# AI Process Optimisation in Chemicals: Financing Savings with Measurement Discipline

*A Stage-Gated Framework for Baselines, Controlled Trials, Verified Cash Benefits and Scalable Industrial Deployment*

**Authored by Chennakeshav Adya, Independent Researcher**

September 2026


## Abstract

*Artificial intelligence can improve chemical-process forecasting, control, maintenance, energy use, yield and quality. The financial case remains fragile when a programme starts with a model accuracy claim and works backwards to savings. Chemical plants operate inside safety, quality, environmental and customer constraints. Production conditions change with feedstock, campaign sequence, grade mix, ambient conditions, equipment health, maintenance, utilities and operator intervention. A simple before-and-after comparison can therefore attribute normal operating variation to the model.*

*This paper develops a stage-gated system for turning a proposed AI process-optimisation use case into an investible and auditable programme. It begins with the physical and economic perimeter, reconstructs the data lineage, defines the baseline and counterfactual, designs controlled trials, and separates prediction from automatic control. It then measures yield, energy, material, reliability and quality effects before translating them into cash. The framework connects technical evidence to accounting treatment, financing structure, vendor contracting, governance and scale-up across sites.*

*Five original figures and five implementation tables show how evidence should move from sensors and operating records to a verified benefits register and financing decision. The International Energy Agency identifies meaningful potential for AI-enabled industrial efficiency while also recording barriers around data, skills, cyber security and implementation. ISO 50015 and the International Performance Measurement and Verification Protocol provide principles for measurement boundaries, baselines and adjustments. NIST guidance informs AI and operational-technology risk controls. Process-safety, environmental, accounting and sustainable-finance sources provide further decision gates.*

*These sources do not establish savings for a particular plant, process, product or model. Every amount, percentage, threshold, timing assumption, score and financing outcome in this paper is a hypothetical analytical assumption used to demonstrate the framework. It is not a company forecast, engineering certification, process-safety assessment, environmental assessment, valuation opinion, accounting conclusion, legal conclusion, tax conclusion or financing commitment. An actual programme requires verified plant data, approved operating procedures, qualified engineering and safety review, cyber-security review, accounting advice and lender-specific diligence.*


JEL Classification: G32, L65, M15, O33, Q40

Keywords: chemicals, artificial intelligence, process optimisation, measurement and verification, industrial energy, yield, operational technology, financing

## 1. Start with an investible savings claim

An AI process-optimisation proposal becomes investible when the sponsor can state what physical result will change, how that result will be measured, what would probably have happened without the intervention, and how the verified change reaches cash. A predicted reduction in energy intensity has little financing value until the team defines the production unit, measurement boundary, baseline period, relevant variables, tariff, quality conditions and evidence owner. The same discipline applies to yield, throughput, off-spec production, maintenance and emissions.

The first investment question should therefore be a claim architecture rather than a model architecture. The claim identifies the process, operating envelope, intervention, affected economic line, measurement frequency and settlement period. It also identifies exclusions. A model that improves reactor conversion while increasing downstream purification cost has not necessarily created value. A model that reduces steam use during a favourable product campaign may be receiving credit for mix. A model that recommends a condition outside an approved safety or quality limit cannot be underwritten as an operating improvement.

The claim should distinguish technical potential, demonstrated operating effect, recurring economic benefit and cash realisation. Technical potential is the result expected from analysis or simulation. Demonstrated effect is the observed change under a controlled protocol. Recurring benefit adjusts for degradation, adoption and operating variability. Cash realisation appears in metered utility cost, purchased material, saleable output, avoided maintenance, working capital or another ledger-supported line. Each level carries different uncertainty and can support a different funding instrument.

Investment approval should state the evidence required at each stage, the maximum capital at risk, the condition for advancing and the condition for stopping. The programme remains a portfolio of options until repeatable evidence exists. This approach protects the plant from a large irreversible commitment and gives vendors, lenders and owners a common language for risk allocation.

**Figure 1. Evidence chain from process signal to financed cash benefit**

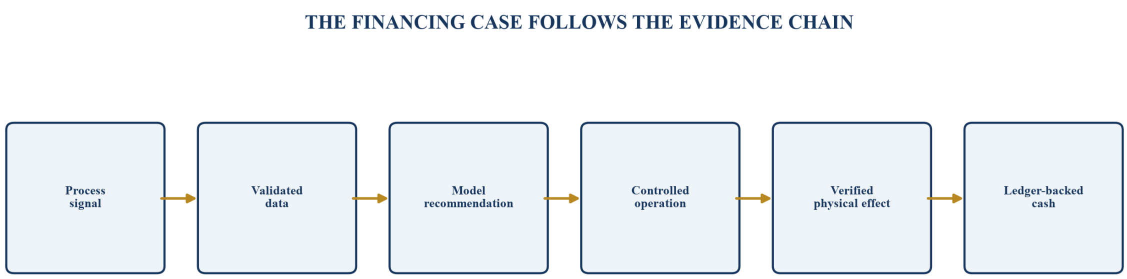


*Each transition requires a named evidence owner and an approved verification test.*

## 2. Fix the physical, economic and decision perimeter

The physical perimeter should identify the equipment, instruments, control loops, utilities, upstream inputs and downstream consequences affected by the use case. A furnace optimiser may change fuel flow, combustion air, excess oxygen, product temperature, emissions, refractory exposure and downstream duty. A reactor-yield model may affect feed ratio, residence time, catalyst use, impurity profile, separation load and inventory. The perimeter should be wide enough to capture material interactions and narrow enough to support credible measurement.

The economic perimeter converts physical effects into value lines. It should include saleable output, raw materials, utilities, consumables, waste treatment, quality losses, rework, maintenance, downtime, working capital, compliance expenditure, software, integration, sensors, cyber controls, training and ongoing model operations. The programme should avoid adding benefits that share the same physical driver. Yield improvement and raw-material reduction may be two descriptions of one material-balance effect.

The decision perimeter defines authority. It identifies who may recommend a set-point change, who may approve it, which control system may execute it, and what happens when data quality, model confidence, safety status or network availability falls below the approved threshold. This perimeter should align with the site's management of change, process-safety, quality and cyber-security procedures. A model can remain valuable in advisory mode when automatic control would create an unacceptable assurance burden.

The three perimeters should be approved together. A narrow technical scope paired with a broad financial claim creates over-attribution. A broad technical scope without decision rights creates an implementation gap. A broad automated-control scope without qualified safety and cyber review creates operational exposure. The signed perimeter document becomes the basis for data access, trial design, M&V, contracting and capital release.

