constraints the Climate Alpha Platform must satisfy. It is written as if a model-risk committee, regulator, or institutional reviewer is reading first, before any implementation detail.

3.1. Prudential model-risk guidance (SR 11-7–type regimes)

Supervisors such as the U.S. Federal Reserve (SR 11-7) treat any quantitative tool that materially informs decisions as a "model": risk-measurement engines, pricing models, optimisation systems, and supporting components. SR 11-7 establishes three pillars: model development/implementation, model use, and model validation. A climate-alpha stack running at institutional scale is in scope.

3.1.1. Model definition and inventory requirements

A "model" is defined as a quantitative method with three elements: input data, processing (methodology), and reporting outputs. Under SR 11-7 this implies:

- The Signal Engine, Geospatial Risk Layer, and Constructor are separate models (or sub-models of a model family).

- Each must have its own inventory record: name, ID, owner, purpose, scope, usage, and implementation platforms.
- The combined “Climate Alpha Platform” is documented as a model family whose outputs (recommended weights, risk metrics, audit artefacts) feed portfolio decisions.

Architectural consequence: each module must be explicitly parameterised, versioned, and addressable in the inventory; the platform is not allowed to be one monolithic opaque engine.

3.1.2. Conceptual soundness and developmental evidence

SR 11-7 expects “sound design, theory, and logic” and “empirical evidence supporting the model’s assumptions, design choices, and limitations.” That maps into:

- Mathematical specification:
- Explicit equations for residualization, signal estimation, hazard aggregation, optimization objectives, and constraints.
- Formal statements about how factor neutrality and climate constraints are enforced.
- Justification of design choices:
- Why use residual alpha rather than raw returns as the target.
- Why specific hazard datasets, scenario sets, and climate metrics were chosen.
- Why QUBO/annealing plus MILP parity is used instead of pure MILP, heuristic screening, or heuristic portfolio heuristics.
- Developmental evidence (captured in the proof-of-value protocol):
- Backtests over multiple regimes (pre- and post-crisis, pre- and post-policy events).
- Sensitivity to hyper-parameters, penalty coefficients, and scenario sets.
- Performance comparisons against benchmark models (factor-only, ESG-score tilts, simple carbon-budget optimisers).

Architectural consequence: the theory sections of the paper are not just academic; they are the conceptual-soundness layer of the SR 11-7 dossier. Every algorithm presented must be stable, traceable, and justified in risk-management language.

3.1.3. Ongoing monitoring and performance thresholds

Model-risk regimes treat climate-alpha and optimisation as live models requiring continuous monitoring:

- Statistical monitoring:
- Out-of-sample tracking of residual alpha, IR, TE, MaxDD, and constraint violations.
- Detection of model drift: degradation relative to baselines or to thresholds defined in the proof-of-value protocol.
- Data and assumption monitoring:
- Alerts when hazard rasters, emissions datasets, or scenario definitions change version.

- Re-evaluation when regulatory definitions of "green," "significant harm," or risk metrics shift.

Architectural consequence: the platform must include monitoring hooks, not just single-run optimisation. Backtests and live runs must produce metrics and log series that a monitoring framework can subscribe to (e.g., a dedicated table keyed by model version and date).

3.1.4. Independent validation and effective challenge

Validation staff need to re-implement or replicate key parts of the stack without developer intervention:

- Reconstruct signals. Given factors, climate features, and return histories, validators must be able to compute residuals, re-fit models, and compare with production $\Delta\mu$.
- Reconstruct optimisation. Given signals, covariance matrices, constraints, and solver parameters, validators must be able to re-solve the QP/LP/MILP and match or challenge proposed weights.
- Challenge key assumptions. Validators must be able to:
- Swap hazard datasets or climate feature sets.
- Replace the QUBO path with pure MILP or alternative heuristics.
- Re-run ablations (Base vs +Climate vs +Geo vs +Constructor).

Architectural consequence:

- The specification must be implementation-agnostic. Nothing in the conceptual layer may rely on proprietary solvers or undocumented vendor code.
- All modules must surface enough intermediate state (residuals, features, matrices, penalties) to allow full off-platform replication.

3.2. Climate-related disclosures and sustainable-finance regulation

The second constraint set comes from climate-disclosure and sustainable-finance frameworks: TCFD, IFRS S2, SFDR, ESRS, EU Taxonomy, and analogous national regimes. Their common theme: climate risk, metrics, and targets must be described quantitatively, with transparent methodologies.

3.2.1. TCFD and IFRS S2

TCFD established the four-pillar structure (governance, strategy, risk management, metrics and targets). IFRS S2 embeds this into financial reporting for listed entities:

- Governance: who is responsible for climate-risk models.
- Strategy: how climate scenarios and physical/transition risks affect business models and portfolios.
- Risk management: processes used to identify, assess, and manage climate-related risks.

- Metrics and targets: quantitative KPIs (e.g., emissions, financed emissions, climate VaR, alignment metrics). In practice, this means a climate-alpha platform must output:
- Portfolio-level climate metrics: aggregated physical-risk scores, emissions intensities, alignment metrics.
- Scenario-analysis outputs: changes in portfolio risk and value under specified climate scenarios and time horizons.
- Methodology descriptions: for hazard aggregation, scenario shocks, exposure calculations, and optimisation under climate constraints. Architectural consequence: the platform's Geospatial Risk Layer and Constructor cannot be internal black boxes. Their methodologies must be expressible in the form of tables and narrative that can flow directly into TCFD/IFRS S2 disclosures.

3.2.2. EU SFDR, ESRS, and taxonomy alignment

SFDR requires funds to classify themselves (Article 6/8/9) based on sustainability claims and to report Principal Adverse Impacts (PAIs) such as GHG emissions, fossil-fuel exposure, and environmental damage. ESRS (European Sustainability Reporting Standards) deepens disclosures. The EU Taxonomy defines substantive contribution and "do no significant harm" criteria for green activities. For the platform this implies:

- Objective-and-constraint mapping:
- Article 8/9 or "taxonomy-aligned" objectives must map to explicit optimisation objectives or constraints (e.g., minimum share of taxonomy-aligned revenue, maximum portfolio-weighted carbon intensity, explicit exclusion lists).
- PAIs must be computed consistently from the same data spine used for optimisation; back-of-envelope calculations from separate spreadsheets are unacceptable.
- Consistency checks:
- The climate-alpha strategy cannot simultaneously claim to maximise taxonomy alignment and hold names that breach defined DNSH (Do No Significant Harm) thresholds unless a documented rationale exists.
- The optimisation engine must be able to show, for each run, whether PAIs and taxonomy

thresholds are satisfied and how close the portfolio is to limits. Architectural consequence: PAIs, taxonomy-alignment flags, and DNSH checks become first-class features in the platform's data model and constraint set. They are not separate reporting layers; they are part of the optimisation problem.

3.2.3. Client and stakeholder expectations

Large asset owners, sovereign funds, and regulators increasingly expect:

- Drill-down capability from portfolio-level climate outputs to issuer, facility, and hazard-level contributors.

- Scenario narratives that link quantitative shifts to underlying events (e.g., what a 2°C vs 3°C scenario means for particular facilities).
- Evidence of robustness: performance under multiple data vendors and scenario sets, sensitivity to parameter choices. Architectural consequence:
- The platform must store intermediate hazard exposures, issuer aggregations, and scenario deltas in a way that supports drill-down and alternative views.
- The methodology must explicitly accommodate multiple scenario families and vendor datasets without structural changes.

3.3. AI-specific frameworks and explainability obligations

As soon as machine-learning or complex optimisation is used, AI-specific frameworks apply, most notably the EU AI Act and internal AI-governance policies at large institutions.

3.3.1. Classification as high-risk AI system

An AI system is likely "high-risk" if:

- It materially influences financial decisions that can harm customers or affect financial stability.
- It operates in a context covered by sectoral regulation (e.g., banking, insurance, investment services). A climate-alpha optimiser that drives portfolio allocation in regulated entities satisfies both criteria. High-risk systems must meet requirements on:
- Risk management: hazards and mitigations identified, monitored, and updated.
- Data governance: training and input data appropriate, documented, and quality-controlled.
- Technical documentation: detailed descriptions of model architecture, training/test splits, performance metrics, limitations.
- Record-keeping: logging of all inputs and outputs per inference.
- Transparency and provision of information: how the system's outputs should be interpreted and used.
- Human oversight: mechanisms to prevent or mitigate harm.
- Accuracy, robustness, and cybersecurity: performance evaluation and resilience.

Architectural consequence:

- The Signal Engine, Geospatial Layer, and Constructor must each have AI-system documentation aligned with these headings.
- Logging and lineage are mandatory, not optional.
- The console and API layers must provide oversight and override capabilities: risk managers can freeze updates, adjust constraints, or revert to deterministic baselines.

3.3.2. Explainability at three abstraction levels

Explainability must align with the cognitive level of the user:

12. Per-asset level (micro):
13. For any asset i, the platform must show:
14. Residual alpha estimate $\Delta\mu\_i$.
15. Contribution of each feature (climate, geospatial, news, traditional) to $\Delta\mu\_i$ (e.g., via SHAP).
16. Facility-level exposures driving those features (e.g., which plants lie in flood zones).
17. Portfolio level (meso):
18. For each portfolio run, the platform must expose:
19. Objective and constraints used.
20. Whether each constraint is binding, and associated dual values.
21. Sensitivity of objective value to constraint relaxations.
22. Breakdown of risk and alpha contributions by sector, region, factor, and climate characteristic.
23. Model-family level (macro):
24. Documentation explaining:
25. Why residualisation is chosen.
26. How hazard data is sourced, validated, and updated.
27. How QUBO and MILP structures ensure constraint satisfaction and tractability.

Architectural consequence: model internals must be designed with natural explanatory handles: linear structures, well-defined metrics, and local contributions. End-to-end deep nets without interpretable intermediate structure are incompatible with this requirement.

3.4. Data lineage, reproducibility, and cryptographic guarantees

Climate-finance disputes and supervisory reviews will question the integrity of data and code used to construct portfolios. A defensible system needs more than internal logs; it needs structured lineage and tamper-evident artefacts.

3.4.1. Lineage semantics

For every data object (field, vector, matrix, table), the platform attaches:

- source: vendor or internal system identifier.
- as_of_date: date the data is meant to represent.
- ingestion_timestamp: when the data was pulled into the system.
- schema_version: identifier for the column layout and units.
- processing_hash: hash of the transformation code and configuration that produced the object.

Transformations are treated as first-class nodes in the lineage graph. For example:

- A facility-level hazard exposure e_{i,f,s,y} records:
- hazard raster source (e.g., Aqueduct v4, ERA5),
- facility geometry source,
- intersection algorithm version and parameters.
- A portfolio weight vector w records:
- IDs of input signal vectors and covariance matrices,
- objective and constraint definitions,
- solver type (QP, LP, MILP, simulated bifurcation),
- random seeds, tolerance parameters, and iteration limits.

Architectural consequence: lineage is not a monolithic audit log but a structured DAG of data and transformation nodes, enabling precise reconstruction of any object's provenance.

3.4.2. Replay and independent reconstruction

Given a run_id, an auditor must be able to:

- Retrieve immutable snapshots of all inputs and configuration used at that time.
- Re-execute signal estimation and optimisation with the same code version or with a validator implementation.
- Confirm that outputs (signals, weights, metrics) match recorded ones up to numeric tolerances.

This demands:

- Strict point-in-time storage rather than "as-latest" views.
- Deterministic or seed-controlled algorithms for any stochastic steps (simulation, annealing).
- Version-locked dependencies (libraries, solver versions).

Architectural consequence: the system's storage and code-deployment patterns must enforce immutability and explicit versioning for anything referenced in an audit run.

3.4.3. Cryptographic attestation

For high-stakes contexts (regulatory inquiries, litigation), internal assurances are not enough. The platform therefore:

- Computes cryptographic hashes of lineage bundles, configuration files, and key outputs (weights, risk metrics, reports) for each run_id.
- Uses post-quantum key-encapsulation and signature schemes (ML-KEM, ML-DSA) to sign these hashes.
- Stores signatures and public verification keys in a controlled registry.

This yields:

- Non-repudiation: the institution can prove that a given audit pack corresponds to a particular historical run and has not been tampered with.
- Third-party verification: regulators or auditors can verify signatures without privileged access to internal systems. Architectural consequence: cryptographic signing is not a cosmetic add-on; it sits on the critical path of audit-pack generation and is integrated into /audit endpoints.

