

# AI for Wealth Management: Hyper-Personalised Advice and Portfolio Construction

*A controlled human-AI operating model for private banks, EAMs, multi-family offices and GCC family offices*

**Prepared by Matchpoint Partners**

Author attribution pending CK approval

August 2026

## ABSTRACT

**Background.** Private-wealth advice depends on fragmented client, entity, portfolio, liability and governance data. Artificial intelligence can assist synthesis and explanation while suitability and investment authority remain accountable professional decisions. **Objective.** This paper develops a controlled AI-assisted wealth-advice and portfolio-construction framework for A6 private banks, EAMs and multi-family offices and A7 GCC family offices and UHNW asset owners. **Approach.** The analysis reviews 40 primary, authoritative and clearly labelled research sources available through 1 August 2026. **Findings.** A governed client truth record, deterministic policy, reproducible portfolio analytics, source-grounded language assistance and meaningful human review form one controlled workflow. **Implications.** Advice quality, review time, overrides, complaints, drift and cost per accepted outcome should be measured before scale. All worked inputs are unverified illustrative management assumptions; attributed Matchpoint or client revenue, cash cost reduction and loss reduction remain USD 0 until approved observed evidence exists.

## Highlights

**Truth record.** Client, entity, portfolio and liability facts retain source, date, status and owner.

**Architecture.** Suitability and portfolio mathematics remain separate from language generation.

**Human decision.** Authorised advisers and investment committees retain judgement, approval and action.

**Evaluation.** Data, constraints, calculations, explanations and human outcomes are tested separately.

**Value.** Accepted outcomes, realised capacity and approved attribution precede commercial claims.

**JEL Classification:** G11, G23, G24, D14, O33

**Keywords:** artificial intelligence, wealth management, hyper-personalisation, private bank, external asset manager, multi-family office, family office, UHNW, suitability, portfolio construction, robo-advice, risk capacity, household balance sheet, human review

*Research paper only. This document is not investment, legal, regulatory, accounting, audit, tax, privacy, cybersecurity, employment, valuation or technology advice. All worked inputs, thresholds, times, rates and economics are unverified illustrative management assumptions.*

## 1. INTRODUCTION

Wealth advice becomes genuinely personal when the portfolio reflects the client's full decision context: objectives, liabilities, liquidity, tax position, currencies, legal entities, family governance, existing holdings, private assets, concentration, preferences and capacity for loss. A short risk questionnaire and a model portfolio capture only part of that context. Artificial intelligence can help organise the information, detect changes, prepare scenarios and explain choices. The investment institution still needs a controlled method that converts information into suitable advice and an authorised portfolio.


This paper develops that method for two Matchpoint Partners ideal customer profiles. A6 covers private banks, external asset managers and multi-family offices that advise or manage assets for many clients. A7 covers GCC family offices and ultra-high-net-worth asset owners that allocate across family entities, operating businesses, liquid markets and private investments. Both profiles face fragmented data, repeated manual preparation, multi-jurisdictional constraints and demanding explanation requirements. Their decision rights and operating economics differ.

The available empirical evidence supports bounded claims. D'Acunto, Prabhala and Rossi found diversification and behavioural effects after adoption of a wealth-management robo-adviser, with outcomes varying by the investor's starting portfolio [26]. Scherer and Lehner analysed 151,200 profile recommendations from a major US robo-adviser and found that goals and horizons influenced equity allocations while more complex hedging demands were largely absent [27]. A 2026 field experiment found that one graphical explanation format did not increase adherence to algorithmic recommendations [28]. A ten-week portfolio experiment found that default implementation produced stronger outcomes than advice alone [29]. These results show that design, starting condition and implementation route matter. They do not establish a universal performance, revenue or productivity effect.

The regulatory evidence supports the same operating boundary. ADGM's digital-investment-management guidance addresses suitability, disclosures, algorithm governance, testing, human resources, cyber controls and outsourcing [1]. ESMA expects firms using AI in investment services to comply with MiFID II organisational and conduct requirements and to act in clients' best interests [8]. SEC and FINRA materials preserve adviser, care, supervision and communication obligations when technology is used [9-12]. IOSCO guidance assigns senior accountability, testing, competence and third-party oversight [13,14].

The paper advances five propositions.

| Proposition                                                                    | Decision implication                                                                                                          |
|--------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| P1. Personalisation begins with a governed client record.                      | Advice cannot exceed the completeness, currency, rights and reconciliation of the underlying facts.                           |
| P2. Suitability, portfolio construction and explanation are separate services. | Deterministic rules and optimisation produce controlled proposals; AI assists intake, synthesis, scenarios and communication. |
| P3. A household or family is a balance sheet, not a questionnaire score.       | Liabilities, currencies, legal entities, private assets, operating-business exposure and governance enter the constraint set. |
| P4. The adviser remains the decision owner.                                    | Human review must be informed, documented and able to change or reject the proposal.                                          |
| P5. Scale is earned through accepted outcomes.                                 | Advice quality, review time, overrides, complaints, drift and realised economics are measured before expansion.               |



**Figure 1. From product-led advice to controlled hyper-personalisation**

Source: Matchpoint wealth-advice operating framework; regulatory and empirical context [1-17,25-37].

All worked cases, thresholds, volumes, times, rates and economics in this paper are explicitly unverified illustrative management assumptions. Attributed Matchpoint or client revenue, cash cost reduction and loss reduction remain USD 0 until approved observed evidence exists.

## 2. SCOPE, DEFINITIONS AND EVIDENCE BOUNDARIES

### 2.1 Scope

The paper covers AI-assisted wealth advice and portfolio construction for advisory and discretionary relationships. It addresses client discovery, data consolidation, suitability, risk capacity, investment-policy design, strategic asset allocation, instrument selection, scenario analysis, tax and currency constraints, portfolio explanation, review preparation, monitoring and change detection.

The method applies to private banks, external asset managers, multi-family offices, single-family offices and UHNW investment organisations. Applicable licensing, fiduciary, conduct, data, outsourcing and cross-border rules depend on the entity, activity, client, product, jurisdiction and delivery channel. Qualified legal, compliance, tax and investment professionals determine applicability and approve implementation.

The paper does not prescribe a security, manager, asset allocation or expected return. It does not provide investment, legal, regulatory, tax, accounting, privacy, cybersecurity or technology advice. It does not treat a language model as a portfolio optimiser, suitability engine, fiduciary or autonomous adviser.

### 2.2 Working definitions

| Term                      | Operational definition used in this paper                                                                                                                                |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hyper-personalised advice | Advice built from a governed client or household record, explicit objectives and constraints, approved methods, current portfolio evidence and recorded human judgement. |
| Client truth record       | Versioned set of client facts, sources, dates, permissions, confidence, conflicts and accountable owners used for advice.                                                |
| Suitability               | The applicable assessment that links client circumstances and objectives to a recommendation or managed action.                                                          |
| Risk tolerance            | Willingness to accept uncertain outcomes, measured with a documented method and reviewed for behavioural instability.                                                    |
| Risk capacity             | Financial ability to absorb loss without impairing required objectives, liabilities or liquidity.                                                                        |
| Investment policy         | Approved objectives, constraints, strategic ranges, governance, benchmarks, liquidity rules, rebalancing and exceptions.                                                 |
| Portfolio proposal        | A non-executed allocation or action prepared for review under the investment policy.                                                                                     |

| Term                    | Operational definition used in this paper                                                                                       |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| Explanation packet      | Source-linked record of facts, constraints, scenarios, trade-offs, proposed action, limitations and approvals.                  |
| Override                | A documented human change to a system proposal, with reason, authority and consequence.                                         |
| Accepted advice outcome | A complete, suitable and approved advice or portfolio packet that passes defined quality and control gates.                     |
| Drift                   | Material change in client facts, portfolio exposure, model behaviour, data or market conditions relative to the approved state. |

## 2.3 Evidence hierarchy

The research gives priority to binding law, official rules, regulator guidance, standards bodies and original empirical studies. A source is used within its stated population and method.

| Evidence class                   | Examples                                                                              | Permitted use                                                                            |
|----------------------------------|---------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------|
| Binding law or regulation        | UAE personal-data law, DIFC data law, EU AI Act, GDPR                                 | Determine obligations after applicability is confirmed.                                  |
| Regulator guidance               | ADGM digital investment management, ESMA statement, SEC and FINRA materials           | Translate supervisory expectations within stated scope.                                  |
| International regulatory report  | IOSCO, IMF, OECD                                                                      | Identify common risks, market evidence and policy direction.                             |
| Consensus standard               | ISO/IEC 42001 and ISO/IEC 23894                                                       | Structure management and risk processes; legal compliance remains a separate conclusion. |
| Voluntary framework              | NIST AI RMF and GenAI Profile                                                         | Define governance, mapping, measurement and management outcomes.                         |
| Peer-reviewed empirical research | RFS, Economic Journal, Journal of Empirical Finance                                   | Establish observed results for the studied design and population.                        |
| Preprint                         | 2026 algorithmic-transparency field study                                             | Inform hypotheses with explicit preprint status and limitations.                         |
| Local production evidence        | Accepted packets, cycle time, overrides, complaints, incidents and realised economics | Support the institution's own approval and scaling decision.                             |

Sources were reviewed through 1 August 2026. NIST states that AI RMF 1.0 is being revised [18]. Regulatory and product status can change. The obligations register should record source date, current status, next review date and accountable owner.

