MatchPoint Research Working Paper No. 2026/T12

Synthetic Data and Simulation for Risk Management

A three-gate framework for family-office portfolios and private-credit risk

Chennakeshav (CK) Adya
Independent Researcher
August 2026

ABSTRACT

Background. Family offices and private-credit funds must assess concentrated, illiquid and path-dependent risks from limited and heterogeneous observations. **Objective.** This paper develops a governed method for using synthetic data, simulation and stress testing in those decisions. **Approach.** The analysis reviews 26 primary, regulatory, central-bank and peer-reviewed sources available through 1 August 2026 and translates them into a three-gate operating standard. **Findings.** Synthetic records require separate purpose-specific utility, privacy and tail-adequacy tests. Similarity does not establish anonymity, decision utility or stress adequacy. The proposed architecture keeps real data within an approved enclave, versions every generator and scenario, requires independent validation, and binds each release to a purpose, recipient and expiry. **Implications.** Family offices can begin with liquidity, commitments, concentration and operational testing. Private-credit funds can begin with covenant surveillance, correlated borrower stress, recovery timing, refinancing and fund liquidity. Attributed revenue, cost reduction and loss reduction remain zero until approved observed evidence exists.

Highlights

Unit. The accountable unit is a versioned, accepted risk packet.

Utility. Fitness is tested against the named downstream decision on held-out real data.

Privacy. Synthetic provenance does not establish anonymity or disclosure safety.

Tails. Extreme dependence, transitions, duration and reverse thresholds receive distinct challenge.

Authority. Data, model, risk, credit, investment, legal and release decisions remain human-owned.

JEL Classification: C15, C53, C58, G11, G17, G23, G32, L86, M15, O33

Keywords: synthetic data, simulation, stress testing, model risk, differential privacy, family office, private credit, liquidity risk, tail risk, data governance

1. INTRODUCTION

Risk management has a data problem before it has a technology problem. A family-office investment committee often has concentrated exposures, uneven manager reporting, private-company information, short internal histories and a limited number of realised losses. A private-credit fund may have detailed borrower files, monthly covenants, collateral evidence and servicing data, yet only a small population of comparable defaults. The events that matter most are consequently the events observed least often.

Synthetic data and simulation can expand the set of questions that a risk team is able to test. Synthetic data creates artificial records from a statistical, rules-based or generative process. Simulation creates paths or scenarios from stated dynamics and assumptions. Stress testing applies adverse conditions to positions, counterparties, cash flows or operating systems. Each tool can support model development, control testing, data sharing, scenario rehearsal and resilience analysis. None of them converts an unobserved tail event into an observed fact.

The regulatory evidence supports controlled experimentation. The Financial Conduct Authority's Synthetic Data Expert Group identified data augmentation and bias mitigation, model testing and validation, and data sharing for fraud controls as three areas of potential financial-services value [1]. The FCA also states that the group's report was collectively authored and does not represent FCA views or imply compliance [1]. The FCA's 2026 anti-money-laundering project used a fully synthetic transaction dataset based on real UK retail-banking data and supplemented it with realistic synthetic machine-learning scenarios [3]. Project Aurora at the Bank for International Settlements used synthetic transaction data to compare machine-learning configurations across bank, national and cross-border settings; its report states that real-world data remains important for assessing feasibility and impact [12].

Privacy requires a separate test. NIST states that many synthetic-data techniques do not satisfy differential privacy and that purpose-built differentially private analyses can be more accurate than releasing a general synthetic dataset [6]. NIST also states that no privacy-preserving release can be guaranteed valid for every analysis [7]. The UK Information Commissioner's Office describes a utility-privacy trade-off: higher fidelity can raise disclosure risk, and source-data bias can be reproduced [8]. Empirical work has found that synthetic-data privacy and utility can be difficult to predict [24], while the TAPAS project provides a threat-modelled set of attacks for testing disclosure risk [25]. The label **synthetic** is therefore a description of provenance, not a privacy certificate.

Model and stress-test governance also remains intact. The Prudential Regulation Authority's model-risk principles cover model identification, governance, development and use, independent validation, and mitigants [14]. The Federal Reserve's 2026 guidance uses a risk-based and tailored approach, generally expects validation before use, and retains oversight of vendor models [15]. Basel stress-testing principles cover governance, objectives, methodology, resources and documentation, including sensitivity, scenario and reverse stress testing [16]. These publications concern institutions within their respective scopes. This paper uses them as control references and does not determine the legal or supervisory position of a particular family office or private-credit firm.

This paper develops an operating framework for **Family-Office CIOs and Heads of Alternatives (A2)** and **Private Credit, Direct-Lending and Special-Situations Funds (A3)**. It addresses six questions:

1. Which risk-data gaps can synthetic data or simulation usefully address?
2. How should a team distinguish artificial records, simulated paths and stress scenarios?
3. Which utility, privacy and tail tests should a risk committee require?
4. What architecture protects real data and preserves scenario lineage?
5. How should the two ICPs build scenario libraries that reflect their different exposures?
6. How should productivity and economics be measured without converting illustrative assumptions into reported benefits?

The central proposition is testable: **synthetic data becomes decision-useful only when it passes purpose-specific utility, privacy and tail-adequacy gates inside a human-owned risk process.** A dataset may be sufficiently useful for system testing and still be unsuitable for credit calibration. A generator may preserve averages and still erase the joint defaults, liquidity spirals or concentration losses that drive an investment decision. The proposed control model treats each release as a governed risk object with an owner, purpose, source boundary, generator version, privacy test, utility test, tail test and expiry date.

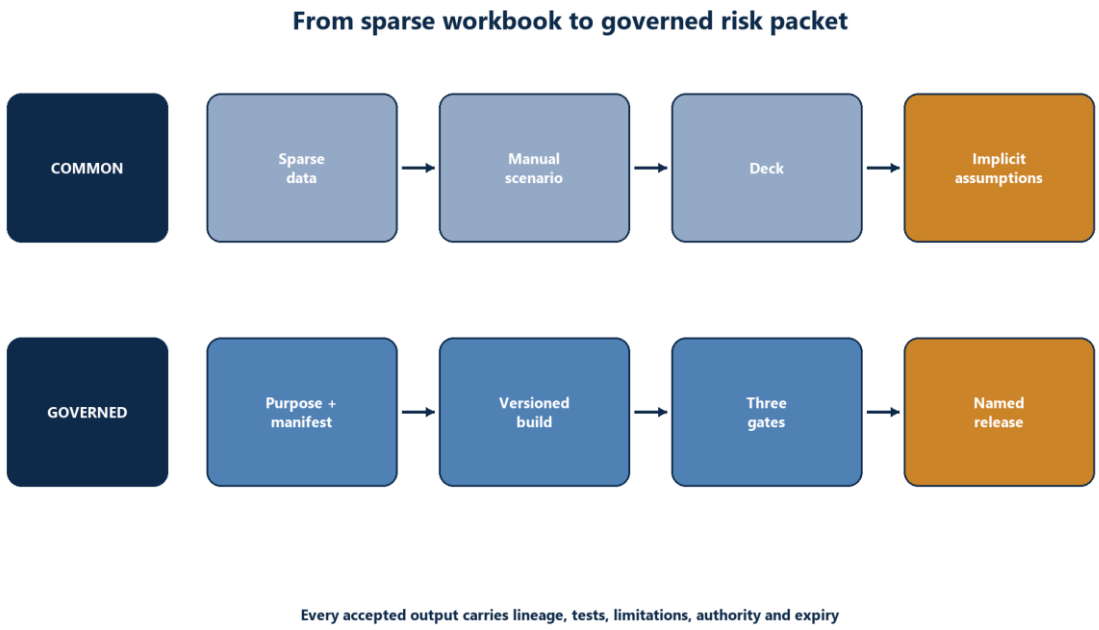

Figure 1. Risk workflow before and after governed synthetic data and simulation
Source: Matchpoint framework.

2. DEFINITIONS, SCOPE AND EVIDENCE BOUNDARIES

2.1 Three related instruments

The three instruments answer different questions. Synthetic data asks, "What artificial records can reproduce specified properties of an approved source population?" Simulation asks, "What paths arise under a stated model, parameter set and dependency structure?" Stress testing asks, "What happens under a deliberately adverse condition, including a condition chosen by reverse engineering a failure threshold?" A single programme may use all three, but their evidence labels should remain distinct.

