

A new model of Automation of Intelligence¹ at Scale— transparent, collaborative, and anchored in human judgment.

Introduction: The Shift Underway

Asset managers have long deployed technology—from algorithmic trading systems to risk models—to automate calculations, monitor risk, and scale reporting.

Now a new type of technology — **AI agents** — is taking this evolution a step further.

Instead of simply running calculations or generating reports, these systems can **think, reason, and collaborate** with professionals in real time. They don't replace human judgment — they amplify it. Imagine working alongside a digital partner that understands your goals, tracks markets continuously, and helps turn insights into action — instantly and transparently.

This is what we call **Automation of Intelligence at Scale**: technology that extends human expertise across every client relationship, portfolio, and decision.

Why it matters

Today, asset managers and advisors face growing challenges:

- Clients expect faster, more personalized service.

- Regulations demand greater transparency.
- Teams must handle increasing dynamic risk and complexity — with limited time.

AI agents address those challenges by becoming **digital collaborators** — intelligent assistants that handle research, monitoring, reporting, and client insights — so professionals can focus on what truly matters: **judgment, strategy, and relationships**.

Early results are striking:

- Teams can deliver **3–5× more personalized portfolios**,
- **Client retention** can rise by up to **40%**,
- And advisors can engage more frequently and meaningfully, supported by intelligent automation that works around the clock.

. Beyond productivity, agents enhance **relationship quality, transparency, and institutional resilience**.

This white paper explains **what agentic AI is, how automated intelligence changes workflows end-to-end, and why it introduces transparency and performance rather than risk**.


¹ At Finbuddygroup.com, AI stand for **Automation of Intelligence at Scale** — the fusion of embedding reasoning, contextual awareness, and strategic intent redefines how wealth management operates at scale.

1. From Automation to Collaboration

Traditional automation followed rules. It worked well for repetitive tasks — but it couldn't adapt, reason, or understand context.

AI agents represent a new paradigm.


They interpret objectives, plan actions, and adapt based on feedback or conditions. They can automate



intelligence — *that is the perception and reasoning processes that connect, risk, data and strategy into informed, goal-directed workflows.*

Where traditional AI produced predictions, agents deliver **continuous collaboration** — interpreting a goal such as *‘limit exposure to cyclical equities while maintaining beta neutrality’* or *‘modify portfolio weights according to new guidelines, compliant with set risk limits and client preference’*. They break down these goals, executing single tasks, and explaining results in natural language.

This evolution marks a shift from *using tools* to *working with digital counterparts*.



Discretionary workflows are now progressively automated. Furthermore, the goal definition has been progressively shifted towards more high level specifications, orchestrating workflows of different agents.

Portfolio managers and analysts remain the ultimate decision-makers, but they now collaborate with intelligent assistants that handle procedural and analytical complexity.

The objective is amplification, not replacement — technology extending the reach of human expertise².

² <https://www.cyberark.com>

2. Defining AI Agents

AI agents are autonomous software systems that extend human intelligence across the full investment process — from data gathering and analysis to execution and reporting — under continuous human oversight. They have evolved from simple chat-bots. The World Economic Forum notes³ that agentic AI agents “perceive, reason, act and learn”. They function as digital “co-worker” on investment teams. They can:

- **Perceive:** Agents continuously integrate market data, client portfolios, and research sources — detecting relevant shifts such as rising inflation signals or sector rotations, allowing speed and precision in fragmented data.
- **Reason** They interpret this information in context, combining firm knowledge and policy constraints to propose actionable insights — for example, suggesting a defensive portfolio tilt. They increase productivity and consistency.
- **Act:** Execute decisions and operational tasks — orchestrating actions, tools, triggering simulations, reducing operational cost and latency.
- **Communicate:** Every action is explained in natural language, allowing advisors to ask “*Why was this change proposed?*” and

receive a transparent, auditable answer. They ensure compliance and client trust.

Architecture Overview


- **Semantic Data Mesh:** Connects market data, client portfolios, internal research, CRM systems, and compliance frameworks into a single contextual data fabric.
*Outcome: **Creates a unified, intelligent data "memory."** It moves beyond simple connections to establish **meaning and context** for every data point*
- **Intelligence Core — Adaptive Reasoning Engine:** Combines multiple large-language models. Finbuddy refined the core to include proprietary data, core financial knowledge, risk management, compliance.
Outcome: Reliability of results and investment specific Intelligence anchored in data and expertise.
- **Agents:** Agents reason across objectives, risk limits, and client goals, turning instructions into executable strategies.
Outcome: continuous, explainable decision support aligned with firm policy.
- **Interaction Layer:** Professionals interact via chat or visual interface. Every recommendation is transparent, traceable, and editable.
Outcome: governance and oversight

³ <https://www.weforum.org/>

stay with the human team while execution scales automatically.

