

The Treasury Copilot: Liquidity, FX, Covenants and Counterparty Risk in One Control Loop

Combining Cash Positioning, Forecast Uncertainty, Covenant Headroom, Hedge Exposures and Counterparty Limits with Authorised Human Execution

Authored by Chennakeshav Adya, Independent Researcher

September 2026

Abstract

Corporate treasury decisions are connected even when systems and teams treat them separately. A delayed receipt changes liquidity, foreign-exchange exposure and borrowing. A hedge consumes counterparty capacity and may create collateral or settlement cash flows. A refinancing decision changes covenant headroom, interest exposure and concentration. Fragmented evidence can therefore cause a locally reasonable action to create a wider balance-sheet problem.

This paper develops a controlled treasury-copilot framework that connects cash positioning, probabilistic forecasting, covenant headroom, foreign-exchange and interest-rate exposure, hedge policy, counterparty limits and settlement. The copilot assembles evidence, detects exceptions, runs approved scenarios and prepares recommendations. Authorised people retain execution, override and escalation. The framework separates observed balances, forecast assumptions, contractual terms, model outputs and decisions so that each recommendation can be reproduced and challenged.

Five original figures and five decision tables present the treasury control loop, cash-forecast fan, covenant-headroom bridge, hedge decision matrix and counterparty heat map. A worked example uses a hypothetical multinational company and illustrative assumptions. All amounts, percentages, probabilities, scores, limits and scenarios are management assumptions for analytical design. The paper does not provide accounting, audit, legal, regulatory, tax, investment, treasury, banking, valuation, security, privacy or technology advice and does not recommend any transaction, hedge, counterparty, system or model.

JEL Classification: G32, G33, F31, C53, M15

Keywords: corporate treasury, artificial intelligence, liquidity, cash forecasting, foreign exchange, covenants, hedging, counterparty risk, model governance, human authorisation

1. Treat treasury as one connected decision system

Treasury protects the organisation's ability to meet obligations while controlling financing, market and counterparty risk. Cash balances, receivables, payables, debt, derivatives, guarantees and bank relationships interact across legal entities and currencies. A decision made inside one workstream can change several other exposures before the next reporting cycle.

A copilot can create a common evidence and recommendation layer. It can reconcile bank data, ledgers, forecasts, facilities, covenants, market data and limits. It can identify missing information, rank exceptions and calculate approved scenarios. Its output should be a decision packet with sources, assumptions, confidence, controls and required authority.

The system should not collapse evidence and judgement into one score. An observed bank balance differs from a forecast receipt. A modelled covenant ratio differs from a legally

interpreted covenant. A proposed hedge differs from an authorised trade. The control loop preserves these distinctions while showing their interaction.

The connected view should distinguish measurement frequency from decision frequency. Bank balances can refresh many times a day, covenant forecasts may update weekly, and counterparty credit evidence may follow a formal review cycle. The system can recompute an indicator when data change, while the governance process determines whether the change requires action. This prevents routine noise from overwhelming approvers and prevents a slowly refreshed control from being mistaken for a current fact.

Economic value can arise through better timing and avoided error. Earlier recognition of a shortfall can preserve financing alternatives. Better exposure evidence can reduce unnecessary hedging. Consistent counterparty aggregation can avoid excess concentration. These benefits require a documented comparison with the decision process that would otherwise have occurred. They should be measured net of technology, data, assurance, operating and change costs.

2. Define the operating perimeter

The perimeter should identify legal entities, accounts, currencies, instruments, facilities, counterparties, payment systems, forecasts and users. It should state which records are authoritative and which are derived. Exclusions, manual accounts, restricted cash, trapped cash and unconnected subsidiaries require visible treatment.

Time horizons need separate control. Intraday liquidity supports payment completion. Daily and weekly horizons govern cash concentration and facilities. Monthly and quarterly horizons support covenant, funding and hedge decisions. Longer horizons inform refinancing and strategic liquidity. A single forecast frequency cannot serve all decisions equally well.

Materiality should be defined for both individual items and connected effects. A small forecast variance can become material when it occurs near a covenant threshold or payment cut-off. A large balance can be less decision-relevant when it is restricted and stable. The perimeter should therefore connect amount, timing, accessibility and consequence rather than rely on one absolute threshold.

Ownership can follow the three-lines model where appropriate. Treasury and finance operate the process and own decisions within delegation. Risk, compliance and control functions provide challenge and monitoring within their mandates. Internal audit can assess design and effectiveness independently. The allocation depends on the organisation, and the system should represent the approved structure instead of inventing authority.

The perimeter should also identify jurisdictions, data restrictions, delegated authorities and contractual constraints. Cash can be visible without being transferable. A facility can be committed while draw conditions remain unsatisfied. A hedge can be economically useful while accounting, documentation or counterparty capacity limits its feasibility.

3. Build the treasury control loop

The loop begins with observed positions and contractual obligations. It adds forecast distributions, exposure calculations, covenant tests and counterparty capacity. Approved policy then defines eligible actions. Human review authorises execution. Confirmations, settlement and realised outcomes return evidence to the next cycle.


Every stage needs an owner, timestamp and status. Data that arrive after the decision cut-off should not silently replace the version used for approval. Model changes should carry version, validation and effective date. Overrides should record the reason, authority and effect on risk.

The loop should include explicit hold states. A recommendation can be technically complete while a legal interpretation, bank confirmation or counterparty review is pending. A hold prevents downstream execution and records the missing evidence. Release from hold requires the named owner and evidence, creating a clear alternative to informal messages or assumed consent.