## 2.4 Research method

The method maps four questions: what an A6 or A7 decision owner needs; what empirical evidence establishes about robo-advice and personalisation; what applicable authorities expect; and what production evidence would justify deployment. Reported statistics retain their study population. Scenario values are labelled unverified. Vendor performance and market-size claims are excluded from the core business case unless supported by a primary or authoritative source.

## 3. A6 AND A7 DECISION MAP

### 3.1 A6 private banks, EAMs and multi-family offices

An A6 institution serves a portfolio of households or legal entities. The operating challenge is repeated delivery under a common control framework while preserving client-specific circumstances, restrictions and explanation. The firm must reconcile front-office judgement, investment policy, product governance, custody data, research, compliance, tax input and client communication.

| A6 decision                 | Personalisation requirement                                                                                 | Required evidence                                       |
|-----------------------------|-------------------------------------------------------------------------------------------------------------|---------------------------------------------------------|
| Onboard or refresh a client | Identity, entities, beneficial ownership, objectives, experience, liquidity, horizon, risk and restrictions | Dated source record, unresolved items and approval.     |
| Set investment policy       | Objectives, capacity for loss, currencies, liabilities, strategic ranges and prohibited exposures           | Approved policy, scenario record and client acceptance. |

| A6 decision         | Personalisation requirement                                                                  | Required evidence                                                    |
|---------------------|----------------------------------------------------------------------------------------------|----------------------------------------------------------------------|
| Construct portfolio | Existing holdings, tax lots, product eligibility, fees, concentration, liquidity and mandate | Constraint file, eligible universe, optimisation version and review. |
| Recommend a change  | Trigger, alternatives, costs, conflicts and portfolio-level effect                           | Before-and-after analysis, rationale, disclosure and sign-off.       |
| Review relationship | Fact changes, drift, performance, cash flows, restrictions and client understanding          | Review packet, exceptions, decisions and next actions.               |

Scale creates a control requirement. A reusable platform may assemble data and analysis. Each client still requires an individual decision record. Segment labels, next-best-action scores and marketing propensity cannot replace suitability.

### 3.2 A7 GCC family offices and UHNW asset owners

An A7 family-office portfolio often extends beyond custody accounts. It may include operating-company equity, real estate, private funds, direct investments, shareholder loans, trusts, foundations, holding companies, art, insurance, family liabilities and several currencies. Valuations and liquidity can be irregular. Family governance can restrict transfers or require approvals. A portfolio optimiser that sees only marketable securities can therefore recommend a locally efficient and globally unsuitable allocation.


| A7 decision                | Personalisation requirement                                                                                    | Required evidence                                                      |
|----------------------------|----------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------|
| Define family capital      | Entities, ownership, control, beneficiaries, liabilities and excluded assets                                   | Consolidation map, legal-owner confirmation and valuation date.        |
| Set liquidity architecture | Operating needs, distributions, capital calls, debt service and reserves                                       | Cash-flow ladder, stress periods and minimum buffers.                  |
| Allocate risk              | Operating-business concentration, regional and currency exposure, private-asset pacing and public-market hedge | Look-through exposure, scenario set and investment-committee decision. |
| Assess opportunity         | Strategic fit, liquidity, governance, tax, concentration and opportunity cost                                  | Portfolio contribution, diligence, alternatives and approval.          |
| Communicate across family  | Beneficiary objectives, governance rights, confidentiality and decision cadence                                | Approved audience, explanation packet and recorded decisions.          |

The family principal or investment committee can retain broad discretion. That discretion benefits from a consistent evidence architecture. The system should expose missing information, valuation staleness and competing objectives instead of collapsing them into one risk score.

### 3.3 Shared decision contract

Every AI-assisted workflow should begin with a decision contract.

| Field                  | Required record                                                                                         |
|------------------------|---------------------------------------------------------------------------------------------------------|
| Decision               | Advice, allocation, rebalance, liquidity action or review supported.                                    |
| Accountable owner      | Adviser, CIO, investment committee or authorised delegate.                                              |
| Client or family scope | Legal entities, accounts, beneficiaries and excluded assets.                                            |
| Inputs                 | Authoritative data sources, dates, rights, reconciliation and known gaps.                               |
| Method                 | Rules, optimiser, capital-market assumptions, scenarios and model versions.                             |
| AI authority           | Permitted synthesis, drafting and monitoring; prohibited approval, execution and unsupported invention. |
| Review                 | Required professional checks, alternatives and override route.                                          |
| Output                 | Explanation packet, decision record, client communication and retention.                                |

**A6 and A7 decision perimeter****Figure 2. A6 and A7 decision perimeter**

Source: Matchpoint ICP and decision-rights framework.

## 4. WHAT THE EMPIRICAL AND MARKET EVIDENCE SUPPORTS

### 4.1 Portfolio effects vary by starting condition

D'Acunto, Prabhala and Rossi studied the introduction of a wealth-management robo-adviser that tailored portfolios to holdings and preferences [26]. Adoption improved diversification for initially under-diversified investors. Already diversified investors changed differently. Adopters also showed reduced disposition, trend-chasing and rank effects while trading activity and attention changed. The operational lesson is a conditional baseline: value depends on the client's starting portfolio, behaviour and use of the recommendation.

This finding prevents a universal claim that automation improves every portfolio. A pilot should stratify by initial diversification, asset complexity, advice channel and client type. It should measure portfolio-level effects after fees and implementation, not only recommendation acceptance.

### 4.2 Commercial recommendations remain simplified

Scherer and Lehner web-scraped recommendations for 151,200 investor types from a major US robo-adviser [27]. Goals and time horizons influenced equity allocation. The authors found that the observed recommendations largely ignored Merton-type hedging demands. This is a direct warning for UHNW and family-office use: accessible digital recommendations can simplify household-finance theory and omit important background risks.

A controlled institution can use simplicity intentionally. The investment policy should state which risks are represented, approximated or excluded. Operating-company exposure, human capital, tax, liabilities, illiquid assets and currency obligations should enter either the optimiser or a documented human overlay.

### 4.3 Explanation does not automatically create adherence

Boulu-Reshef, Direr and Louafi ran a randomised controlled trial within a French robo-adviser platform [28]. The study, posted as a 2026 preprint, found no evidence that one profile-based graphical explanation increased adherence or engagement. Device and platform experience appeared more important in some behaviours. The result does not show that explanation is ineffective. It shows that one explanation format cannot be assumed to change behaviour.

The institution should test whether an explanation helps the intended user identify facts, trade-offs, costs and downside. Comprehension, correction and informed override are stronger targets than simple adherence. A client who correctly rejects a recommendation after seeing an omitted constraint may demonstrate a successful explanation control.

### 4.4 Implementation design changes outcomes

The 2026 Economic Journal study used a ten-week portfolio-choice experiment [29]. Robo-advice supported continued participation, rebalancing and portfolios closer to a utility-maximising benchmark. Default implementation performed better than advice alone. This result is design-specific and creates a conduct question. A default can improve implementation while increasing the importance of suitability, disclosure, conflict, monitoring and reversibility.

Back, Morana and Spann used two consequential experiments and found that robo-advice availability reduced the disposition effect [30]. Human-like social design elements reduced advice seeking in their design. The lesson is to test interface features for decision quality. Anthropomorphic language and conversational fluency should not be treated as trust evidence.

#### 4.5 Personalisation can improve diversification within bounded designs

A case-based recommender study evaluated portfolio recommendations with 1,172 real users and reported improved yield in most experimental settings while meeting stated risk profiles [31]. A 2026 natural experiment using account-level data from a major Chinese commercial bank reported a 7.6 per cent reduction in concentration after personalised recommendations, alongside higher investment income and lower overall risk [32]. These studies differ in market, product, population and design. They support evaluation of diversification and risk-capacity alignment. They do not establish the same effect for a GCC private bank or family office.

Human-robo collaboration research identifies heterogeneous user responses and the importance of performance information and framing [33]. The institution should therefore record advice source, explanation version, adviser interaction and override. Aggregated acceptance rates can hide important differences by client sophistication, channel and starting condition.


#### 4.6 Current financial-sector context

The Bank of England and FCA reported in their 2024 survey that 75 per cent of respondent firms were using AI and another 10 per cent planned to use it over three years; one third of use cases were third-party implementations [25]. The figures describe the respondent population. They establish material adoption and dependency exposure in that survey.

The DFSA's 2025 survey covered 661 authorised firms with an 88 per cent participation rate. It reported 52 per cent using AI, 60 per cent with some form of governance structure and 21 per cent lacking clear accountability or oversight even where AI use was critical [4]. These DIFC-specific findings support a governance gate for A6 institutions.

The IMF's 2024 Global Financial Stability Report describes opportunities in risk management and market liquidity alongside opacity, operational dependency, cyber risk and correlated behaviour [15]. OECD work describes the potential for differentiated, personalised recommendations and highlights data, bias, explainability, market-conduct and stability concerns [16,17]. None of these reports establish local portfolio alpha or realised commercial value.