Instrument	Operational definition	Typical output	Appropriate initial use	Principal control question
Rules-based synthetic data	Artificial records constructed from explicit business rules and distributions	Test borrowers, transactions, covenants or holdings	System, workflow and control testing	Do rules cover invalid, boundary and exception states?
Statistical synthetic data	Artificial records sampled from fitted distributions or graphical models	Tabular portfolio or loan records	Exploratory analysis and selected sharing	Which marginal and joint properties are preserved?
Generative synthetic data	Artificial records produced by a learned generator such as a GAN or related model	High-dimensional tabular or time-series records	Model development, augmentation and controlled research	What utility and privacy failures arise from fitting and sampling?
Monte Carlo simulation	Repeated paths drawn from an explicit stochastic model	Returns, cash flows, defaults, recoveries or liquidity paths	Distributional risk analysis	Are assumptions, dependencies and parameters decision-relevant?
Scenario simulation	Coherent path built from linked macro, market and operating assumptions	Multi-period asset and liability outcomes	Planning and portfolio resilience	Are scenario transmissions internally consistent?

Instrument	Operational definition	Typical output	Appropriate initial use	Principal control question
Stress test	Adverse sensitivity, scenario or reverse stress	Loss, breach, liquidity or solvency outcome	Risk appetite and contingency planning	Is the severity adequate and are management actions credible?

Synthetic population means the records generated for an approved purpose. **Source population** means the real or previously approved data used to fit, calibrate or validate the synthetic process. **Fidelity** means similarity on the properties relevant to the intended task. **Privacy** means resistance to the specified disclosure threats; it is not inferred from visual similarity. **Tail adequacy** means the ability of the model or scenario set to represent the adverse dependence, concentration, transition and liquidity behaviour material to the decision.

2.2 Evidence hierarchy

This paper uses four evidence tiers. Tier 1 consists of legislation, regulation and supervisory standards. Tier 2 consists of official regulator, central-bank and public-authority research. Tier 3 consists of peer-reviewed or conference research with a stated method. Tier 4 consists of proposed Matchpoint frameworks and explicitly labelled illustrative management assumptions. A source can support only the claim tested or stated within its scope.

The evidence cut-off is 1 August 2026. The paper does not represent that every cited control is legally required for every target firm. Legal, regulatory, privacy and model-governance applicability depends on jurisdiction, entity, activity, data, client and contractual facts. The authorship and economic scenarios in this paper remain unverified until the named author and Matchpoint management approve them.

2.3 What the external evidence supports

Evidence	Verified contribution	Boundary retained in this paper
FCA Synthetic Data Expert Group [1]	Financial-services use cases and practical considerations across augmentation, validation and sharing	Collective expert report; no FCA endorsement or compliance conclusion
FCA AML project [3]	Fully synthetic retail-banking transaction data plus synthetic scenarios for AML research	Project-specific research; no demonstrated family-office or private-credit outcome
NIST SDNist [4]	Benchmark utility and privacy evaluation for synthetic data	Tool metrics require purpose-specific interpretation
NIST differential-privacy guidance [5-7]	Formal privacy-loss evaluation and limits of universal utility	Differential privacy is technically demanding and can reduce accuracy
ICO PET guidance [8]	Data minimisation and sharing potential, with utility, disclosure and bias considerations	Guidance is under review following UK legislative change
BIS Project Aurora [12]	Synthetic transaction scenarios for comparing collaborative AML configurations	Real-world data remains necessary for feasibility and impact assessment
PRA and Federal Reserve model-risk guidance [14,15]	Inventory, tiering, validation, monitoring, governance and mitigants	Applies according to stated institutional and jurisdictional scope
Basel, EBA and NGFS stress references [16-18]	Governance, methodology, scenario and reverse-stress disciplines	Scenario sets do not exhaust every tail or tipping point
IMF and FSB private-credit work [19,20]	Data gaps, valuation opacity, borrower quality, leverage, concentration, liquidity and interconnectedness	System-level analysis does not estimate a named fund's loss
TimeGAN and CTGAN [21,22]	Methods for temporal and mixed-type tabular generation	Experimental results do not establish suitability for a specific risk decision

Evidence	Verified contribution	Boundary retained in this paper
PATE-GAN [23]	Generator design with differential-privacy guarantees	Privacy parameters and task utility require validation
Stadler et al. and TAPAS [24,25]	Empirical privacy-utility limitations and adversarial privacy auditing	Tested generators, datasets and threat models define the result boundary

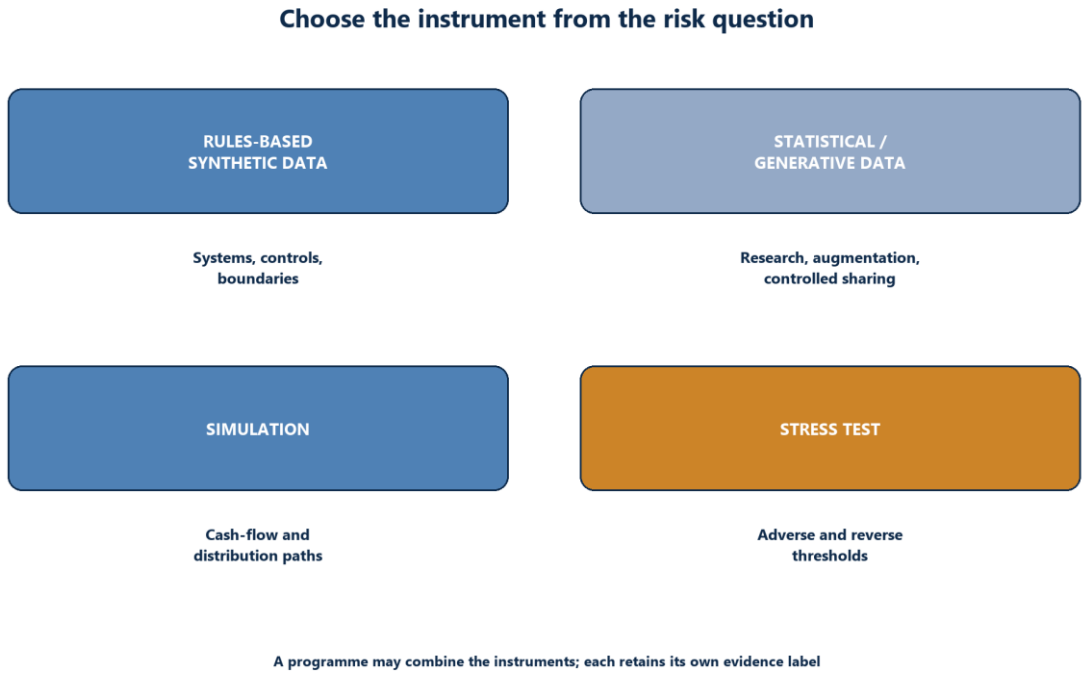

Figure 2. Synthetic-data and simulation methods mapped to risk uses
Source: Matchpoint synthesis of [1-7,12,21-26].

3. THE RISK-DATA PROBLEM FOR A2 AND A3

3.1 Family offices: concentration, heterogeneity and limited history

A family office may combine public securities, private funds, direct companies, real estate, private credit, operating businesses and liquidity reserves. The reporting frequency, valuation convention, currency, legal structure and look-through coverage vary across those assets. Internal loss history can be short because the institution is young, the portfolio has changed, or the rare event has not occurred. The same family can create correlated exposures through geography, sector, sponsor, bank, currency, operating business and personal guarantees.

Synthetic data can help build test populations for aggregation, exposure mapping, capital-call forecasting and operational-control rehearsal. Simulation can explore cash-flow sequencing, NAV shocks, FX moves, commitment pacing and distributions. Stress testing can challenge the family's ability to fund obligations while protecting strategic assets. The family-office use case is strongest when the scenario library starts with actual holdings, commitments, legal terms and liquidity rules, while keeping source data within its approved enclave.

A2 risk question	Source limitation	Suitable instrument	Required acceptance evidence
Can commitments be funded during a two-year distribution drought?	Few comparable historical pacing cycles	Scenario simulation plus reverse liquidity stress	Contracted commitments, call rules, liquid resources, time-to-cash and management-action assumptions

A2 risk question	Source limitation	Suitable instrument	Required acceptance evidence
How much hidden sponsor or sector concentration exists?	Incomplete look-through and inconsistent labels	Synthetic records for mapping tests; deterministic aggregation on real data	Coverage ratio, entity resolution, source dates and unresolved exposures
What happens when public, private and operating assets fall together?	Sparse joint-tail observations	Dependency stress and Monte Carlo sensitivity	Correlation or copula rationale, tail dependence tests and alternative structures
Can reporting and control systems process unusual structures?	Production data lacks edge cases	Rules-based synthetic test data	Boundary cases, expected outputs and defect log
How could heirs, advisers or vehicles gain inappropriate data access?	Real access incidents are rare and sensitive	Synthetic identity and permission scenarios	Threat model, least-privilege tests and incident evidence

3.2 Private credit: sparse defaults, stale valuations and nonlinear recovery

The IMF's April 2024 work identifies data gaps and risks associated with borrower quality, valuations, fund liquidity, leverage and interconnectedness in private credit [19]. The Financial Stability Board's May 2026 report estimates the market at USD 1.5 trillion to USD 2 trillion at end-2024 and highlights data gaps, valuation opacity, credit quality, interconnections, leverage, concentration and liquidity [20]. The FSB states that the market has not been tested at its current scale through a prolonged downturn [20]. These findings describe the market and identified vulnerabilities; they do not determine the performance of a named manager or loan.