Control design should also address conflicting recommendations. A liquidity model may favour drawing a facility while a covenant model shows reduced headroom. A counterparty engine may favour reallocating a hedge while settlement operations cannot support the alternative before cut-off. The copilot should present the conflict and feasible combinations, then route the decision to the authorised level.

The copilot can shorten the loop by bringing exceptions to the appropriate person. It should also preserve enough context for challenge. Faster presentation has economic value when it reduces idle cash, emergency borrowing, missed hedges or payment failure. Speed without control can amplify error.

Figure 1. The controlled treasury decision loop



Evidence moves through forecasting, constraint testing, authorised action and realised-outcome review without giving the model execution authority.

Table 1. Minimum treasury decision record

Field	Required content	Evidence	Control use
position	cash, debt, derivatives and restrictions	bank, ledger and contract records	starting state
forecast	distribution by entity, currency and bucket	driver-linked assumptions and error history	uncertainty
exposure	amount, timing, offset and sensitivity	contracts, orders and approved forecast	risk measurement
constraints	covenants, liquidity, limits and permissions	agreements, policy and delegated authority	feasibility
proposal	action, amount, tenor, price and alternatives	scenario and market evidence	decision
control	preparer, checker, approver and execution route	workflow and access records	accountability
outcome	confirmation, settlement, realised cash and variance	independent records	learning

Each recommendation should remain reproducible from controlled evidence and policy.

4. Establish authoritative data layers

Treasury data should be classified by authority and latency. Bank statements and payment-system records can evidence external cash. The general ledger provides accounting classification. Treasury-management systems can maintain positions and deals. Enterprise systems contribute orders, invoices, payroll and tax. Contracts determine legal obligations and rights.

The copilot should preserve lineage from output to source. Reconciliations need tolerance, timing and ownership. A balance difference caused by cut-off differs from an unexplained break. Manual adjustments require reason, evidence and expiry. Stale data should remain visibly stale rather than be presented as current.

Master data is part of the control environment. Entity, account, currency, counterparty, instrument and covenant identifiers should be consistent. Duplicate or ambiguous identifiers can create false offsets, limit breaches or missed obligations. Changes need approval and effective dating.

5. Separate restricted, available and transferable cash

Gross cash is rarely equal to available liquidity. Accounts can be restricted by regulation, security, joint-venture arrangements, debt terms, payment controls or local operating needs. Tax, payroll and customer money can have specific treatment. Transferability can depend on documentation, timing and approvals.

The cash position should therefore show gross, restricted, operational reserve, committed outflow and available amount. It should also show the legal entity and currency in which liquidity resides. Group-level aggregation without transfer constraints can overstate the ability to meet a local obligation.

The copilot can flag cash that appears idle and test a permitted concentration route. The recommendation should include transfer cost, tax and legal review status, cut-off, counterparty and resulting entity positions. Authorised treasury staff decide whether the transfer is appropriate.

6. Build a driver-based cash forecast

A driver-based forecast connects receipts and payments to operating evidence. Customer receipts depend on invoices, disputes, terms and behaviour. Supplier payments depend on approved obligations and scheduling. Payroll, tax, capex, debt service and dividends have different certainty and governance.

Forecast rows should carry source, owner, expected date, currency, amount, confidence and update history. Rules can infer timing from behaviour, while material items retain accountable review. Aggregate statistical accuracy cannot replace scrutiny of a single payment whose failure would breach a covenant or interrupt operations.

Forecast logic should distinguish committed, highly probable, probable and contingent flows using defined criteria. The categories support scenarios and liquidity buffers. They should not be used as accounting conclusions unless qualified professionals confirm the relevant requirements.

7. Represent forecast uncertainty as a distribution

A point forecast can conceal the range that matters for liquidity. The model should estimate a distribution by time bucket, entity and currency. Scenario paths can combine receipt delays, payment acceleration, FX movement, rate change, facility draw and collateral calls.


The distribution should reflect correlations where evidence supports them. A market shock can delay customers, weaken currency and tighten bank capacity together. Treating these effects as independent can understate downside. Correlation assumptions require history, rationale and sensitivity.

The board or treasury committee should agree the decision percentile and buffer policy. A higher percentile reduces shortfall probability while consuming more liquidity. The selected level depends on consequence, funding reliability, risk appetite and contingency options.

Distribution design should separate ordinary volatility from discrete events. Customer-payment behaviour can support a statistical range when sufficient controlled history exists. A refinancing failure, sanctions event, cyber incident or market closure may require a structured scenario because history is sparse or irrelevant. Combining both approaches gives management a central distribution and named severe pathways without pretending that every tail event has a stable probability.

Back-testing should compare realised outcomes with the forecast distribution, not only the median. If outcomes fall outside the stated range too often, the model is miscalibrated or the environment has changed. If ranges are consistently very wide, the output may offer little decision value. Calibration, sharpness and decision usefulness should therefore be reviewed together.

Figure 2. Cash forecast fan with liquidity threshold



The paths and values are hypothetical assumptions that illustrate how uncertainty widens with the decision horizon.

Table 2. Forecast evidence and response by time horizon

Horizon	Primary evidence	Main uncertainty	Typical control response
intraday	payment queue, balances and settlement cut-offs	timing and operational failure	reserve, sequence and escalate
one to seven days	confirmed receipts, payroll, tax and debt service	customer and cut-off variation	position cash and facilities
two to thirteen weeks	invoices, procurement, capex and financing plans	behaviour and operating change	buffers, hedges and funding
three to twelve months	budget, pipeline, maturities and scenarios	demand, margin and market movements	refinancing and policy review
beyond twelve months	strategic plan and capital programme	structural and model uncertainty	capacity and contingency design

Accuracy, control and action should be matched to the decision horizon.