#### Evidence-to-decision ladder for wealth AI



**Figure 3. Evidence-to-decision ladder for wealth AI**

Source: Matchpoint evidence framework; empirical context [26-33,37].

### 5. THE PERSONALISATION DECISION MODEL

#### 5.1 Separate facts, policy, mathematics and language

The operating model uses four layers.

| Layer       | Function                                                                           | Control rule                                                           |
|-------------|------------------------------------------------------------------------------------|------------------------------------------------------------------------|
| Facts       | Client, household, entity, portfolio, market and product data                      | Preserve source, date, rights, confidence and reconciliation.          |
| Policy      | Eligibility, suitability, strategic ranges, liquidity, tax and mandate constraints | Approved rules execute deterministically and version explicitly.       |
| Mathematics | Optimisation, scenario analysis, attribution and risk calculation                  | Reproducible methods, assumptions, tolerances and independent testing. |
| Language    | Intake summary, questions, scenario explanation and draft communication            | Source-grounded generation, no authority to approve or execute.        |

The separation makes errors diagnosable. A wrong client currency is a fact problem. A prohibited instrument entering the eligible universe is a policy problem. An unstable allocation is a mathematics problem. A fluent but unsupported rationale is a language problem. Each requires a different owner and test.

## 5.2 Personalisation ladder

The institution should state its current level rather than use the term hyper-personalisation as a marketing claim.

| Level             | Description                                                                                | Minimum evidence                                                |
|-------------------|--------------------------------------------------------------------------------------------|-----------------------------------------------------------------|
| 0 Generic         | Same content or portfolio for a broad audience                                             | Product and communication approval.                             |
| 1 Segment         | Rules vary by broad client segment                                                         | Segment definition and eligibility tests.                       |
| 2 Profile         | Goals, horizon and risk questionnaire affect recommendation                                | Validated questionnaire, scoring and mapping.                   |
| 3 Portfolio-aware | Existing holdings, costs, tax lots, concentration and liquidity enter advice               | Reconciled holdings and portfolio-level comparison.             |
| 4 Household-aware | Liabilities, entities, currencies, private assets and external exposures enter constraints | Consolidated truth record and staleness controls.               |
| 5 Adaptive        | Approved change events trigger re-evaluation and human review                              | Monitoring, drift thresholds, review and client-consent design. |

Progression requires evidence. An A6 institution may operate different levels by mandate or client. An A7 family office may reach household awareness for strategic allocation while keeping private-asset valuations as explicit ranges.

## 5.3 Objective and constraint hierarchy

The portfolio problem should be written as a hierarchy before optimisation.

1. Preserve legal, mandate and client prohibitions.
2. Meet near-term cash, debt-service and capital-call requirements.
3. Protect required minimum wealth or spending objectives under approved stress tests.
4. Remain within strategic risk, concentration, currency and liquidity ranges.
5. Improve the selected objective, such as expected utility, tracking, after-tax outcome or liability coverage.
6. Minimise avoidable costs, turnover and implementation burden.

A stylised constrained objective can be written as:

$$\max_w E[R_p] - \lambda \text{Var}(R_p) - \kappa \text{Cost}(w, w_0) - \eta * \text{Tax}(w, w_0)$$

subject to:

$$A w \leq b, \sum(w) = 1, \text{Liquidity}_h \geq \text{RequiredCash}_h, \text{and } \text{ProhibitedExposure} = 0.$$

The symbols and inputs are not universal. The approved method must define expected returns, covariance, liability horizon, transaction costs, tax treatment, illiquidity, leverage, look-through and tolerance. The language model should not invent these values.


## 5.4 Multiple objectives and family governance

Families can hold competing objectives: preserve purchasing power, distribute cash, support an operating business, fund philanthropy, transfer wealth, maintain control and create optionality. The decision record should show whose objective applies, the legal entity, time horizon, priority and conflict-resolution route.

| Objective field | Required value                                                                 |
|-----------------|--------------------------------------------------------------------------------|
| Owner           | Client, beneficiary group, trust, company or family body.                      |
| Measure         | Currency amount, probability, real return, liability coverage or policy range. |
| Horizon         | Date or review interval.                                                       |
| Priority        | Binding, target or preference.                                                 |
| Flexibility     | Permitted shortfall, delay or substitution.                                    |
| Authority       | Person or body able to approve change.                                         |

## 6. CONTROLLED ARCHITECTURE AND TOOL STACK

### 6.1 Reference architecture



**Figure 4. Controlled architecture for AI-assisted wealth advice**

Source: Matchpoint control architecture; regulatory, privacy and risk context [1-25,34-36].

The architecture begins with identity and entitlement. A client or adviser can access only authorised entities, accounts and documents. A data layer preserves raw source, reconciled values and lineage. A policy service applies eligibility, suitability and mandate rules. An analytics layer calculates exposures, risk, scenarios and optimisation. An AI layer can retrieve approved evidence, prepare questions and draft explanations. A validation layer checks schema, citations, numbers, prohibited actions and policy status. An authorised person approves the final advice or portfolio action.

| Component               | Required control                                                                               |
|-------------------------|------------------------------------------------------------------------------------------------|
| Identity and consent    | Client, household and staff access; purpose; delegation; revocation.                           |
| Data ingestion          | Source contracts, field mapping, timestamps, completeness, duplicate and reconciliation rules. |
| Client truth record     | Versioning, provenance, conflicts, confidence, owner and expiry.                               |
| Policy engine           | Deterministic eligibility, suitability, mandate, concentration and authority rules.            |
| Analytics and optimiser | Versioned code, assumptions, constraints, benchmarks, tests and reproducible output.           |
| Retrieval               | Approved source set, permissions, evidence spans and citation.                                 |
| Generative service      | Structured prompts, bounded output, abstention, logging and no execution authority.            |
| Validation              | Calculation reproduction, source checks, policy checks and adverse tests.                      |
| Workflow                | Adviser review, override, client communication, execution separation and retention.            |
| Monitoring              | Data drift, portfolio drift, model change, incident, complaint and control status.             |

## 6.2 The LLM's permitted role

For an initial deployment, a large language model may summarise source-linked client notes; detect missing or inconsistent facts; draft clarifying questions; retrieve approved research; explain an already calculated scenario; compare a proposal with the approved policy; and prepare a draft review packet.

The model should have no autonomous authority to change risk classification, alter policy constraints, select an unapproved product, approve advice, execute a trade, contact a client, move data, change a source record or suppress an exception. Later authority requires a separate approval supported by evaluation, legal analysis, monitoring and reversibility.

## 6.3 Data and model isolation

Client data should remain partitioned by legal and operational entitlement. Retrieval should enforce access before a passage enters model context. Prompts and outputs should inherit the highest applicable classification. Provider use of input or output for training, retention, support access and cross-border processing requires explicit review.

The Bank of England and FCA survey found significant third-party use and concentration in provider categories [25]. IOSCO guidance expects firms to manage third-party relationships and clarify performance, responsibility, rights and remedies [13]. A6 and A7 organisations should document provider, model, hosting, data, retrieval, identity, monitoring and fourth-party dependencies.

## 6.4 Resilience and manual fallback

The advice process must continue or stop safely when a data feed, model, optimiser or provider fails. The fallback states are explicit: use last verified data with a staleness warning; switch to manual preparation; pause recommendation; or escalate to an authorised owner. The system must never silently substitute a model, assumption or data source.

# 7. CLIENT AND HOUSEHOLD DATA MODEL

## 7.1 Client truth record

Personalisation begins with a source hierarchy. Signed client statements, legal documents, custodian records, administrator records, audited accounts, approved valuations and direct confirmations may carry different authority. The record should retain both the normalised field and the source evidence.

| Field group           | Illustrative contents                                                                  | Required controls                                                 |
|-----------------------|----------------------------------------------------------------------------------------|-------------------------------------------------------------------|
| Identity and entities | Client, beneficial owner, company, trust, foundation, account and relationship         | Verified source, legal owner, effective date and entitlement.     |
| Objectives            | Spending, capital preservation, growth, philanthropy, succession and optionality       | Owner, amount, currency, horizon, priority and flexibility.       |
| Liabilities           | Debt service, commitments, tax, distributions and planned expenditure                  | Amount, date, currency, certainty and legal entity.               |
| Holdings              | Securities, funds, private assets, operating companies, cash and collateral            | Source, valuation, look-through, liquidity, encumbrance and date. |
| Risk                  | Tolerance, capacity, required return, drawdown boundary and scenario response          | Method, reviewer, inconsistencies and approval.                   |
| Constraints           | Product, jurisdiction, ethics, Sharia, ESG, tax, leverage, liquidity and concentration | Binding source, scope, owner and expiry.                          |
| Preferences           | Communication, involvement, exclusions and implementation choices                      | Consent, priority and review date.                                |
| Governance            | Delegation, committee, signatory, reserved matter and beneficiary rights               | Legal confirmation and authority matrix.                          |

## 7.2 Provenance and confidence

Every decision-relevant field should record source, date, owner and status. A practical status set is verified, client-stated, third-party supplied, management-estimated, conflicted, stale or missing. A value labelled management-estimated must not appear as verified in an explanation or optimiser input.