Private-credit risk is shaped by path dependence. A borrower can remain current while covenant headroom erodes, liquidity falls, add-backs grow, collateral values weaken and refinancing options contract. Amendments, payment-in-kind interest, delayed reporting and sponsor support can change the timing and visibility of stress. A generator trained mainly on performing loans can produce realistic-looking borrowers while under-representing the transition sequence that matters for loss estimation.

A3 risk question	Source limitation	Suitable instrument	Required acceptance evidence
Which covenant combinations precede a breach?	Few breaches and changing documentation	Synthetic sequence augmentation for model development	Train-test separation, rare-state coverage and real-data holdout performance
How do default, recovery and workout duration interact?	Realised workouts are sparse and heterogeneous	Rules-based plus stochastic simulation	Legal seniority, collateral, jurisdiction, cure and recovery timing assumptions
Can the fund meet redemptions, capital calls and asset funding?	Cash-flow timing is irregular	Multi-period liquidity simulation and reverse stress	Facility terms, investor liquidity, unfunded commitments and sale haircuts
How would sponsor, sector or arranger concentration transmit?	Entity relationships are incomplete	Graph simulation and concentration stress	Resolved entities, exposure coverage and dependency rationale
Does the surveillance pipeline detect unusual deterioration?	Production examples lack labelled edge cases	Synthetic control-testing records	Expected alerts, false-positive and false-negative results

3.3 The missing-tail problem

A model can reproduce central distributions and fail at the tail. Marginal default rates do not determine joint defaults. Average recovery does not determine the cash-flow consequence of a long workout. Pairwise correlations do not capture a common refinancing shock. A scenario that uses a normal distribution for convenience can suppress

skewness, fat tails and regime changes. A synthetic dataset that balances default classes can improve classifier training while distorting base rates if the deployment calibration is not restored.

The design response is to separate three objectives:

- **Representation:** reproduce approved statistical properties of the source data.
- **Decision utility:** preserve the answer to a named risk question on held-out real data.
- **Stress adequacy:** cover severe dependencies and paths that may be weakly represented or absent in the source data.

No single fidelity score proves all three. NIST's utility guidance calls for a suite of metrics because a release cannot be guaranteed valid for every analysis [7]. NGFS states that users should apply its climate scenarios through their own risk frameworks and notes that tail risks may be more severe than represented, while some tipping points are absent [18]. The same discipline applies to institution-specific financial scenarios.

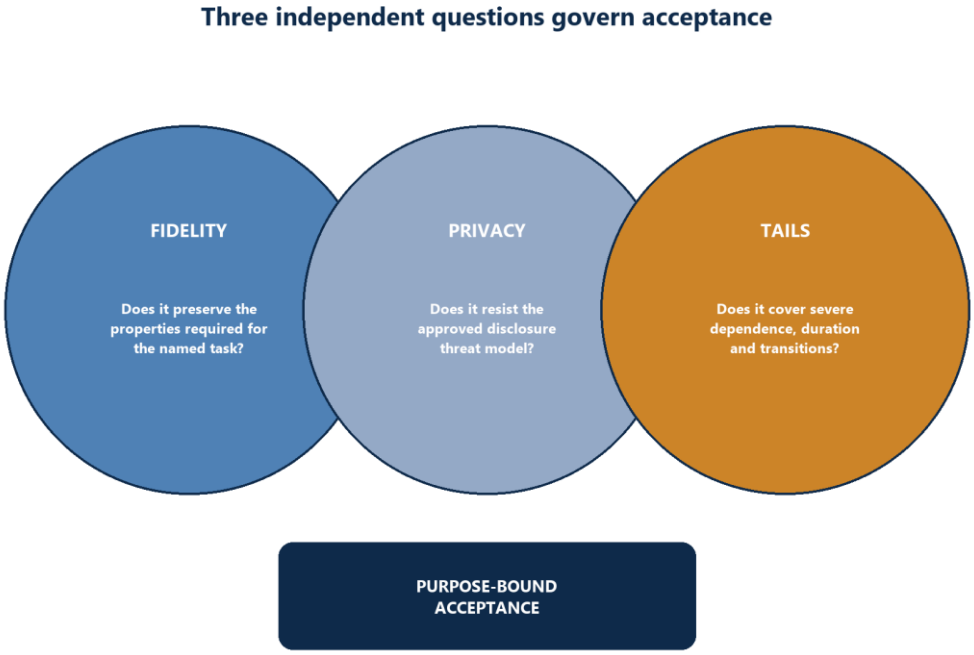

Figure 3. The fidelity, privacy and tail-adequacy risk cube
Source: Matchpoint framework; validation references [4-8,14-18,24,25].

4. GENERATION AND SIMULATION METHODS

4.1 Rules before models

Rules-based generation is often the right first instrument for control testing. A team can enumerate valid and invalid combinations of currency, covenant, legal entity, cash-flow sign, date, collateral status and permission. The expected result is deterministic, so a defect has an interpretable cause. Rules-based populations also expose data definitions that a learned generator might otherwise absorb without challenge.

Statistical generators model explicit distributions and dependencies. They can be appropriate when the data is modest, the relationships are interpretable and the intended use is narrow. Bayesian networks, copulas, mixture models and hierarchical models can incorporate domain constraints and uncertainty. Their limitations should be visible: a fitted dependency is historical, an independence assumption is substantive, and an estimated tail from few observations is unstable.

Generative models can represent complex mixed-type or sequential data. CTGAN addresses mixed continuous and discrete tabular fields and imbalance in categorical variables [22]. TimeGAN combines supervised and adversarial objectives in a learned embedding to preserve temporal dynamics [21]. PATE-GAN adapts private aggregation of teacher ensembles to provide a differential-privacy guarantee for its generator [23]. These designs offer method choices; they do not remove the need for risk-purpose validation.

Method	Strength	Material limitation	Initial risk use
Explicit rules and combinatorial generation	Interpretable, reproducible edge cases	Coverage depends on expert enumeration	Pipeline, policy and reconciliation testing
Parametric distribution	Transparent assumptions and simple sensitivity	Misspecification and weak tail representation	Cash-flow, rate and spread sensitivities
Copula or dependency model	Separates marginals and dependence	Tail choice and calibration are consequential	Joint asset, default or liquidity stress
Bayesian network	Interpretable conditional structure	Structure may omit latent drivers	Borrower and control dependencies
CTGAN-style tabular generator [22]	Mixed types and imbalanced categories	Mode collapse, memorisation and unstable rare-state fidelity	Controlled tabular augmentation
TimeGAN-style sequence generator [21]	Temporal dynamics and predictive similarity	Sequence realism may not preserve decision tails	Covenant and cash-flow sequence research
Differentially private generator [5,6,23]	Quantified privacy-loss design	Privacy budget, implementation and utility loss require expert validation	Selected sharing or research releases
Agent-based simulation	Heterogeneous actors and feedback	Behavioural rules and calibration can dominate results	Market, redemption or refinancing feedback

4.2 Simulation design

A risk simulation begins with a decision, not a distribution. The owner specifies the action that the output may influence, the time horizon, the positions, the risk factors, the state transitions, the management actions and the acceptance thresholds. Parameters then receive an evidence label: observed, externally sourced, calibrated, expert-judged or illustrative. A single scenario can mix these labels, and the final output should disclose them.

Scenario class	Construction	Example for A2	Example for A3
Historical replay	Apply an observed period or event	Public-market drawdown plus FX move	Spread widening and reduced refinancing access
Hypothetical coherent scenario	Link adverse macro, market and operating assumptions	Distribution drought plus property markdown	EBITDA decline, base-rate path and covenant breach
Sensitivity	Change one input or narrow group	Commitment-call speed	Recovery haircut or exit multiple
Monte Carlo	Sample repeated paths from a stated model	Multi-asset cash-flow-at-risk	Default, recovery and funding paths
Reverse stress	Solve for conditions that breach an outcome	Minimum liquid-reserve breach	Facility, NAV or redemption failure
Operational simulation	Generate transactions, users or failures	Permission, payment and reporting exception	Surveillance, covenant and servicing exception
Climate or transition scenario	Map macro and sector pathways into exposures	Family operating business and real assets	Borrower cash flow, collateral and sector migration

4.3 Dependency, regime and management action

Three modelling choices dominate many risk outcomes. First, dependency can rise during stress. A family office may experience falling market assets, lower private distributions and higher operating-business liquidity needs at the same time. A lender may face correlated borrower stress, lower recoveries and slower exits. Second, regimes can change the parameters themselves: refinancing availability, rates, valuation multiples and covenant behaviour can move to a different state. Third, management actions can reduce or amplify loss, and their timing and feasibility should be tested.