8. Measure forecast error as a managed risk

Forecast error should be measured by horizon, entity, currency, flow type and direction. Mean absolute error can show size. Bias can reveal consistent optimism or conservatism. Tail error can show shortfall risk. Service measures can show whether treasury had enough warning to act.

An error is useful only when its cause is classified. Timing, amount, missing item, duplicate, cancellation, FX translation and data latency need separate treatment. The owner should know which errors are controllable and which reflect genuine uncertainty.

Model performance should influence buffers and confidence. A flow with weak history can receive a wider distribution or stronger review. Improved accuracy should be demonstrated out of sample and across operating conditions before the organisation reduces liquidity protection.

9. Connect the forecast to executable liquidity

The forecast should map each downside path to actions that are legally and operationally available. Actions can include cash concentration, deposit maturity, facility draw, payment rescheduling, receivables action, asset sale, capital call or contingency funding. Each has lead time, cost, conditions and capacity.

Committed facilities should be tested for draw conditions, representations, covenants, utilisation limits and notice periods. Uncommitted lines are potential capacity rather than assured liquidity. Market issuance can become unavailable when it is most needed.

The copilot can rank actions by cost, speed, capacity and residual risk. It should expose assumptions and conflicts. A facility draw can protect cash while consuming covenant headroom or counterparty concentration. The authorised decision should consider the combined effect.

10. Maintain a liquidity ladder

The liquidity ladder aligns inflows, outflows, buffers and actions by time bucket. It should show opening cash, contractual flows, forecast flows, stress adjustments and closing available liquidity. It should remain available at group, entity and currency levels.

Buckets should match operating reality and cut-offs. A monthly bucket can conceal a large payment due before a receipt. Intraday peaks can matter where payment systems or collateral require gross liquidity. The method should avoid netting flows that cannot legally or operationally offset.

Thresholds need early-warning levels and escalation. A forecast breach at a longer horizon may trigger refinancing preparation. A near-term breach may trigger contingency actions. The record should show who can approve each response.

11. Translate debt agreements into controlled covenant logic

Covenants begin with executed agreements, amendments and qualified legal interpretation. Definitions can differ from accounting measures and between facilities. EBITDA adjustments, debt, cash, permitted acquisitions, distributions and cure rights require exact treatment.

The covenant engine should store definitions, test dates, thresholds, baskets, cure mechanics, reporting requirements and evidence. Formula changes need four-eyes review. A natural-language model can assist extraction, while qualified people confirm meaning and implementation.

Outputs should distinguish legal result, management forecast and model sensitivity. A ratio that appears compliant under one interpretation can still require review. The copilot should link every component to its source and highlight unresolved terms.

12. Build the covenant-headroom bridge

Headroom should be explained from current reported position to the next test and downside cases. The bridge can include trading, working capital, acquisitions, disposals, capex, debt draw, FX translation, interest, permitted adjustments and cure actions.


Sensitivity should show both numerator and denominator effects. An adverse event can reduce earnings while increasing debt. Currency movements can change translated debt and earnings

differently. Interest-rate changes can affect coverage and cash.

The bridge supports earlier intervention. Management can identify which operating or financing actions consume headroom and which preserve it. It should also show lead time for waivers, equity cures, asset sales or refinancing, without assuming availability.

Definitions should be tested against transaction proposals before documents are signed. An acquisition can affect leverage through purchase consideration, assumed debt, earn-outs, integration cost and permitted adjustments. A disposal can release cash while reducing earnings. The bridge should show these effects across test dates and identify which components require lender consent or qualified interpretation.

Forecast headroom should include reporting and compliance lead time. A ratio may remain above the legal threshold while approaching an internal point at which management needs to prepare certificates, engage lenders or restrict discretionary action. Internal triggers should be approved, linked to actions and monitored separately from the contractual covenant. This avoids describing an internal warning as a legal breach.



Values are hypothetical assumptions and do not represent an actual borrower, facility or forecast.

13. Forecast covenant headroom probabilistically

A single base-case ratio is insufficient when cash, earnings, rates and FX are uncertain. The model can generate a distribution of covenant outcomes using approved driver ranges and correlations. It should identify the probability and timing of early-warning and legal thresholds.

Simulation does not convert legal uncertainty into mathematics. Unresolved definitions should appear as separate interpretations or explicit limitations. The decision packet should show which source of uncertainty dominates the result.

Management can use the distribution to decide when to preserve liquidity, reduce commitments, engage lenders or prepare alternatives. The action threshold should account for lead time and consequence, not only the contractual boundary.

14. Connect capital commitments to headroom

Capex, acquisitions, dividends, buybacks and restructuring can change liquidity and covenants together. Approval papers should show the combined effect across forecast periods. Committed expenditure and contingent consideration require separate visibility.

The copilot can test a proposed transaction against facility definitions, baskets and forecasts. It can surface conditions and missing evidence. Final legal, accounting and financing conclusions require authorised specialists.

The decision should preserve a record of the approved case and conditions. If price, funding or operating assumptions change, the commitment should return to the relevant gate. This prevents a transaction from consuming more headroom than the approved analysis contemplated.