**Table 1. Evidence map for an AI process-optimisation programme**

| Evidence family | Minimum control | Decision supported | Primary owner |
|---|---|---|---|
| process and laboratory data | tag lineage, calibration, timestamp and quality flag | technical baseline | operations and laboratory |
| production context | grade, campaign, feedstock, rate, shutdown and operator action | counterfactual adjustment | production planning |
| safety and environmental status | alarm, trip, bypass, permit and incident record | operating envelope | process safety and EHS |
| economic data | invoice, tariff, standard cost, sales and ledger reconciliation | cash conversion | finance |
| model records | version, feature set, training window, test result and override | model governance | digital and engineering |

*Evidence should be traceable to a system, owner, time period and approved use.*

## 3. Build a data lineage that survives diligence

Chemical-process data often appear abundant while remaining unsuitable for an investment claim. A historian may contain thousands of tags, but tag names, units, scaling, calibration, sampling frequency, compression and bad-quality handling can change over time. Laboratory results may be timestamped at collection, receipt or release. Production quantities may come from tank gauges, flow meters, batch tickets or reconciled accounting records. The programme should document how each input was created and transformed.

Lineage starts at the instrument or source record and ends at the feature used by the model or the value used by the benefits calculation. For each material tag, the team should record sensor type, physical location, engineering unit, calibration history, historian path, aggregation, missing-value treatment and known breaks. Derived variables should carry a reproducible formula and version. Manual entries should identify the person, authority and reason for the change.

Time alignment requires special attention. Residence time, laboratory delay, batch boundaries, recycles and inventory buffers can separate an operating action from its measured output. A naive same-hour comparison can associate an input change with the wrong product result. The engineering team should define the process lag and test it against material flow, batch genealogy and controlled events. Clock changes, daylight-saving treatment and server synchronization should also be checked in multi-system data.

The diligence pack should include a tag dictionary, lineage diagram, data-quality report, version-controlled extraction logic and reconciliation to selected source records. This pack allows an independent reviewer to reproduce the dataset and identify where judgment entered. Reproducibility matters because the model, the M&V calculation and the financing claim may otherwise rely on three different versions of the same operating history.

## 4. Establish the baseline before the model changes behaviour

A baseline should represent the relationship between production conditions and the outcome of interest before the intervention. It is not simply the average of a convenient period. The chosen window should capture normal operating modes, relevant seasonal or campaign variation, sufficient stable production and known disruptions. The team should document why excluded periods do not represent the intended operating case and retain an audit trail of every exclusion.

Baseline variables may include product grade, production rate, feedstock composition, catalyst age, ambient temperature, utility conditions, equipment availability, turnaround status and operator actions. The choice should follow physical causality and data quality. Adding many variables can improve statistical fit while weakening interpretability and increasing the risk of overfitting. Omitting a material driver can credit the programme for a favourable change unrelated to the model.

ISO 50015 provides general principles for measuring and verifying organisational energy performance. The International Performance Measurement and Verification Protocol emphasizes consistent comparison, measurement boundaries and adjustments for changed conditions. Their central logic applies beyond energy: the intervention effect is observed relative to an agreed baseline adjusted for relevant independent variables and non-routine events.

The baseline should be signed by engineering, operations, finance and the independent verifier before the trial begins. The sign-off should cover the period, boundary, variables, data treatment, model form, uncertainty approach and change-control process. A baseline rewritten after results are known creates selection bias and weakens any performance-linked payment or financing claim.

**Figure 2. Baseline and counterfactual architecture**

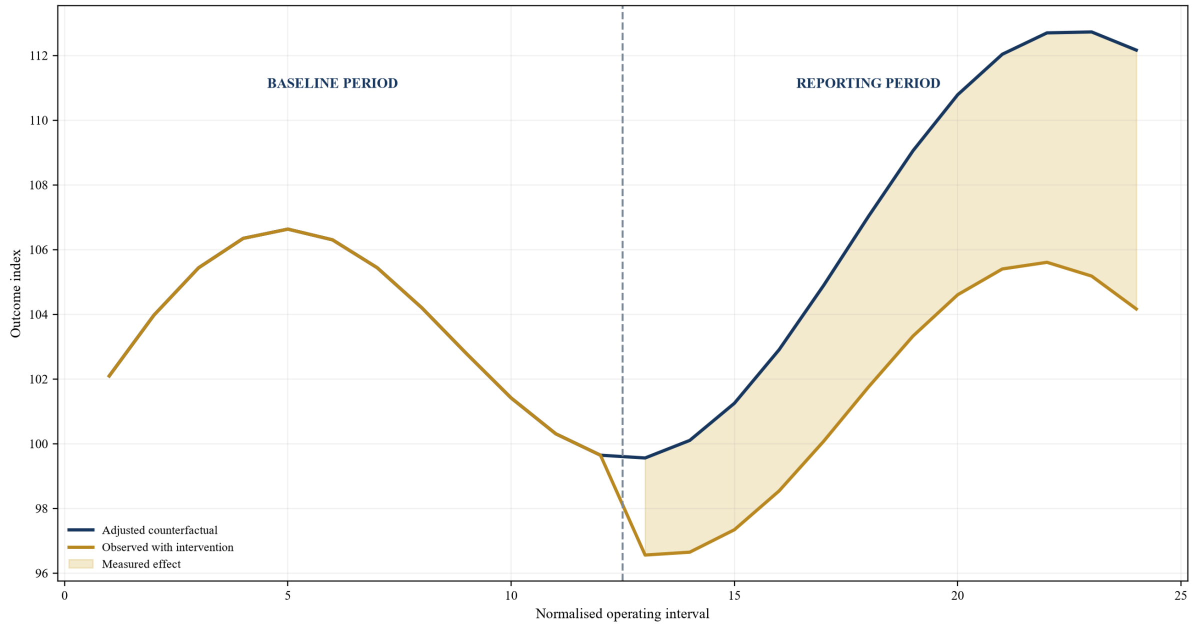


*The counterfactual estimates the outcome under observed conditions without the intervention; it is not a guaranteed forecast.*

## 5. Define the counterfactual and adjustment rules

Savings and improvements are counterfactual claims. The plant can measure what occurred; it cannot directly measure the same period without the intervention. The M&V plan should therefore specify how the expected outcome will be calculated under actual reporting-period conditions. This may use a regression, first-principles model, matched operating periods, a control line, a stepped-wedge rollout or another method suited to the process and decision.