The scenario registry should record every management action separately. A planned asset sale needs a buyer, time-to-cash, price haircut and decision authority. A facility draw needs available capacity, covenant compliance and lender behaviour. A capital call needs a contractual right and investor ability to fund. A borrower amendment needs consent, economics, reporting and a revised recovery path. An action that has not been operationally rehearsed should not receive full credit in a severe scenario.

5. THE THREE-GATE VALIDATION STANDARD

5.1 Gate one: purpose-specific utility

Utility is defined by the task. A synthetic population for interface testing needs valid schemas, boundary values and expected exceptions. A population for model development needs performance on untouched real data. A population for portfolio analysis needs exposure, dependency and cash-flow properties relevant to that analysis. Similarity charts are diagnostic evidence and do not establish decision utility by themselves.

NIST's SDNist tool evaluates utility and privacy using a range of benchmark metrics [4]. NIST's utility guidance states that there is no universal utility measure and recommends a suite of metrics [7]. This paper applies a layered test that starts with structure and ends with the actual decision task.

Utility layer	Example test	Acceptance principle
Schema	Types, ranges, nulls, keys, referential integrity	All hard rules pass
Marginal	Quantiles, category shares, zero mass, missingness	Tolerance set before generation
Pairwise	Correlations, conditional rates and cross-tabs	Material relationships preserved within tolerance
Multivariate	Classifier two-sample test, propensity separation	Differences investigated, not hidden by an aggregate score
Temporal	Autocorrelation, transition matrix, duration and event ordering	Relevant dynamics preserved
Rare state	Breach, default, recovery, fraud or concentration representation	Coverage and base-rate treatment documented
Downstream task	Train-synthetic/test-real or query agreement	Held-out real-data criterion passes
Decision reproduction	Ranking, limit, liquidity or breach conclusion	Decision result and uncertainty are acceptable to owner

Train-synthetic/test-real testing is particularly valuable for model development because it evaluates the intended use on untouched real data. Query agreement is appropriate where analysts will run defined aggregations. Ranking agreement may be required where the output prioritises borrowers or exposures. A generator should fail utility validation if it produces an attractive global score while materially changing the decision cohort.

5.2 Gate two: privacy and disclosure risk

Synthetic data can reveal information through exact or near-exact records, attribute inference, membership inference, model memorisation or linkage with external data. Stadler, Oprisanu and Troncoso found that the evaluated synthetic-data approaches either failed to prevent inference attacks or failed to retain utility, and that the

trade-off was difficult to predict [24]. TAPAS provides a framework for defining attacker knowledge and goals and running relevant attacks [25].

Differential privacy provides a formal method for bounding the influence of an individual record. NIST SP 800-226 describes how to evaluate differential-privacy guarantees and warns that implementation is difficult and can be easy to get wrong [5]. NIST recommends validated libraries and appropriate expertise [5]. Differential privacy also has a budget: repeated releases or analyses consume privacy loss, and a nominal parameter has meaning only with its adjacency definition, accounting method, implementation and threat model.

Privacy test	Question	Evidence retained
Exact-match and near-neighbour	Did the release reproduce a source record or distinctive combination?	Distance method, thresholds and flagged records
Membership inference	Can an attacker infer whether a person or entity was in the training data?	Threat model, attack advantage and vulnerable cohorts
Attribute inference	Can known fields reveal a sensitive unknown field?	Attacker knowledge and measured success
Rare-record exposure	Are unique or under-represented records disproportionately vulnerable?	Cohort results and suppression or redesign decision
Linkage	Can external data reconnect a record to a person, borrower or vehicle?	Linkage sources, match rate and residual risk
Differential-privacy review	Is the claimed guarantee correctly specified and implemented?	Epsilon, delta, adjacency, accountant, library and expert review
Release composition	What cumulative privacy loss arises across versions and recipients?	Release ledger and budget accounting
Human and contractual review	Is the intended sharing lawful and permitted?	Legal basis, agreement, recipient, purpose and retention terms

The EDPB states that whether an AI model is anonymous requires a case-by-case assessment and should involve a very low likelihood of identifying people or extracting personal data through queries [9]. The EU AI Act requires data governance and management practices for training, validation and test datasets used by high-risk AI systems [10]. The UAE federal data-protection portal describes requirements covering processing, consent, security, rights and cross-border transfers under Federal Decree-Law No. 45 of 2021 [11]. Applicability requires legal assessment. A synthetic release should remain inside the same approval process as other sensitive data until privacy and legal owners approve the exact purpose and recipient.

5.3 Gate three: tail and stress adequacy

Tail validation tests whether the synthetic or simulated population preserves the events and dependencies that drive adverse outcomes. It is distinct from utility because average prediction or query accuracy may remain high while tail behaviour fails. It is distinct from privacy because a differentially private generator can still omit the adverse states required for risk management.

Tail test	Family-office example	Private-credit example
Extreme quantiles	Liquid-asset drawdown and call burden	EBITDA decline and recovery delay
Joint exceedance	Public drawdown plus distribution drought	Sector defaults plus lower collateral values
Tail dependence	Currency, property and operating-business stress	Borrower, sponsor and refinancing stress
Transition coverage	From normal liquidity to reserve breach	From covenant pressure to amendment, default and workout
Duration	Time below liquidity threshold	Time in breach, default and recovery

Tail test	Family-office example	Private-credit example
Concentration	Largest family, sponsor, bank or geography	Largest borrower, sponsor, industry or arranger
Reverse threshold	Scenario that exhausts liquid reserve	Scenario that breaches facility, NAV or redemption limit
Alternative model	Different copula, regime or parameter set	Different default correlation or recovery process

Basel's stress principles call for documented objectives, governance, policies, methodology, resources and challenge [16]. EBA guidance provides a common stress-testing framework [17]. NGFS asks users to apply its scenarios through their own risk frameworks and states that some tail risks and tipping points may be absent [18]. The tail gate therefore requires an institution-specific overlay, alternative models and a record of what remains outside the scenario set.

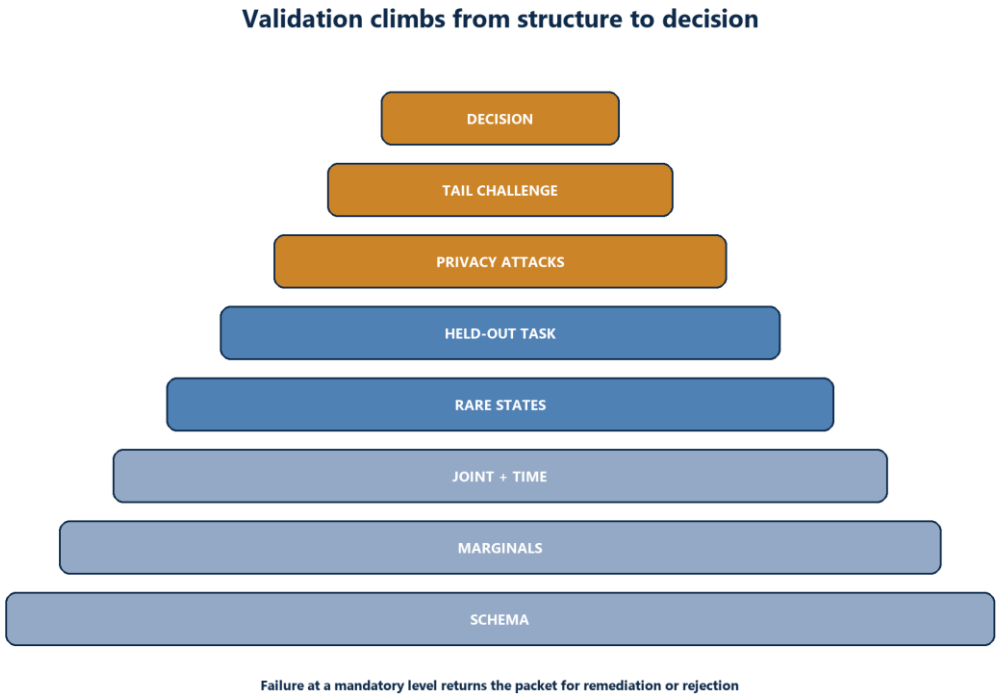

Figure 4. Validation ladder from schema to decision and tail challenge

Source: Matchpoint validation framework.

5.4 Acceptance decision

Each use case receives one of four decisions: approved for the named purpose; approved with restrictions; return for remediation; rejected. Approval attaches to a version, purpose, population, recipient class and expiry date. A population approved for development may remain prohibited for calibration or external sharing. Any change to source population, generator, privacy mechanism, scenario logic, material feature or recipient triggers a change assessment.