Conflicting sources require an exception. The system can display both values and ask for resolution. It should not select the more recent or convenient value without an approved rule. A missing value can trigger a conservative policy or pause. The chosen response depends on consequence.

### 7.3 Staleness model

Staleness should depend on the field. Listed holdings may require daily or event-driven updates. Private valuations can be quarterly or less frequent. Family objectives may remain stable until a life or governance event. Legal authority changes only when the underlying instrument changes.

| Field                         | Example review trigger                                                  |
|-------------------------------|-------------------------------------------------------------------------|
| Cash and listed holdings      | New feed, failed reconciliation or advice date.                         |
| Private-asset valuation       | New report, material company event or approved expiry.                  |
| Liabilities and capital calls | Notice, drawdown, repayment or forecast change.                         |
| Risk capacity                 | Material wealth, liability, income or family-event change.              |
| Risk tolerance                | Periodic review, behaviour inconsistent with profile or client request. |
| Tax or residency              | Confirmed change or planned move.                                       |
| Authority                     | New legal document, delegation or reserved-matter decision.             |

### 7.4 Consent, purpose and minimisation

The UAE personal-data framework applies to personal-data processing within its scope and establishes governance, security, rights and cross-border requirements [6]. DIFC Data Protection Regulation 10 addresses personal data processed through autonomous and semi-autonomous systems [7]. GDPR and EDPB guidance can apply to relevant EU processing and automated decision-making contexts [23,24]. Applicability requires qualified legal analysis.

The operating control is purpose-specific. The institution should state why each field is collected, where it is used, who can see it, how long it is retained, which providers receive it and how the client can correct it. Behavioural data collected for communication design should not silently become an investment-risk input.

### 7.5 Household consolidation and look-through

Consolidation should preserve legal separation. The household view is analytical and does not merge ownership, custody or authority. Look-through exposure should record the source and method. An operating company may create sector, country, currency, leverage and income dependence. A private fund may expose the portfolio to several strategies and unfunded commitments. Unavailable look-through data remains an explicit unknown.

For A7 portfolios, a useful exposure cube covers legal entity, asset, manager, strategy, sector, geography, currency, liquidity bucket, valuation date, leverage and decision right. The cube supports concentration and cash-flow decisions while retaining ownership boundaries.

## 8. SUITABILITY AND PORTFOLIO CONSTRUCTION

### 8.1 Suitability as an evidence chain

ADGM's digital-investment guidance places suitability, disclosures and algorithm governance within the operating framework for digital investment managers [1]. The DFSA's client-classification and suitability review identified inadequate documentation and over-reliance on tick-box approaches in the reviewed firms [3]. ESMA's 2024 statement identifies bias, data quality, opacity, overreliance, privacy and security risks when AI supports investment services [8]. SEC care guidance expects an understanding of the investment and the specific retail investor's profile, including financial situation, needs, assets, debts, tax status, horizon, liquidity, tolerance and experience [10].

The practical evidence chain is:

`source fact -> client objective -> risk capacity -> applicable constraint -> considered alternatives -> portfolio effect -> adviser judgement -> client communication -> approval`.

Each link should be reconstructable. A risk score without its answers, date, method, conflicts and reviewer is insufficient.

## 8.2 Risk tolerance and capacity

Risk tolerance is behaviourally unstable and can be affected by framing, recent returns and question design. Risk capacity is balance-sheet dependent. The system should preserve both and escalate conflict. A client can state a high willingness to take risk while holding binding near-term liabilities. The resulting portfolio must respect capacity and applicable suitability standards.

| Test                      | Purpose                                                       | Evidence                                                        |
|---------------------------|---------------------------------------------------------------|-----------------------------------------------------------------|
| Questionnaire consistency | Identify contradictory answers and response patterns          | Item-level answers, scoring version and flagged contradictions. |
| Loss scenario             | Test response to a defined monetary and percentage decline    | Scenario, client response and adviser discussion.               |
| Required-return test      | Identify objectives that demand unrealistic or excessive risk | Cash-flow model, assumptions and alternatives.                  |
| Liquidity stress          | Protect known and contingent obligations                      | Cash ladder, haircuts and stress window.                        |
| Concentration test        | Capture existing wealth and income dependence                 | Look-through exposure and limit rationale.                      |
| Behaviour review          | Compare stated tolerance with observed actions                | Source, context and adviser interpretation.                     |

## 8.3 Investment-policy construction

The investment policy translates client facts into binding rules and ranges. It should identify base currency, objectives, benchmark, strategic allocation, permitted and prohibited assets, liquidity, concentration, leverage, currency, derivatives, private-market pacing, ESG or Sharia requirements, rebalancing, valuation, delegation and review.

An AI assistant can draft a policy from structured approved fields and identify missing sections. The adviser or investment committee owns the policy. Source-linked drafting allows the reviewer to trace each restriction to the client instruction, legal document or approved judgement.

## 8.4 Portfolio optimisation

Markowitz formalised mean-variance portfolio selection [38]. Merton developed continuous-time portfolio choice with consumption and wealth dynamics [39]. Black and Litterman presented an approach to combine equilibrium returns with investor views [40]. These methods provide foundations and remain sensitive to assumptions, estimation error, objectives and constraints.

The optimiser should be a tested analytical service. Its input pack includes the eligible universe, return and risk assumptions, currency, fees, taxes, liquidity, turnover, asset ranges and the current portfolio. Its output pack includes the objective value, constraints, binding limits, expected characteristics, scenario results, turnover and alternatives.

The institution should test stability. Small changes in expected returns or covariance can create large allocation changes. Controls can include robust estimation, shrinkage, regularisation, resampling, minimum trade sizes, turnover penalties and strategic ranges. The selected technique requires independent review and documented rationale.

## 8.5 Private assets and illiquidity

Private investments require cash-flow and valuation treatment beyond a listed-asset optimiser. The portfolio should record unfunded commitments, expected calls and distributions, liquidity horizon, valuation date, leverage, currency, sector and concentration. Scenario analysis should include delayed exits, reduced distributions, accelerated calls and valuation haircuts.

For A7 families, the operating business can dominate economic exposure. Treating it as absent creates a distorted public-market allocation. A range-based exposure and scenario overlay can be more honest than a false point estimate. The investment committee should approve how incomplete private exposures enter the decision.

## 8.6 Tax, currency and entity constraints

Advice should distinguish pre-tax optimisation from after-tax implementation. Tax lots, wrappers, legal ownership, residence, withholding, realised gains and succession structures can change the preferred action. Qualified tax advice is required. The system can preserve inputs, calculate approved scenarios and flag conflicts.

Currency should be linked to liabilities and objectives. Reporting currency alone is insufficient. The allocation record should show asset currency, economic currency, hedging, liability currency and the policy for residual exposure.

## 8.7 Product and manager selection

The portfolio allocation and the product choice are separate decisions. Product selection should consider mandate fit, liquidity, fees, tracking, counterparty, credit, operational risk, tax, legal terms, capacity and conflicts. A model should retrieve only from the approved eligible universe and current due-diligence records. Marketing language and popularity cannot substitute for product governance.

| Gate           | Required evidence                                                           |
|----------------|-----------------------------------------------------------------------------|
| Eligibility    | Jurisdiction, client type, mandate, platform and policy.                    |
| Investment fit | Exposure, role, benchmark, risk, liquidity and scenario behaviour.          |
| Cost           | Explicit fee, embedded cost, transaction, custody, tax and performance fee. |
| Operational    | Dealing, valuation, custody, reporting, service and continuity.             |
| Conflict       | Revenue, rebate, proprietary product, placement or related party.           |
| Approval       | Product committee or authorised owner, conditions and expiry.               |

## 9. ADVISER WORKFLOW AND EXPLANATION

### 9.1 Before-and-after workflow

The existing workflow often moves through meetings, notes, spreadsheets, risk systems, investment proposals, compliance checks and manually assembled review documents. The controlled workflow retains professional discussion and creates one evidence path.

1. Ingest authorised source records and client notes.
2. Reconcile holdings, entities, liabilities and restrictions.
3. Identify missing, stale or conflicting facts.
4. Apply deterministic suitability and policy rules.
5. Calculate portfolio, alternatives, costs and scenarios.
6. Draft a source-linked explanation packet.
7. Obtain adviser and required control review.
8. Discuss with the client or decision body.
9. Record acceptance, override, deferral or rejection.
10. Route approved instructions through separated execution controls.