15. Create an exposure inventory before hedging

Foreign-exchange and interest-rate risk should begin with a controlled inventory of contractual, highly probable and balance-sheet exposures. Each item needs currency, amount, timing, entity, source, confidence and offset. Forecast exposures should remain distinct from recognised items.

Natural offsets can reduce risk when timing, entity and currency align and when policy permits netting. A receivable and payable in the same currency may not offset if their dates or legal entities differ. Intercompany funding can create additional exposures and documentation requirements.

The inventory should avoid hedging gross flows that are already economically offset. It should also avoid relying on a forecast net amount when the underlying components can diverge. Exposure quality determines hedge quality.

16. Measure risk in cash and accounting terms

Treasury should show the cash effect of market movements, not only a mark-to-market measure. FX changes can affect receipts, payments, debt service, collateral and covenant ratios. Rate changes can alter interest cost, derivative value and refinancing capacity.

Accounting treatment can change earnings and equity volatility. IFRS 9 hedge accounting requires formal designation, documentation and effectiveness requirements for qualifying relationships. Treasury, finance and auditors should align the operational hedge with the applicable accounting analysis.

Risk measures can include sensitivity, cash-flow-at-risk, value-at-risk and scenario loss. Methods, assumptions, limitations and use should be documented. IFRS 7 disclosures also emphasise the nature and extent of financial-instrument risks and how management addresses them.

17. Define the hedge objective

A hedge should state the exposure, risk, horizon and management objective. Objectives can include protecting a budget rate, preserving cash, limiting covenant volatility or reducing refinancing sensitivity. Different objectives can produce different instruments and hedge ratios.

The policy should define eligible instruments, counterparties, tenors, limits, delegation and prohibited activity. It should address forecast uncertainty, over-hedging, optionality, rollover and

unwind. The recommendation should show compliance before execution.

Hedge effectiveness should be assessed economically and, where relevant, for accounting. A low-cost instrument can be unsuitable if it creates liquidity, basis, settlement or documentation risk outside the approved objective.

18. Construct the hedge decision matrix

The decision matrix combines exposure confidence, downside consequence, liquidity effect, accounting treatment, counterparty capacity and instrument characteristics. It should present alternatives and the unhedged case. Pricing should include spread, premium, collateral, credit and operational cost where relevant.

Forecast exposures can support layered hedging when policy permits. Nearer and more certain flows may receive higher coverage. Longer or uncertain exposures may use lower ratios or options. The method should respond to evidence rather than rely on a static percentage.


The copilot can calculate approved alternatives and detect policy exceptions. An authorised dealer or committee selects the action. Trade execution should remain segregated from recommendation generation, confirmation and settlement.

Transaction-cost comparison should use executable evidence. Indicative prices can support exploration, while authorisation should use current market inputs and approved dealing procedures. The analysis should reflect bid-offer spread, option premium, credit valuation, collateral, break cost and operational burden where material. Price alone does not determine suitability.

Optionality should be valued against the protected decision. An option can preserve participation in favourable movements and limit adverse cash outcomes, while its premium consumes liquidity. A forward can provide certainty with different termination and collateral consequences. The matrix should show the cash profile across relevant scenarios and avoid treating instrument complexity as evidence of sophistication.

Layered hedging needs a rule for changing confidence. As an exposure moves from forecast to order to invoice, its evidence strengthens and the eligible hedge ratio may change. The process should also define what happens when dates shift. Rollover decisions can create cost and accounting effects that belong in the initial policy design.

Figure 4. Hedge decision matrix across evidence and consequence



The categories illustrate a governance method; they are not instrument recommendations or policy limits.

Table 3. Hedge recommendation record

Dimension	Required question	Evidence	Exception response
exposure	what amount, currency, date and confidence?	contract, order or forecast	validate or reduce scope
objective	which cash or risk outcome is protected?	approved policy and budget	escalate unclear purpose
alternatives	unhedged, natural offset, forward, option or structure?	comparable scenarios and pricing	obtain qualified review
liquidity	what premium, margin, collateral or settlement cash arises?	terms and stress tests	reserve or redesign
accounting	what designation and volatility can arise?	finance and audit analysis	preserve separate conclusion
counterparty	is limit, documentation and settlement capacity available?	limits, agreements and credit evidence	allocate or defer
authority	who prepares, checks, approves and executes?	delegation and workflow	block unauthorised action

A recommendation should show economic purpose, constraints and alternatives before authorisation.

19. Control over-hedging and forecast change

Forecast transactions can fall, move or disappear. The hedge can then exceed the remaining exposure. The system should compare hedge notional and maturity against updated exposure and policy at every cycle.

Exceptions require a defined response. Treasury can rebalance, roll, unwind or retain an instrument only within authority and after considering cost, accounting and risk. The copilot should calculate alternatives and preserve the reason for action.

Forecast change should also update performance analysis. A hedge should not be judged only by standalone gain or loss. Its outcome should be considered against the exposure and approved objective.

20. Incorporate settlement risk

FX settlement risk arises when one party delivers currency and the counter-currency is not received. The 2025 BIS Triennial Survey reported that payment-versus-payment eliminated settlement risk for a portion of surveyed settlement, while substantial volumes used methods that mitigated or retained exposure.

The recommendation should identify settlement method, timing, currency pair, cut-off and gross exposure. Payment-versus-payment, netting and other mitigants have eligibility and operational requirements. Treasury should understand the actual route used for each trade.

Settlement events affect intraday liquidity and counterparty exposure. Confirmation, matching, payment release and receipt should be monitored independently from front-office execution. Failed or delayed settlement needs immediate escalation and fallback.