Routine adjustments account for variables expected to change, such as output, grade, feedstock, weather or operating hours. Non-routine adjustments address events such as a major equipment replacement, changed product specification, new utility system, sustained market curtailment or altered operating policy. The plan should define who identifies such events, the evidence required, the calculation method and the dispute process. Uncontrolled retrospective adjustment can absorb poor performance or manufacture savings.

Model uncertainty and measurement uncertainty should be reported separately where practical. A highly accurate meter does not resolve an uncertain counterfactual. A sophisticated model does not repair biased or poorly calibrated input data. The benefits report should disclose the applicable interval, confidence or sensitivity analysis, material assumptions and the effect of alternative specifications. A single precise number can misrepresent a wide decision range.

The programme should also define persistence. An initial improvement may decline when feedstock changes, equipment fouls, operators revert, the model drifts or the production schedule moves outside the training envelope. Verification should continue across agreed operating modes and time periods. Financing should rely on the persistent case, with a reserve or haircut for adoption, downtime, degradation and uncertainty.

**Table 2. Measurement-and-verification design choices**

| Situation | Possible design | Principal strength | Principal limitation |
|---|---|---|---|
| isolated equipment with reliable meters | retrofit isolation | direct physical boundary | interaction outside boundary may be missed |
| whole line with multiple interactions | adjusted facility or line model | captures combined effect | counterfactual can be model-sensitive |
| repeated batches or parallel trains | matched or randomized trial | strong comparison | production constraints may limit randomisation |
| staged site rollout | stepped-wedge design | operationally practical comparison | time effects require control |
| new process without stable history | calibrated engineering simulation | supports early decision | depends heavily on assumptions and calibration |

*The method should match the process, available evidence and intended financial use.*

## 6. Prioritise use cases by evidence-adjusted value

Use-case selection should compare expected value, evidence readiness, operational controllability, implementation cost, safety criticality and time to verification. A high theoretical saving with weak instrumentation and a narrow operating window may rank below a moderate opportunity with reliable data, frequent repetition and clear cash conversion. The portfolio should include quick learning cases and strategically important cases without treating every model as a production deployment.

Yield optimisation can create value where material balance, product quality and saleable output are well measured. Energy optimisation can be attractive where tariffs, meters and process-normalisation variables are available. Predictive maintenance can be valuable when failure modes, work orders, downtime costs and intervention decisions are consistently recorded. Quality prediction can reduce laboratory delay or off-spec production, but it needs verified batch genealogy and a clear authority boundary.

Each use case should have a falsifiable hypothesis. For example: within a defined product grade and production-rate range, advisory recommendations will reduce normalized steam per tonne while maintaining specified product quality, emissions and equipment constraints. The hypothesis states the outcome, boundary, mode and non-negotiable conditions. It allows the team to design a trial that can fail cleanly rather than redefining success after deployment.

A portfolio score should be a decision aid, not a substitute for judgment. Scores depend on assumptions and can create false comparability. The approval paper should show the underlying evidence, major dependencies, downside and reason for selection. Safety-critical or permit-sensitive use cases require qualified review regardless of a high economic score.

**Table 3. Evidence-adjusted use-case screen**

| Dimension | Question | Strong evidence | Escalation trigger |
|---|---|---|---|
| economic materiality | which cash line changes and by how much | reconciled volume, price and cost driver | value depends on unverified allocation |
| data readiness | can the dataset be reproduced | calibrated tags and controlled lineage | missing context or unstable definitions |
| controllability | can an approved action change the outcome | clear lever and authority | recommendation has no executable owner |
| verification | can the counterfactual be defended | signed baseline and trial design | retrospective baseline selection |
| operating risk | can the case remain within safe and compliant limits | approved envelope and fallback | safety or permit implication unresolved |

*Scores are illustrative; the decision should retain the underlying evidence and constraints.*

## 7. Put process safety ahead of optimisation

An optimiser should operate only inside a qualified process-safety and operating envelope. The envelope should reflect design limits, safe operating limits, alarm and trip settings, material compatibility, relief capacity, environmental constraints, product specifications and equipment integrity. Model recommendations should never be treated as an alternative to a safety instrumented function, protective layer or approved operating procedure.

IEC 61511 addresses the lifecycle for safety instrumented systems in the process sector. An AI model outside the safety system can still affect the demand placed on protective layers by moving the process closer to a constraint or increasing the frequency of excursions. Process-safety specialists should assess the proposed use, failure modes, independence, common-cause exposure, operator response and management-of-change requirement.

The control philosophy should define advisory, supervisory and closed-loop modes. Advisory mode presents a recommendation for operator review. Supervisory mode may adjust a target through an approved control layer. Closed-loop mode executes within defined constraints. Progression should depend on evidence, hazard review, human factors, cyber controls, fallback performance and operator competence. A technically capable model may remain in advisory mode when the assurance case for automation is incomplete.

Fallback must be deterministic and tested. Loss of model service, stale data, abnormal operation, low confidence or a cyber event should return the process to an approved control strategy without unsafe transition. The programme should record every recommendation, acceptance, rejection, override and outcome. These records support learning, incident review and the evidence needed for any later change in authority.

## 8. Design a controlled trial that operations can run

A controlled trial should balance causal evidence with plant practicality. The team should define eligible operating periods, intervention and comparison conditions, washout or lag, minimum duration, required observations, stop conditions and quality review. Trial design should be approved before outcome data are examined. Production planning, customer commitments and maintenance windows should be integrated so that the comparison does not depend on unrecorded operational concessions.

Randomised crossover trials can be strong where the process permits repeated switching between advisory-on and advisory-off conditions without material disruption. Parallel trains can support contemporaneous comparison if their equipment, feed and operating history are sufficiently similar. Batch processes may support matched campaigns. A stepped-wedge rollout can introduce the model across lines or shifts over time, allowing later groups to act as temporary controls.

The trial should retain operator agency and document behavioural effects. Operators may pay more attention during a pilot, maintenance may respond faster, or production scheduling may favour the trial. These changes can create real value while weakening attribution to the model. The benefits register should separate model recommendation, operating action and supporting management response. This separation helps determine what must be funded and sustained during scale-up.

Trial governance should include daily safety and data checks, a weekly technical review and a pre-defined decision board. Stop conditions may include safety or quality excursion, unstable control, unreliable data, cyber concern, inadequate adoption or insufficient evidence. A stopped trial can be a successful risk decision. Capital should advance only when the evidence supports the next operating mode and financial claim.