- **High Availability Fabric:** Finbuddy is built on trading-grade infrastructure, guaranteeing reliability, speed, and fault-tolerant agent execution in production.

*Outcome: assurance of consistency and **uninterrupted and reliable operation.***



In the new architecture, humans define intent — goals, constraints, and strategic context.

Each agent operates autonomously within clear boundaries yet remains continuously open to supervision and adjustment. Over time, agents learn from feedback, refining workflows and surfacing new insights as conditions evolve.

The Measurable Shift

Independent research underscores the scale of the transformation.

KPMG observes that agentic AI in wealth management frees advisors from routine monitoring, allowing greater focus on

strategy and client engagement.

McKinsey estimates that, combined with generative AI, agentic systems can affect 25–40 % of an asset manager’s cost base. Firms deploying them report 50–70 % time savings on data-driven tasks, enabling the same teams to expand client coverage 3× to 5× — with client touchpoints increasing monthly and retention improving by around 40 %.

Wealth Agents

FinBuddy’s agentic architecture is composed of specialised digital entities. In the following few agents are briefly introduced.

Analytical agents — turning data into strategy

- **Research Agent** interprets key disclosures, generates structured insights. For private markets it can review hundreds of pages for due diligence and prepare client ready evaluation documents. It can automate weeks of work and revisions in minutes.
- **Risk Agent** It tackles hard to identify risks like Central Banks tone shift, emerging themes. Transforms high level risk description into deliverable constraints and portfolios. Generates dynamic, qualitative stress models in near real time.

Portfolio Agents — turning strategy into execution

- **Portfolio Agent** generates cross asset portfolios compliant to client mandate and risk limits. Process high level specifications into concrete portfolios.
- **Compliance Agent** Verifies every adjustment against regulatory and internal frameworks, ensuring automated governance at scale.

Client Agents — turning intelligence into engagement

- **Client Agent:** From gathering requirements to summarises key changes in the strategy in plain language. It can generate ad-hoc reports and can answer directly on most immediate questions from clients.
- **Advisory Agent** understands client requirements, manages ongoing communication, and maintains relationships through personalized insights and timely follow-ups.
- **Planning Agent** supports comprehensive financial and wealth planning, integrating investment, real estate, and long-term objectives into coherent, adaptive strategies.

3. End-to-End Intelligence: The Agentic Workflow

Today's workflows are fragmented: research in one system, risk in another, operations in a third.

Today's asset management stacks are shifting from isolated tools toward **multi-agent AI systems**. New cutting edge architecture lets specialized AI agents work together. As IBM⁴ explains, agentic workflows break complex processes into adaptive steps: autonomous agents “reason, plan and use tools to execute complex tasks efficiently.” Agentic design unifies these elements through **cooperating digital specialists** that share context.

In asset management, the model operates **end-to-end** — spanning client specification and onboarding, tailored service design, portfolio management, and ongoing reporting with continuous client engagement.

This ecosystem transforms linear processes into **adaptive collaboration** and feedback. Humans remain the architects of intent; agents handle orchestration and execution. The result is faster insight, cleaner communication, and no operational gaps.

⁴ <http://ibm.com>

4. Use cases in wealth management

The following examples highlight how intelligent agents create faster insight, personalized service, and scalable efficiency. — not an exhaustive list.

Adaptive Risk & Research Intelligence

- **Adaptive risk models**

Traditional factor risk models rely on predefined, static style factors which may fail to capture emerging sources of risk as market conditions evolve. Agentic systems add adaptability, by analysing and reasoning on new emerging risk such as inflation, supply-chain, geopolitical exposure, tariff sensitivity, or geopolitical dependencies. These signals can be integrated dynamically into the firm's risk model, — enhancing foresight and deepening client engagement.

- **Adaptive Macro Research Synthesis**

Agents can reason across macroeconomic releases, policy statements, central-bank communications, to extract early signals that shape portfolio positioning. For example, they can detect a rising focus on “persistent inflation” in Federal Reserve speeches or “fiscal discipline” in European policy documents — indicators of shifting regimes. By connecting these signals to market variables such as yield-curve movements or sector sensitivities, the

agent translates macro narratives into actionable portfolio insights.