### 9.2 Explanation packet

An explanation should help the recipient understand the decision. It should state facts used, relevant constraints, recommended change, considered alternatives, costs, key scenarios, uncertainty, conflicts, limitations and required action.

| Section             | Question answered                                                             |
|---------------------|-------------------------------------------------------------------------------|
| Client context      | Which dated facts and objectives drove the analysis?                          |
| Current portfolio   | What exposures, concentrations, liquidity and costs exist?                    |
| Proposed action     | What changes, in which entities or accounts, are proposed?                    |
| Rationale           | Which objectives and constraints support the proposal?                        |
| Alternatives        | Which reasonable alternatives were considered and why were they not selected? |
| Downside            | What could happen under defined adverse scenarios?                            |
| Costs and conflicts | What direct and indirect costs or institutional interests apply?              |
| Limitations         | Which data, estimates, model boundaries or uncertainties remain?              |
| Authority           | Who reviewed, approved and can challenge the decision?                        |

The explanation should not expose a model's private chain of thought. It should expose the decision evidence, method, material assumptions, alternatives and accountable judgement.

### 9.3 Adviser review

Meaningful review requires time, competence, source access and authority to reject. A review button alone does not establish oversight. The adviser should confirm client facts, policy, portfolio calculations, alternatives, costs, conflicts, language and action. Material overrides require a coded reason and free-text rationale.

| Override reason   | Example                                                        |
|-------------------|----------------------------------------------------------------|
| Missing fact      | External holding or planned liability omitted.                 |
| Stale data        | Private valuation no longer decision-useful.                   |
| Client preference | Binding preference not represented in the policy engine.       |
| Model limitation  | Illiquidity or tax interaction outside approved scope.         |
| Market judgement  | Approved tactical or implementation judgement.                 |
| Conduct concern   | Recommendation or communication could disadvantage the client. |

Override analysis is a learning signal. Persistent overrides can reveal data, policy, optimiser or workflow defects. A low override rate is not automatically good; automation bias can suppress challenge.

### 9.4 Client comprehension and challenge

The client should be able to correct material facts, ask for alternatives, understand costs and decline the recommendation. Explanation testing should measure comprehension and correction, not persuasion. The 2026 transparency field study supports direct testing of formats and channels [28].

For a family investment committee, the packet can include a decision dashboard, appendix of assumptions and entity-level actions. Different audiences can receive different depth while the underlying facts and calculations remain consistent.

### 9.5 Communications control

FINRA states that technology-neutral rules continue to apply to generative AI, including supervision and communications standards [12]. SEC enforcement against misleading AI claims demonstrates the importance of accurate descriptions of system use [11]. Client communications should identify the service accurately, avoid unsupported performance or capability claims, preserve required disclosures and record the approved version.

## 10. GOVERNANCE, PRIVACY, SECURITY AND THIRD PARTIES

### 10.1 Senior accountability

IOSCO's 2021 guidance recommends designated senior-management responsibility, testing and monitoring, appropriate competence, third-party oversight, disclosure and data controls [13]. ADGM expects board and senior-management oversight of algorithm design, performance, deployment and security for digital investment managers [1]. An A6 institution should assign an accountable executive, advice owner, investment-method owner, data owner, compliance owner, privacy owner, security owner and independent validator.

An A7 family office can use a proportionate structure. The CIO or investment committee owns investment policy and portfolio decisions. A named technology owner operates the service. External counsel, tax advisers, cybersecurity specialists and investment consultants provide bounded expertise. Delegation and reserved matters remain documented.

### 10.2 AI and model register

The register should identify each system, component and use. A language model embedded in a customer-relationship system, portfolio platform or office application remains part of the perimeter when it can process client information or influence advice.

| Register field  | Required content                                                 |
|-----------------|------------------------------------------------------------------|
| Identity        | System, provider, model, version, owner and status.              |
| Purpose         | User, process, client population, output and excluded use.       |
| Decision effect | Facts, policy, portfolio, communication or execution influenced. |
| Data            | Sources, sensitivity, location, rights, retention and transfer.  |
| Authority       | Read, draft, recommend, approve, communicate or execute.         |

| Register field | Required content                                                     |
|----------------|----------------------------------------------------------------------|
| Evidence       | Design, evaluation, validation, privacy, security and legal reviews. |
| Operation      | Monitoring, changes, incidents, complaints and exceptions.           |
| Exit           | Fallback, portability, deletion, transition and termination.         |

### 10.3 Risk framework

NIST AI RMF organises outcomes under Govern, Map, Measure and Manage [18]. The NIST Generative AI Profile identifies risks including confabulation, harmful bias, data privacy, information integrity, information security and value-chain dependencies [19]. ISO/IEC 42001 specifies an AI management system [20]. ISO/IEC 23894 provides AI risk-management guidance [21]. These sources can structure controls. Their use does not establish legal compliance or certification.

The risk assessment should cover unsuitable advice, client harm, discrimination, privacy, confidentiality, cyberattack, inaccurate data, model instability, unsupported language, conflicts, third-party failure, concentration, unauthorised action and inability to reconstruct the decision.

### 10.4 Privacy and automated processing

The institution should map controller and processor roles, legal basis, purpose, data category, access, transfer, retention, rights and incident obligations. Where automated decision-making or profiling rules apply, qualified legal owners should determine the required information, safeguards, human intervention and challenge route [7,23,24].

Data minimisation matters because hyper-personalisation can invite excessive collection. A data item should enter the record only when its decision purpose and control are defined. Sensitive family information needs particularly restrictive access, logging and disclosure.

### 10.5 Security

The service should use least privilege, client and entity isolation, encryption, secret management, secure integration, output validation, logging, incident response and provider controls. Retrieval content can contain malicious or misleading instructions. Documents and web content are data, not authority. The system should keep tool permissions separate from retrieved text and validate every proposed action against deterministic policy.

### 10.6 Third-party due diligence

| Area        | Due-diligence evidence                                                          |
|-------------|---------------------------------------------------------------------------------|
| Service     | Purpose, architecture, model, sub-processors and support.                       |
| Data        | Training use, retention, deletion, residency, transfer and access.              |
| Control     | Identity, isolation, encryption, logging, change and incident.                  |
| Performance | Relevant evaluation, limitations, service levels and monitoring.                |
| Change      | Notice, version control, regression evidence and right to object.               |
| Resilience  | Continuity, recovery, concentration, portability and exit.                      |
| Contract    | Responsibility, audit, incident, confidentiality, IP, remedies and termination. |

### 10.7 Conflicts and optimisation objective

An optimisation function can create a conduct conflict if it includes institutional revenue, engagement or product-placement objectives without appropriate governance. The decision contract should record the optimisation objective and prohibited objectives. Revenue propensity should remain outside suitability and portfolio construction unless qualified owners determine a lawful, fair and disclosed use.

The SEC's care guidance and fiduciary interpretation require client interests and a reasonable understanding of objectives [10,36]. The system must preserve considered alternatives, product costs and institutional conflicts.

## 11. EVALUATION AND MONITORING

### 11.1 Evaluation unit

The evaluation unit is the configured workflow for a defined client population and advice task. Provider benchmark scores do not establish fitness. Tests should include data ingestion, policy rules, calculations, retrieval, generation, human review, communication and failure handling.



**Figure 5. Wealth-advice evaluation scorecard**

Source: Matchpoint evaluation framework.

### 11.2 Representative evaluation set

The set should cover A6 and A7 cases with different complexity: new and established clients; advisory and discretionary mandates; listed-only and mixed liquid-private portfolios; single and multiple entities; several currencies; conflicting tolerance and capacity; missing or stale data; product restrictions; tax lots; concentrated operating businesses; capital calls; and market stress.

Each case needs an approved reference outcome or adjudication process. Experts should document acceptable alternatives. Suitability and portfolio decisions can permit more than one valid answer. Evaluation should distinguish an unacceptable breach from a reasonable professional alternative.

### 11.3 Metric stack

| Layer      | Metric                          | Definition                                                     |
|------------|---------------------------------|----------------------------------------------------------------|
| Data       | Critical-field completeness     | Required current fields present with acceptable source status. |
| Data       | Reconciliation exceptions       | Unresolved differences by source and consequence.              |
| Policy     | Constraint-breach rate          | Proposals violating an approved binding rule.                  |
| Portfolio  | Calculation reproduction        | Independent reproduction within approved tolerance.            |
| Portfolio  | Allocation stability            | Change under defined input perturbations.                      |
| Portfolio  | Scenario coverage               | Required shocks calculated and explained.                      |
| Retrieval  | Evidence precision and recall   | Relevant approved evidence returned for defined questions.     |
| Generation | Grounded material statements    | Material statements supported by cited source or calculation.  |
| Human      | First-pass acceptance           | Packets accepted without material correction.                  |
| Human      | Override pattern                | Rate, reason, consequence and reviewer.                        |
| Client     | Comprehension                   | Correct responses to defined decision questions.               |
| Conduct    | Complaint or unsuitable outcome | Confirmed issue under approved taxonomy.                       |

| Layer     | Metric                    | Definition                                        |
|-----------|---------------------------|---------------------------------------------------|
| Economics | Cost per accepted outcome | Full controlled cost divided by accepted packets. |

### 11.4 Hard gates

Certain failures should block production regardless of average score: wrong client or entity; unauthorised data access; binding-constraint breach; unsupported product; calculation outside tolerance; missing material cost or conflict; unsupported material claim; prohibited action; failed approval; or inability to reconstruct the decision.

An average accuracy percentage can conceal a severe suitability or privacy failure. The gate should use consequence-weighted errors and separate thresholds by task.