**Figure 3. Controlled trial inside the approved operating envelope**

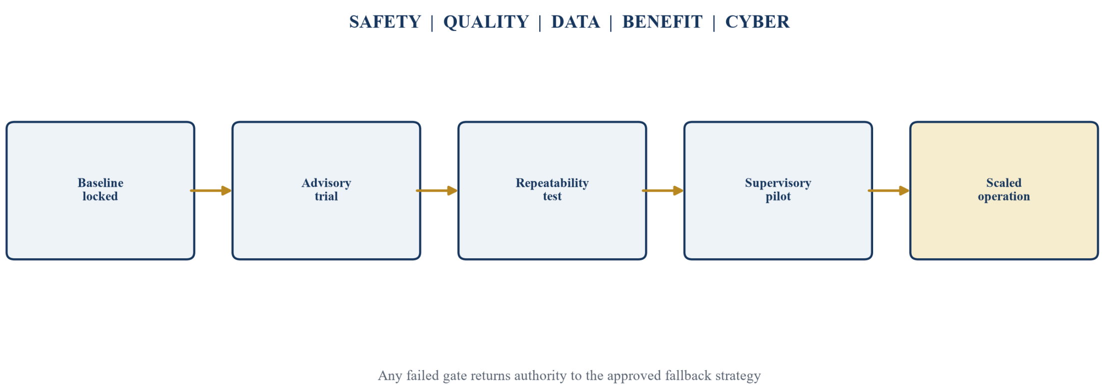


*The trial advances only when safety, quality, data and benefit gates remain satisfied.*

## 9. Separate prediction, recommendation and control

A prediction estimates a future state, a recommendation proposes an action, and control changes the process. These functions have different evidence and risk requirements. A model may accurately predict a quality result while recommending an uneconomic or unsafe response. A recommendation may be economically sound while remaining unsuitable for automatic execution because the control layer, fallback or operator interface is incomplete.

The model specification should state the decision it supports, prediction horizon, input availability, target definition, loss function and operating envelope. Offline accuracy should be assessed by relevant mode and cost, not only an aggregate score. Rare but material conditions deserve separate analysis. A low average error can coexist with poor performance during startup, grade transition, catalyst ageing or high-rate operation.

Recommendation logic should incorporate constraints and the economics of action. The value of a predicted deviation depends on whether the plant can intervene in time, the cost of intervention, the risk of a false alarm and the effect on other objectives. The programme should record both accepted and rejected recommendations so that adoption and operator judgment are visible. Silent filtering of difficult cases can exaggerate performance.

Control authority should advance through explicit management of change. The team should verify interfaces, permissions, rate limits, constraint handling, failover, logging and operator override. A model update should not change plant behaviour without version approval and regression testing. The financial case should reflect the operating mode actually approved rather than the automation level imagined in the business case.

## 10. Secure operational technology and model interfaces

AI deployment can create new data paths between historians, laboratories, maintenance systems, cloud platforms, vendor services and operational technology. NIST SP 800-82 Rev. 3 addresses security for operational technology while recognizing performance, reliability and safety requirements. The programme should begin with an architecture and data-flow map that identifies trust boundaries, protocols, identities, privileges, remote access and the consequence of compromise.

Read-only extraction can still create risk if credentials, jump hosts, shared services or vendor connections are poorly controlled. Write access to a control layer requires stronger assurance. The design should apply least privilege, network segmentation, monitored access, controlled software deployment, secure configuration, backup, recovery and incident response suited to the site. Security controls should be tested with operations because an IT control that interrupts availability can create process risk.

Model supply-chain risk includes code, libraries, containers, pretrained components, data services and vendor updates. The owner should maintain an inventory, approved versions, vulnerability response, rollback capability and contractual access to material records. Sensitive process data, product information and operating know-how require clear rights, retention, location and confidentiality controls.

The business case should fund cyber and resilience work as part of the production asset. Treating these costs as optional overhead overstates returns and delays deployment. The financing model should also include the cost and downtime of patching, testing, backups, disaster recovery and vendor exit. A resilient fallback can preserve production value even when the AI service is unavailable.

## 11. Govern model validity, drift and change

The NIST AI Risk Management Framework organizes AI risk activity around governance, mapping, measurement and management. For a chemical process, model governance should connect those functions to the plant's existing engineering, quality, process-safety, cyber and capital-control systems. The AI model is one component of an operating decision and should inherit the controls required by the consequence of that decision.

The model record should include purpose, owner, approved users, version, training period, features, target, limitations, validation results, operating envelope and fallback. Performance should be monitored by relevant process mode. Drift can arise from sensor recalibration, equipment work, catalyst changes, feedstock, recipe, product specification or operating policy. A stable statistical distribution does not prove that the physical relationship remains valid.

Change control should distinguish data refresh, parameter tuning, feature change, model replacement, interface change and authority change. Each category needs a defined test and approver. A vendor's routine model update should not bypass the site's management-of-change process. The production system should retain the ability to reproduce the prior version and compare recommendations on a controlled dataset.

Retirement is part of governance. A model may become uneconomic, unsupported, duplicative or unsafe within the revised process. The owner should define triggers, archive evidence, remove privileges and return to the approved fallback. The benefits register should stop recognising recurring benefit when adoption or performance no longer meets the verification rule.

## 12. Measure yield without losing the material balance

Yield improvement should be defined in saleable product terms and reconciled through the material balance. Gross production can rise while saleable yield falls because impurities, off-spec material, rework, purge, moisture or inventory measurement changes. The programme should define feed basis, product basis, normalisation, quality acceptance, recycle treatment and the timing boundary between input and output.

The baseline should account for product grade, feedstock composition, catalyst condition, rate and campaign sequence. Laboratory and production records should be linked through batch or lot genealogy. Tank movements and work-in-progress can shift apparent yield between periods. Finance should reconcile the physical effect to inventory and cost accounting before value is recognised.

The cash bridge can distinguish increased saleable volume, avoided raw material, avoided disposal and incremental variable conversion cost. Market value should reflect the actual product and sales constraint. Extra output has limited near-term value when the plant is demand constrained, storage constrained or required to discount marginal volume. Avoided input may produce faster cash if purchases and inventory policy respond.

The trial should monitor quality, emissions, energy, run length and downstream load alongside yield. A local improvement that increases purification duty, fouling or customer complaints may destroy value outside the model boundary. Persistent yield value requires stable results across representative modes and a documented operating practice that can be repeated without exceptional attention.