6. CONTROLLED ARCHITECTURE AND WORKFLOW

6.1 Architecture

The controlled architecture keeps real data inside an approved boundary and moves only approved artefacts between stages. The source enclave contains holdings, loan, borrower, transaction, document and outcome data with permissions and retention controls. A curation layer resolves identifiers, versions, lineage and quality exceptions.

Generation and simulation occur in a controlled environment. Validation is independent of the developer where materiality requires it. The release registry records the approved purpose, dataset, recipient and privacy budget. Decision systems consume only approved versions, while human authorities own capital, credit, disclosure and release decisions.

Layer	Function	Minimum control	Evidence object
Source enclave	Hold real and sensitive data	Least privilege, encryption, retention and access logging	Source manifest and access log
Curation	Resolve entities, definitions, quality and versions	Data owner, rules and exception queue	Curated-version record
Generator or simulator	Produce artificial records or paths	Versioned code, parameters, seeds and environment	Generation manifest
Privacy evaluation	Test disclosure threats and formal claims	Independent threat model and attack suite	Privacy report
Utility evaluation	Test purpose-specific performance	Predefined holdout and tolerances	Utility report
Tail challenge	Test severe dependencies and reverse thresholds	Alternative models and risk-owner challenge	Tail report
Release registry	Control purpose, recipient and expiry	Approval, fingerprint, budget and revocation	Release certificate
Decision application	Use approved output for named task	Version pinning and human review	Decision packet
Monitoring	Detect drift, misuse, leakage and invalidation	Alerts, periodic review and incident response	Monitoring and incident log

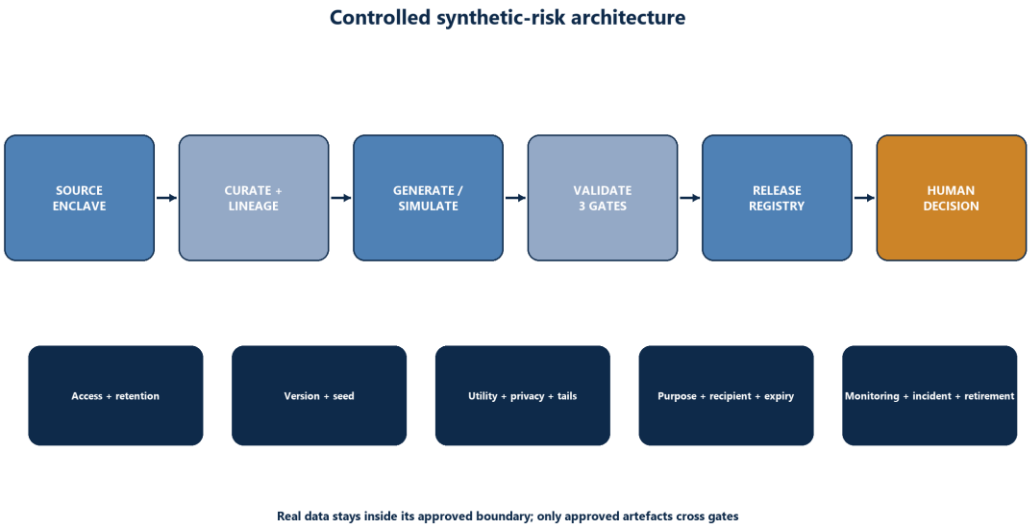

Figure 5. Controlled architecture for synthetic risk data and simulation
Source: Matchpoint architecture; governance references [5,8-18].

6.2 Before-and-after workflow

The common manual pattern moves from a sparse workbook to a hand-built scenario and then directly into a committee deck. Assumptions, versions and exceptions may be dispersed across email and files. The governed workflow creates a stable chain: purpose registration; approved source manifest; curated dataset; versioned

generator or simulator; three-gate validation; independent challenge; approved release; decision use; monitoring and expiry.

The productivity opportunity comes from reusable evidence. Once definitions, entity mappings, validation tests and scenario cards are versioned, a team can rerun them with fewer manual reconciliations. Productivity remains quality-adjusted. A faster scenario that hides an assumption or omits an exposure is a control failure.

Workflow stage	Common manual state	Governed state	Release condition
Request	Broad instruction to create a dataset or stress	Named decision, owner, purpose and prohibited uses	Complete use-case card
Data	Spreadsheet extracts and local copies	Approved manifest inside source enclave	Data-owner approval
Build	Analyst code and undocumented assumptions	Versioned pipeline, parameters and seeds	Reproducible generation
Validate	Visual similarity and spot checks	Utility, privacy and tail gates	All mandatory criteria pass
Challenge	Informal review	Independent validation and risk-owner challenge	Exceptions resolved or accepted
Release	File copied to user	Fingerprinted, purpose-bound and time-limited release	Named release authority
Use	Output enters model or deck	Version-pinned decision packet with caveats	Human decision and sign-off
Monitor	Rebuild when problem appears	Drift, incident, expiry and privacy-budget monitoring	Continued-use review

6.3 Tool stack as control categories

The framework does not prescribe a vendor. A team needs secure data storage, data-quality and lineage controls, a reproducible modelling environment, a generator or simulator appropriate to the data, privacy-test tooling, statistical and downstream-task validation, model registry, workflow approvals, monitoring and immutable logging. Vendor components remain subject to security, data-use, subprocessor, retention, portability and exit review.

NIST recommends using validated differential-privacy libraries because implementations can fail in subtle ways [5]. The PRA and Federal Reserve guidance both retain oversight of third-party or vendor models within their stated scopes [14,15]. A contract or vendor score should therefore enter the evidence pack and cannot substitute for institutional validation.

7. ICP-SPECIFIC SCENARIO LIBRARIES

7.1 A2 family-office library

The family-office library starts with liquidity and concentration because these risks span asset classes and can force decisions. Each scenario card records the balance-sheet date, reporting coverage, look-through gaps, currency basis, commitment rules, valuation convention, liquidity tier and management actions. Synthetic records can expand edge cases for operational testing; capital decisions use the approved real exposure base with simulated paths.

Scenario	Core shocks	Decision output	Important limitation
Distribution drought	Private-fund distributions fall; calls continue	Minimum liquidity and funding sequence	Historical pacing may not represent a new regime
Family concentration event	Operating business, property and public sector exposure decline together	Liquidity, collateral and strategic-asset trade-offs	Cross-holding and guarantee data may be incomplete

Scenario	Core shocks	Decision output	Important limitation
Currency and rate shock	Base rates rise; key currency weakens; hedges reprice	Cash-flow-at-risk and hedge capacity	Hedge liquidity and counterparty behaviour require separate stress
GP-led extension wave	Exits delay; NAVs adjust slowly; continuation requests rise	Commitment pacing and governance workload	Valuation lag may conceal economic loss
Bank or custodian interruption	Access, settlement or facility availability is impaired	Operational liquidity and fallback capability	Event duration and legal access differ by account
Succession and permission event	Authority or data access changes suddenly	Control continuity and decision rights	Legal succession facts need qualified advice

7.2 A3 private-credit library

The private-credit library follows the borrower from early deterioration through restructuring and recovery. It includes deal terms, covenant definitions, amendment history, sponsor, sector, currency, geography, collateral, facility funding, investor liquidity and valuation. Synthetic sequence augmentation may help develop surveillance tools; calibration and risk limits remain anchored to real and externally validated evidence.

Scenario	Core shocks	Decision output	Important limitation
Rates remain high	Interest expense and cash conversion weaken	Coverage, liquidity and amendment need	Borrower hedging and sponsor support vary
Sector revenue shock	Revenue falls; working capital absorbs cash	Breach timing, liquidity need and downside case	Revenue elasticity requires borrower evidence
Refinancing closure	Exit debt unavailable or materially more expensive	Maturity wall and extension requirement	Lender behaviour is scenario-dependent
Collateral and recovery stress	Collateral value falls; workout duration rises	Recovery distribution and liquidity path	Legal priority and jurisdiction are decisive
Correlated sponsor stress	Several borrowers rely on the same sponsor or capital source	Concentration and support capacity	Sponsor exposure may be incomplete
Investor-liquidity pressure	Redemptions or distributions rise while asset exits slow	Fund liquidity and facility capacity	Vehicle terms and gates must be modelled exactly
Operational surveillance failure	Data arrives late or mapping suppresses an alert	Detection delay and control remediation	Synthetic alerts test the system, not true borrower incidence