3.5. Consolidated design constraints

Collecting the previous subsections, the regulatory and institutional context imposes the following hard constraints on the Climate Alpha Platform:

28. Modular, inventory-ready architecture. Signal, geospatial, and construction components must be separable models in the SR 11-7 sense, each with clear mathematical specifications and governance.
29. Full transparency from data to decision. For any asset and any portfolio, the platform must explain: which data were used, how they were transformed, which models operated on them, and how this produced signals and weights.
30. Residual, factor-clean definition of climate alpha. To satisfy effective challenge and avoid mis-selling, any "alpha" claim must be demonstrably orthogonal to standard risk premia, with formal tests and ablations to prove it.
31. Facility-to-portfolio climate mapping integrated into optimisation. Disclosure and SFDR/ESRS regimes require portfolio-level climate metrics grounded in facility-level hazards and emissions. These metrics must be embedded as constraints and objectives, not computed post-hoc.
32. Explainable optimisation with deterministic parity. Any stochastic or heuristic search (e.g., quantum-inspired QUBO solvers) must be backed by an exact MILP/QP formulation that produces constraint-clean solutions and exposes duals and binding sets.
33. Comprehensive lineage, replay, and tamper-evidence. Every run must be reconstructible from immutable snapshots, and audit artefacts must be cryptographically bound to their historical context.
34. Human-override and governance hooks. Risk and investment committees must be able to inspect, override, or freeze model outputs, integrate challenger models, and feed results into TCFD/IFRS/ SFDR reporting. These constraints are not optional; they are the boundary conditions within which the architecture and mathematics in subsequent sections are constructed.

4. Data Spine and Lineage

This section fixes the data model the platform assumes. Everything else—signals, geospatial risk mapping, optimisation, governance—is built on top of this spine and lineage semantics.

4.1. Data domains and objects

The data spine is a set of versioned, point-in-time tables and rasters that cover all inputs required by the Signal Engine, Geospatial Risk Layer, and Constructor. Core domains:

35. Market and pricing data
36. Daily asset-level OHLCV and returns.
37. Corporate-action streams (splits, dividends, rights issues).
38. Benchmark index membership histories.
39. FX rates and risk-free curves for excess-return computation.
40. Factor and fundamental data
41. Exposures to style and risk factors (market, size, value, profitability, investment, momentum, quality, etc.).
42. Industry classifications (GICS/ICB/NAICS).
43. Issuer-level financials (balance sheet, P&L, cashflow) with point-in-time revisions.
44. Pre-computed factor returns (e.g., Fama–French) for validation and monitoring.
45. Climate and environmental data
46. Emissions data: scopes 1, 2, and 3; intensity metrics (tCO_2e / revenue, /EVIC, /MWh, etc.).
47. Targets and transition plans: SBTi flags, net-zero commitments, decarbonisation trajectories.
48. EU Taxonomy, SFDR PAI, and other sustainability indicators as structured fields.
49. Facility registry and geospatial data
50. Facility-level polygons and point geometries for plants, offices, warehouses, critical suppliers.
51. Metadata: facility type, capacity, sector, ownership, operational status, start/end dates.
52. Jurisdiction attributes: country, sub-national region, regulatory regime tags.
53. Hazard rasters and climate scenarios
54. Tiled rasters for physical hazards: flood, drought/low-flow, heat stress, windstorm, wildfire, coastal inundation.
55. Time-stamped layers per scenario (e.g., NGFS "Orderly", "Disorderly", "Hot House World") and time horizons (e.g., 2030, 2050, 2100).
56. Metadata: source dataset (e.g., WRI Aqueduct 4.0, ERA5, NASA FIRMS), scenario labels, units, resolution, coordinate reference system.
57. Climate-news and NLP features
58. Document store of climate-relevant news with timestamps, issuers, and extracted entities.
59. Derived sentiment and topic scores per issuer/date.
60. Derived features and signals
61. Engineered features for signal estimation (joined factors + climate + news).
62. Residual returns, estimated $\Delta\mu$ vectors, SHAP attributions, posterior variances.

63. Optimisation artefacts
64. Covariance matrices, scenario stress matrices, constraint matrices.
65. Weights, constraint duals, binding-set indicators, and risk decompositions per run.

Each domain is stored as a separate set of tables or rasters with explicit keys and versioning. No "magic join" is assumed; joins are defined and versioned explicitly as transformations.

4.2. Point-in-time schema and temporal semantics

The spine is bitemporal: it tracks both when data is valid and when it was ingested/known. This is required for unbiased backtesting and audit replay. Each row in a core table carries:

- as_of_date: the date the data describes economically (e.g., fiscal period end, disclosure effective date).
- ingestion_timestamp: the exact time the row entered the system.
- valid_from / valid_to: optional additional fields for effective intervals (e.g., facility operational period).

Example: a price table row:

- asset_id
- date (trading date)
- px_open, px_high, px_low, px_close, volume
- as_of_date = date
- ingestion_timestamp

Example: an emissions table row:

- issuer_id
- scope_1_tCO2e, scope_2_tCO2e, scope_3_tCO2e
- intensity_tCO2e_per_revenue
- as_of_date = 2023-12-31 (company's reporting year)
- ingestion_timestamp = 2024-04-15T10:37:00Z

Backtests at time t are only allowed to use data whose ingestion timestamp $\leq$ t and as_of_date $\leq$ t to enforce strict no-look-ahead. Facilities and hazard rasters are treated similarly:

- Facility geometry rows include valid_from and valid_to representing operational periods.
- Hazard rasters include scenario_id, horizon_year, and generation_timestamp.

This is non-negotiable. Without such temporal semantics, any reported historical performance is contaminated by restatement or look-ahead bias.

4.3. Facility registry and hazard rasters

The Geospatial Risk Layer depends on a clean mapping between assets (issuers and securities), facilities, and hazard rasters.

4.3.1. Facility registry structure

Core fields for each facility record:

- facility_id (surrogate key)
- issuer_id and, optionally, ultimate_parent_id
- geom (POLYGON or POINT, in a canonical CRS such as EPSG:4326)
- facility_type (plant, mine, warehouse, office, data centre, etc.)
- capacity (MW, tonnes/year, square meters, as applicable)
- status (operational, under construction, decommissioned)
- valid_from, valid_to

Additional optional fields:

- supply_chain_role (tier-1 supplier, logistics hub, etc.)
- jurisdiction (country, region codes)
- regulatory_tags (carbon market participation, zoning restrictions)

Facilities are linked to issuers via issuer_id; multiple facilities per issuer are allowed.

4.3.2. Hazard raster catalogue

Each hazard dataset is catalogued with metadata:

- hazard_id (e.g., aqueduct_flood_4_0)
- hazard_type (riverine_flood, coastal_flood, wildfire, heat, drought, wind, surge)
- scenario_id (e.g., NGFS_orderly_2C, NGFS_hot_house_3C)
- horizon_year (e.g., 2030, 2050, 2100)
- resolution (spatial resolution in degrees or metres)
- crs (coordinate reference system)
- units (e.g., annual damage probability, expected depth in metres, index score)
- generation_timestamp (when this layer was created/ingested)

Hazard rasters themselves are stored in tiled formats with indexes enabling efficient spatial queries (e.g., via PostGIS, raster stores, or cloud object storage with overviews).

4.3.3. Facility–hazard intersection semantics

4.4. Feature engineering on the spine

66. Climate transition-risk features
67. Emissions levels and intensities, by scope.
68. Targets encoded as distances from required path (e.g., overshoot vs SBTi-aligned trajectory).
69. Classification into hard-to-abate vs easy-to-abate sectors.
70. Climate-news features
71. Rolling sentiment scores, controversy indicators, litigation mentions.
72. Topic distributions over climate-related themes (policy, technology, regulation, activism).
73. Interaction and cross terms
74. Products of physical and transition features (e.g., high physical flood risk and high emissions intensity).
75. Interaction between climate features and financial variables (e.g., leverage × hazard).

Feature pipelines are defined as DAGs of operations: filter, join, aggregate, transform, normalise. Each node has its own code hash and configuration, and produces outputs that carry updated lineage fields. Normalisation conventions:

- Continuous features are z-scored within (region, sector) at each date t.
- Categorical variables are one-hot encoded with explicit reference categories.
- Winsorisation is applied to extreme values before z-scoring, with cut levels documented (e.g., 0.5%/ 99.5% tails). These conventions must be fixed and versioned. Changing them produces a new feature-pipeline version and therefore a new signal-engine model version.

4.5. Lineage model and /audit artefacts

Lineage is explicit, not approximate. Every data object and transformation is represented in a graph.

4.5.1. Lineage metadata per object

Each stored object (table, raster, feature vector, signal vector, weight vector, report) carries at minimum:

- object_id: unique identifier.
- object_type: e.g., raw_price_table, hazard_raster, feature_matrix, signal_vector, weights, report_pdf.
- source: original external source(s), if any.
- as_of_date / horizon_year for economic interpretation.
- ingestion_timestamp / generation_timestamp.
- schema_hash: hash of column names, types, units, and, for rasters, resolution and CRS.
- processing_hash: hash of the transformation code and configuration that created the object.

- parents: list of object_ids that were inputs to this object.

This gives a directed acyclic graph (DAG) of provenance.

4.5.2. Run-level lineage bundles

Each optimisation or scoring run has a run_id. The platform materialises a lineage bundle for the run:

- List of all input objects (features, covariances, constraints) with full metadata.
- Model versions for each component (Signal Engine version, Geospatial layer version, Constructor version).
- Solver configurations (solver type, tolerance, random seed, max iterations).
- Outputs (signals, weights, risk metrics, SHAP attributions).

Bundles are serialised to JSON (machine-readable) and also summarised in human-readable PDFs for governance packs.

4.5.3. /audit endpoint behaviour

The /audit?run_id= endpoint returns:

- The full lineage bundle.
- Links to stored reports and logs.
- Verification status of cryptographic signatures.

This is designed to be callable by risk, internal audit, or external regulators, not just portfolio managers. Ensuring that this endpoint can always reconstruct lineage is a hard requirement.

4.6. Cryptographic integrity and PQC signing

Given the sensitivity of climate-finance decisions and the probability of legal/regulatory disputes, the platform binds lineage and key outputs with post-quantum cryptographic signatures.

4.6.1. Hashing and signing process

For each run_id the platform:

76. Serialises the lineage bundle and core outputs (weights, risk metrics, constraints, key plots) into a canonical byte representation.
77. Computes a cryptographic hash H over this representation (e.g., SHA-3).
78. Uses an ML-DSA private key to sign H, producing signature \sigma.
79. Optionally uses ML-KEM to encrypt or encapsulate symmetric keys if confidentiality is required in external transmission. Stored artefacts:

    - hash: hexadecimal encoding of H.
    - signature: \sigma.

- public_key_id: identifier of the public key required to verify the signature.

4.6.2. Verification semantics

Verification consists of:

- Recomputing H from the stored lineage bundle and outputs.
- Checking the signature \sigma against H using the ML-DSA public key.

If verification passes, the bundle is proven to be unaltered since signing. If verification fails, the artefact is either corrupted or maliciously modified.

4.6.3. Integration into KPIs and controls

Lineage signing and verification feed into platform KPIs and controls:

- Correctness KPI: proportion of runs where lineage signatures verify successfully; target is 100%.
- Change control: any modification to past bundles is impossible without invalidating signatures, enforcing immutability of historical evidence.
- External audits: regulators or clients can verify signatures offline, independent of internal systems'

integrity. This closes the loop: the data spine provides the raw material for signals and optimisation; the lineage graph and PQC signing provide a defensible chain of custody from raw data to final portfolios.

5. Methodology

This section fixes the mathematical structure and algorithms of the Climate Alpha Platform. It covers: (i) the factor-residual baseline; (ii) climate-aware signal estimation; (iii) geospatial risk aggregation; (iv) portfolio construction under financial and climate constraints; (v) discrete selection as QUBO with quantum-inspired solvers and classical parity; and (vi) explainability.

5.1. Notation and universes

Let: All variables are point-in-time consistent with the data spine: for a given t, only information whose ingestion timestamp is $\leq$ t is used.

5.2. Factor residualization and definition of climate alpha

The platform defines "climate alpha" as expected excess returns orthogonal to standard risk premia. This explicitly removes the most common failure mode of ESG strategies—repackaged value/quality/ sector exposure.

5.2.1. Cross-sectional factor model

On each date t, the platform fits a cross-sectional factor model (for instance, FF-5 + industry dummies):

5.2.2. Residual alpha and its estimation

5.3. Climate-aware signal estimation

5.3.2. Bayesian temporal smoothing

5.3.3. Multiple-testing and signal-quality control

Across many assets and dates, naive t-stat thresholding overstates signal strength. The platform uses:

- Deflated Sharpe/IR adjustments to correct for selection bias and multiple testing.
- obson–Korkie/Memmel tests to compare IR of climate-augmented portfolios vs factor-only

baselines. Climate features and model specifications are only promoted to production if they demonstrate statistically robust, out-of-sample improvements over the baseline, after transaction costs.

5.4. Geospatial risk mapping (method level)

The Geospatial Risk Layer translates facility geometries and hazard rasters into issuer- and portfolio-level exposure vectors. It is used both to construct features for the Signal Engine and to enforce constraints in optimisation.

5.4.1. Facility-level hazard computation

5.4.2. Issuer-level hazard aggregation

5.4.3. Portfolio-level footprints

Given weights wt, portfolio exposure under scenario s, horizon y, and hazard k is: Climate constraints in optimisation typically enforce:

5.5. Portfolio construction: objectives and constraints

The Constructor solves an optimisation problem that trades off climate-alpha signals versus risk and implementation costs, under factor and footprint constraints.

5.5.1. Mean–variance with factor and climate constraints

Baseline mean–variance problem at time t:

5.5.2. ES/CVaR formulation

5.5.3. Tracking-error-constrained climate alpha

5.6. Discrete selection as QUBO and quantum-inspired optimisation

To handle cardinality and discrete inclusion decisions, the platform formulates a QUBO over binary variables and solves it via quantum-inspired methods, then refines results with convex optimisation.