### 11.5 Monitoring

Monitoring should cover data, portfolio, method, provider, human and client outcomes.

| Signal                                    | Triggered review                                                                  |
|-------------------------------------------|-----------------------------------------------------------------------------------|
| Data feed failure or reconciliation break | Pause affected calculations and resolve source.                                   |
| Material client or family event           | Refresh truth record, policy and suitability.                                     |
| Portfolio range or liquidity breach       | Review action under investment policy.                                            |
| Model or provider change                  | Regression test and approval before production.                                   |
| Rising overrides or rework                | Diagnose data, policy, method or interface.                                       |
| Complaint or unsuitable outcome           | Incident, client remediation and control review.                                  |
| Security or privacy event                 | Incident response, containment and required notification.                         |
| Performance drift                         | Review assumptions, attribution and communication; avoid automatic model chasing. |

### 11.6 Independent challenge

Independent reviewers should have access to methods, cases, source records and failure logs. They should reproduce calculations, challenge data and assumptions, test prohibited actions, evaluate subgroup outcomes and inspect whether human review is meaningful. Conditions and residual risks should have owners and expiry dates.

## 12. PRODUCTIVITY AND ECONOMIC MEASUREMENT

### 12.1 Accepted outcome

The economic unit is an accepted advice or review packet. A generated draft has no recognised value until it passes data, suitability, portfolio, control and human gates. Time saved on drafting can be consumed by reconciliation, correction, supervision, client discussion or incident handling.

`Cost per accepted outcome = Total controlled workflow cost / Accepted outcomes`

Total controlled workflow cost includes staff, data, licences, integration, compute, evaluation, validation, security, privacy, compliance, vendor oversight, change, monitoring, incident and retirement cost.

### 12.2 Productivity bridge

| Stage                        | Baseline measure               | Assisted measure                     | Quality guardrail                        |
|------------------------------|--------------------------------|--------------------------------------|------------------------------------------|
| Client discovery preparation | Staff minutes per packet       | Staff and compute minutes            | Critical facts and questions complete.   |
| Data reconciliation          | Exceptions and elapsed time    | Exceptions and elapsed time          | No unresolved material difference.       |
| Portfolio analysis           | Analyst minutes and rework     | Analyst and compute minutes          | Reproduced calculations and constraints. |
| Explanation drafting         | Adviser and support minutes    | Adviser, support and compute minutes | Grounded material statements.            |
| Review and approval          | Review minutes and corrections | Review minutes and corrections       | Suitable, complete and approved.         |
| Client meeting               | Preparation and follow-up      | Preparation and follow-up            | Client comprehension and correction.     |

The productivity rate should use equivalent case mix and output standard. Measurement should report median and tail time, not only a mean. Cases that fail or require extensive rework remain in the denominator.

### 12.3 Capacity and revenue

Released capacity has economic value only when the institution changes an observable outcome. The adviser can spend more time with clients, serve additional qualified relationships, shorten review cycles or reduce external support. Recognition requires time records, accepted output, capacity release and a realised operational or commercial event.

Realised capacity value = Verified hours released *Approved realised-use factor* Approved value per realised hour

Revenue attribution requires a documented causal rule linking the intervention to an acquired, retained or expanded client relationship and collected revenue. Pipeline, proposal value, client engagement and assets under advisement are intermediate indicators. They are not collected revenue.

### 12.4 Risk and loss

Avoided loss is difficult to prove. A scenario can support risk appetite and control investment. It should remain a management estimate until an approved methodology and observed counterfactual evidence exist.

Quality costs can be measured directly: material corrections, delayed advice, client remediation, trade cancellation, complaint handling, incident response and external review. A productivity case should subtract these costs.

### 12.5 Unverified illustrative business-case formulae

All variables in this subsection are unverified illustrative management assumptions.

Gross time benefit = (Baseline minutes - Assisted minutes) *Eligible volume* Loaded minute cost


Quality-adjusted benefit = Gross time benefit *Acceptance factor* Realisation factor

Net benefit = Quality-adjusted benefit + Approved realised revenue + Approved observed loss reduction - Full incremental cost

ROI = Net benefit / Full incremental cost

No numeric scenario is presented as a forecast. Attributed Matchpoint or client revenue, cash cost reduction and loss reduction remain USD 0.

#### Evidence-gated value bridge



**Figure 6. Evidence-gated value bridge**

Source: Matchpoint business-case framework; all scenario values require approved observed evidence.

## 13. ILLUSTRATIVE OPERATING CASES

### 13.1 A6 annual-review workflow

The following case is an unverified illustrative management scenario. It does not describe an observed Matchpoint or client implementation.

An A6 multi-family office prepares annual reviews for clients with several custodians. The proposed pilot covers data consolidation, fact-change questions, portfolio diagnostics, approved scenarios and a draft explanation packet. Advice, product selection, client communication and execution remain with authorised professionals.

| Scenario input                           | Unverified illustrative assumption |
|------------------------------------------|------------------------------------|
| Pilot relationships                      | 24                                 |
| Custody sources per relationship         | 3                                  |
| Baseline staff time per accepted packet  | 11.0 hours                         |
| Assisted staff time target               | 7.5 hours                          |
| Required critical-field completeness     | 100%                               |
| Binding-constraint tolerance             | 0 breaches                         |
| Material unsupported statement tolerance | 0                                  |
| Minimum first-pass acceptance target     | 85%                                |

These are management assumptions for measurement design. They are not benchmark evidence. The pilot should record actual case complexity, time by role, rework, overrides, data exceptions, client corrections and accepted outcomes. It should stop if a hard gate fails or if reviewers lack enough time to challenge the output.

### 13.2 A6 decision gate

The pilot advances when the institution verifies data entitlement and reconciliation; reproduces calculations; records zero binding-constraint breaches in the approved test set and pilot; establishes grounded explanations; demonstrates meaningful adviser review; and observes a quality-adjusted cost per accepted outcome within the approved threshold.

An average time reduction cannot override a suitability, access or material-explanation failure. Expansion should add one controlled dimension at a time, such as a new client segment, data source, product type or workflow authority.

### 13.3 A7 family-capital allocation workflow

The following case is also an unverified illustrative management scenario.

An A7 family office maintains listed investments, several private funds, direct company interests, real estate and an operating-company stake. The proposed pilot creates an exposure cube, a twelve-quarter liquidity ladder, strategic-allocation scenarios and an investment-committee packet. Private valuations remain dated and range-based. The optimiser produces alternatives only for the approved liquid portfolio and uses committee-approved overlays for illiquid exposures.

| Scenario input                           | Unverified illustrative assumption       |
|------------------------------------------|------------------------------------------|
| Legal entities in analytical perimeter   | 9                                        |
| Custody and administrator sources        | 6                                        |
| Private funds and direct investments     | 28                                       |
| Liquidity forecast horizon               | 12 quarters                              |
| Strategic scenarios                      | Base, downside and severe downside       |
| Operating-company valuation              | Approved range, not point estimate       |
| Investment-committee decisions           | Human only                               |
| Attributed revenue, cost or loss outcome | USD 0 pending approved observed evidence |

### 13.4 A7 decision gate

The family office should confirm legal ownership and authority; reconcile listed assets and cash; label private valuations and look-through gaps; document capital-call assumptions; test liquidity under stress; show the public portfolio with and without operating-business exposure; and record committee decisions and overrides.


The output should expose uncertainty. A wide valuation range, missing manager exposure or disputed family objective remains visible. The system should not create false precision to make the allocation appear complete.

### 13.5 Comparison of operating cases

| Dimension              | A6 institution                                     | A7 family office                                             |
|------------------------|----------------------------------------------------|--------------------------------------------------------------|
| Unit of scale          | Repeated client relationship                       | Consolidated family-capital decision                         |
| Primary owner          | Adviser and investment committee                   | CIO, principal or family investment committee                |
| Main data challenge    | Multi-client consistency and entitlement           | Entity, private-asset and operating-business consolidation   |
| Main conduct challenge | Suitability, conflicts and communications at scale | Governance, competing objectives and incomplete valuations   |
| Main economic measure  | Cost and cycle time per accepted packet            | Decision quality, liquidity readiness and committee capacity |
| Initial AI authority   | Synthesis, questions and explanation draft         | Consolidation support, scenarios and committee draft         |

## 14. IMPLEMENTATION ROADMAP

### Ninety-day adoption roadmap for controlled wealth advice



**Figure 7. Ninety-day adoption roadmap for controlled wealth advice**

Source: Matchpoint implementation framework.

#### 14.1 Days 0-15: authority and scope

Name the accountable executive and advice or investment owner. Select one client population and one decision. Define prohibited actions, data boundary, review, fallback and hard gates. Reconcile existing AI features in portfolio, CRM, office and communication tools.

#### 14.2 Days 16-30: truth record and baseline

Define the client and household schema, source hierarchy, staleness, reconciliation and exception rules. Sample representative cases and measure current time, rework, data breaks, review and quality. Obtain privacy, security, legal and compliance analysis for the bounded pilot.

#### 14.3 Days 31-45: policy and analytics

Version suitability rules, eligible universe, portfolio constraints, assumptions, scenarios and calculation methods. Create an independent reproduction test. Build the explanation template and evidence requirements.