## 13. Measure energy as a normalised process outcome

Energy savings should be measured against useful production and relevant operating conditions. Steam, electricity, fuel, compressed air, refrigeration and cooling-water effects may cross several meters and cost centres. The measurement boundary should follow the affected utilities and interactive effects. A furnace improvement can reduce fuel while changing electricity, oxygen, emissions or downstream separation duty. Whole-line measurement may therefore be more credible than a single-equipment claim.

The baseline should use energy-performance indicators suited to the process. Energy per tonne may be adequate within a stable grade and rate range. A multivariable baseline may be needed when grade, feedstock, ambient conditions, throughput or utility properties materially affect consumption. The team should retain raw meter data, tariff structure, correction factors, calibration and evidence of non-routine events.

Cash value depends on tariff and purchasing mechanics. Avoided kilowatt-hours may have a different value by time, demand band, contract, self-generation dispatch or take-or-pay commitment. Fuel savings may change emissions cost, tax or certificate exposure. The finance bridge should calculate the avoided cash under the applicable contract and identify any saving that remains an accounting allocation rather than a cash reduction.

The IEA identifies potential for AI-enabled industrial efficiency alongside barriers around data, digital infrastructure, skills and security. Site underwriting should remain based on measured plant evidence. A global potential estimate does not establish the result for a particular process. The financing case should include the energy used by data collection, computing and model operation where material.

## 14. Quantify raw-material, consumable and waste effects

Raw-material optimisation can create value through recipe control, feed-ratio stability, impurity management, catalyst use and reduced giveaway. The programme should define the material basis and distinguish purchase quantity, process consumption, inventory movement and standard-cost allocation. A reduction recorded in the process model may not reach cash until purchasing, minimum-order quantities and safety stock change.

Consumables such as catalysts, solvents, filters, treatment chemicals and packaging require life-cycle treatment. A model may extend catalyst life while changing activity, selectivity or regeneration cost. It may reduce solvent use while increasing recovery energy. The benefits calculation should cover the full consequence within the approved boundary and avoid annualising a short trial beyond the demonstrated operating cycle.

Waste value includes avoided material loss, handling, storage, treatment, transport, disposal and regulatory cost. It may also include recovered by-product value. Quantity and classification should be supported by manifests, meters, laboratory data and production reconciliation.

Environmental and legal treatment depends on the actual material and jurisdiction, and qualified advisers should confirm applicable requirements.

The cash bridge should state the response mechanism. Reduced theoretical dosage has little near-term cash value if purchase orders remain unchanged or obsolete stock increases. Procurement and inventory owners should commit to an action linked to the verified operating effect. The benefits register should record both physical saving and realised purchase or disposal cash so that working-capital timing is visible.

## 15. Measure reliability and maintenance value carefully

Predictive-maintenance models often report avoided failure value, which is inherently counterfactual and vulnerable to overstatement. A warning does not prove that a failure would have occurred, and a maintenance action may itself create cost or downtime. The programme should define the failure mode, detectable degradation, intervention threshold, maintenance response and evidence that the action changed the outcome.

Validation can use historical back-testing, seeded or simulated faults where safe, prospective shadow mode, and controlled operational deployment. Precision and recall should be translated into operational consequences. False positives consume inspection and maintenance capacity. False negatives leave exposure. Lead time matters only when the plant can plan and execute an effective response before the failure window closes.

Value may include avoided repair, reduced collateral damage, shorter outage, improved maintenance scheduling, lower spare-part urgency and protected production. Each component should use approved cost and production assumptions without overlap. Avoided downtime should reflect the realistic constraint, inventory buffer, market demand and recovery profile. The model should not value every predicted failure at peak lost-margin assumptions.

Persistence requires feedback from work orders, inspection findings, parts condition and final failure classification. Maintenance data are often inconsistent, so the programme may need a disciplined coding and closure process. The recurring case should include model operations, sensor maintenance, analyst review and the cost of investigating alerts. Financing should use a conservative demonstrated case until sufficient event evidence accumulates.

## 16. Protect product quality and customer commitments

Quality prediction can reduce laboratory delay, transition loss, rework and off-spec production. The programme should distinguish a soft sensor used for early warning from an approved quality-release method. A model should not replace a contractual or regulatory test without the required validation, quality approval and customer acceptance. The operating procedure should state which decision the model may support.

Ground-truth quality data require controlled sampling, laboratory method, chain of custody, result timing and linkage to the correct production lot. Measurement error and sample representativeness can place a ceiling on model performance. The dataset should preserve retests, invalid results and specification changes. Removing difficult observations can create an attractive model that fails under real production conditions.

The economic effect should separate reduced transition material, lower giveaway, fewer customer claims, shorter release time, inventory reduction and increased throughput. These benefits can share a driver and should be reconciled. Faster release may reduce working capital even if production cost is unchanged. Reduced giveaway can increase saleable output while requiring no additional market volume when product is sold to specification rather than fixed formulation.

Customer risk should remain visible. A small increase in variability can outweigh an average yield improvement if it threatens a critical specification or qualification. Trial approval should identify protected quality attributes, escalation, quarantine and communication rules. The persistent value case should include customer complaints, returns and claim experience, not only in-process model metrics.

## 17. Translate physical improvement into verified cash

The savings bridge should begin with the verified physical effect and move through price, timing and operational response to cash. It should show gross physical value, interactive costs, implementation run cost, adoption, persistence, market or purchasing constraint, tax or accounting timing where relevant, and the resulting cash benefit. Every line should identify the source system and accountable owner.

Yield value may enter through additional sales, avoided input or reduced waste. Energy value should use the actual tariff and contract. Maintenance value should use completed work and avoided-cost logic approved by finance. Working-capital value should distinguish a release of cash from an earnings improvement. The bridge should avoid treating cost allocation changes as external cash savings unless spending changes.

Benefits should be reported in a waterfall from technical potential to verified cash. The technical case may remain useful for portfolio prioritisation. The financed case should rely on demonstrated, repeatable and conservatively valued results. Haircuts can reflect measurement uncertainty, adoption, downtime, model degradation, limited operating coverage and execution delay. These deductions should be transparent rather than hidden in a high discount rate.

Finance should reconcile the benefits register to utility invoices, purchase orders, production and sales records, inventory, maintenance work orders and the general ledger. Timing differences and accounting treatment should be explained. Independent verification can strengthen performance-linked contracting or financing, provided that the verifier's scope, data access, methodology and conflict controls are defined.