5.6.1. Binary selection variables and energy function

5.6.2. Simulated bifurcation / annealing-class solvers

To approximate the minimum of E(z), the platform uses simulated bifurcation or similar dynamical algorithms:

- Introduce continuous pseudo-spin variables x_i $[-1,1]$ that will be mapped to $z_i = (1+x_i)/2$.
- Define Hamiltonian dynamics with potential energy mirroring QUBO, plus a control parameter that is

slowly changed ("bifurcated") to steer the system toward low-energy states.

- Time-discretise the equations of motion and simulate until convergence.

Key properties:

- The algorithm explores the high-dimensional energy surface more globally than greedy local search.
- It is fully classical and implementable on CPUs/GPUs; it is "quantum-inspired" only in its analogy to

adiabatic processes. Alternative solvers (simulated annealing, tabu search, genetic algorithms) can be plugged into the same QUBO definition, as long as they are seeded and logged for reproducibility.

5.6.3. Classical MILP/QP parity

Heuristic QUBO solvers are not acceptable as sole decision engines in a regulated context. The platform therefore defines a parity MILP/QP whose feasible set and objective correspond as closely as possible to the QUBO formulation:

5.7. Explainability layer

Explainability is handled at three levels: per-asset feature attribution, portfolio-level constraint diagnostics, and hazard-to-exposure tracing.

5.7.1. Kernel-SHAP for climate-alpha signals

5.7.2. Constraint duals and binding-set diagnostics

5.7.3. Hazard-to-exposure-to-portfolio path

For each portfolio run, the audit pack can include:

- A list of issuers contributing most to a given hazard exposure.
- Maps showing facilities in risk hotspots.
- Quantitative decomposition of exposure by sector, region, and hazard type.

This completes the "story" from raw climate data to financial decisions in a way that is verifiable and regulator-compatible.

5.8. Consistency, robustness, and ablations

This methodology section, together with the data spine and context, defines a complete, implementation-ready blueprint for a governed climate-alpha stack: residual, factor-clean signals; geospatially grounded exposures; and optimisation with quantum-inspired exploration, convex refinement, and full explainability.

6. Governance, Security, and Compliance

This section fixes how the platform is governed, documented, secured, and made compliant with supervisory expectations. It expands the "SR 11-7 Pack / EU AI Act Alignment / Security / PQC Signatures" bullets into a complete operating regime.

6.1. SR 11-7 model-risk framework

The platform is treated as a model family under SR 11-7 and equivalent model-risk policies. The family contains three primary constituent models:

- Model A: Signal Engine (climate-aware residual alpha).
- Model B: Geospatial Risk Layer (facility-to-issuer-to-portfolio hazard mapping).
- Model C: Constructor (optimisation stack: QUBO + convex refinement + MILP/QP parity).

Each model and the family as a whole must satisfy SR 11-7 requirements: model inventory, conceptual soundness, developmental evidence, monitoring, and independent validation.

6.1.1. Model inventory and classification

For each of A, B, C, and for the combined Climate Alpha Platform, the inventory entry includes:

- Model ID and name – e.g., "CA-A-XAI-SignalEngine-v1".
- Purpose and scope – what decisions the model supports (e.g., signal ranking for equity universes,

hazard footprint computation, portfolio weights for climate-aware mandates).

- Model type – statistical/ML (Signal Engine), deterministic numerical (Geospatial), optimisation (Constructor).
- Owner and accountable executive – named individuals or roles; ownership cannot be anonymous or collective.
- Usage classification – materiality rating (e.g., "material for investment decisions in climate-labelled funds"). This inventory ensures that any supervisory or internal review can immediately see where climate-alpha models sit in the broader risk architecture.

6.1.2. Conceptual-soundness documentation

For each model, a technical note is maintained; these notes are essentially the formal exposition already given in Sections 4–5, re-packaged into SR 11-7 vocabulary:

- Signal Engine note sets out:
- The factor-residualisation operator Mt its algebraic properties, and why climate alpha is defined on rt+1 instead of yt+t
- The elastic-net regression, the Bayesian smoothing layer, prior choices, and the rationale for using cross-sectional rather than time-series models.
- The criteria for feature inclusion, normalization, winsorisation, and interaction terms.
- Geospatial note sets out:
- Hazard raster definitions, facility geometry assumptions, intersection rules, and aggregation weights.
- Limitations and edge cases (e.g., missing facility locations, low-resolution rasters).
- Constructor note sets out:
- Mean–variance and ES/CVaR formulations, with KKT conditions and interpretation of duals.
- QUBO energy definition, penalty terms, and equivalence to constrained optimisation in the limit of large penalties.
- MILP/QP parity formulation and conditions under which cardinality and inclusion constraints are enforced exactly. Each note has a Limitations subsection explicitly listing model boundary conditions: missing data patterns, assumed stationarity windows, hazard-model epistemic uncertainty, solver tolerances, and failure modes.

6.1.3. Developmental evidence and validation packs

SR 11-7 requires evidence that models behave as intended. The platform's Proof-of-Value Protocol constitutes the core developmental evidence:

- Design: large/mid universes per region; 2005–2020 training/parameterisation; 2021–2025 walk-forward; strict factor neutrality; realistic costs; liquidity and borrow limits; turnover caps.
- Ablations: Base → +Climate Signals → +Geospatial → +Quantum-Inspired Construction.

- Metrics: residual alpha vs multi-factor, IR, TE, MaxDD, hit-ratio, signal decay, capacity, cost-adjusted net performance.
- Gates: IR uplift +0.15–0.30, MaxDD reduction 10–20%, stability across rebalancing periods.

The validation pack references this protocol and includes:

- Backtest result summaries (no full code; these are summaries, consistent with SSRN working-paper norms).
- Sensitivity tests: changing hazard datasets, excluding climate features, altering penalty coefficients in QUBO, switching to alternative optimisers.
- Stress applications: how portfolios behave under adverse climate scenarios and macro shocks.

6.1.4. Monitoring, thresholds, and change control

Monitoring is defined at two levels:

80. Statistical performance:

- Ongoing tracking of realised residual alpha vs factor-only baselines, IR, TE, and MaxDD.
- Thresholds consistent with PoV gates; breaches trigger review.

81. Data/assumption monitoring:

- Alerts when hazard models (e.g., Aqueduct 4.0 to 5.0) or emissions methodologies change; this triggers a re-run of validations and possibly retuning of signals and constraints. Change control:
- Any model change (code, feature set, solver configuration) increments the model version and logs a change record with rationale, risk impact, and sign-offs.
- Material changes require re-validation and a short addendum to the SR 11-7 pack.

Independent validation (by a separate team or external reviewer) re-implements key components (residualiser, climate regression, geospatial aggregation, optimisation) on sample data and checks consistency with platform outputs.

6.2. EU AI Act and AI-governance alignment

The platform is designed to meet requirements of AI regulations such as the EU AI Act for high-risk AI systems.

6.2.1. AI-system boundary and classification

The AI system boundary includes:

- The Signal Engine when implemented using ML (regularised regression + possible non-linear extensions).
- The optimisation routines when they incorporate heuristic or stochastic solvers (QUBO + simulated bifurcation / annealing-class algorithms). The platform documentation explicitly classifies these as high-risk AI components because they materially influence financial decisions in regulated environments.

6.2.2. Technical documentation (Annex IV–style)

For each AI component, an Annex-style technical file is maintained (mapping to EU AI Annex IV):

- System overview: purpose, intended use, and limitations.
- Data description: training and validation data domains, sampling frames, and handling of bias.
- Model description: architecture (linear + regularisation, state-space smoothing), loss functions, hyperparameters, and reasoning for choices.
- Training and validation: procedures for cross-validation, stability selection, and performance metrics.
- Monitoring and update strategy: triggers for retraining or deprecation.
- Known limitations and risks: identified model risks, external dependencies, potential fairness or bias issues in climate exposure, and overfitting concerns. This documentation is structurally aligned with what EU AI expects for high-risk AI systems.

6.2.3. Logging, traceability, and risk management

The system logs at the run level (per /scores or /weights call) and at the data-object level (as per the lineage model). For each AI inference or optimisation run, logs include:

- Input object IDs (feature matrices, covariances, constraint sets).
- Model versions (Signal Engine, Geospatial, Constructor).
- Random seeds and solver configurations (for any stochastic algorithm).
- Output object IDs (signals, weights, risk metrics, SHAP attributions).

Risk management for AI components is captured in a Model Risk Register, listing:

- Specific AI-related hazards (e.g., non-stationarity of climate relationships, undue reliance on a single hazard dataset, black-box drift if non-linear models are introduced later).
- Mitigations (residualization, dual-vendor data ingestion, ablations, human overrides, etc.).

6.2.4. Human oversight mechanisms

The platform exposes oversight hooks via the console and APIs. Key controls:

- Pre-trade override:
- Risk managers can modify constraint sets (tighten caps, adjust turnover or TE limits), freeze climate constraints, or set manual hard limits on exposures.
- The system records overrides as structured events tied to run_ids.
- Model freeze:
- A governance flag can freeze specific model versions: the system refuses to run new optimisations with frozen models until a review is completed.
- Fallback modes:
  - If AI components are deemed unreliable (e.g., due to data shocks), the platform can be configured to fall back to:
  - Factor-only signals plus climate constraints based on static footprints.
  - Deterministic optimisations (pure QP/LP/MILP) without QUBO or stochastic exploration.

Every override and mode change is logged and enters the audit trail.

6.3. Security architecture

The platform's security model addresses confidentiality, integrity, and availability, with specific attention to regulatory expectations for financial services.

6.3.1. Data in transit and at rest

- In transit: all client–platform and internal service–service communications use TLS (modern versions, e.g., TLS 1.2+), with strong cipher suites and certificate pinning where appropriate. At rest:
- Data at rest in databases and object stores is encrypted with per-tenant keys.
- Keys are managed by a KMS (Key Management System); rotation policies are enforced.
- Access control uses least-privilege principles; access is role-based and audited.

Data residency:

- Tenants can be mapped to specific storage regions to satisfy data-residency laws; the data spine is deployed per-region, not globally mirrored by default.

6.3.2. Access control and segregation of duties

Access is segmented across roles:

- Research / model developers: read access to anonymised or synthetic data samples and model code repositories; no direct access to production secrets or live client data.
- Production operators: access to deployment configuration and monitoring; no permission to modify model code (except via controlled releases).
- Risk/compliance: read access to lineage, audit packs, and logs; limited ability to change model-state flags (e.g., freeze). Segregation of duties reduces the attack surface and aligns with expectations in SOC2-like audits (noted in the roadmap).

6.3.3. Logging and tamper-evident event trails

Security logging covers:

- Authentication and authorisation events.
- Access to sensitive objects (raw hazard data, facility registries, client portfolios).
- Changes to model configuration, constraint sets, and override flags.
- Invocation of /audit and other governance endpoints.

Tamper-evidence:

- Logs are written to append-only stores (e.g., WORM volumes or log services with append-only guarantees).
- Hash-chains or other cryptographic anchoring schemes can be used to detect deletion or modification of log entries. This ensures that security events, like modelling events, are forensically traceable.

6.4. Post-quantum signatures and trust guarantees

The platform uses post-quantum cryptography (PQC) to sign lineage bundles and key reports, making them tamper-evident and resilient to future cryptanalytic advances.

6.4.1. PQC scheme selection and key management

- ML-KEM (Kyber) is used for key establishment and secure key distribution.
- ML-DSA (Dilithium) is used for digital signatures on lineage bundles and critical outputs.

Keys are managed as follows:

- A master signing key pair per platform deployment, with public keys distributed to auditors and regulators as needed.
- Rotation policy for signing keys; older public keys remain available for verification of historical runs.
- Key material stored in hardened HSMs or KMS with strict access controls.

6.4.2. Signing of lineage bundles and reports

For each run_id, the system:

82. Serializes the lineage bundle (inputs, configurations, outputs) into a canonical format.
83. Computes a cryptographic hash H (e.g., SHA-3) of the serialized bundle.
84. Signs H with the ML-DSA private key to produce signature \sigma.
85. Stores (run_id, H, σ, public_key_id) alongside the bundle.

Similar procedures are applied to SR 11-7 dossiers, EU AI transparency notes, and any externally shared audit pack.

6.4.3. Verification and external assurance

Verification is straightforward: •

- Recompute the hash H' of the retrieved lineage bundle.
- Verify \sigma against H' using the associated ML-DSA public key.

If verification passes, the bundle is confirmed as identical to the originally signed artefact. If it fails, either the content or the signature has been altered or corrupted. Applications:

- Internal audit: quick integrity checks during reviews or forensics.
- Regulatory inspections: regulators can independently validate that materials provided to them correspond to original run outputs and have not been edited.
- Client due diligence: institutional clients can verify the authenticity of performance and methodology reports. This mechanism is integral to the platform's "quantum-resilient trust" positioning: even if classical cryptographic schemes degrade, signed artefacts remain verifiable.