21. Create a counterparty evidence register

Counterparty limits should combine approved credit appetite, current exposure, potential future exposure, settlement exposure, deposits, derivatives, facilities and operational dependency. Legal entities within a banking group require correct aggregation according to policy.

Credit ratings can inform the register but should not be the only evidence. Market signals, financial information, country exposure, collateral, documentation, service incidents and concentration matter. Evidence should carry source date and owner.

The copilot can consolidate exposures and compare them with limits. It should identify stale credit reviews, missing agreements and proposed actions that consume capacity. Limit approval and exception authority remain human governance decisions.

22. Calculate exposure consistently

Current exposure can include deposits, positive derivative value, unsettled trades and other claims. Potential future exposure depends on instrument, maturity and market movement. Settlement exposure can be gross when currencies are delivered at different times.

Netting and collateral should be recognised only when enforceability and operational conditions are supported. A contractual provision that has not been reviewed or implemented should not automatically reduce exposure. Wrong-way risk requires attention when counterparty quality deteriorates with the underlying exposure.

Methods should be documented and independently reviewed. Changes can move reported utilisation materially. The dashboard should distinguish calculation change from economic change.

23. Visualise concentration and capacity

A counterparty heat map can show exposure against limit and credit or operational condition. It should include near-term settlement peaks and proposed transactions, not only current end-of-day balances. Group and country concentration can sit above individual limits.

Colour should trigger action through defined thresholds. Amber can require review or reallocation. Red can block new exposure and escalate remediation. Overrides need authority, duration and compensating controls.

The map should not compress every dimension into one colour without evidence. Users need the underlying amount, limit, tenor, instrument and source. A visual supports triage and does not replace the credit decision.

Capacity should be reserved for expected operational needs. A counterparty can have unused credit limit while the remaining capacity is needed for payroll accounts, trade finance, guarantees or existing derivative settlement. The recommendation engine should identify earmarked capacity and avoid allocating the same headroom to multiple proposed actions.

Country and group concentration can matter even when every individual legal-entity limit is respected. Common ownership, clearing bank, payment infrastructure, currency or sovereign exposure can create correlated failure. Aggregation rules should reflect approved policy and available evidence. The heat map should allow users to move from individual banks to these higher-level dependencies.

Limit breaches can arise from market movements after a valid trade. The response should distinguish passive breach from unauthorised new exposure, while preserving escalation and remediation. Automatic blocking should apply to new recommendations according to policy, and existing risk should be managed through an authorised plan rather than an uncontrolled forced transaction.

Figure 5. Illustrative counterparty exposure and capacity heat map



Counterparty labels, limits, scores and exposures are hypothetical assumptions for framework illustration.

Table 4. Counterparty capacity decision rules

Condition	Evidence	Copilot response	Human decision
within limit with current review	exposure, limit and current credit evidence	present available capacity	approve within delegation
approaching early warning	utilisation, settlement peak or concentration	identify alternatives	allocate or reduce
limit breach forecast	proposed action or market movement	block recommendation and escalate	exception or remediation
stale or missing evidence	expired review, agreement or netting opinion	remove unsupported capacity	validate before use
operational incident	payment, confirmation or service failure	flag dependency and fallback	suspend or continue with controls
stress concentration	correlated banks, country or currency	show aggregate exposure	diversify or accept explicitly

Thresholds are management choices and require approved limits, definitions and escalation.

24. Link counterparty risk to liquidity

Counterparty deterioration can create both credit and liquidity effects. Deposits can become inaccessible. Derivative replacement can require cash. Facility capacity can weaken. A settlement failure can interrupt payment chains. The scenario engine should combine these pathways.

Liquidity buffers should consider concentration. Large balances at one bank can create an apparent surplus and a practical dependency. Operational diversification also matters because accounts, portals, payment files and mandates need to work before a contingency.

The copilot can show the cash effect of removing one counterparty or service. Treasury should test whether alternative banks, facilities and payment routes are genuinely ready.

25. Integrate interest-rate risk and debt service

Fixed and floating debt, derivatives, deposits and forecast funding create interest-rate sensitivity. The system should show cash interest by period, reset dates, benchmarks, margins, floors and hedges. Refinancing assumptions should remain visible.

Rate scenarios should feed liquidity and covenant forecasts. Higher interest can reduce coverage, cash and headroom simultaneously. Derivative collateral or break costs can add liquidity exposure. A hedge can reduce one sensitivity while creating another.

Debt-service evidence should reconcile to executed terms and current positions. Proposed refinancing requires range and conditions rather than an assumed closing. Treasury can prepare alternatives before maturities enter a constrained window.

26. Treat facilities as conditional instruments

A committed amount is not the same as drawable cash. Conditions precedent, representations, utilisation procedures, collateral, covenants and material-event provisions can affect access. The register should show legal status and operational readiness.

Draw tests should be refreshed before a liquidity need. Signatories, notices, bank details and cut-offs should be current. A contingency plan that depends on an untested facility process has execution risk.

The copilot can pre-populate evidence and flag unresolved conditions. It should not provide a legal conclusion. Authorised treasury and legal teams confirm availability and execute.

27. Design the recommendation hierarchy

Recommendations can be classified as information, review, prepare, approve, execute or escalate. The class determines evidence, user role and urgency. A missing forecast item can require review. A projected shortfall can require preparation. A limit breach can block execution.

Every recommendation should state the trigger, evidence, alternatives, affected limits, expected cash effect, expiry and authority. Users should be able to reject or modify it with a recorded reason. Silent dismissal weakens learning.