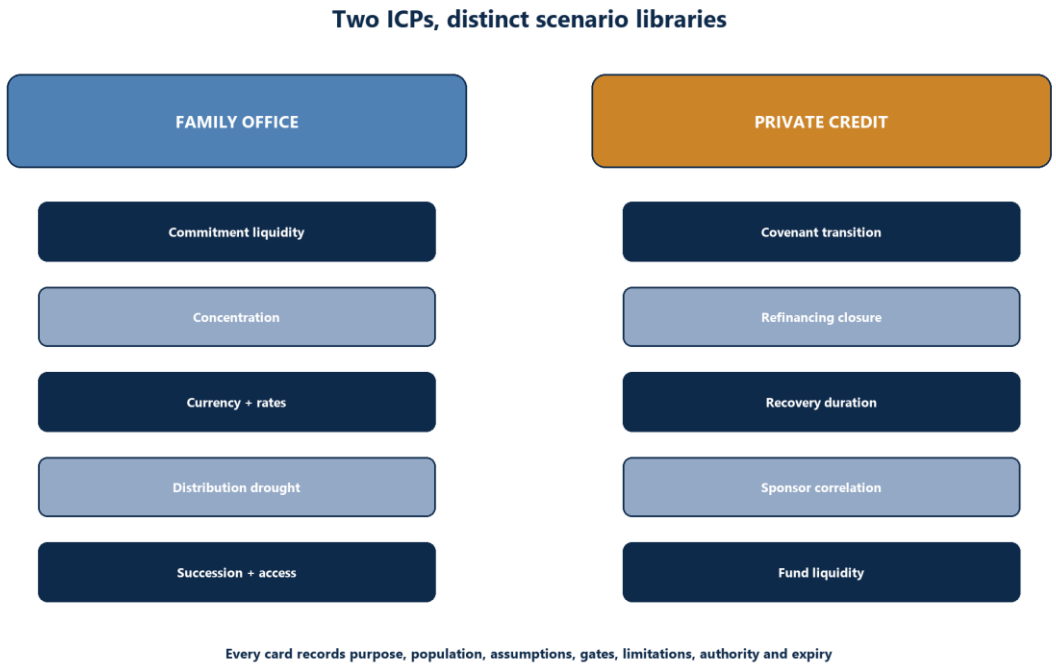

Figure 6. Scenario library for family offices and private-credit funds
Source: Matchpoint scenario framework; risk context [18-20].

7.3 Scenario-card standard

Every scenario card contains: identifier and version; owner and challenge owner; decision question; positions and population; horizon and step; starting date; risk factors; dependencies; parameter evidence labels; synthetic or simulated components; management actions; utility criteria; privacy criteria; tail criteria; outputs; limitations; approval; expiry; and linked decisions. The standard permits comparison across reruns and makes a committee aware when a changed result comes from changed data, model, assumption or action.

8. MEASUREMENT AND ILLUSTRATIVE ECONOMICS

8.1 Productivity measurement

The productivity unit is an **accepted risk packet**. A packet can be a validated synthetic dataset, a scenario run, a stress report, a control-test population or a committee-ready risk exhibit. Acceptance requires the registered purpose, source manifest, reproducible build, three-gate validation, resolved exceptions and named approval. Rows generated, scenarios run and model speed are operating metrics; they are not the value unit.

Metric	Definition	Decision use
Purpose-complete rate	Packets with complete use-case card / attempted packets	Intake quality
Reproducibility rate	Builds reproduced from manifest / builds tested	Engineering control
Utility-pass rate	Packets passing all mandatory utility criteria / tested packets	Fitness for purpose
Privacy-pass rate	Packets within approved disclosure thresholds / tested packets	Release control
Tail-pass rate	Packets passing severe-state and alternative-model criteria / tested packets	Risk adequacy
First-pass acceptance	Packets accepted without material correction / reviewed packets	Workflow quality

Metric	Definition	Decision use
Cycle time	Elapsed time from approved request to accepted packet	Speed
Review burden	Reviewer and remediation hours / accepted packet	Hidden labour
Reuse rate	New decisions using an approved reusable component / eligible decisions	Platform value
Decision lead time	Time from data cut to committee-ready risk packet	Organisational responsiveness
Incidents	Unapproved access, release, use or material defect	Control outcome

The primary productivity measure is:

Quality-adjusted accepted risk packets per paid hour = accepted packets x quality weight / (builder time + validator time + reviewer time + remediation time).

The comparison uses the approved baseline workflow for matched packet classes. All failed builds, rejected releases, privacy failures, tail failures and manual fallbacks remain in the denominator. A team should report distributions and packet classes because complex stress work and simple control-test data have different effort profiles.

8.2 Illustrative operating model

The table below is an **unverified illustrative management scenario**. It is a design example and is not a forecast, budget, proposal, achieved result or recommendation. Values require replacement with approved observed data before management use.

Input	Illustrative value	Status
Professionals contributing to risk-data and scenario work	8	Unverified illustrative management assumption
Paid hours per professional per year	1,760	Unverified illustrative management assumption
Total paid hours	14,080	Arithmetic from illustrative assumptions
Blended fully loaded value per hour	AED 625	Unverified illustrative management assumption
Annual recurring platform and control cost	AED 450,000	Unverified illustrative management assumption
Baseline accepted packets	160	Unverified illustrative management assumption
Baseline hours per accepted packet	52	Unverified illustrative management assumption
Baseline first-pass acceptance	62%	Unverified illustrative management assumption
Attributed revenue	USD 0	No approved observed evidence
Attributed cost reduction	USD 0	No approved observed evidence
Quantified loss reduction	USD 0	No approved observed evidence

Three illustrative productivity cases show the measurement logic. Case L raises first-pass acceptance modestly and reduces hours per packet to 48. Case M uses 43 hours and 75% first-pass acceptance. Case H uses 38 hours and 82% first-pass acceptance. These are unverified scenarios. Finance has not approved a capacity value, cost reduction,

revenue effect or avoided-loss value. The programme therefore reports operational changes separately and retains economic attribution at zero.

Scenario	Hours per accepted packet	First-pass acceptance	Annual accepted packets at 8,320 packet hours	Gross hours released versus baseline output	Approved economic attribution
Baseline	52	62%	160	0	USD 0
Low case	48	68%	173	640	USD 0
Medium case	43	75%	193	1,440	USD 0
High case	38	82%	219	2,240	USD 0

The packet-hour pool of 8,320 is an unverified illustrative allocation derived from 160 baseline packets multiplied by 52 hours. Gross hours released compares the hours needed to produce the baseline 160 packets under each illustrative case. It does not establish labour removal, cost reduction or cash benefit. Realised capacity requires evidence that the organisation redeployed, avoided or removed the time. Economic attribution requires finance approval, a period, a counterfactual and an allocation method.

8.3 Value bridge

The value bridge has five gates: measured task effect; accepted capacity release; observed redeployment; approved cost or revenue connection; realised cash outcome. The first three can support operational management. The final two require finance evidence. Risk reduction remains especially difficult to monetise because a prevented event is unobserved and controls can overlap. The dashboard can report defects, threshold breaches, incidents and decision lead time without claiming avoided loss.

9. GOVERNANCE AND CONTROLS

9.1 Ownership

Model governance begins with inventory and tiering. The inventory covers generators, simulators, privacy mechanisms, transformations, scenario models, vendor tools and material spreadsheets. Tiering reflects decision materiality, data sensitivity, external release, complexity, replaceability and detectability of failure. A high-tier generator used for external data sharing receives a different validation and approval path from rules-based data used in a closed test system.

Role	Core accountability
Board or delegated risk committee	Risk appetite, material use approval and challenge
CIO or Chief Risk Officer	Portfolio-risk purpose, limits and decision integration
Credit or investment owner	Use-case definition, real-world relevance and decision ownership
Data owner	Source authority, quality, permissions, retention and lineage
Model developer	Reproducible build, documentation, testing and remediation
Independent validator	Utility, privacy, tail, conceptual and implementation challenge
Privacy and legal owners	Data-protection, sharing, contract and jurisdiction assessment
Information security	Enclave, identity, access, logging, incident and vendor controls
Release authority	Purpose, recipient, expiry and release approval
Internal audit	Independent assessment of design and operating effectiveness

The developer cannot approve a material external release alone. The business owner cannot waive a failed privacy gate alone. The risk committee cannot interpret a model result without the disclosed data and scenario limitations. Named accountability creates an escalation route when evidence conflicts.

9.2 Minimum controls

Control	Failure addressed	Evidence
Use-case registration	Unbounded or repurposed use	Approved purpose and prohibited uses
Data manifest	Unknown population or unlawful source	Source, fields, dates, owner and permissions
Environment isolation	Real-data leakage	Architecture, access and network logs
Reproducible build	Untraceable change	Code, dependencies, parameters, seeds and fingerprint
Independent validation	Developer confirmation bias	Validation report and issue log
Privacy threat model	Assumed anonymity	Attacker, knowledge, goal and test suite
Tail challenge	Realistic averages with missing stress	Alternative models, reverse stress and severe states
Release certificate	Dataset copied beyond purpose	Recipient, purpose, version, expiry and revocation
Output marking	Synthetic record mistaken for fact	Machine and human-readable provenance label
Monitoring	Drift or invalid continued use	Thresholds, alerts and review record
Incident response	Delayed containment	Owner, playbook, notification and corrective action
Retirement	Stale versions remain in decisions	Withdrawal, archive and consumer notification

9.3 Validation frequency and change control

Validation occurs before first use, after a material change and periodically according to tier. A material change can include new source population, new jurisdiction, new purpose, generator version, changed privacy parameter, altered dependency, new recipient, changed decision threshold or significant performance drift. Emergency use remains documented, time-limited and subject to retrospective review.