6.5. Governance roles, RACI, and operating cycle

Governance is expressed as a RACI matrix (Responsible, Accountable, Consulted, Informed) across roles and processes:

- Model Developer (MD):
- Responsible for implementing and maintaining model code and documentation.
- Consulted on change requests and validation feedback.
- Model Owner (MO):
- Accountable for model performance, appropriateness, and alignment with business objectives.
- Approves material changes and sign-offs for SR 11-7 and AI documentation.
- Model Validation (MV):
- Responsible for independent review and challenge of models; cannot report to MD.
- Approves initial go-live and material changes.

- Risk/Compliance (RC):
- Accountable for ensuring regulatory compliance; sets constraints and risk appetite.
- Consulted on model scope, disclosure mapping, SFDR/IFRS/S2 alignment.
- IT/Security (IT):
- Responsible for deployment security, access control, encryption, and log integrity.
- Operations (Ops):
- Responsible for running daily processes, monitoring run health and SLOs.

Operating cycle:

86. Design: MD drafts model specs and initial prototypes; RC provides policy constraints; MV defines validation expectations.
87. Build/Calibrate: MD implements; IT sets up environments; Ops configures pipelines using the data spine.
88. Validate: MV runs independent tests, ablations, and challenge models; findings documented; MD corrects; MO decides on readiness.
89. Deploy: IT/ Ops deploy model version; RC confirms constraints; SR 11-7 and AI documents are filed in inventory.
90. Monitor: Ops tracks KPIs/SLOs (latency, constraint violations, coverage, performance); RC monitors risk metrics; MD monitors model drift.
91. Review/Change: On breach of thresholds or scheduled reviews, MV re-opens validation; MD proposes changes; cycle repeats. This operating rhythm is required to keep the platform within supervisory expectations while allowing iteration.

6.6. Documentation and disclosure mapping

Finally, governance and compliance structures must feed into external-facing documents and disclosures.

- SR 11-7 dossier:
- Collates the conceptual-soundness notes, developmental evidence, monitoring framework, and validation reports into a coherent pack per model family.
- EU AI transparency note:
- Summarizes AI components (Signal Engine and Constructor), their purpose, training data, performance, oversight, and limitations in language aligned with AI regulation.
- IFRS S2 / TCFD-aligned climate disclosures:
- Uses outputs from the Geospatial Risk Layer and Constructor to populate climate-related risk metrics, scenario analyses, and descriptions of how climate is integrated into investment processes.
- SFDR / ESRS climate annexes:
- Trace PAIs, taxonomy alignment, and DNSH conditions back through the same lineage used for

optimisation, ensuring consistency between product behavior and disclosure claims. All these documents pull directly from the same model specifications and lineage bundles described earlier; there is no separate “marketing model”. Governance, security, and compliance are thus structurally fused with the platform’s mathematical and data architecture.

7. Proof-of-Value Protocol

This section specifies the empirical protocol used to test the Climate Alpha Platform. It is an experimental design that an institution can implement to verify or falsify the claimed performance.

7.1. Objectives and evaluation philosophy

Single objective: Determine whether the governed climate-alpha stack delivers statistically robust, factor-neutral, cost-adjusted improvement in risk-adjusted performance vs a factor-neutral baseline, under realistic constraints. Secondary requirements:

- Uplift must survive transaction costs, liquidity limits, and borrow/capacity constraints.
- Uplift must be orthogonal to standard factors by construction (residualisation) and by ex-post tests (factor regressions).
- Uplift must be robust across regions, time periods, hazard datasets, and scenario sets. Protocol is strictly walk-forward. No in-period parameter refitting using future information.

7.2. Universe construction and data windows

Regions and universes.

- Regions:
- North America.
- Europe.
- Developed Asia.
- Optional: EM ex-China, EM including China.
- For each region, define a liquid large/mid-cap universe:
- Rank by free-float market cap; take top 1,000–1,500 names per region.
- Impose minimum ADV and listing-age filters.
- Exclude penny stocks, persistent illiquid names, and securities under restructuring or prolonged suspension. Time structure.
- Training window (design/calibration): 2005–2020 inclusive. Used for:
- Setting factor model specification and residualiser parameters.
- Designing feature sets and hyperparameter ranges for the Signal Engine.
- Choosing penalty magnitudes for QUBO and tolerances for MILP/QP.
- Roughly calibrating climate-constraint ranges and turnover bands.
- Evaluation window (true out-of-sample): 2021–2025.

- Walk-forward only.
- No parameter tuning based on any data after 2020.

Rebalancing frequency.

- Primary: monthly or quarterly.
- At each rebalance date t:
- Construct features, signals, and constraints using only information with ingestion timestamp $\leq t$.
- Hold portfolio until next t+1 (no intramonth trading in the protocol).

This setup forces the backtest to approximate realistic institutional workflows.

7.3. Strategy ladder (ablation design)

The protocol evaluates an ablation ladder, isolating incremental contribution of each architectural component.

92. Base – Factor-Neutral Baseline

- Signals: standard factor-based (value, size, momentum, profitability, investment, quality, low-vol) derived from the data spine; no climate features.
- Residualization: returns are projected off the same multi-factor model described in Section 5; alpha is defined on residuals but learned only from non-climate features.
- Portfolio construction:
- Mean–variance or ES/CVaR optimisation.
- Strict factor-neutrality relative to benchmark.
- Liquidity, turnover, sector/country caps enforced.
- No climate constraints, no geospatial inputs, no QUBO.

93. +Climate Signals

- Signals: extend feature matrix \boldsymbol{\phi}_t with:
- Emissions (Scopes 1/2/3), intensities.
- Transition indicators (targets, SBTi flags, fossil share, etc.).
- Basic physical-risk proxies (country-level or coarse hazard metrics).
- Model: residual returns r_{i,t+1} regressed on extended \boldsymbol{\phi}_t via elastic-net + Bayesian smoothing.
- Construction: unchanged from Base (mean–variance or ES/CVaR, factor-neutral, no explicit climate constraints).

94. +Geospatial Layer
95. +Quantum-Inspired Construction

- Signals: identical to +Geospatial configuration; no change in Signal Engine.
- Construction:
- Use QUBO formulation with binary variables for inclusion; run simulated bifurcation / annealing to propose candidate subsets.
- Run full MILP/QP parity problem with these subsets as warm starts; final portfolio is from the deterministic solver.
- Enforce all constraints (factor, climate, liquidity, turnover, sector/country) exactly.

7.4. Costs, liquidity, and realistic constraints

7.5. Metrics and definitions

Performance and risk are measured with a consistent metric set.

7.6. Statistical tests and multiple-testing controls

The protocol explicitly handles statistical significance and multiple-testing.

96. Jobson–Korkie / Memmel tests.

- Compare IR of Base vs each ladder rung.
- Null hypothesis: IR\text{Base} = IR\text{Variant}.
- Use Memmel's correction for sample-estimated statistics and overlapping periods.

97. Deflated Sharpe / Deflated IR.

- Estimate the effective number of trials (number of materially different model and hyperparameter variants explored).
- Compute Deflated Sharpe (DSR) or Deflated IR for each variant to adjust for data-mining and non-normality.
- Require climate variants to exceed DSR thresholds, not just naïve IR thresholds.

98. Bootstrap confidence intervals.

- Use block bootstrap on return series to generate empirical distributions for IR, alpha, MaxDD.
- Examine 95% confidence intervals for uplift vs Base; require that intervals exclude zero uplift for claimed improvements.

99. Temporal robustness.

- Split evaluation period into at least two subwindows (e.g., 2021–2022, 2023–2025).
- Check that uplift persists in both; reject configurations where uplift is highly concentrated in a single subperiod.

7.7. Decision gates

The protocol defines explicit gates that must be passed before claiming PoV success. Primary gates (hard conditions).

100. IR uplift:

- +Climate +Geospatial +Constructor must produce IR uplift of +0.15 to +0.30 vs Base in at least two major regions, on net returns.
- Uplift must be statistically significant under Jobson–Korkie/Memmel and DSR.

101. MaxDD improvement:

- MaxDD reduction of 10–20% vs Base, or no worse than Base under equal TE.

102. Persistence:

- IR uplift and MaxDD improvement must hold under small perturbations of rebalancing schedule (e.g., shifting rebalance day within the month) and over subperiods.

103. Constraint cleanliness:

- 100% of evaluation runs have zero hard constraint violations in parity MILP/QP solutions. Secondary gates (strong preferences).
- Hit-ratio > 50% overall and not materially worse in down-markets.
- Capacity such that at least 50% of IR uplift remains at institutionally relevant AUM levels.

Strategies failing primary gates are not climate-alpha platforms under this architecture; they remain prototypes or research curiosities.

7.8. "No-cheat" constraints

To keep the protocol defensible, several practices are explicitly banned:

104. No look-ahead / restated data.

- Only point-in-time data with ingestion timestamps ≤ t is used for decisions at t.
- Restated fundamentals or emissions can appear only when they were actually available in real time.

105. No ex-post horizon tuning.

- Holding periods and rebalance frequencies are fixed before evaluation.
- No retrospective adjustment to align with regimes where performance was accidentally better.
-

106. No uncontrolled model shopping.

- Number of materially distinct configurations tried is logged.
- DSR or other corrections explicitly use this number when assessing significance.

107. No hidden factor tilts.

- Ex-post multi-factor regressions verify factor betas near-zero relative to Base.
- If residual factor tilts remain, they are either brought under additional constraints or acknowledged as non-climate sources of return.

7.9. Artefacts and reproducibility

The protocol requires production of artefacts capable of supporting internal challenge and external audit. For each strategy, region, and ladder rung:

- Lineage bundles (PQC-signed):
- Exact universes, feature definitions, hazard datasets, parameter sets, solver configs, model versions.
- Hashes of all data inputs and code versions.
- SR 11-7 model pack:
- Conceptual-soundness summary.
- Empirical results with statistical tests and limitations.
- AI technical file:
- Description of AI components, training/validation data subsets, monitoring scheme, and known limitations.
- Replication packs:
- Either anonymized data + code sufficient for internal validation teams to reproduce results, or synthetic data verifying methodological correctness. All artefacts are created automatically as part of the PoV runs and stored through the lineage and signing mechanisms defined earlier.

8. Implementation Blueprint

This section specifies how to implement the Climate Alpha Platform as a working system. It assumes a modern data stack and a quantitative engineering team, and focuses on components, data flows, APIs, and service levels, not on re-deriving the mathematics. The blueprint is implementation-agnostic with respect to cloud provider, programming language, and vendor risk model; it defines logical components and interfaces that can be mapped onto any reasonable institutional stack. The implementation objectives are threefold. First, define a minimal but complete set of services that realise the platform: a bitemporal data spine, the Signal Engine, the Geospatial Risk Layer, the Constructor, and the Governance/Lineage layer. Second, specify the main data contracts between these services so that they can be developed and tested independently while still composing into a coherent system.

Third, fix basic operational characteristics such as latency targets, availability expectations, and a pragmatic deployment roadmap so the platform can run as part of production investment workflows rather than remaining a theoretical construct. At system level, the architecture decomposes into six logical layers. The Data Spine Layer provides a bitemporal tabular store for market, factor, fundamental, emissions, transition, taxonomy, and news data, a geospatial store for facility geometries and hazard rasters, and an object store for artefacts such as lineage bundles, model versions, audit packs, and logs. The Feature Layer consists of ETL or ELT jobs that transform raw spine tables into feature matrices Φ_t for each date and universe, including residual-return series, climate features, hazard-derived exposures, and interaction terms. The Signal Engine Layer implements residualisation, regularised regressions, Bayesian smoothing, and SHAP-style explainability, and outputs climate-aware $\Delta\mu$ signals with associated uncertainty measures. The Geospatial Risk Layer implements the facility registry, hazard–facility intersection, issuer-level exposure aggregation, and portfolio-footprint computation. The Constructor Layer hosts deterministic convex optimisers (mean–variance, ES/CVaR, TE-constrained QP/SOCP) and the discrete selection stack (QUBO plus simulated bifurcation or annealing, followed by MILP/QP parity refinement). The Governance, API, and Console Layer provides lineage and signing services, SR 11-7 and AI documentation generators, public-facing APIs (such as /scores, /weights, /exposures, /audit), and a web console for portfolio, risk, and model-governance users. In deployment, these logical layers can be realised as separate services or as coarser-grained components, but their responsibilities and interfaces remain as defined here.