Urgency should reflect cut-offs and lead times. The system can rank items by consequence and time to act. It should avoid using opaque priority scores without explanation.

28. Preserve human execution authority

Payment and trading authority should remain within approved mandates and segregation of duties. The copilot can draft instructions or tickets only where policy permits. Creation, approval, release, confirmation and reconciliation should remain appropriately separated.

High-risk actions may require dual approval, call-back or independent verification. Access should use least privilege, strong authentication and monitored service accounts. Changes to beneficiary, counterparty and settlement data require special control.

Emergency procedures should be defined before an incident. A model or system outage cannot suspend the organisation's ability to meet critical obligations. Manual fallback should be tested and proportionate.

29. Govern model and data risk

Forecast, exposure and recommendation models need inventory, ownership, purpose, limitations, validation and monitoring. Materiality determines the depth of governance. Models that influence payments, hedges or funding decisions require stronger control than descriptive summaries.

Validation should test conceptual design, data, implementation, performance, stability and use. Generative components require controls for unsupported output, prompt injection, confidential data and changing provider behaviour. Source-linked output and deterministic calculations can reduce ambiguity.

Users need training in appropriate reliance. The interface should show uncertainty and exceptions. A confident narrative should never obscure missing data or an unresolved constraint.

Model inventory should identify components that calculate, classify, retrieve or generate. Deterministic interest and covenant calculations have different risks from machine-learning forecasts or language-model document extraction. A combined application can contain several models and non-model rules. Governance should be proportionate to each component and to the

consequence of the final use.

Independent validation should use representative data and operational workflows. Back-tests can assess forecast calibration. Benchmarking can compare calculations with controlled alternatives. Sensitivity can expose unstable assumptions. Outcome analysis can show whether users follow, override or misunderstand recommendations. Findings need severity, owner, remediation date and approved use restriction.

Change control should cover data mappings, prompts, retrieval sources, model versions, policy rules and interface logic. A vendor model can change without a conventional software release. The organisation needs monitoring and contractual notification where feasible. Material change should return the affected use case to an appropriate validation and approval gate.

30. Manage third-party and concentration risk

Treasury technology can depend on banks, market-data providers, cloud platforms, model providers and integration vendors. A shared dependency can affect several processes at once. The FSB has identified third-party concentration, cyber risk, model risk and data governance as relevant AI-related vulnerabilities.

Contracts should address service, security, data use, audit, incident, change, portability, subcontractors, termination and transition. Critical services require contingency and exit planning. Evidence should reflect the actual implementation.

Concentration analysis should cover technology as well as financial counterparties. Multiple treasury applications can still depend on one cloud, identity provider or data feed. Resilience testing should include these common points.

31. Secure the treasury-copilot environment

Treasury data contain balances, counterparties, payment details, funding terms and market positions. Access, encryption, logging, environment separation and secure development should match the sensitivity. Approved data should not be entered into unapproved public tools.

The system should defend against manipulation of instructions, documents and prompts. External files can contain malicious content. Source ingestion needs scanning, isolation and validation. Material output should be checked against authoritative records.

Incident plans should address fraudulent payment requests, deepfakes, credential compromise, data leakage and model misuse. Escalation should connect treasury, security, legal, compliance and banking partners.

32. Build an immutable evidence trail

The decision record should preserve inputs, model version, assumptions, policy version, recommendation, users, approvals, execution details, confirmations and outcomes. It should support reconstruction after an error, audit or dispute.

Logs need integrity, retention and access control. A user should not be able to alter the evidence that supported an earlier approval. Corrections should create a new version and retain history.

The trail also supports improvement. Management can identify recurring overrides, weak forecasts, late decisions and control breaks. Learning should update methods through approved change rather than silent tuning.

33. Design the daily treasury dashboard

The dashboard should begin with decisions and exceptions. It can show available liquidity, forecast distribution, threshold breaches, covenant headroom, unhedged exposure, counterparty capacity, settlement status and required approvals. Each item should link to evidence.

Users need role-specific views. A dealer needs executable exposure and limits. The treasurer needs consolidated risk and decisions. Finance needs accounting and cash evidence. The board needs material trends, downside and governance rather than transaction detail.

The dashboard should show data freshness and model status. A green measure built on stale data is misleading. Missing sources and unresolved reconciliations belong in the decision view.

34. Establish the weekly control forum

A weekly forum can review forecast error, liquidity scenarios, covenant paths, hedge status, counterparty concentration and model exceptions. Participants should include treasury, finance, risk and relevant business owners. Legal, tax, accounting and technology specialists join when required.

The forum should decide actions, owners and dates. It should distinguish routine delegation from exceptions needing committee or board authority. Unresolved matters remain visible until closed.

Minutes and evidence should connect to the system record. The process can be short when data and authority are clear. Long meetings often reflect missing ownership or reconciliation.

35. Report to the board through decision ranges

Board reporting should show available liquidity under approved scenarios, covenant headroom, funding maturities, material market exposure, counterparty concentration and control incidents. Ranges and triggers are more useful than false precision.

The report should explain changes from the prior period. It should identify which drivers are observed, forecast or conditional. Proposed transactions should show their incremental effect on liquidity and headroom.

Management recommendations require clear authority and next evidence. The board retains oversight of risk appetite, funding capacity and material exceptions. The copilot provides a controlled information layer.