Portfolio Design & Thematic Innovation

- **Thematic portfolios**

Agents can design and manage thematic portfolios end-to-end — from identifying ideas, market narratives to construct fully weighted allocations. The process is a transparent, fully documented and client-ready, auditable process, transforming a new idea into a fully implemented portfolio in a fraction of the time.

- **Client-Aware Portfolio Intelligence**

When risk or opportunity is identified for a client, the agent can suggest tactical tilts — such as increasing infrastructure exposure during fiscal-expansion phases, while dynamically managing duration as rate expectations evolve. They also generate plain-language rationales, cutting review cycles from days to minutes.

Client Advisory & Engagement

- **Advisory:** Agents assist advisors in capturing client objectives, design tailored financial and wealth plans, and maintaining dialogue through proactive insights and automated follow-ups. Every interaction feels personal, yet is powered by transparent, data-driven intelligence.

Private Markets & Due Diligence

Agents can automate intelligence for due diligence reviewing hundreds of pages of legal documents, financial statements and operational data. They summarise key

exposures, highlighting sensitivities, generate scenarios and produce client-ready reports. They compare evaluation with historical funding rounds and cooperate with other agents for risk and compliance — accelerating complex reviews while maintaining rigor.

5. A New Interaction Model: Conversing with Intelligence

The familiar dashboard which was a source of inefficiency both in the learning curve and the custom design, is evolving into a **conversation interface**. Text is the new interface. We take a step further, Adopting a dynamic hybrid approach of text and GUI creation which is the core of our workflow generation. Our platform offers high-level agents that can be customised for specific task or orchestrated in custom agentic workflow and run in a high availability environment, from on demand to long running jobs, to complex triggering execution. Workflows lifespans range from minutes to years.

Professionals can now ask high level instructions:

“Show me factor drift in European equities since last rebalance.”


“Run a 20% lower-duration simulation and prepare a client summary.”

Agents interpret such requests, reason across data, and produce immediate

answers — or seek clarification when uncertain.

Multiple agents can coordinate: *“Ask the Research Agent to include inflation-sensitive sectors.”*

Complex rules of triggering can be configured within the Finbuddy framework, which allows to run risk checks, internal validations, client relationships and interaction and more.



The experience feels less like querying a system and more like **collaborating with a colleague**. This collaboration model redefines efficiency. Instead of humans manually steering every process, professionals set intent and boundaries while agents execute, learn, and report back. Like any partnership, agentic collaboration demands a period of adaptation, as professionals adjust to and learn new ways of interacting with intelligent systems. It is a shift from *human-in-the-loop* to *human-above-the-loop* — where oversight becomes strategic rather than operational. The result is faster execution, consistent governance, and a structure where intelligence scales without diluting **accountability**. This conversational model reduces


friction, encourages transparency, and ensures that decision-making remains intuitive and explainable.

Humans stay in command; technology simply extends their reach.

6. A Magnifying Glass, Not a Black Box

Governance, Transparency, and the Human Loop

Finance demands explainability.



Agentic platforms are therefore designed

not as opaque engines, but as **transparent frameworks of reasoning**.

Each agent's logic, data sources, and contextual assumptions are **visible and traceable**.

Our solution is not a black box, but a magnifying glass. Professionals can inspect *why* a portfolio change was proposed, *which constraints* were applied, and *what confidence level* guided the result.

Transparency isn't just good practice; it's regulatory expectation.

The FCA and the UK government's AI strategy explicitly call for "*appropriate transparency and explainability*" in automated systems. Similarly, UK GDPR requires firms to inform clients about any automated profiling or AI-driven decisions and to give "meaningful information about the logic involved". These rules reinforce advisors' long-standing fiduciary obligation to keep clients informed.

A recent Bank of England survey⁵ found **81% of financial firms** using AI now employ explainability tools, and **84%** have named a senior accountability lead for AI projects. Transparency and human governance are rapidly becoming industry standards.

Built-In Safeguards

- **Human approval:** Any material portfolio change proposed by the AI must be reviewed and signed off by an advisor. This ensures AI output is

⁵ bankofengland.co.uk

vetted against human judgment and client mandates.