**Figure 4. From technical potential to verified cash**

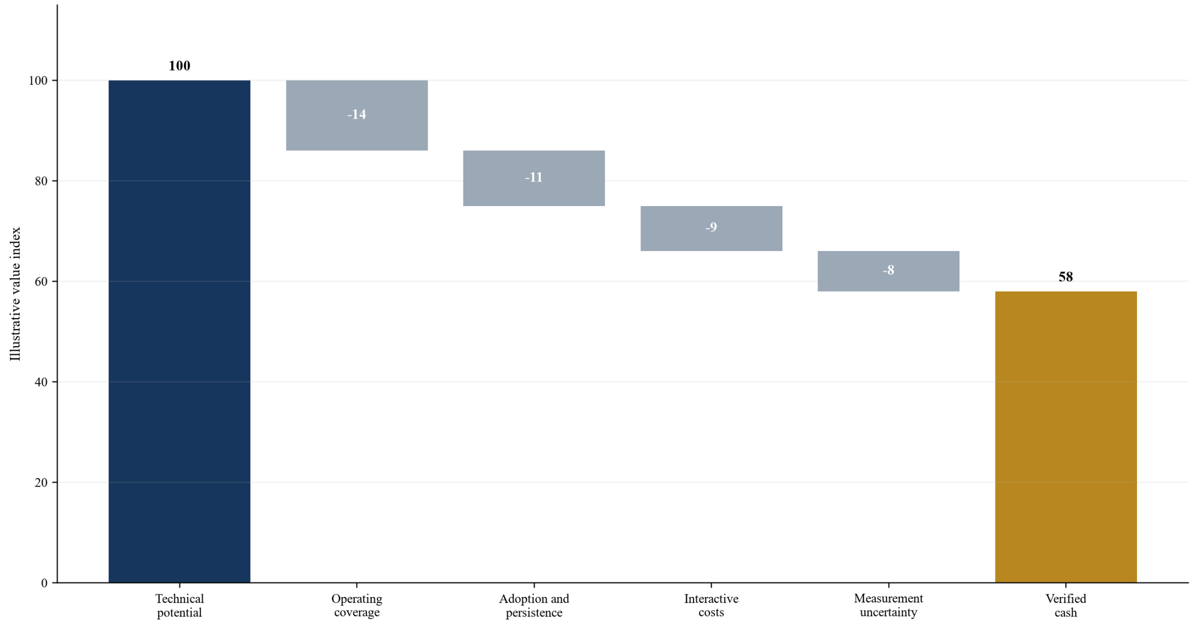


*Values are illustrative indices and do not represent a plant forecast.*

## 18. Decide accounting treatment before promising returns

AI programme expenditure can include research, data preparation, software, sensors, integration, control-system changes, equipment modifications, testing, training and ongoing operation. Accounting treatment depends on the facts, applicable standards and the nature of each component. IAS 38 distinguishes research from development and sets recognition criteria for intangible assets. IAS 16 addresses property, plant and equipment. Cloud and service arrangements can require separate analysis.

The investment case should not assume capitalisation merely to improve reported earnings. Finance should map each cost to the applicable policy and obtain qualified accounting advice. Internally generated data, models and configuration may lack separability or demonstrable future economic benefits at early stages. Hardware and directly attributable installation may follow a different path. Maintenance, retraining and support may remain period expense.

Impairment and useful-life considerations should enter the economics. Model performance can degrade, a vendor can discontinue support, the process can change, or a superior method can make the asset obsolete. The owner should identify the unit of account, dependencies, expected life, residual value, testing indicators and decommissioning cost. A long accounting life should not substitute for evidence of operating persistence.

The benefits report should distinguish EBITDA, operating cash, capital expenditure, working capital and balance-sheet treatment. This prevents the programme from mixing a cost reclassification with an economic saving. Lenders and boards can then assess debt service and return on invested capital using reconciled definitions.

## 19. Match financing structure to evidence maturity

Early discovery and baseline work carry technology and data risk. They are commonly suited to corporate operating budget, innovation budget, vendor-funded discovery or a small capped pilot. Milestone funding can release capital after data-readiness, trial, repeatability and production-assurance gates. This limits loss when the use case lacks evidence and preserves the option to scale a successful case.

Performance-linked vendor economics can align part of payment with verified results. The contract needs an agreed baseline, measurement boundary, adjustment rules, data access, verification, exclusions, cap, floor, payment period and dispute process. The vendor should not control both the model and the final savings calculation without independent governance. Payment should also reflect the operating mode and the plant's responsibility for adoption.

Asset finance or project-style debt may become relevant when the programme includes identifiable sensors, control upgrades, energy equipment or a portfolio of repeatable projects with stable cash evidence. General corporate debt can fund scale-up when the borrower retains performance risk. Sustainability-linked debt can reference material corporate KPIs, while use-of-proceeds instruments depend on eligible expenditure and applicable principles. Each structure requires instrument-specific review.

The funding model should include contingency, cyber, integration, training, model operations, verification and decommissioning. Debt service should rely on a downside cash case rather than technical potential. The lender should see the evidence chain, contractual risk allocation, residual operating risk and fallback. Financing does not convert an uncertain use case into a certain saving.

**Table 4. Financing structure by evidence maturity**

| Evidence stage | Possible funding route | Risk retained by owner | Key financing gate |
|---|---|---|---|
| discovery and baseline | operating or innovation budget | data and use-case risk | reproducible data and approved hypothesis |
| controlled pilot | milestone capital or vendor co-funding | trial and adoption risk | signed M&V plan and stop conditions |
| repeatable site deployment | corporate capex, lease or performance contract | persistence and integration risk | independently verified operating effect |
| multi-site scale | corporate debt or portfolio facility | portfolio and execution risk | standard design, controls and cash evidence |
| material sustainability programme | eligible green or sustainability-linked instrument | KPI, disclosure and verification risk | instrument-specific alignment and external review |

*Instrument suitability depends on borrower, jurisdiction, documentation and verified project facts.*

## 20. Design performance-linked financing with credible KPIs

A financing KPI should be material to the borrower, measurable, comparable over time and resistant to operational gaming. A narrow pilot metric may be useful for vendor payment and unsuitable for a corporate sustainability-linked instrument. The KPI definition should state perimeter, unit, baseline, target, observation period, adjustment, assurance, reporting and consequence.

The Sustainability-Linked Loan Principles and associated guidance emphasize core components including KPI selection, performance targets, loan characteristics, reporting and verification. The Sustainability-Linked Bond Principles similarly address material quantitative KPIs, targets, reporting and external verification. Transaction counsel and arrangers should confirm the current requirements and market expectations for the chosen instrument.