#### 14.4 Days 46-60: offline evaluation

Run representative historical and synthetic cases without client or execution impact. Test missing data, contradictory facts, stale valuations, constraint conflicts, provider failure, malicious content, unsupported questions and adverse markets. Record every error and adjudicate acceptable alternatives.

### 14.5 Days 61-75: shadow operation

Run the assisted workflow beside the existing process. Reviewers remain blind to assisted output where the evaluation design requires it. Compare completeness, calculations, recommendations, explanations, time, rework and overrides. No assisted output reaches the client or execution route.

### 14.6 Days 76-90: restricted pilot gate

Approve a restricted pilot only after hard gates, data rights, validation, supervision, monitoring and fallback pass. Define maximum population, permitted functions, named reviewers, daily or weekly issue review and stop authority. Record collected economic evidence without forecasting unobserved benefits.

### 14.7 Months four to twelve

Expand one dimension after evidence. Add data sources, cases, products or clients through change control and regression testing. Test provider exit and manual continuity. Analyse overrides, subgroup outcomes, complaints and accepted-outcome economics. Obtain independent assurance before material authority expands.

| Stage    | Evidence gate                                                  | Decision              |
|----------|----------------------------------------------------------------|-----------------------|
| Scope    | Decision contract and authority approved                       | Design.               |
| Baseline | Representative current process measured                        | Build.                |
| Offline  | Hard gates and performance thresholds pass                     | Shadow.               |
| Shadow   | Equivalent or better accepted outcomes with meaningful review  | Restricted pilot.     |
| Pilot    | Stable quality, conduct, control and economics                 | Controlled expansion. |
| Scale    | Portfolio-wide dependency, resilience and assurance acceptable | Broader production.   |

## 15. LIMITATIONS, RESEARCH AGENDA AND CONCLUSION

The cited empirical studies use different countries, products, investor populations, interfaces and methods. Several focus on retail robo-advice rather than private banks, EAMs, multi-family offices or GCC family offices. The 2026 transparency paper was a preprint at the evidence cut-off. Portfolio and behavioural effects do not establish adviser productivity, collected revenue, local suitability or performance for a specific institution.

The regulatory synthesis is cross-jurisdictional. Applicability depends on entity, activity, client and location. Standards and voluntary frameworks support control design and do not establish legal compliance. AI capabilities, provider terms and regulatory timetables can change after the evidence cut-off.

Future research should test household-aware advice with private assets and liabilities; compare explanation formats for comprehension and correction; measure adviser overrides and automation bias; evaluate after-tax and multi-currency implementation; test whether data completeness improves suitable outcomes; and report incidents, complaints and subgroup outcomes. Institutions should publish methods and bounded results where confidentiality permits.

The practical conclusion is that hyper-personalised wealth advice is a controlled decision system. A governed truth record captures the client or family. Approved policy and mathematics generate reproducible alternatives. AI assists synthesis, questions and explanation. An accountable professional owns suitability, judgement and action. Production evidence determines scale.

Attributed Matchpoint or client revenue, cash cost reduction and loss reduction remain USD 0 until approved observed evidence supports attribution.

### DECLARATIONS

Author attribution. Prepared by Matchpoint Partners. Named-person author attribution is pending CK approval.

Evidence cut-off. Sources were reviewed through 1 August 2026.

Scenario status. Every worked case, threshold, volume, time, rate and economic input labelled unverified is an illustrative management assumption. It is not observed client or Matchpoint evidence.

Financial attribution. Attributed Matchpoint or client revenue, cash cost reduction and loss reduction remain USD 0 until approved observed evidence exists.

Professional boundaries. Research paper only. This document is not investment, legal, regulatory, accounting, audit, tax, privacy, cybersecurity, valuation or technology advice.

Conflicts and funding. No external funding or client-specific data was used. Matchpoint Partners has a commercial interest in advisory work related to the subject.

## REFERENCES

- [1] Financial Services Regulatory Authority of Abu Dhabi Global Market (2025). Supplementary Guidance: Digital Investment Management. <https://assets.adgm.com/download/assets/supplementary-guidance-authorisation-of-digital-investment-management- robo-advisory-activities.pdf/f2ac6f36606c11ef946beccc50de2f68>
- [2] Central Bank of the UAE, Securities and Commodities Authority, Dubai Financial Services Authority and Financial Services Regulatory Authority (2021). Guidelines for Financial Institutions Adopting Enabling Technologies. <https://www.adgm.com/media/announcements/uae-regulatory-authorities-jointly-issue-guidelines-for-financial-institutions>
- [3] Dubai Financial Services Authority (2018). Findings of Client Classification and Suitability Review. <https://www.dfsa.ae/news/dfs-a-publishes-findings-client-classification-and-suitability-review>
- [4] Dubai Financial Services Authority (2025). Artificial Intelligence Survey 2025. <https://www.dfsa.ae/news/new-dfs-a-survey-generative-ai-adoption-has-nearly-tripled-within-difc-last-12-months-governance-continues-develop>
- [5] Central Bank of the United Arab Emirates (2026). Guidance Note on the Consumer Protection and Responsible Adoption and Use of Artificial Intelligence and Machine Learning by Licensed Financial Institutions in the U.A.E. <https://rulebook.centralbank.ae/en/rulebook/guidance-note-consumer-protection-and-responsible-adoption-and-use-artificial-intelligence>
- [6] United Arab Emirates Government (2021). Federal Decree-Law No. 45 of 2021 Regarding the Protection of Personal Data. <https://u.ae/en/about-the-uae/digital-uae/data/data-protection-laws>
- [7] Dubai International Financial Centre (current at 2026). Data Protection Regulations; Regulation 10: Personal Data Processed through Autonomous and Semi-Autonomous Systems. <https://assets.difc.com/v1/media/edge/images/dubaiintern0078-difcexperie96c5-production-3253/media/project/difcexperiences/difc/difcwebsite/documents/laws--regulations/data-protection-regulation.pdf>
- [8] European Securities and Markets Authority (2024). Public Statement on Artificial Intelligence and the Provision of Investment Services. <https://www.esma.europa.eu/press-news/esma-news/esma-provides-guidance-firms-using-artificial-intelligence-investment-services>
- [9] US Securities and Exchange Commission (2017). Robo-Advisers: Guidance Update and Investor Bulletin. <https://www.sec.gov/newsroom/press-releases/2017-52>
- [10] US Securities and Exchange Commission Staff (2023). Standards of Conduct for Broker-Dealers and Investment Advisers Care Obligations. <https://www.sec.gov/about/divisions-offices/division-trading-markets/broker-dealers/staff-bulletin-standards-conduct-broker-dealers-investment-advisers-care-obligations>
- [11] US Securities and Exchange Commission (2024). SEC Charges Two Investment Advisers with Making False and Misleading Statements About Their Use of Artificial Intelligence. <https://www.sec.gov/newsroom/press-releases/2024-36>
- [12] Financial Industry Regulatory Authority (2024). Regulatory Notice 24-09: Regulatory Obligations When Using Generative Artificial Intelligence and Large Language Models. <https://www.finra.org/rules-guidance/notices/24-09>
- [13] International Organization of Securities Commissions (2021). The Use of Artificial Intelligence and Machine Learning by Market Intermediaries and Asset Managers. <https://www.iosco.org/library/pubdocs/pdf/IOSCOPD684.pdf>
- [14] International Organization of Securities Commissions (2025). Artificial Intelligence in Capital Markets: Use Cases, Risks, and Challenges. <https://www.iosco.org/library/pubdocs/pdf/IOSCOPD788.pdf>
- [15] International Monetary Fund (2024). Global Financial Stability Report, Chapter 3: Advances in Artificial Intelligence: Implications for Capital Market Activities. <https://www.imf.org/en/publications/gfsr/issues/2024/10/22/global-financial-stability-report-october-2024>
- [16] Organisation for Economic Co-operation and Development (2023). Generative Artificial Intelligence in Finance. [https://www.oecd.org/content/dam/oecd/en/publications/reports/2023/12/generative-artificial-intelligence-in-finance\\_37bb17c6/ac7149cc-en.pdf](https://www.oecd.org/content/dam/oecd/en/publications/reports/2023/12/generative-artificial-intelligence-in-finance_37bb17c6/ac7149cc-en.pdf)
- [17] Organisation for Economic Co-operation and Development (2021). Artificial Intelligence in Finance. [https://www.oecd.org/en/publications/oecd-business-and-finance-outlook-2021\\_ba682899-en/full-report/component-6.html](https://www.oecd.org/en/publications/oecd-business-and-finance-outlook-2021_ba682899-en/full-report/component-6.html)
- [18] National Institute of Standards and Technology (2023). Artificial Intelligence Risk Management Framework 1.0. <https://www.nist.gov/itl/ai-risk-management-framework>
- [19] National Institute of Standards and Technology (2024). Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
- [20] International Organization for Standardization (2023). ISO/IEC 42001:2023, Artificial Intelligence Management Systems. <https://www.iso.org/standard/42001>
- [21] International Organization for Standardization (2023). ISO/IEC 23894:2023, Artificial Intelligence: Guidance on Risk Management. <https://www.iso.org/standard/77304.html>
- [22] European Union (2024). Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:32024R1689>
- [23] European Union (2016). Regulation (EU) 2016/679, General Data Protection Regulation. <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- [24] European Data Protection Board (2018). Guidelines on Automated Individual Decision-Making and Profiling. [https://www.edpb.europa.eu/documents/guideline/automated-decision-making-and-profiling\\_en](https://www.edpb.europa.eu/documents/guideline/automated-decision-making-and-profiling_en)
- [25] Bank of England and Financial Conduct Authority (2024). Artificial Intelligence in UK Financial Services, 2024. <https://www.bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024>
- [26] D'Acunto, F., Prabhala, N. and Rossi, A.G. (2019). The Promises and Pitfalls of Robo-Advising. *Review of Financial Studies*, 32(5), 1983-2020. <https://doi.org/10.1093/rfs/hhz014>
- [27] Scherer, B. and Lehner, S. (2025). What Drives Robo-Advice? *Journal of Empirical Finance*, 80, 101574. <https://doi.org/10.1016/j.jempfin.2024.101574>
- [28] Boulou-Reshef, B., Direr, A. and Louafi, M. (2026). Algorithmic Transparency and Portfolio Choices: Field Evidence. SSRN preprint. <https://doi.org/10.2139/ssrn.6042654>
- [29] Lambrecht, M., Oechssler, J. and Weidenholzer, S. (2026). On the Benefits of Robo-Advice in Financial Markets. *Economic Journal*, accepted manuscript. <https://doi.org/10.1093/ej/ueag076>
- [30] Back, C., Morana, S. and Spann, M. (2023). When Do Robo-Advisors Make Us Better Investors? The Impact of Social Design Elements on Investor Behavior. *Journal of Behavioral and Experimental Economics*, 103, 101984. <https://doi.org/10.1016/j.socec.2023.101984>
- [31] Musto, C., Semeraro, G., Lops, P., de Gemmis, M. and Lekkass, G. (2015). Personalized Finance Advisory through Case-Based Recommender Systems and Diversification Strategies. *Decision Support Systems*, 77, 100-111. <https://doi.org/10.1016/j.dss.2015.06.001>
- [32] Wang, C., Zhang, T. and Zhu, S. (2026). Big Data Recommendations and Portfolio Diversification: Evidence from Account-Level Data. *Finance Research Letters*, 92, 109541. <https://doi.org/10.1016/j.frl.2026.109541>
- [33] Mahmud, H., Islam, N. and Krishnan, S. (2025). Human-Robo-Advisor Collaboration in Decision-Making: Evidence from a Multiphase Mixed Methods Experimental Study. *Decision Support Systems*, 198, 114541. <https://doi.org/10.1016/j.dss.2025.114541>
- [34] Organisation for Economic Co-operation and Development (2024). Regulatory Approaches to Artificial Intelligence in Finance. [https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/09/regulatory-approaches-to-artificial-intelligence-in-finance\\_43d082c3/f1498c02-en.pdf](https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/09/regulatory-approaches-to-artificial-intelligence-in-finance_43d082c3/f1498c02-en.pdf)
- [35] Financial Stability Board (2024). The Financial Stability Implications of Artificial Intelligence. <https://www.fsb.org/2024/11/the-financial-stability-implications-of-artificial-intelligence/>
- [36] US Securities and Exchange Commission (2019). Commission Interpretation Regarding Standard of Conduct for Investment Advisers. <https://www.sec.gov/files/rules/interp/2019/ia-5248.pdf>
- [37] Foerster, S., Linnainmaa, J.T., Melzer, B.T. and Previtero, A. (2017). Retail Financial Advice: Does One Size Fit All? *Journal of Finance*, 72(4), 1441-1482. <https://doi.org/10.1111/jofi.12514>
- [38] Markowitz, H. (1952). Portfolio Selection. *Journal of Finance*, 7(1), 77-91. <https://doi.org/10.1111/j.1540-6261.1952.tb01525.x>
- [39] Merton, R.C. (1969). Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case. *Review of Economics and Statistics*, 51(3), 247-257. <https://doi.org/10.2307/1926560>
- [40] Black, F. and Litterman, R. (1992). Global Portfolio Optimization. *Financial Analysts Journal*, 48(5), 28-43. <https://doi.org/10.2469/faj.v48.n5.28>