PRA SS1/23 and Federal Reserve SR 26-2 both emphasise risk-based governance, validation, monitoring and controls within their respective scopes [14,15]. Basel and EBA stress guidance support regular challenge of methodology, scenarios, assumptions and results [16,17]. This paper applies those disciplines through a single evidence packet so that model, data, privacy and stress decisions can be challenged together.

10. ADOPTION ROADMAP

10.1 Six gated stages

The proposed roadmap begins with bounded operational uses and advances only when evidence passes. The durations below are proposed planning ranges and remain subject to team size, data state, governance and scope.

Stage	Proposed period	Deliverable	Exit criterion
0. Charter	Weeks 0-2	Owner, inventory, purpose, legal and data boundary	Governance approves one bounded use case
1. Rules and baseline	Weeks 2-6	Definitions, real-data baseline, rules-based test population	Data quality and baseline accepted
2. Controlled prototype	Weeks 6-10	Versioned generator or simulator inside enclave	Reproducible build and no critical control failure

Stage	Proposed period	Deliverable	Exit criterion
3. Three-gate validation	Weeks 10-14	Utility, privacy and tail reports	Mandatory thresholds pass
4. Shadow decisions	Weeks 14-18	Output compared with current process; no decision reliance	Independent validator and owner accept evidence
5. Restricted production	Weeks 18-26	Approved use, release registry and monitoring	Stable operating evidence and review approval
6. Scale or retire	After week 26	Expanded use or controlled withdrawal	New purpose passes a fresh assessment

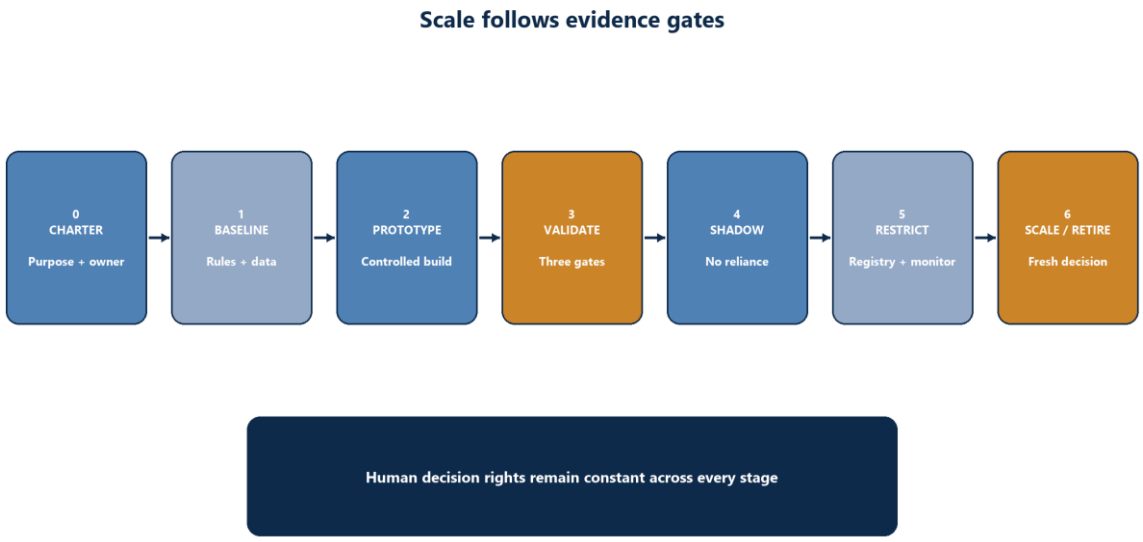

Figure 7. Gated adoption roadmap and decision rights
Source: Matchpoint adoption framework.

10.2 Ninety-day starting programme

During days 1-30, the team inventories the current risk-data and scenario process, selects one decision, records the source boundary and establishes the baseline. Suitable starting cases include synthetic operational-test data for a covenant pipeline, a family-office commitment-liquidity simulator, or a private-credit surveillance shadow model. External sharing and capital decisions remain outside the first prototype unless separately approved.

During days 31-60, the team builds a reproducible pipeline inside the approved environment. It implements schema and lineage controls, creates a holdout, defines the privacy threat model, selects tail scenarios and records all parameter evidence labels. The team deliberately tests failure states: rare records, missing data, new categories, regime changes, access violations and invalid combinations.

During days 61-90, independent validation runs the three gates. The output enters a shadow decision alongside the approved current process. The team measures accepted packets, cycle time, review burden, defects and incidents. Management receives operational evidence and a list of unresolved limitations. Revenue, cost reduction and avoided loss remain zero until observed and approved evidence exists.

11. LIMITATIONS AND FURTHER RESEARCH

This paper is a control and operating framework. It does not test a named family-office portfolio, fund, loan book, generator, privacy mechanism or vendor. It does not estimate market returns, default probability, recovery, liquidity loss, operating cost or regulatory capital. Its illustrative economics are unverified management assumptions.

The cited empirical generator studies use particular datasets, models and metrics. Their results cannot establish performance for another population. Differential-privacy guarantees depend on formal definitions, parameters, accounting and implementation. Adversarial privacy tests provide empirical evidence under specified threat models and cannot enumerate every attacker. Tail scenarios remain incomplete by construction, especially where structural breaks, feedback, legal process, geopolitical events or management behaviour are weakly observed.

Private-market source data can contain valuation lag, survivor selection, changing documentation, missing look-through and inconsistent outcomes. Synthetic data can reproduce these defects. Scenario precision can exceed parameter evidence and create false confidence. A committee should receive sensitivity ranges, alternative models and explicit unknowns.

Further research should compare generator families on private-market decision tasks using secure multi-institution evaluation; develop benchmark rare-state and transition metrics for covenant and liquidity data; test privacy risk for entity and relationship graphs; evaluate how differential privacy affects tail estimates; and study how scenario libraries change actual committee decisions. Any multi-institution study would require legal, confidentiality, competition, data-protection and governance approval.

12. CONCLUSION

Synthetic data and simulation can make risk work more testable, reusable and comprehensive. The value comes from a controlled chain that begins with an actual decision and ends with an accepted evidence packet. The method does not receive authority from realism, volume or technical sophistication.

For family offices, the highest-value initial questions concern liquidity, commitment pacing, concentration, look-through and operational continuity. For private-credit funds, they concern covenant transitions, correlated borrower stress, recovery timing, refinancing, fund liquidity and surveillance controls. Rules-based synthetic data is well suited to system and edge-case testing. Statistical and generative methods can support research and development where held-out utility, privacy and tail tests pass. Simulation and stress testing remain assumption-dependent decision aids.

The governing standard is three-dimensional: fidelity to the properties required for the task; privacy against a stated threat model; and adequacy for adverse paths and dependencies. Approval is versioned, purpose-bound, recipient-bound and time-limited. Human professionals retain data, model, risk, credit, investment, legal and release authority.

The economic discipline is equally clear. Rows generated and scenarios run are operating counts. Accepted risk packets per paid hour measure workflow productivity. Realised capacity, cost reduction, revenue and avoided loss require separate evidence and approval. At launch, **Attributed revenue: USD 0. Attributed cost reduction: USD 0. Quantified loss reduction: USD 0.**

DECLARATIONS

Author: Chennakeshav (CK) Adya. Author status: unverified pending named-author approval.
Funding: No external funding has been identified for this paper.
Conflicts of interest: The author is associated with Matchpoint Partners. This paper describes a proposed advisory and operating framework relevant to Matchpoint's AI and technology advisory practice.
Data availability: No client, borrower or portfolio dataset was used. All quantitative operating scenarios are unverified illustrative management assumptions.
Use of artificial intelligence: AI-assisted research and drafting were used under human review. Primary and authoritative sources were checked individually. The named author remains responsible for final approval.
Disclaimer: Research paper only. This document is not investment, legal, regulatory, tax, accounting, valuation, cybersecurity, data-protection, credit, model-risk or technology-procurement advice.