An AI programme can contribute to an energy- or emissions-intensity KPI, but the financing framework should avoid attributing the full corporate result to the model. Production mix, acquisitions, closures, renewable procurement and other initiatives may also affect the KPI. The borrower should disclose the strategic levers, their expected contribution, governance and limitations at the level required by the instrument.

The verification protocol should identify data systems, assurance scope, fallback when data are unavailable, restatement and the treatment of acquisitions or structural change. The financing model should test the economic consequence of missing a target and the cost of ongoing reporting and assurance. KPI integrity protects both funding credibility and the broader programme.

## 21. Contract vendors around evidence, access and exit

The vendor contract should define deliverables across discovery, data engineering, model development, trial, production deployment, support and exit. Acceptance should use evidence and operating gates rather than a broad statement that the model is accurate. The customer should know which artefacts, source materials, configurations, interfaces and documentation it receives at each milestone.

Data and intellectual-property terms should cover pre-existing IP, plant data, derived data, features, model artefacts, improvements, confidential know-how, permitted use, location, retention and deletion. The plant should retain access to records needed for safety, compliance, verification, dispute and transition. A vendor should not be able to remove the evidence supporting a financing claim when the service ends.

Service levels should reflect operating consequence. Availability, response, model latency, data freshness, incident handling, patching, backup, recovery and change notice may matter. The contract should state the approved interface and prohibit unapproved write access or model updates. Audit and security rights should match the architecture and risk.

Performance payment needs a complete M&V annex. It should define baseline, counterfactual, routine and non-routine adjustments, quality and safety constraints, owner actions, exclusions, verification, cap, floor, settlement and dispute resolution. Exit assistance should include data export, model and configuration transfer where licensed, credential removal, secure deletion and a tested return to the fallback operating strategy.

## 22. Release capital through stage gates

A stage-gated programme treats capital as a sequence of learning commitments. Gate zero confirms the economic claim and operating sponsor. Gate one confirms the physical perimeter, data lineage and baseline feasibility. Gate two approves the trial design, safety case, cyber architecture and M&V plan. Gate three confirms repeatable physical effect. Gate four confirms verified cash and production assurance. Gate five approves scale.

Each gate should have a decision paper, evidence pack, accountable approvers and a maximum next-stage exposure. Criteria should be set before results are known. A gate can approve, approve with conditions, repeat, redesign or stop. The programme should preserve negative evidence because failed use cases improve portfolio selection and vendor governance.

Capital release should include the complete cost to reach the next gate. Data remediation, sensor work, integration, process-safety review, cyber controls, operator time, verification and contingency should be visible. Funding only the model can leave the plant with a promising algorithm and no production asset.

The board should distinguish a technical milestone from an economic milestone. Offline validation supports a technical decision. Controlled operation supports an operating decision. Ledger-backed benefit supports a financing decision. Multi-site repeatability supports a scale decision. This hierarchy reduces the risk of financing a demonstration as if it were a recurring cash stream.

**Figure 5. Capital-release gates for industrial AI**

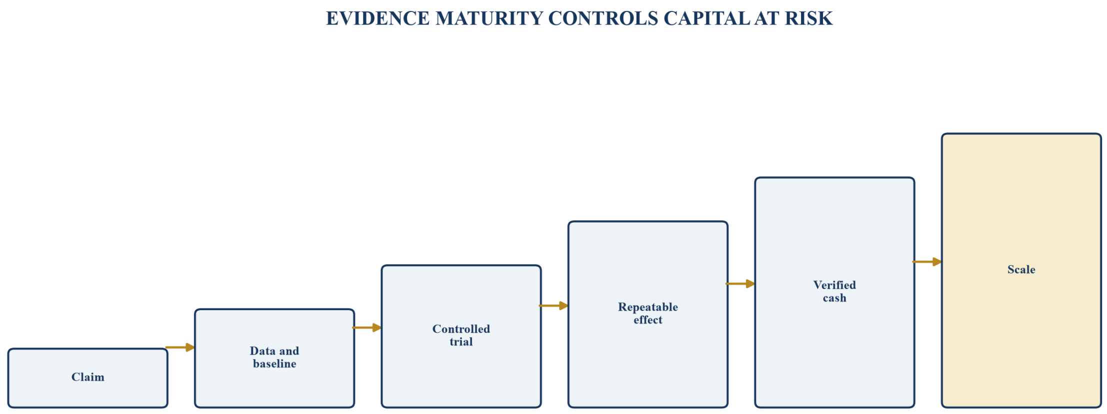


*Capital at risk increases only after evidence and operating assurance strengthen.*

## 23. Operate one benefits register across engineering and finance

The benefits register should be the controlled link between use-case evidence and management reporting. Each entry should identify the use case, physical metric, boundary, baseline, counterfactual, intervention, operating coverage, measurement result, uncertainty, economic conversion, owner, verification status, cash status and expiry or review date. Supporting records should be linked rather than copied into an uncontrolled spreadsheet.

Engineering owns the physical relationship and operating conditions. Operations owns adoption and execution. Finance owns valuation, ledger reconciliation and cash status. Digital teams own model and data records. Process safety, quality, EHS and cyber owners approve the relevant constraints. An independent verifier may review selected claims. Shared ownership should not obscure final accountability for each field.

The register should prevent double counting across use cases and periods. A yield optimiser and quality predictor may both claim reduced off-spec production. A maintenance model and throughput optimiser may both claim avoided downtime. The benefits office should allocate the physical effect once and explain contribution where several interventions interact.

Reporting should include gross potential, verified physical benefit, verified economic benefit and realised cash. It should also show cost, variance, persistence, adoption and unresolved evidence. Benefits should expire or return to review when the process, model, tariff, product mix or measurement basis changes materially.

**Table 5. Benefits-register control fields**

| Control field | Required evidence | Approval | Review trigger |
|---|---|---|---|
| baseline and counterfactual | signed M&V plan and reproducible calculation | engineering and verifier | material process or data change |
| operating effect | trial and production-period records | operations and quality | drift, low adoption or constraint breach |
| economic conversion | price, tariff, cost and ledger bridge | finance | contract, mix or accounting change |
| cash realisation | invoice, purchase, sales, inventory or payment record | finance owner | timing variance or reversal |
| model assurance | version, validation, monitoring and fallback | digital, cyber and engineering | update, incident or authority change |

*The register should retain supporting evidence and a complete change history.*

## 24. Scale across sites without copying assumptions

Scale should standardise the operating system while revalidating the physical and economic case at each site. Plants can differ in equipment, instrumentation, feedstock, product mix, control systems, maintenance maturity, utilities, tariffs, staffing, cyber architecture and regulatory obligations. A model or baseline should not be transferred solely because the process name is similar.