36. Run a hypothetical worked example

Consider a hypothetical group with five operating entities, six currencies, USD 120 million gross cash, USD 85 million available liquidity and two revolving facilities. It has quarterly leverage and interest-cover tests, forecast imports, export receipts and bank exposure across five counterparties. These figures are illustrative assumptions.

The thirteen-week point forecast shows a minimum USD 28 million available liquidity. The distribution shows an illustrative tenth-percentile low of USD 14 million after receipt delays, currency movement and capex timing. The approved early-warning threshold is assumed at USD 20 million.

The copilot identifies that the same downside also reduces covenant headroom and increases utilisation at one hedge counterparty. It prepares three alternatives: defer discretionary capex, draw a facility earlier or adjust the hedge allocation. Each alternative shows cost, capacity, conditions and residual risk.

The scenario model assumes that USD 18 million of the gross cash is restricted, USD 17 million is held as an operating reserve and USD 8 million sits in an entity that cannot transfer it before the payment date. It also assumes that a USD 22 million revolving-facility draw is available after satisfaction of specified conditions. These assumptions are visible and require confirmation.

The exposure inventory identifies an illustrative EUR payable and GBP receivable whose dates do not align. Treating them as a natural offset would leave an intervening liquidity need. The proposed hedge alternatives therefore show the gross timing profile, expected net position and settlement cash. The counterparty view also includes the peak settlement amount rather than only replacement value.

The covenant bridge shows an illustrative 32-point opening headroom index, positive trading contribution and negative effects from working capital, rates, currency and committed investment. The selected downside closes at 20 points, above an assumed 10-point internal warning. No legal conclusion is implied. The example requires the actual agreement and qualified interpretation in a real decision.

37. Apply the control loop to the example

Treasury verifies the largest receipts and discovers that one expected customer payment lacks current evidence. Finance updates the forecast confidence. Legal confirms the facility draw conditions. Risk confirms that the proposed hedge allocation would breach an internal early-warning level at Bank B.

The authorised committee selects an illustrative combination of reduced discretionary capex, a smaller facility draw and hedge allocation across two approved counterparties. The action preserves a modelled minimum liquidity of USD 24 million and keeps covenant headroom above the internal trigger in the selected scenario.

These outputs are assumptions and do not predict an outcome. The example shows how connected evidence changes a decision. A cash-only view would have missed covenant and counterparty effects.

38. Measure realised performance

After the period, the system compares realised cash, forecast distribution, actions and costs. It records whether the receipt delay occurred, whether the facility was drawn, the hedge outcome and the covenant path. Error is attributed to driver and owner.

Performance should include decisions avoided as well as actions taken. Early evidence can prevent an unnecessary draw or unsuitable hedge. The benefit claim needs an agreed

counterfactual and should avoid treating every model alert as value.

Control measures include unauthorised actions, overrides, reconciliation breaks, settlement failures, data latency, model drift and fallback performance. Financial and control outcomes belong together.

39. Implement the framework in ninety days

Days One to Thirty can define perimeter, authority, source systems, decision records and critical thresholds. The team can reconcile accounts, facilities, covenants, exposures and counterparty limits. It can select one entity group for controlled design.

Days Thirty-One to Sixty can build the forecast distribution, covenant bridge, hedge matrix and counterparty view. Model, data, security and workflow controls should be tested. Users can run parallel recommendations without execution authority.

Days Sixty-One to Ninety can operate a bounded live control loop with human approval, independent confirmation and daily reconciliation. Management can measure forecast error, recommendation quality, control exceptions and realised outcomes before widening scope.

The implementation should begin with a narrow set of high-value decisions. One legal-entity group, a small currency set and one facility family can reveal data and authority problems without exposing the whole treasury estate. Expansion can follow evidence that reconciliations, recommendations, approvals and fallback work under representative conditions.

Success criteria should combine timeliness, accuracy, control and economic outcome. Measures can include reconciled account coverage, forecast calibration, warning lead time, policy-compliant recommendations, override quality, settlement completion and attributable financing or liquidity effect. Activity counts, such as alerts generated or documents indexed, do not establish decision value.

The operating model needs named product, treasury, finance, risk, security and data owners. It also needs support and incident arrangements across business hours, currencies and payment cut-offs. A pilot that depends on exceptional effort from a small build team may not represent a sustainable control environment.

Table 5. Ninety-day controlled implementation plan

Period	Build	Evidence gate	Capital or scope decision
days 1 to 15	perimeter, authority and source register	owners and authoritative records confirmed	approve discovery
days 16 to 30	reconciled positions and constraint library	balances, facilities and covenant logic accepted	approve modelling
days 31 to 45	forecast distribution and error history	representative back-test and limitations	approve parallel run
days 46 to 60	hedge and counterparty recommendations	policy, limit and control tests passed	approve bounded users
days 61 to 75	daily loop without autonomous execution	stable evidence, overrides and fallback	approve bounded live use
days 76 to 90	outcome review and replication plan	realised performance and control acceptance	scale, redesign or stop

40. Make connected evidence the operating advantage

Treasury risk becomes harder to manage when cash, FX, covenants and counterparties are viewed in separate cycles. One delayed receipt can change liquidity, debt, hedging and concentration. A connected control loop exposes these interactions before authority is exercised.

The copilot's economic role is to assemble evidence, quantify uncertainty, test constraints and prepare reproducible alternatives. Its governance role is to preserve source, policy, approval and outcome. Authorised people retain judgement and execution.

The strongest implementation begins with controlled data and specific decisions. It measures forecast error, override, settlement and realised cash. Scope expands only when evidence shows that the system improves decisions within acceptable operational, model, security and legal boundaries.