The data spine is the foundation and must be built first, in a small but principled form. The tabular store is bitemporal: core tables such as prices_pt, factors_pt, fundamentals_pt, emissions_pt, taxonomy_pai_pt, and news_features_pt are keyed by asset or issuer and date, and each row carries both an as_of_date and an ingestion_timestamp. This ensures that all queries can be restricted to data that were actually available at decision time, preventing look-ahead. Prices and returns, factor exposures and sector dummies, financial ratios and balance-sheet items, emissions and transition indicators, taxonomy and PAI fields, and news-derived sentiment and topic features are all stored in this form. The geospatial store holds a facility_registry table in a spatial database such as PostGIS, with facility identifiers, issuer identifiers, geometries in a well-defined CRS, facility type, capacity, status, and validity intervals, and a raster store for hazard layers tagged by hazard type, scenario identifier, horizon year, resolution, CRS, and generation timestamp. An object store holds model artefacts, lineage bundles, signed audit packs, and large logs that do not fit naturally in structured tables. Ingestion follows a common pattern: raw vendor and internal files land in a staging area with metadata, are validated against schemas and basic sanity checks, normalised to canonical IDs and units, and then loaded into the _pt tables with the appropriate temporal metadata and schema versions; indexes on asset–date and issuer–date are built to support efficient joins; hazard rasters are tiled and indexed for spatial access. On top of the spine, a Feature Service builds the feature matrices. Conceptually, it exposes a function such as build_features(universe_id, as_of_date, config_id), which first resolves the investable universe at the given date (applying eligibility rules based on exchanges, size, liquidity, and listing age), then joins the relevant point-in-time tables for prices, factors, fundamentals, emissions, taxonomy, PAIs, news, and issuer hazard exposures. It then applies a pipeline defined by config_id that specifies winsorisation rules (for example, percentile caps by region and sector), normalisation (for example, z-scoring within region–sector), and the construction of interaction features (for example, hazard times leverage, emissions times valuation, news sentiment times sector). The service emits a feature matrix Φ_t with N_t rows and P columns, and writes metadata describing the schema, the as_of_date, the generation timestamp, the parent spine objects, and the pipeline version hash. With deterministic transformations and versioned configurations, the triplet (universe_id, as_of_date, config_id) is sufficient to reproduce any feature matrix. The Signal Engine then operates on feature matrices and residual returns. It exposes a function such as compute_signals(feature_matrix_id, returns_config_id, model_id) that retrieves Φ_t by its identifier, computes or fetches residual returns r_{i,t+1} based on the specified residualisation configuration (projecting raw excess returns off the factor and sector model), and fits the signal model indicated by model_id. In the baseline specification, this model is a regularised regression (elastic-net) of residual returns on features, potentially combined with a Bayesian state-space update to smooth parameters over time. The engine produces, for each asset, a climate-aware $\Delta\mu_t$ estimate interpreted as the posterior mean expected residual return, along with an uncertainty measure such as posterior variance or a derived confidence score. It also computes SHAP-style feature attributions per asset and date (or at least for the most material assets), making the signal locally explainable. The engine writes a signal_run record containing the signal_run_id, the linked feature matrix, the residualisation configuration, the model identifier, the as_of_date, and hashes or pointers for all input and output objects. This run identifier is what downstream services use to refer to signals in a controlled, traceable way. The Geospatial Risk Layer is implemented as a dedicated service responsible for hazard-to-exposure mapping and portfolio footprint computation. It must support computing facility-level hazard scores by intersecting facility geometries with hazard rasters for a chosen (hazard type, scenario, horizon) combination, aggregating those scores to issuer exposures using explicit weighting schemes (for example, by capacity, asset value, or revenue share), computing scenario deltas between different hazard or scenario pairs, and finally computing portfolio-level metrics by weighting issuer exposures with portfolio weights. Internally, the service may have methods that compute facility exposures given a hazard set and facility set configuration, return an identifier for facility_hazard_pt tables, then aggregate to a issuer_exposure_id that can be used by both the Feature Service (when adding geospatial features into Φ_t) and the Constructor (when building climate constraints). The key design choice is to treat geospatial outputs as versioned objects with IDs, not as ad-hoc columns created separately by each team; this ensures consistency across signals, constraints, and reporting.

The Constructor is implemented as a service that generates portfolios from universes, signals, risk objects, and constraint profiles. A core function such as construct_portfolio(universe_id, as_of_date, signal_run_id, constraint_profile_id) first resolves the universe and retrieves the $\Delta\mu$

vector and uncertainty estimates for that date and signal run. It then builds risk objects: a covariance matrix Σ_t derived from a risk model or historical returns, tracking-error matrices, or scenario return matrices if ES/CVaR optimisation is used. It then constructs a full constraint set based on the constraint profile: budget and bound constraints; long-only or long/short structure; factor neutrality relative to a benchmark; climate constraints defined via issuer exposure objects (such as caps on portfolio emissions intensity, hazard exposure, or PAI metrics); turnover constraints relative to the last weights_run_id; sector, country, liquidity, and concentration limits. The optimisation sequence typically starts with a continuous convex pre-optimisation, solving a standard mean–variance or ES/CVaR problem under all linear and convex constraints without cardinality constraints. This yields a baseline portfolio used as a sanity check and as a candidate solution. Optionally, the service then constructs a QUBO formulation with binary inclusion variables z_i and an energy function that rewards high $\Delta\mu$, penalises risk, encodes cardinality targets, and soft-penalises factor and climate deviations. A simulated bifurcation or annealing algorithm is run with logged seeds to explore combinatorial selections; the resulting selections are used as warm starts for a MILP/QP parity problem that links binary inclusion to continuous weights via constraints of the form $l_i z_i \leq w_i \leq u_i z_i$ and enforces the exact objective and constraint set of the continuous problem plus cardinality. The MILP/QP solutions are the authoritative portfolios for production, and the Constructor chooses the best feasible solution by objective value. It then emits weights per asset, risk and factor decompositions, climate footprints, PAI and taxonomy metrics, dual variables and binding flags for each constraint, and links all of this into a weights_run_id with a complete lineage bundle. Externally, the platform is exposed through a small, stable API surface and a console. A /scores endpoint accepts a universe, date, model, and feature configuration and triggers feature construction and signal computation, returning a signal_run_id and, optionally, the signals and top SHAP attributions. A /weights endpoint accepts a universe, date, signal_run_id, and constraint profile and triggers the entire construction pipeline, returning a weights_run_id plus weights, risk metrics, climate metrics, and diagnostics. A /exposures endpoint returns factor exposures, climate footprints, PAI and taxonomy metrics, and decomposition by issuer, sector, or region for a given weights_run_id. An /audit endpoint returns the lineage bundle, including model versions, data-object identifiers, solver configurations, hashes, and signatures for any given run. The console is a thin layer over these APIs, providing a portfolio view (weights, active weights, risk, TE, ES, MaxDD estimates), a climate view (maps with facility dots, hazard overlays, footprint breakdowns), a constraint view (constraint values, bounds, duals, binding status), and a governance view (model versions, lineage completeness, signature verification, change history, overrides). Portfolio managers, risk managers, validators, and compliance staff operate primarily through the console, while production workflows integrate the APIs into existing OMS and RMS systems. From an operational perspective, the implementation blueprint fixes indicative service-level objectives suitable for batch and low-frequency portfolio construction. A typical target is that /scores calls for universes of up to around a thousand assets return within a couple of seconds once data are cached, and /weights calls using only convex optimisation complete in a few seconds, while the

QUBO plus MILP/QP parity path can take up to tens of seconds for the same universe size. Availability for core services during trading hours should be at least 99.5 per cent; lineage completeness should be 100 per cent for all signal_run_id and weights_run_id; and production portfolios must exhibit zero hard constraint violations. These values can be tightened or loosened in a particular institution, but they provide concrete baselines. Finally, a staged deployment roadmap aligns with the architecture. In an initial phase, institutions implement the bitemporal data spine and residualiser, including the basic facility registry and at least one emissions and one hazard dataset, and compute factor-clean residual returns. In a second phase, they build the Feature and Signal services for issuer-level climate features and a baseline Constructor implementing QP or ES/CVaR with factor neutrality and simple climate constraints, and begin running the proof-of-value protocol on at least one region and universe. In a third phase, they extend the geospatial layer and discrete constructor: the facility registry and geospatial ingestion pipeline are expanded, facility–hazard intersections and issuer exposures are computed and plumbed into features and constraints, and the QUBO plus MILP/QP parity path is implemented.

In a fourth phase, governance and externalisation are hardened: lineage, logging, and SR 11-7 and AI documentation generation are formalised; post-quantum signing of lineage bundles and audit packs is integrated as needed; and the APIs are exposed to internal clients and integrated into production investment processes. This phased blueprint, combined with the mathematical and governance specifications in earlier sections, is sufficient for a competent quantitative engineering team to implement the Climate Alpha Platform in a regulated institutional environment.

9. Risks and Mitigations

This section enumerates key risks and the mechanisms the architecture uses to contain them.

9.1. Model risk

Risk 1: Spurious climate alpha (factor contamination).

- Issue: Apparent climate alpha may be driven by unrecognised factor bets (sector, value, quality, momentum).
- Mitigation:
- Residualization against a robust multi-factor model before signal estimation; explicit factor-neutral constraints in optimisation.
- Ex-post factor regressions on portfolio returns to verify near-zero factor betas.
- Ablation ladder: climate variants must beat the factor-only Base under identical constraints and costs. Risk 2: Overfitting high-dimensional climate features.
- Issue: Many features (hazards, emissions, news, interactions) vs limited history can lead to unstable coefficients.
- Mitigation:

- Elastic-net regularisation with nested cross-validation and stability selection.
- Bayesian smoothing over time; shrinkage toward priors; penalty for low-confidence signals in optimisation.
- Use of Deflated Sharpe/IR and multiple-testing corrections before promoting features.

Risk 3: Mis-specified geospatial aggregation.

- Issue: Wrong facility weights, coarse rasters, or misaligned scenarios can misrepresent exposures.
- Mitigation:
- Explicit, versioned rules for facility weights (capacity, revenue) and interpolation.
- Validation via independent geospatial checks and sanity tests (e.g., coastal facilities actually near coasts).
- Parallel ingestion of alternative hazard datasets or resolutions for robustness tests.

Risk 4: Optimiser pathologies.

- Issue: Heuristic or annealing-based optimisers can get stuck in poor local minima or violate constraints if misused.
- Mitigation:
- QUBO outputs never used directly; MILP/QP parity is the authoritative portfolio.
- Strict solver tolerances; monitoring of convergence flags; fallback to pure deterministic QP/MILP if heuristics misbehave.
- Logging of all solver states and seeds for forensic reproducibility.

9.2. Data risk

Risk 5: Vendor dependence and methodology drift.

- Issue: Heavy reliance on a single hazard or emissions provider; silent methodology changes can alter exposures.
- Mitigation:
- Dual-vendor ingestion where possible; capability to swap hazard/emissions datasets without architecture changes.
- Schema versioning; ingestion timestamps; alerts on dataset version changes.
- Periodic re-validation when vendors change methodology.

Risk 6: Data quality and coverage gaps.

- Issue: Missing or low-quality facility locations, emissions, or news; biases by sector or region.
- Mitigation:
    - Coverage dashboards showing data completeness by sector/region.

    - Explicit missingness indicators as features; avoid silent imputation that hides gaps.
    - Conservative default behaviour (e.g., excluding names from climate-alpha tilts when data quality is too low).

9.3. Operational and governance risk

Risk 7: Misuse of outputs.

- Issue: Portfolios used in contexts beyond design scope; constraints modified ad hoc without governance.
- Mitigation:
    - Clear scope statements in model documentation; enforcement of allowed use-cases.
    - Configuration management for constraint profiles; change control and approvals for material modifications.
- Audit logs for overrides and manual interventions.

Risk 8: Process failures and outages.

- Issue: Late or failed runs, partial data availability, or solver failures.
- Mitigation:
    - SLO monitoring; automated alerts for missed runs or incomplete pipelines.
    - Fallback modes (e.g., hold previous day's weights) with carefully defined usage conditions.
    - Disaster-recovery and redundancy for data and compute components.

9.4. Regulatory and legal risk

Risk 9: Greenwashing and disclosure inconsistency.

- Issue: Portfolios labelled as climate-aligned or SFDR Article 8/9 but optimisation does not actually enforce stated objectives.
- Mitigation:
    - Direct mapping of SFDR/IFRS/TCFD objectives into explicit constraints and objectives.
    - Use of the same exposure metrics in optimisation and reporting; no parallel "marketing" metrics.
    - Periodic external review of mandates vs portfolios.

Risk 10: Litigation and dispute over data and decisions.

- Issue: Clients or regulators challenge the basis of allocations or risk assessments.
- Mitigation:
    - End-to-end lineage and PQC-signed audit packs for each run.
    - Reproducibility: ability to re-run historical decisions exactly.
    - Clear documentation of limitations and areas of model uncertainty (see next section).

10. Limitations and Scope

The architecture is intentionally incomplete in several dimensions. These are not oversights; they are explicit boundaries of the current work.

1. No empirical performance claims.
2. The paper does not present realized backtests, Sharpe ratios, or p-values.
3. It specifies the Proof-of-Value Protocol that others must run; performance is a contingent outcome, not a claim here.
4. Hazard and climate-model uncertainty.
5. Hazard rasters and climate scenarios are themselves uncertain. The platform treats them as given inputs with known provenance but does not solve their epistemic uncertainty.
6. Alternative scenarios and datasets can be plugged into the same architecture; evaluating their realism is outside scope.
7. Simplified economic transmission channels.
8. The platform models climate exposures primarily via physical hazard intensities, emissions, and transition indicators.
9. It does not explicitly model full macroeconomic feedbacks, general equilibrium effects, or multi-period capital-allocation dynamics.
10. Limited asset-class coverage.
11. Focus is on listed equities (and possibly listed corporate bonds).
12. Sovereign bonds, structured products, private markets, and derivatives are only indirectly addressed, though the architecture could be extended.
13. No real-time intraday behavior.
14. Design assumes daily or lower-frequency signals and rebalancing.
15. Intraday optimisation, execution algorithms, and microstructure effects are outside scope.
16. Vendor and system heterogeneity.
17. The specification is vendor-agnostic and implementation-agnostic.
18. Mapping it into a specific institution's technology stack, legacy systems, and vendor mix requires additional engineering work not covered here.
19. Governance variations across jurisdictions.
20. The paper aligns with SR 11-7, IFRS S2, SFDR/ESRS, and EU AI Act at a generic level.
21. Local interpretations and additional national rules may require adaptation of documentation and control structures. Future work sits exactly on these boundaries: empirical validation, broadening asset-class scope, deeper macro scenario integration, and refined treatment of hazard model uncertainty.