- **Complete audit trails and version control:** The system logs every decision with full provenance — which data sources, model version and rules were used. Under FCA rules (SYSC 4.1.1R), firms must have clear, transparent controls on their information systems⁶, and a complete audit trail meets that requirement directly. All changes can be rolled back or inspected later, so nothing is hidden.
- **Access boundaries:** Role-based permissions ensure only authorized staff can run or modify agents. Boundaries align with internal policies and regulatory requirements, maintaining data privacy and client confidentiality.
- **Explainable outputs:** Advisors can interrogate the model's reasoning in plain English. For example, the user might ask, "Why was stock X added to the portfolio?" prompts a natural-language explanation. Our platform also provides guidance on how to maximize transparency.

The philosophy

Agentic AI doesn't replace trust — it builds it.

It transforms automation into a transparent partner, one that expands

access to knowledge rather than concealing it.

Our platform acts as a **magnifying glass**, not a black box — opening the reasoning process instead of hiding it. They invite scrutiny and dialogue, enabling firms to meet governance and fiduciary obligations with greater confidence than before.



⁶ [fca.org.uk](https://www.fca.org.uk)

7. The Client-Centric Future of Wealth Management at Scale

The wealth manager of the near future will revolve around one principle: **every client experience should feel personally crafted — yet instantly delivered.** The platform enables wealth managers to understand and anticipate each client's evolving goals, translating them into portfolios that adjust dynamically as markets and personal circumstances change.

Agentic platforms transform the traditional operating model by embedding reasoning and adaptability into every step of the client journey. From initial specification and portfolio design to tactical rebalancing and ongoing dialogue, agents coordinate tasks, share context, and surface insight continuously.

This new model redefines the client relationship:

Personalisation scales intelligently.

Agents translate client goals into actionable portfolio strategies, advising on evolving market conditions. What once required weeks of manual preparation — model updates, risk assessments, or report generation — can now occur in near real-time. The result: faster reviews and more responsive engagement.

Human expertise focuses on strategy, not process.

With agents managing data collection, simulation, and documentation, advisors dedicate their time to judgment, relationships, and creative portfolio construction. Automation amplifies rather than replaces it — enabling richer, more interactive client conversations.

Thematic and tactical portfolios

Agentic systems can detect emerging trends and construct thematic or tactical portfolios that reflect clients' evolving interests, fully documented and client-ready.

Efficiency and scalability redefine competitiveness.

By automating intelligence, firms achieve operational scale once limited to the largest institutions and the wealthiest clients. This new efficiency allows wealth managers to remain competitive while responding to ongoing fee pressures.

Intelligence becomes transparent and repeatable.

Agentic workflow turns automation into a transparent partner that delivers **hyper-personalised wealth management at scale**, where every client experience is both bespoke and continuously informed by intelligent insight.

8. Conclusion

AI agents redefine how investment organisations think, act, and interact. They introduce an **operating layer of transparent automated intelligence** — one that connects human expertise with continuous, proactive support.

The transformation is architectural and cultural:

- From **tools to teammates** – agents collaborate with professionals, not replace them.
- From **black boxes to magnifying glasses** – reasoning becomes inspectable, auditable, and open to dialogue.
- From **automation to partnership** — accountability turns efficiency into a collaborative outcome between people and intelligent systems.

The firms that embrace this change will not only work faster — they will **see more clearly**, understand their processes more deeply, and build trust through visibility.

Agentic AI doesn't diminish the human role; it **elevates** it.

The new efficiency enabled by agentic systems leads to a **3–5× increase in bespoke service capacity**, while advisors experience a **40% rise in client responsiveness and retention**. Firms

unlock unprecedented research capabilities. Beyond productivity, agents enhance **relationship quality, transparency, and institutional resilience**. The next generation of asset management will be defined by **people who collaborate fluently with intelligent machines**.

About Finbuddy

Finbuddy⁷⁷ is a fintech company offering the next generation agentic platform fundamentally changing the **economics** of wealth management.

We operationalises Institutional intelligence, elevating and scaling it across every function of the Firm.

Our platform delivers a suite of proprietary, highly specialised agents that can be easily tailored to your data, strategic view, and processes, creating decisive edge across portfolio construction, risk management, alpha generation, and client engagement. Hyper **customised portfolios** and enhanced client experiences and engagement — delivered **at scale**, without eroding margins — are now a practical reality.



⁷⁷ Finbuddygroup.com | For more in depth Q&A with our AI ambassador scan QR code or follow link <https://chatgpt.com/g/g-68e780d5e848819184e5fda11c266c62-finbuddy-ambassador>