A durable system also creates an institutional memory of how treasury decisions perform. Forecast versions reveal persistent bias. Covenant bridges show which commitments repeatedly consume flexibility. Hedge records show whether the protected cash objective was achieved. Counterparty histories show how limits, incidents and operating dependencies evolved. This evidence can improve policy, buffers, delegation and financing strategy over time.

Management should preserve the ability to challenge and retire the copilot. Performance can deteriorate when business models, payment behaviour, banking relationships, market structure or data systems change. A periodic review should confirm continued purpose, source quality, model performance, user behaviour, control effectiveness and vendor viability. Material weakness can lead to restricted use, redesign or withdrawal.

The central discipline is traceability from obligation to evidence, from evidence to scenario, from scenario to authorised decision and from decision to realised outcome. That traceability supports speed because approvers receive a coherent record. It supports resilience because fallback owners understand the position and constraints. It supports accountability because a recommendation never becomes an unexplained instruction.

References

- [1] IFRS Foundation, IFRS 7 Financial Instruments: Disclosures.
<https://www.ifrs.org/issued-standards/list-of-standards/ifrs-7-financial-instruments-disclosures/>
- [2] IFRS Foundation, IFRS 9 Financial Instruments.
<https://www.ifrs.org/issued-standards/list-of-standards/ifrs-9-financial-instruments/>
- [3] IFRS Foundation, IAS 21 The Effects of Changes in Foreign Exchange Rates.
<https://www.ifrs.org/issued-standards/list-of-standards/ias-21-the-effects-of-changes-in-foreign-exchange-rates/>
- [4] IFRS Foundation, IAS 7 Statement of Cash Flows.
<https://www.ifrs.org/issued-standards/list-of-standards/ias-7-statement-of-cash-flows/>
- [5] Bank for International Settlements, Uncovering FX Settlement Risk: New Measures from the 2025 BIS Triennial Survey, 2026.
<https://www.bis.org/publications/uncovering-fx-settlement-risk-new-measures-2025-bis-triennial-survey>
- [6] Committee on Payments and Market Infrastructures, Principles for Financial Market Infrastructures.
<https://www.bis.org/cpmi/publ/d101a.pdf>

- [7] Basel Committee on Banking Supervision, Principles for Sound Liquidity Risk Management and Supervision. <https://www.bis.org/publ/bcbs144.htm>
- [8] Central Bank of the UAE, Measurement of Liquidity Risk. <https://rulebook.centralbank.ae/en/rulebook/b-measurement-liquidity-risk>
- [9] Central Bank of the UAE, Qualitative Requirements for Liquidity Risk Management. <https://rulebook.centralbank.ae/en/rulebook/article-2-qualitative-requirements>
- [10] Central Bank of the UAE, Risk Management Regulation and Standards. <https://rulebook.centralbank.ae/en/rulebook/risk-management-regulation-and-standards>
- [11] Central Bank of the UAE, Risk Management. <https://www.centralbank.ae/en/our-operations/risk-management/>
- [12] Central Bank of the UAE, Guidance Note on Consumer Protection and Responsible Adoption and Use of Artificial Intelligence and Machine Learning, 2026. <https://rulebook.centralbank.ae/en/rulebook/guidance-note-consumer-protection-and-responsible-adoption-and-use-artificial-intelligence>
- [13] Dubai Financial Services Authority, AI Survey 2025: Generative AI Adoption and Governance in the DIFC, 2025. <https://www.dfsa.ae/news/new-dfsa-ai-survey-generative-ai-adoption-has-nearly-tripled-within-difc-last-12-months-governance-continues-develop>
- [14] Financial Stability Board, The Financial Stability Implications of Artificial Intelligence, 2024. <https://www.fsb.org/2024/11/the-financial-stability-implications-of-artificial-intelligence/>
- [15] Financial Stability Board, Enhancing Third-Party Risk Management and Oversight: A Toolkit for Financial Institutions and Financial Authorities, 2023. <https://www.fsb.org/2023/12/enhancing-third-party-risk-management-and-oversight-a-toolkit-for-financial-institutions-and-financial-authorities/>
- [16] Bank for International Settlements, Governance of AI Adoption in Central Banks, 2025. <https://www.bis.org/publ/othp90.htm>
- [17] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework 1.0, 2023. <https://doi.org/10.6028/NIST.AI.100-1>
- [18] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, 2024. <https://doi.org/10.6028/NIST.AI.600-1>
- [19] National Institute of Standards and Technology, Cybersecurity Framework 2.0, 2024. <https://doi.org/10.6028/NIST.CSWP.29>
- [20] US Federal Reserve and Office of the Comptroller of the Currency, Supervisory Guidance on Model Risk Management SR 11-7 and OCC 2011-12. <https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm>
- [21] Bank of England Prudential Regulation Authority, Supervisory Statement SS1/23 Model Risk Management Principles for Banks. <https://www.bankofengland.co.uk/prudential-regulation/publication/2023/may/model-risk-management-principles-for-banks-supervisory-statement>
- [22] UK National Cyber Security Centre, Guidelines for Secure AI System Development, 2023. <https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development>
- [23] SWIFT, ISO 20022. <https://www.swift.com/standards/iso-20022>
- [24] SWIFT, Customer Security Controls Framework. <https://www.swift.com/myswift/customer-security-programme-csp/security-controls>
- [25] International Organization for Standardization, ISO/IEC 42001 Artificial Intelligence Management System. <https://www.iso.org/standard/81230.html>

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.