11. Conclusion

The Climate Alpha Platform is defined as a governed architecture rather than a single trading rule or black-box strategy. It starts from a factor-residual view of returns, pushes climate information into the residual space, and only then estimates expected return differentials. Facility-level hazards are mapped into issuer and portfolio exposures in a way that is mathematically explicit and point-in-time consistent. Portfolio construction then operates on these residual signals and exposures via a quantum-inspired discrete selection layer, followed by deterministic convex refinement, so that the final result is both expressive and constraint-clean. The essential design choices are structural, not cosmetic. Residualization ensures that any "climate alpha" is orthogonal to known risk premia, so the platform cannot quietly rely on value, quality, or sector tilts while marketing them as climate effects. The geospatial layer forces physical risk to be spatially explicit, linking individual facilities to hazard rasters, scenarios, and time horizons, and reusing the same exposures consistently in both signal estimation and optimisation constraints. The constructor deliberately separates name selection from continuous sizing: QUBO-based exploration proposes combinatorial configurations, while MILP/QP parity enforces exact constraints, produces interpretable duals, and supports full audibility. Governance and lineage are not add-ons; they are core design primitives. Every inference and optimisation run is bound to a lineage bundle, a model version set, and a signed audit pack. Model-risk expectations (SR 11-7–type regimes), climate-disclosure standards (TCFD/IFRS S2, SFDR/ESRS), and AI-governance requirements (EU AI Act–style) are built into the architecture through explicit model inventories, technical files, reproducible PoV protocols, and post-quantum signatures. The result is a platform that can be interrogated by risk committees and supervisors at the same level of depth as internal risk models, rather than an opaque ESG overlay. Deliberately, the paper stops before empirical performance claims. It does not present Sharpe ratios or production backtests. Instead, it specifies a Proof-of-Value Protocol—universe definitions, data windows, cost models, constraint sets, metrics, and statistical tests—that any institution can implement on its own data. Performance becomes an empirical outcome of that protocol, not a marketing promise embedded in this document. This separation makes the contribution falsifiable: if an implementation fails the PoV gates, the architecture is exposed to critique; if it passes, the uplift is attributable to a well-defined stack rather than ad-hoc tweaks. For practitioners, the document can be read directly as a build plan. Implement the bitemporal data spine and lineage semantics; construct the residualizer and Signal Engine with explicit factor baselines; integrate the geospatial layer for facility-to-portfolio climate mapping; layer the constructor with QUBO exploration and MILP/QP parity; and wrap the whole stack in SR 11-7 / AI-compliant governance with signed audit artefacts. Once that implementation is in place, the PoV protocol defines whether the resulting system delivers the targeted information-ratio uplift and drawdown improvements under realistic constraints. If it does, the institution acquires a climate-aware investing capability that is simultaneously technically robust, spatially explicit, and regulator-ready; if it does not, the failure is diagnosable at each layer of the architecture, rather than hidden in a monolithic black box.

12. References

1. Federal Reserve Board. “SR 11-7: Supervisory Guidance on Model Risk Management.” (2011). https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm
2. IFRS Foundation. “IFRS S2 Climate-related Disclosures (Issued Standard, Part A).” (2023). https:// www.ifrs.org/content/dam/ifrs/publications/pdf-standards issb/english/2023/issued/part-a/issb-2023-a-ifrs-s2-climate-related-disclosures.pdf
3. IFRS Foundation. “IFRS S2 — Navigator page.” (access). https://www.ifrs.org/issued standards/ifrs-sustainability-standards-navigator/ifrs-s2-climate-related-disclosures/
4. IFRS Foundation. “Greenhouse Gas Emissions Disclosure requirements applying IFRS S2 (Educational Material).” (May 2025). https://www.ifrs.org/content/dam/ifrs/supporting-implementation/ifrs-s2/ghg-ifrs-s2-educational-material.pdf
5. EFRAG. “ESRS E1 Climate change — Delegated Act Annex.” (2023). https://www.efrag.org/AssetsDownloadassetUrl=%2Fsites%2Fwebpublishing%2FSiteAssets%2FESRS%2520E1%2520Delegated-act-2023-5303-annex-1_en.pdf
6. European Commission. “Sustainability-related disclosure in the financial services sector (SFDR).” (portal). https://finance.ec.europa.eu/sustainable-finance/disclosures/sustainability-related-disclosure-financial-services-sector_en
7. EUR-Lex. “Commission Delegated Regulation (EU) 2022/1288 (SFDR RTS).” (2023 consolidated). https://eur-lex.europa.eu/eli/reg_del/2022/1288/2023-02-20/eng
8. WRI. “Aqueduct Water Risk Atlas.” (tool). https://www.wri.org/applications/aqueduct/water-risk-atlas/
9. WRI. “Aqueduct 4.0 Current and Future Global Maps Data.” (2023). https://www.wri.org/data/aqueduct-global-maps-40-data
10. WRI. “Aqueduct 4.0 — Data & Methodology (GitHub).” https://github.com/wri/Aqueduct40
11. Copernicus C3S/ECMWF. “ERA5 hourly data on single levels (dataset).” https://cds.climate.copernicus.eu/ datasets/reanalysis-era5-single-levels
12. ECMWF. “ERA5 — dataset overview.” https://www.ecmwf.int/en/forecasts/dataset/ecmwf-reanalysis-v5
13. NASA Earthdata. “FIRMS — Fire Information for Resource Management System.” https:// firms.modaps.eosdis.nasa.gov/
14. NASA/GSFC MODIS. “MCD64A1 Burned Area Product.” https://ladsweb.modaps.eosdis.nasa.gov/ missions-and-measurements/products/MCD64A1
15. Google Earth Engine Catalog. “MODIS MCD64A1.061 Burned Area Monthly Global 500 m.” https:// developers.google.com/earth-engine/datasets/catalog/MODIS_061_MCD64A1

16. EPA. “Greenhouse Gas Reporting Program (GHGRP) — FLIGHT facility-level data.” https://www.epa.gov/ ghgreporting
17. EPA. “GHGRP Data Sets.” https://www.epa.gov/ghgreporting/data-sets
18. GHG Protocol. “Corporate Accounting and Reporting Standard.” https://ghgprotocol.org/sites/default/files/ standards/ghg-protocol-revised.pdf
19. GHG Protocol. “Corporate Standard (overview page).” https://ghgprotocol.org/corporate-standard
20. Kenneth R. French. “Data Library (Factors & Portfolios).” https://mba.tuck.dartmouth.edu/pages/faculty/ ken.french/data_library.html
21. Kenneth R. French. “Description of Fama/French 5 Factors (2x3).” https://mba.tuck.dartmouth.edu/pages/ faculty/ken.french/Data_Library/f-f_5_factors_2x3.html
22. Markowitz, H. “Portfolio Selection.” Journal of Finance (1952). PDF (Cowles Foundation). https:// cowles.yale.edu/sites/default/files/files/pub/d03/d0317.pdf
23. Rockafellar, R.T., Uryasev, S. “Optimization of Conditional Value-at-Risk.” Journal of Risk (2000). https:// www.ise.ufl.edu/uryasev/files/2011/11/2000-CVaR1-JR.pdf
24. Almgren, R., Chriss, N. “Optimal Execution of Portfolio Transactions.” Journal of Risk (2001; working 2000). https://www.math.nyu.edu/faculty/chriss/optliq.pdf
25. Lundberg, S.M., Lee, S.-I. “A Unified Approach to Interpreting Model Predictions (SHAP).” NeurIPS 2017. https://arxiv.org/pdf/1705.07874.pdf
26. Goto, H. et al. “Combinatorial optimization by simulating adiabatic bifurcations in nonlinear Hamiltonian systems (Simulated Bifurcation).” Nature Scientific Reports (2019). https://www.nature.com/articles/ s41598-019-49172-3.pdf
27. Venturelli, D. et al. “Quantum optimization for real-world portfolio allocation problems.” (QUBO/Ising, review/use case). https://link.springer.com/article/10.1007/s42484-019-00013-7
28. NIST. “FIPS 203 — Module-Lattice-Based Key-Encapsulation Mechanism (ML-KEM).” https:// csrc.nist.gov/pubs/fips/203/final
29. NIST. “FIPS 204 — Module-Lattice-Based Digital Signature Algorithm (ML-DSA).” https://csrc.nist.gov/ pubs/fips/204/final
30. TCFD. “Final Recommendations of the Task Force on Climate-related Financial Disclosures.” (2017). https:// www.fsb-tcfd.org/wp-content/uploads/2017/06/FINAL-TCFD-Report-062817.pdf
31. TCFD. “Implementing the Recommendations — Annex.” (2017). https://www.fsb-tcfd.org/wp-content/ uploads/2017/06/FINAL-TCFD-Annex-062817.pdf

Appendix A. Governance Anchors and Publication-Grade Credibility Controls

This appendix adds governance and credibility scaffolding expected in regulated environments and high-quality research dissemination. It is compatible with supervisory model-risk frameworks and research reproducibility expectations.

- Model Risk Governance: Align lifecycle controls (development, validation, governance) with supervisory guidance (SR 11-7 / OCC MRM).
- Reproducibility: Define point-in-time datasets, parameter registries, code versioning, and independent replication steps.
- Data Lineage: Every derived dataset MUST be traceable to raw sources with timestamps and transformation hashes.
- Disclosure: Document intended use, limitations, and failure modes; separate hypothesis generation from confirmatory claims.

Appendix B. Formal Derivations and Correctness Proofs

B.1 Residualisation as an orthogonal projection

Let X be the cross-sectional factor exposure matrix ($N \times K$) and r the vector of asset excess returns ($N \times 1$). Define $\text{beta} = (X^T X)^{-1} X^T r$ and residual $e = r - X\,\text{beta}$. Then $X^T e = 0$ because $X^T e = X^T r - X^T X\,\text{beta} = 0$. Therefore e lies in the orthogonal complement of the column space of X. This provides exact factor neutrality (with respect to X) by construction.

B.2 QUBO mapping for cardinality-constrained selection

With binary z in $\{0,1\}^N$ and target K selected names, enforce cardinality via penalty $\text{lambda}(1^T z - K)^2$. The resulting objective is quadratic in z and can be solved by quantum-inspired heuristics or classical QUBO solvers. A deterministic MILP/MIQP parity solver SHOULD be used as a correctness backstop for small-to-mid universes.

Appendix C. Worked Numerical Examples (Step-by-Step Calculations)

C.1 Cross-sectional residualisation (OLS) — worked example

Given exposures X (intercept + 2 factors) and excess returns r:

Asset	1 (Intercept)	Factor 1	Factor 2
i=1	1.000	0.200	-0.100
i=2	1.000	-0.100	0.050
i=3	1.000	0.400	0.100
i=4	1.000	0.000	-0.200
Asset		r_i	
i=1		0.0120	
i=2		-0.0080	
i=3		0.0100	
i=4		0.0020	
Coefficient		Value	
Intercept		-0.002545	
Factor 1		0.042655	
Factor 2		-0.032345	
Asset		Residual e_i	
i=1		0.002779	
i=2		0.000428	
i=3		-0.001283	
i=4		-0.001924	
Orthogonality component		X^T e	
Intercept		0.00000000	
Factor 1		-0.00000000	
Factor 2		-0.00000000	

C.2 CVaR (Expected Shortfall) — worked example

Scenario	Loss L_s
s=1	0.0100
s=2	0.0200
s=3	0.0150
s=4	0.0300
s=5	0.0050
s=6	0.0180
s=7	0.0400
s=8	0.0120
s=9	0.0220
s=10	0.0700

VaR_95 = 0.0565; CVaR_95 = 0.0700 (mean of tail losses >= VaR).

Appendix D. Loophole Analysis and Specification Hardening

This appendix enumerates common failure modes in climate-aware strategies and prescribes MUST-level controls: factor leakage, look-ahead bias, coverage drift, constraint gaming, scenario misuse, and governance gaps. Each item is written to be testable.

Factor leakage

- MUST enforce residualisation; MUST monitor post-trade factor exposures within tolerance bands; MUST log and explain any breach.

Look-ahead / leakage

- MUST use point-in-time data; MUST attest ingestion times; MUST run leakage tests on every pipeline change.

Coverage drift

- MUST publish coverage KPIs for facilities, hazards, and emissions; MUST degrade gracefully via documented fallbacks.

Constraint gaming

- MUST expose binding constraints and dual variables; SHOULD provide deterministic MIQP parity for small universes.

Scenario misuse

- MUST document scenario sources and intended interpretation; MUST separate physical vs transition scenarios; MUST audit scenario parameter changes.