REFERENCES

[1] Financial Conduct Authority. (2024). Report: Using Synthetic Data in Financial Services. https://www.fca.org.uk/publications/corporate-documents/report-using-synthetic-data-financial-services
[2] Financial Conduct Authority. (2023). FS23/1: Feedback Statement on Synthetic Data Call for Input. https://www.fca.org.uk/publications/feedback-statements/fs23-1-feedback-statement-synthetic-data-call-for-input
[3] Financial Conduct Authority. (2026). Research Note: Synthetic Data and Anti-Money Laundering Project Report. https://www.fca.org.uk/publications/research-notes/research-note-synthetic-data-anti-money-laundering-project-report
[4] National Institute of Standards and Technology. (2025). SDNist Synthetic Data Report Tool v1.4. https://www.nist.gov/services-resources/software/sdnist-synthetic-data-report-tool
[5] National Institute of Standards and Technology. (2025). NIST SP 800-226: Guidelines for Evaluating Differential Privacy Guarantees. https://doi.org/10.6028/NIST.SP.800-226
[6] National Institute of Standards and Technology. (2021). Differentially Private Synthetic Data. https://www.nist.gov/blogs/cybersecurity-insights/differentially-private-synthetic-data
[7] National Institute of Standards and Technology. (2021). Utility Metrics for Differential Privacy: No One-Size-Fits-All. https://www.nist.gov/blogs/cybersecurity-insights/utility-metrics-differential-privacy-no-one-size-fits-all
[8] Information Commissioner's Office. (2023). Privacy-enhancing technologies. Guidance marked under review following the Data (Use and Access) Act. https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/data-sharing/privacy-enhancing-technologies
[9] European Data Protection Board. (2024). Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models. https://www.edpb.europa.eu/documents/opinion-of-the-board-art-64/opinion-282024-on-certain-data-protection-aspects-related-to_en
[10] European Union. (2024). Regulation (EU) 2024/1689, Article 10: Data and data governance. https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng
[11] UAE Government. (2022). Data protection laws. Federal Decree-Law No. 45 of 2021. https://u.ae/en/about-the-uae/digital-uae/data/data-protection-laws
[12] Bank for International Settlements. (2023). Project Aurora: the power of data, technology and collaboration to combat money laundering across institutions and borders. https://www.bis.org/publ/othp66.pdf
[13] Bank for International Settlements. (2024). Artificial intelligence and the economy: implications for central banks. Annual Economic Report 2024, Chapter III. https://www.bis.org/publ/arpdf/ar2024e3.htm
[14] Prudential Regulation Authority. (2026). SS1/23: Model risk management principles for banks. Current version April 2026. https://www.bankofengland.co.uk/prudential-regulation/publication/2023/may/model-risk-management-principles-for-banks-ss
[15] Board of Governors of the Federal Reserve System. (2026). SR 26-2: Supervisory Guidance on Model Risk Management. https://www.federalreserve.gov/supervisionreg/srletters/SR2602.htm
[16] Basel Committee on Banking Supervision. (2026). RMA30: Stress testing principles. https://www.bis.org/basel_consolidated_guidelines/chapter/RMA/30.htm
[17] European Banking Authority. (2018). EBA/GL/2018/04: Guidelines on institutions' stress testing. https://eba.europa.eu/activities/single-rulebook/regulatory-activities/supervisory-review-and-evaluation-process-srep-and-pillar-2/guidelines-stress-testing
[18] Network for Greening the Financial System. (2026). NGFS Climate Scenarios for central banks and supervisors, Phase V. https://www.ngfs.net/en/publications-and-statistics/publications/ngfs-climate-scenarios-central-banks-and-supervisors-phase-v
[19] International Monetary Fund. (2024). Global Financial Stability Report, April 2024, Chapter 2: The Rise and Risks of Private Credit. https://www.elibrary.imf.org/display/book/9798400257704/CH002.xml
[20] Financial Stability Board. (2026). Report on Vulnerabilities in Private Credit. https://www.fsb.org/2026/05/report-on-vulnerabilities-in-private-credit/
[21] Yoon, J., Jarrett, D., and van der Schaar, M. (2019). Time-series Generative Adversarial Networks. Advances in Neural Information Processing Systems 32. https://proceedings.neurips.cc/paper/2019/hash/c9efe5f26cd17ba6216bbe2a7d26d490-Abstract.html
[22] Xu, L., Skoularidou, M., Cuesta-Infante, A., and Veeramachaneni, K. (2019). Modeling Tabular Data using Conditional GAN. Advances in Neural Information Processing Systems 32, 7333-7343. https://arxiv.org/abs/1907.00503
[23] Jordon, J., Yoon, J., and van der Schaar, M. (2019). PATE-GAN: Generating Synthetic Data with Differential Privacy Guarantees. International Conference on Learning Representations. https://openreview.net/forum?id=S1zk9iRqF7
[24] Stadler, T., Oprisanu, B., and Troncoso, C. (2022). Synthetic Data: Anonymisation Groundhog Day. 31st USENIX Security Symposium, 1451-1468. https://www.usenix.org/conference/usenixsecurity22/presentation/stadler
[25] Houssiau, F., Jordon, J., Cohen, S. N., Daniel, O., Elliott, A., Geddes, J., Mole, C., Rangel-Smith, C., and Szpruch, L. (2022). TAPAS: a Toolbox for Adversarial Privacy Auditing of Synthetic Data. NeurIPS 2022 Workshop on Synthetic Data for Empowering ML Research. https://arxiv.org/abs/2211.06550
[26] Patki, N., Wedge, R., and Veeramachaneni, K. (2016). The Synthetic Data Vault. IEEE International Conference on Data Science and Advanced Analytics. https://dai.lids.mit.edu/wp-content/uploads/2018/03/SDV.pdf

APPENDIX A. USE-CASE AND RELEASE CHARTER

Field	Required entry
Use-case identifier	Stable identifier and version
Decision	Exact decision, test or research activity supported
Owner	Named business or risk owner
Source population	Dataset, date, entity, fields and lineage
Purpose	Approved use and recipient class
Prohibited uses	Decisions, recipients and environments excluded
Method	Rules, statistical generator, generative model or simulation
Parameters	Values, evidence labels and approval
Utility gate	Metrics, holdout, thresholds and results
Privacy gate	Threat model, attacks, formal guarantee and results
Tail gate	Severe states, dependencies, alternatives and results
Independent validation	Validator, date, issues and decision
Release	Fingerprint, recipient, environment, expiry and revocation
Monitoring	Drift, incident, review and retirement triggers

APPENDIX B. VALIDATION ACCEPTANCE CHECKLIST

1. The decision, purpose, owner and prohibited uses are explicit.
2. The real source population remains inside the approved environment.
3. Entity, date, unit, currency and missing-data definitions are versioned.
4. Training, validation and final real-data holdout populations are separated.
5. The generator or simulator is reproducible from code, parameters, dependencies and seeds.
6. Schema, marginal, joint, temporal, rare-state and downstream-task tests are complete.
7. The privacy threat model states attacker knowledge, access and goal.
8. Exact-match, near-neighbour, membership, attribute, rare-record and linkage tests are complete where relevant.
9. Any differential-privacy claim records epsilon, delta, adjacency, accountant, composition and expert validation.
10. Extreme quantiles, joint exceedance, duration, transition, concentration and reverse-stress tests are complete.
11. Alternative dependency, regime and parameter structures have been challenged.
12. Management actions have authority, capacity, timing, cost and operational evidence.
13. Independent validation issues are resolved or explicitly accepted by the authorised owner.
14. The release certificate binds version, purpose, recipient, environment and expiry.
15. The output is marked as synthetic or simulated and cannot be mistaken for observed fact.
16. Monitoring, incident, invalidation, revocation and retirement processes are live.
17. Decision materials disclose source gaps, scenario omissions and remaining uncertainty.
18. Revenue, cost reduction and loss reduction remain zero without approved observed evidence.

APPENDIX C. ILLUSTRATIVE SCENARIO CARD

Status: unverified illustrative example.
Identifier: A3-LIQ-001. Decision: assess whether a private-credit vehicle can meet obligations during a refinancing and recovery slowdown. Horizon: 24 months, monthly steps. Starting population: approved real fund positions and contractual obligations inside the source enclave. Synthetic component: rules-based future borrower transitions used for operational and model challenge. Simulation component: default timing, recovery amount, recovery duration, facility availability and investor cash flows. Stress component: correlated sponsor and sector pressure, refinancing closure and higher sale haircuts.
Mandatory outputs: monthly available liquidity; facility headroom; unfunded asset obligations; investor payments; defaults; recoveries; concentration; threshold-breach month; required management action; and unresolved data gaps.
Utility acceptance: exact contractual cash flows reconcile; held-out real-data queries meet predefined tolerances; transition coverage includes all approved states. Privacy acceptance: release remains internal; no source record is reproduced; rare borrowers pass the approved attack suite. Tail acceptance: reverse stress identifies the first liquidity breach; at least two alternative dependency models and a longer recovery-duration case are reported.
Authority: fund risk owner accepts the scenario purpose; data owner approves the population; independent validator signs the three gates; authorised committee retains the liquidity and investment decision.