**APPENDIX A. CLIENT TRUTH RECORD TEMPLATE**

| Field           | Required entry                                                                                     |
|-----------------|----------------------------------------------------------------------------------------------------|
| Record identity | Client, household, legal entity, account and version.                                              |
| Fact            | Normalised value, unit and applicable entity.                                                      |
| Source          | Document, system, person, date and exact evidence.                                                 |
| Status          | Verified, client-stated, third-party supplied, management-estimated, conflicted, stale or missing. |
| Purpose         | Advice or portfolio decisions that use the fact.                                                   |
| Rights          | Consent, access, transfer, retention and correction route.                                         |
| Owner           | Person accountable for accuracy and refresh.                                                       |
| Expiry          | Date or event that requires review.                                                                |

**APPENDIX B. ADVICE PACKET CHECKLIST**

| Gate        | Reviewer evidence                                                       |
|-------------|-------------------------------------------------------------------------|
| Identity    | Correct client, entity, account and mandate.                            |
| Facts       | Critical fields current, sourced and reconciled.                        |
| Objectives  | Owner, amount, currency, horizon and priority confirmed.                |
| Suitability | Tolerance, capacity, experience, liquidity and conflicts assessed.      |
| Policy      | Current approved constraints and eligible universe applied.             |
| Portfolio   | Calculations reproduced; alternatives, costs and scenarios reviewed.    |
| Product     | Eligibility, due diligence, cost and conflict current.                  |
| Explanation | Material statements grounded; limitations and uncertainty visible.      |
| Approval    | Named adviser, control owner and client or committee action recorded.   |
| Execution   | Approved instruction separated from analysis and checked independently. |

**APPENDIX C. PORTFOLIO INPUT AND OUTPUT RECORD**

| Group             | Minimum fields                                                                            |
|-------------------|-------------------------------------------------------------------------------------------|
| Current portfolio | Holdings, tax lots, price, currency, cost, liquidity, entity and source date.             |
| Assumptions       | Expected return, risk, correlation, inflation, currency, fee and tax treatment.           |
| Constraints       | Asset ranges, concentration, liquidity, turnover, leverage, product and entity rules.     |
| Objective         | Function, risk-aversion or priority parameters and rationale.                             |
| Alternatives      | Current, recommended, minimum-change and approved comparator.                             |
| Scenarios         | Base, downside, severe downside and client-specific shocks.                               |
| Output            | Allocation, trades, expected characteristics, binding constraints, costs and limitations. |
| Reproduction      | Code or engine version, timestamp, input hash and independent result.                     |

**APPENDIX D. MINIMUM MONITORING DASHBOARD**

| Metric                      | Definition                                                        |
|-----------------------------|-------------------------------------------------------------------|
| Active relationships        | Unique client or family records within approved scope.            |
| Current truth records       | Records with every critical field within approved staleness.      |
| Reconciliation exceptions   | Open data differences by severity and age.                        |
| Binding-constraint breaches | Proposals or actions that violated an approved binding rule.      |
| Accepted packets            | Packets passing every required gate.                              |
| First-pass acceptance       | Accepted without material correction divided by reviewed packets. |

| Metric                             | Definition                                                                  |
|------------------------------------|-----------------------------------------------------------------------------|
| Material corrections               | Changes affecting suitability, allocation, cost, conflict or client action. |
| Overrides                          | Rate and reason by adviser, client segment and system version.              |
| Complaints and incidents           | Confirmed outcomes under approved taxonomies.                               |
| Cost per accepted packet           | Full controlled cost divided by accepted packets.                           |
| Attributed revenue or loss outcome | Approved observed value; USD 0 until evidence is approved.                  |

## APPENDIX E. RED-FLAG REGISTER

| Red flag                                                 | Required response                                               |
|----------------------------------------------------------|-----------------------------------------------------------------|
| Wrong client, entity or account                          | Stop, contain, investigate access and notify required owners.   |
| Missing or stale critical fact                           | Pause affected advice or apply approved conservative policy.    |
| Tolerance and capacity conflict                          | Adviser discussion and recorded resolution.                     |
| Unreconciled holding or liability                        | Resolve source or state exclusion and consequence.              |
| Binding constraint breach                                | Reject proposal and diagnose policy or integration failure.     |
| Large allocation instability                             | Review assumptions, method and strategic ranges.                |
| Unsupported material explanation                         | Reject communication and correct retrieval or generation.       |
| Product outside eligible universe                        | Reject and investigate retrieval, policy and approval controls. |
| Unrecorded adviser override                              | Pause action and complete decision record.                      |
| Provider or model change                                 | Regression test and approval before use.                        |
| Client data in unauthorised context                      | Security and privacy incident process.                          |
| Claimed revenue, saving or avoided loss without approval | Set attributed value to USD 0 and correct reporting.            |