Appendix E. Figures Extracted from the Original Manuscript

Given a facility geometry $G_{i,f}$ and a hazard raster h_k defined over a grid $\mathcal{S}$, the platform defines:

- A binary mask $l_{f,s} \in \{0,1\}$ indicating whether grid cell s intersects $G_{i,f}$.
- Facility-level hazard exposure:

$$e_{i,f,k} = \frac{\sum_{s \in \mathcal{S}} h_k(s)\, l_{f,s}}{\sum_{s \in \mathcal{S}} l_{f,s}}$$

if the denominator is non-zero; otherwise a defined fallback (nearest-neighbour interpolation or flagged missing).

Extensions:

- **Weighted means** using facility capacity or asset value if those represent better impact proxies than area.
- **Non-linear transformations** (e.g., thresholds, percentiles) to capture tail risks.

Intersection is deterministic for a given set of geometries, rasters, and interpolation rules. The transformation code (and its version) is part of lineage (next subsection).

Figure E1: Extracted figure from original PDF (page 14).

The Signal Engine consumes a feature matrix Φ_t for each date t. That matrix is built from the spine via a sequence of defined transformations.

Feature classes:

1. **Financial and factor features**
   - Standard factor exposures (size, value, momentum, profitability, investment, quality, low-volatility).
   - Leverage, profitability ratios, growth metrics, payout ratios.
   - Sector and region fixed effects (dummies).
2. **Climate physical-risk features**
   - Facility-level $e_{i,f,k}$ aggregated to issuer level:

$$E_{i,k} = \sum_f \omega_{i,f} e_{i,f,k}, \quad \sum_f \omega_{i,f} = 1$$

where $\omega_{i,f}$ are weights (capacity, revenue share, asset value).

- Scenario-delta exposures: $\Delta E_{i,k}^{(s_1,s_2)} = E_{i,k}^{(s_2)} - E_{i,k}^{(s_1)}$.
- Peer-relative measures: z-scores of $E_{i,k}$ within sector/region.

Figure E2: Extracted figure from original PDF (page 15).

- $t \in \{1, \ldots, T\}$ index discrete rebalancing dates.
- $i \in \{1, \ldots, N_t\}$ index assets existent and investable at t.
- $y_{i,t+1}$: excess return of asset i from t to $t+1$, in portfolio currency.
- $\mathbf{y}t + 1 \in \mathbb{R}^{N_t}$: *vector stacking* $yi, t+1$.
- $\mathbf{w}_t \in \mathbb{R}^{N_t}$: portfolio weights after rebalancing at t.
- $\mathbf{b}_t \in \mathbb{R}^{N_t}$: benchmark weights at t.
- $\mathbf{X}_t \in \mathbb{R}^{N_t \times K}$: factor-exposure matrix at t (market, size, value, profitability, investment, sector, etc.).
- $\boldsymbol{\phi}_t \in \mathbb{R}^{N_t \times P}$: climate/geospatial/news feature matrix at t.
- $\boldsymbol{\Sigma}_t \in \mathbb{R}^{N_t \times N_t}$: covariance matrix of returns over a chosen horizon (e.g., one month, one quarter).

Figure E3: Extracted figure from original PDF (page 17).

$$\mathbf{y}_{t+1} = \mathbf{X}_t \boldsymbol{\beta} t + \boldsymbol{\varepsilon} t + 1,$$

where:

- $\boldsymbol{\beta}_t \in \mathbb{R}^K$ are factor prices at t.
- $\boldsymbol{\varepsilon}_{t+1} \in \mathbb{R}^{N_t}$ are residuals (idiosyncratic returns).

Figure E4: Extracted figure from original PDF (page 17).

Define the **baseline residual alpha** of asset i across a backtest window $t = 1, \ldots, T$:

$$\hat{\alpha} i = \frac{1}{T} \sum t = 1^T r_{i,t+1}.$$

To assess its significance, the platform uses a HAC (Newey–West) estimator of the variance of $\hat{\alpha} i$*, allowing for serial correlation in* $\{ri, t\}$. The t-statistic is:

$$t_i = \frac{\hat{\alpha}_i}{\widehat{\mathrm{se}}(\hat{\alpha}_i)}.$$

In production, residual alpha is not used directly for ranking; instead, it provides a baseline for **signal quality** checks and ablations (climate-augmented signals must beat this baseline after costs). ecoquant2.pdf

Figure E5: Extracted figure from original PDF (page 18).

A standard OLS estimate is:

$$\widehat{\beta}_t = (\mathbf{X}_t^\top \mathbf{X}_t)^{-1}\mathbf{X}t^\top \mathbf{y}t + 1.$$

Define the **projection matrix** onto the factor space:

$$\mathbf{P}_t = \mathbf{X}_t(\mathbf{X}_t^\top \mathbf{X}_t)^{-1}\mathbf{X}_t^\top,$$

and the residualiser:

$$\mathbf{M}_t = \mathbf{I} - \mathbf{P}_t.$$

Then the residual return vector is:

$$\mathbf{r}_{t+1} \equiv \mathbf{M}t\mathbf{y}t + 1.$$

Properties:

- Idempotence and symmetry: $\mathbf{M}_t^2 = \mathbf{M}_t$, $\mathbf{M}_t^\top = \mathbf{M}_t$.
- Orthogonality: $\mathbf{X}t^\top \mathbf{r}t + 1 = 0$.

Thus, any later expectation or prediction based on $\mathbf{r}_{t+1}$ is, by construction, factor-neutral.

Figure E6: Extracted figure from original PDF (page 18).

On each date t, the platform estimates a cross-sectional mapping from climate/geospatial/news features to future residual returns. This is the "Signal Engine". ecoquant2.pdf

5.3.1. Linear regularised climate-alpha model
For each asset i:

$$r_{i,t+1} = \phi_{i,t}^{\top}\gamma_t \eta_{i,t+1},$$

where:

- $\phi_{i,t} \in \mathbb{R}^P$: feature vector (climate exposures, emissions, targets, hazard deltas, news, interactions).
- $\gamma_t \in \mathbb{R}^P$: time-$t$ feature prices for climate-linked characteristics.
- $\eta_{i,t+1}$: residual noise after climate features.

The platform estimates γ_t via elastic-net regularisation:

$$\widehat{\gamma}t = \arg\min \gamma \in \mathbb{R}^P \left\{ \sum_{i=1}^{N_t} \left(r_{i,t+1} - \phi_{i,t}^{\top}\gamma\right)^2 \lambda_1 \|\gamma\|_1 \lambda_2 \|\gamma\|_2^2 \right\},$$

with penalties $\lambda_1, \lambda_2 \geq 0$ chosen via nested cross-validation and stability selection.

The one-step-ahead **climate-alpha signal** is then:

$$\widehat{\Delta\mu} i, t = \phi i, t^{\top} \widehat{\gamma}_t.$$

This is a prediction of **residual** excess return, conditional on climate and related features, and is factor-neutral by construction because the target is $\mathbf{r}t+1$, *not* $\mathbf{y}t+1$.

Figure E7: Extracted figure from original PDF (page 19).

To avoid over-reacting to noise in any single cross-section, the platform embeds γ_t in a simple Bayesian state-space:

- Prior: $\gamma t \sim \mathcal{N}(\mathbf{m}t|t-1, \mathbf{V}_{t|t-1})$.
- Observation model: $\mathbf{r}_{t+1} = \mathbf{\Phi} t \gamma t + \epsilon t + 1$,
  where $\mathbf{\Phi} t$ stacks $\phi i, t$ row-wise and $\epsilon t + 1 \sim \mathcal{N}(0, \mathbf{R})$.

Posterior update (Bayesian linear regression):

$$\mathbf{V}t|t = \left(\mathbf{V}t|t-1^{-1}\mathbf{\Phi}_t^\top \mathbf{R}^{-1}\mathbf{\Phi}_t\right)^{-1},$$

$$\mathbf{m}t|t = \mathbf{V}t|t\left(\mathbf{V}t|t-1^{-1}\mathbf{m}t|t-1\mathbf{\Phi}t^\top \mathbf{R}^{-1}\mathbf{r}t+1\right).$$

The posterior mean $\mathbf{m}_{t|t}$ is used as a smoothed estimator of γt*, and the posterior variance* $\mathbf{V}t|t$ quantifies parameter uncertainty. This is exploited in portfolio construction by down-weighting signals with large posterior uncertainty.

The resulting **posterior signal** is:

$$\widetilde{\Delta\mu} i, t = \phi i, t^\top \mathbf{m}_{t|t},$$

with uncertainty measure

$$\mathrm{Var}(\widetilde{\Delta\mu} i, t) = \phi i, t^\top \mathbf{V}t|t\phi i, t.$$

The constructor can transform these into **confidence-scaled signals**: for example,

$$\Delta\mu^{\mathrm{eff}} i, t = \widetilde{\Delta\mu} i, t \cdot \left(1 + \kappa \cdot \frac{\widetilde{\Delta\mu} i, t}{\sqrt{\mathrm{Var}(\widetilde{\Delta\mu} i, t)}}\right)^{-1},$$

for some $\kappa > 0$, shrinking aggressive forecasts with low statistical support.

Figure E8: Extracted figure from original PDF (page 20).

Given:

- Facility geometry $G_{i,f}$ for facility f of issuer i.
- Hazard raster $h^{(k,s,y)}(\mathbf{s})$ for hazard type k, scenario s, horizon y, defined on grid cells $c \in \mathcal{C}$ with centroids $\mathbf{s}_c$.

Define indicator $l_{i,f,c} = 1$ if cell c intersects $G_{i,f}$, 0 otherwise. Then:

$$e_{i,f}^{(k,s,y)} \begin{cases} \dfrac{\sum_{c\in\mathcal{C}} h^{(k,s,y)}(\mathbf{s}c)\, li, f, c\, w_{i,f,c}}{\sum_{c\in\mathcal{C}} l_{i,f,c}\, w_{i,f,c}} & \text{if } \sum_c l_{i,f,c} w_{i,f,c} > 0, \\ \text{NA} & \text{otherwise,} \end{cases}$$

where $w_{i,f,c}$ are cell weights (e.g., area, asset value, or capacity).

Interpolation fallback: if no cells intersect directly (small facilities), nearest-neighbour or kernel interpolation is used with explicit radius r, documented in lineage.

Figure E9: Extracted figure from original PDF (page 21).

Issuer-level exposures aggregate facility hazards:

$$E_i^{(k,s,y)} = \sum_f \omega_{i,f} e_{i,f}^{(k,s,y)}, \quad \sum_f \omega_{i,f} = 1,$$

with $\omega_{i,f}$ proportional to facility capacity, revenue share, or asset value. Different weighting schemes can be defined as separate exposure families.

Scenario deltas compare scenarios s_1, s_2 or horizons y_1, y_2:

$$\Delta E_i^{(k;s_1 \to s_2, y_1 \to y_2)} = E_i^{(k,s_2,y_2)} - E_i^{(k,s_1,y_1)}.$$

These deltas represent sensitivity to worsened physical risk and can either enter the feature matrix Φ or define constraints.

Figure E10: Extracted figure from original PDF (page 21).

$$\mathcal{E}^{(k,s,y)}(\mathbf{w}t) = \sum i w_{i,t} E_i^{(k,s,y)}.$$

Figure E11: Extracted figure from original PDF (page 21).

$$\mathcal{E}^{(k,s,y)}(\mathbf{w}_t) \leq \bar{\mathcal{E}}^{(k,s,y)},$$

for specified k, s, y and caps $\bar{\mathcal{E}}^{(k,s,y)}$ defined by risk policy or product design.

Because both features and constraints are built on the same $E_i^{(k,s,y)}$, the system is self-consistent: the feature “high flood risk” and the portfolio constraint “limit portfolio flood exposure” refer to identical underlying calculations.

Figure E12: Extracted figure from original PDF (page 21).

When tail risk is more relevant than variance, the platform uses ES (CVaR) optimisation (Rockafellar–Uryasev). Let $\mathbf{R}$ be a matrix of scenario returns (rows = scenarios, columns = assets). Portfolio scenario returns: $\mathbf{r}^{\text{port}} = \mathbf{R}\mathbf{w}$. CVaR at level α is minimised by:

$$\begin{aligned} &\min_{\mathbf{w},\xi,\mathbf{u}} \xi + \frac{1}{(1-\alpha)S} \sum_{s=1}^{S} u_s \\ &\text{s.t.} u_s \geq -r_s^{\text{port}} - \xi, \quad u_s \geq 0, \quad s = 1, \ldots, S, \\ &\quad \text{same constraints on } \mathbf{w} \text{ as above.} \end{aligned}$$

Objective can be augmented with a linear reward $+\gamma \mathbf{w}^\top \boldsymbol{\Delta\mu}_t$. This yields a linear program (LP) plus linear constraints, which is efficiently solvable.
ecoquant2.pdf

Figure E13: Extracted figure from original PDF (page 22).

An alternative is to maximise residual alpha subject to a tracking-error bound against benchmark $\mathbf{b}_t$:

- Define active weights $\mathbf{u} = \mathbf{w} - \mathbf{b}_t$.
- Define tracking-error covariance $\boldsymbol{\Sigma}^{\mathrm{TE}} = \boldsymbol{\Sigma} - \boldsymbol{\Sigma}\mathbf{b}\mathbf{b}^\top\boldsymbol{\Sigma}$ or similar.