The scale package should separate reusable components from site-specific components. Reusable elements may include data contracts, model-development standards, trial protocol, governance, cyber pattern, M&V template, vendor clauses and benefits-register fields. Site-specific elements include tag mapping, calibration, process lags, safety envelope, baseline, tariff, customer constraints and approved control authority.

Portfolio funding can release capital by cohort. A lead site proves the end-to-end evidence chain. A second site tests transferability. A wider cohort uses standard deployment with local gates. The portfolio should track failure and adaptation cost, not only successful replication. This allows the board to estimate scale economics from observed deployment effort.

Central governance can maintain standards, shared infrastructure, vendor management and portfolio reporting. Local plant leadership retains accountability for safe operation and verified benefit. The scaling decision should show where central economies of scale outweigh local integration and assurance cost. A slower controlled rollout can preserve more value than a rapid deployment that creates unreliable claims or operational resistance.

## 25. Use a final investment-committee checklist

The investment committee should receive the signed claim, physical and economic perimeter, data-lineage report, baseline, counterfactual, trial design, process-safety and cyber approvals, model-validation record, M&V plan, benefits bridge, accounting analysis, financing proposal, vendor contract and scale plan. Open items should have a consequence, owner and decision date.

The committee should confirm that the programme can distinguish model effect from product mix, feedstock, weather, maintenance, operator behaviour and other changes. The physical effect should reconcile to quality, material balance, energy, reliability and compliance constraints. The cash bridge should reconcile to the applicable ledger or contract without double counting.

Financing should match evidence maturity and downside liquidity. Performance-linked terms require an agreed calculation and independent dispute path. Debt service should be supportable under a conservative persistent case. The owner should retain a tested fallback, vendor exit and decommissioning plan. Cyber, safety, integration, verification and operating costs should be funded.

Approval should state the maximum capital, next evidence gate, approved operating mode, prohibited actions, target review date and conditions for stopping. The final record should explain the residual technical, operating, cyber, measurement, adoption and financing risks that the owner accepts. A defensible programme finances verified process improvement rather than an accuracy score.

## References


[1] International Energy Agency. Key Questions on Energy and AI, 2026. https://www.iea.org/reports/key-questions-on-energy-and-ai

[2] International Energy Agency. Energy and AI, 2025. https://www.iea.org/reports/energy-and-ai

[3] International Energy Agency. AI for energy optimisation and innovation, 2025. https://www.iea.org/reports/energy-and-ai/ai-for-energy-optimisation-and-innovation

[4] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework 1.0. https://doi.org/10.6028/NIST.AI.100-1

[5] National Institute of Standards and Technology. AI Risk Management Framework resources. https://www.nist.gov/itl/ai-risk-management-framework

[6] National Institute of Standards and Technology. Guide to Operational Technology Security, SP 800-82 Rev. 3. https://doi.org/10.6028/NIST.SP.800-82r3

[7] National Institute of Standards and Technology. Cybersecurity Framework 2.0. https://www.nist.gov/cyberframework

[8] International Organization for Standardization. ISO 50015:2014, Measurement and verification of energy performance. https://www.iso.org/standard/60043.html

[9] International Organization for Standardization. ISO 50001:2018, Energy management systems. https://www.iso.org/standard/69426.html

[10] International Organization for Standardization. ISO 50006:2023, Energy performance indicators and energy baselines. https://www.iso.org/standard/79367.html

[11] Efficiency Valuation Organization. International Performance Measurement and Verification Protocol Core Concepts 2022. https://evo-world.org/en/products-services-mainmenu-en/protocols/ipmvp

[12] International Electrotechnical Commission. IEC 61511-1:2016 and Amendment 1:2017, functional safety for the process industry sector. https://webstore.iec.ch/en/publication/24241

[13] International Finance Corporation. Environmental, Health and Safety Guidelines: Chemicals. https://www.ifc.org/en/insights-reports/general-environmental-health-and-safety-guidelines/ehs-guidelines-chemicals

[14] International Finance Corporation. General Environmental, Health and Safety Guidelines. https://www.ifc.org/en/insights-reports/2000/general-environmental-health-and-safety-guidelines

[15] IFRS Foundation. IAS 38 Intangible Assets. https://www.ifrs.org/issued-standards/list-of-standards/ias-38-intangible-assets/

[16] IFRS Foundation. IAS 16 Property, Plant and Equipment. https://www.ifrs.org/issued-standards/list-of-standards/ias-16-property-plant-and-equipment/

[17] IFRS Foundation. IAS 36 Impairment of Assets. https://www.ifrs.org/issued-standards/list-of-standards/ias-36-impairment-of-assets/

[18] IFRS Foundation. Intangible Assets project, 2026. https://www.ifrs.org/projects/work-plan/intangible-assets/

[19] Loan Market Association, Asia Pacific Loan Market Association and Loan Syndications and Trading Association. Sustainability-Linked Loan Principles, 2025. https://www.lma.eu.com/sustainable-lending/resources

[20] Loan Market Association, Asia Pacific Loan Market Association and Loan Syndications and Trading Association. Guidance on Sustainability-Linked Loan Principles, 26 March 2025. https://www.lma.eu.com/application/files/3517/4298/0872/Guidance_on_Sustainability-Linked_Loan_Principles_-_26_March_2025.pdf

[21] International Capital Market Association. Sustainability-Linked Bond Principles, June 2024. https://www.icmagroup.org/assets/documents/Sustainable-finance/2024-updates/Sustainability-Linked-Bond-Principles-June-2024.pdf

[22] International Capital Market Association. The Principles Guidance Handbook, November 2024. https://www.icmagroup.org/assets/documents/Sustainable-finance/2024-updates/The-Principles-Guidance-Handbook-November-2024-041124.pdf

[23] European Commission. AI Act regulatory framework. https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

[24] European Commission Joint Research Centre. Best Available Techniques reference documents for the chemical industry. https://eippcb.jrc.ec.europa.eu/reference

[25] United States Occupational Safety and Health Administration. Process Safety Management. https://www.osha.gov/process-safety-management

[26] United States Chemical Safety and Hazard Investigation Board. Recommendations and investigations. https://www.csb.gov/recommendations/

[27] Organisation for Economic Co-operation and Development. AI Principles. https://oecd.ai/en/ai-principles

## About the Author

Authored by Chennakeshav Adya
Independent Researcher

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds, bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.