Problem:

$$
\begin{aligned}
&\max_{\mathbf{u}} \mathbf{u}^\top \boldsymbol{\Delta}\boldsymbol{\mu}_t \\
&\ \text{s.t.} \sqrt{\mathbf{u}^\top \boldsymbol{\Sigma}^{\mathrm{TE}} \mathbf{u}} \leq \tau, \\
&\qquad \mathbf{1}^\top(\mathbf{b}_t + \mathbf{u}) = 1, \\
&\qquad \mathbf{l} \leq \mathbf{b}_t + \mathbf{u} \leq \mathbf{u}, \\
&\qquad \mathbf{F}t^\top \mathbf{u} = \mathbf{0}, \\
&\qquad \mathcal{E}^{(k,s,y)}(\mathbf{b}t + \mathbf{u}) \leq \bar{\mathcal{E}}^{(k,s,y)}, \\
&\qquad \|\mathbf{u} - (\mathbf{w}t - 1 - \mathbf{b}t - 1)\|_1 \leq \Gamma.
\end{aligned}
$$

This yields a convex problem (second-order cone program if TE constraint kept in norm form) and gives direct control over active risk.

Figure E14: Extracted figure from original PDF (page 22).

$$
\begin{aligned}
&\max_{\mathbf{w}} \mathbf{w}^\top \boldsymbol{\Delta}\boldsymbol{\mu}_t \frac{\lambda}{2} \mathbf{w}^\top \boldsymbol{\Sigma}_t \mathbf{w} \\
&\ \text{s.t.} \mathbf{1}^\top \mathbf{w} = 1 \quad (\text{budget}), \\
&\qquad \mathbf{l} \leq \mathbf{w} \leq \mathbf{u} \quad (\text{bounds}), \\
&\qquad \mathbf{F}_t^\top (\mathbf{w} - \mathbf{b}t) = \mathbf{0} \quad (\text{factor neutrality}), \\
&\qquad \mathcal{E}^{(k,s,y)}(\mathbf{w}) \leq \bar{\mathcal{E}}^{(k,s,y)} \quad \forall (k,s,y) \in \mathcal{K} \quad (\text{climate caps}), \\
&\qquad \|\mathbf{w} - \mathbf{w}t - 1\|_1 \leq \Gamma \quad (\text{turnover}), \\
&\qquad \mathbf{L}\mathbf{w} \leq \mathbf{c} \quad (\text{liquidity/borrow, etc.}).
\end{aligned}
$$

Figure E15: Extracted figure from original PDF (page 22).

Introduce binary vector $\mathbf{z} \in \{0,1\}^{N_t}$:

- $z_i = 1$ if asset i is eligible for non-zero weight.
- $z_i = 0$ otherwise.

Initial approximations treat weights as $w_i \approx z_i / K$, where K is target number of holdings. The QUBO energy encodes alpha, risk, factor neutrality, and climate caps:

$$
\begin{aligned}
E(\mathbf{z}) = & -\sum_i \Delta\mu_{i,t} z_i \lambda \sum_{i,j} \Sigma_{ij,t} z_i z_j \\
& + A\left(\sum_i z_i - K\right)^2 B \left\|\mathbf{F}_t^\top \mathbf{z}\right\| 2^2 \\
& + C \left\|\mathcal{E}(\mathbf{z}) - \bar{\mathcal{E}}\right\| +^2 \mathbf{q}^\top \mathbf{z} c.
\end{aligned}
$$

Interpretation:

- First term rewards inclusion of high-alpha names.
- Second term penalises variance of equal-weighted portfolio.
- Penalty A: enforces cardinality $\sum z_i \approx K$.
- Penalty B: penalises factor exposures (enforces approximate neutrality).
- Penalty C: penalises climate-constraint violations (using positive-part norm).
- $\mathbf{q}$: linear penalty/bonus vector incorporating illiquidity, borrow costs, and other preferences.
- c: constant offset.

All quadratic and linear terms are collected into standard QUBO form $E(\mathbf{z}) = \mathbf{z}^\top Q \mathbf{z} + \mathbf{c}^\top \mathbf{z} + c$.

Figure E16: Extracted figure from original PDF (page 23).

- Binary variables z_i with cardinality and inclusion constraints.
- Continuous weights w_i bounded by z_i:

$$
l_i z_i \le w_i \le u_i z_i,
$$

ensuring $w_i = 0$ when $z_i = 0$.

- Objective equivalent to mean–variance or CVaR formulation with same $\Delta\mu, \Sigma, \mathcal{E}$.
- Exact factor and climate constraints; turnover and liquidity constraints as in Section 5.5.

MILP/QP solution:

$$
\max_{\mathbf{w},\mathbf{z}} \quad \mathbf{w}^\top \boldsymbol{\Delta\mu}_t \frac{\lambda}{2} \mathbf{w}^\top \boldsymbol{\Sigma}_t \mathbf{w}, \quad \text{s.t. constraints as above.}
$$

QUBO-derived $\mathbf{z}$ are used as **warm starts** for MILP. If the QUBO path finds a good approximate selection, MILP fine-tunes weights and guarantees constraint compliance.

Guarantees:

- Any portfolio used in production must come from the **parity MILP/QP** or an equivalent convex program, not directly from heuristic QUBO; QUBO is admissible as a pre-selector only.
- MILP/QP solvers expose dual variables for continuous constraints, allowing full diagnostic reporting (binding sets, shadow prices).

Figure E17: Extracted figure from original PDF (page 24).

For each date t, model $g_t : \phi i, t \mapsto \widetilde{\Delta\mu} i, t$. Kernel-SHAP approximates Shapley values of each feature k for each asset i:

$$\widetilde{\Delta\mu} i, t \approx \phi 0, t + \sum_{k=1}^{P} s_{i,t,k},$$

where:

- $\phi_{0,t}$ is a baseline value (e.g., average prediction under reference distribution).
- $s_{i,t,k}$ is the SHAP value for feature k on asset i.

These values satisfy, approximately, local accuracy and additivity properties. The platform stores top-K positive and negative contributors for each asset:

- "Drivers" list, e.g., high drought exposure in region X, low emissions intensity, improving news sentiment.
- Confidence information derived from model variance and SHAP variance across background samples.

Figure E18: Extracted figure from original PDF (page 24).

From QP/LP/MILP solutions, the platform extracts dual variables:

- For equality constraints $A\mathbf{w} = \mathbf{b}$, duals $\boldsymbol{\nu}$.
- For inequality constraints $G\mathbf{w} \leq \mathbf{h}$, duals $\boldsymbol{\lambda} \geq 0$.

Dual magnitudes indicate shadow prices:

- High $|\nu_j|$: small relaxation of constraint j would materially improve objective.
- High λ_k: inequality k is binding and impactful.

Binding sets are flagged explicitly:

- Factor-neutrality constraints that bind.
- Climate caps that bind (e.g., flood exposure at cap).
- Turnover, liquidity, or sector caps that bind.

These diagnostics let risk managers understand **why** particular names were chosen or excluded and which constraints drove the final shape of the portfolio. ecoquant2.pdf

Figure E19: Extracted figure from original PDF (page 25).

Explainability for climate risk is completed by tracing paths:

- From portfolio-level footprint metric $\mathcal{E}^{(k,s,y)}(\mathbf{w})$ back to per-issuer exposures $E_i^{(k,s,y)}$.
- From E_i back to facility exposures $e_{i,f}^{(k,s,y)}$ and their geometries.
- From facility exposures back to raster cells $h^{(k,s,y)}(\mathbf{s}_c)$.

Figure E20: Extracted figure from original PDF (page 25).

The methodology is designed to support strict internal consistency and empirical validation. ecoquant2.pdf

- **Internal consistency:** the same hazard exposures $E_i^{(k,s,y)}$ drive both signals (as features in ϕ) and constraints ($\mathcal{E}^{(k,s,y)}(\mathbf{w})$). No duplicated or contradictory climate metrics.
- **Robustness checks:** multiple feature sets, hazard datasets, and scenario choices are run through the same residualisation and construction pipeline; performance and constraint satisfaction are compared.
- **Ablation ladder:**
  - Base (factor-only signals, standard optimisation).
  - +Climate Signals (add climate features to Signal Engine).
  - +Geospatial Layer (add facility-based hazard features and climate constraints).
  - +Quantum-Inspired Construction (add QUBO + parity MILP).

Only additions that demonstrably improve residual alpha and drawdown metrics under the same cost and constraint assumptions are retained.

Figure E21: Extracted figure from original PDF (page 25).

- Signals: add facility-derived hazard exposures and hazard deltas $E_i^{(k,s,y)}$, $\Delta E_i^{(k;s_1 \to s_2, y_1 \to y_2)}$ into Φ_t.
- Constraints:
  - Introduce portfolio-level linear climate caps $\mathcal{E}^{(k,s,y)}(\mathbf{w}_t) \leq \bar{\mathcal{E}}^{(k,s,y)}$.
  - Caps set relative to benchmark or absolute risk targets.
- Construction: still classical (QP/LP/SOCP), no QUBO yet.

Figure E22: Extracted figure from original PDF (page 33).

Backtests are run under **implementation-grade** constraints. ecoquant2.pdf

Transaction costs.

- Per-period net return:

$$R_t^{\text{net}} = R_t^{\text{gross}} - \sum_i \left(c^{\text{lin}} i, t |\Delta w i, t| + c^{\text{imp}} i, t (\Delta w i, t)^2 \right).$$

- $c_{i,t}^{\text{lin}}$: linear spread/fee component (bps per unit turnover).
- $c_{i,t}^{\text{imp}}$: convex impact component increasing with trade size vs ADV and volatility.

Liquidity and capacity constraints.

- Per-name trade limit:

$$|\Delta w_{i,t}| \cdot \text{AUM} \leq \theta_{\text{ADV}} \cdot \text{ADV}_{i,t}$$

for θ_{ADV} typically around 10–20%.

- Maximum position size as fraction of free float (e.g., < 2–5%).
- Exclusion of names failing minimum ADV over a lookback window.

Shorting and leverage.

- Long-only version: $0 \leq w_{i,t} \leq u_i$, $\sum_i w_{i,t} = 1$.
- Long–short version (if considered):
  - Gross exposure cap: $\sum_i |w_{i,t}| \leq G_{\max}$.
  - Net exposure band: $|\sum_i w_{i,t} - 1| \leq \delta$.
  - Borrow availability and cost model for shorts.

Figure E23: Extracted figure from original PDF (page 34).

Turnover control.

- L1 turnover constraint:

$$\|\mathbf{w}t - \mathbf{w}t - 1\|_1 \leq \Gamma$$

with Γ calibrated to match realistic annual turnover (e.g., 100–150% annual).

Figure E24: Extracted figure from original PDF (page 35).

Risk and exposure constraints.

- Factor-neutrality vs benchmark:
  - Active factor exposures constrained to a small band around zero.
- TE or volatility bands vs benchmark for TE-constrained variants.
- Sector and country caps (e.g., ±x% active vs benchmark).

These assumptions are fixed across ladder rungs; only climate signals and constructors change.

Figure E25: Extracted figure from original PDF (page 35).

1. **Factor-clean alpha.**

   - Regress net portfolio return on factor returns:

$$R_t^{\text{net}} - R_t^{\text{rf}} = \alpha + \sum_k \beta_k f_{k,t} + \epsilon_t$$

- Use multi-factor set (Fama–French 5 + momentum + quality + low-vol).
- $\hat{\alpha}$ (annualised) and Newey–West t-stat are primary factor-clean performance measures.

2. **Information Ratio (IR).**

$$\text{IR} = \frac{\overline{a}}{\sqrt{\widehat{\text{Var}}(a)}}$$

- $\overline{a}$: mean active return.
- Variance estimated with appropriate scaling to annual IR (e.g., $\sqrt{12}$ for monthly).

3. **Tracking Error (TE).**

$$\text{TE} = \sqrt{\widehat{\text{Var}}(a)} \quad \text{(annualised)}.$$

Figure E26: Extracted figure from original PDF (page 35).

4. **Maximum Drawdown (MaxDD).**

   - From cumulative net value series V_t:

$$\mathrm{MaxDD} = \max_t \left(\frac{\max_{u \le t} V_u - V_t}{\max_{u \le t} V_u} \right).$$

5. **Hit-ratio.**

   - Overall: $\frac{1}{T}\sum_t \mathbf{1}\{a_t > 0\}$.
   - Conditional: hit-ratio in up-market months (benchmark > 0) and down-market months (benchmark < 0).

6. **Signal decay.**

   - Construct decile portfolios sorted on $\widetilde{\Delta\mu}_t$.
   - Track performance of each decile over 1/3/6/12-month holding periods (non-overlapping or overlapping).
   - Estimate how alpha decays with horizon.

7. **Capacity.**

   - Scale trades and positions by AUM multiples and re-run cost model.
   - Compute IR and MaxDD at each AUM level; capacity is the highest AUM with acceptable IR degradation (e.g., ≤ 20–30% drop in IR vs base AUM).

Figure E27: Extracted figure from original PDF (page 36).