# Practical AI for Asset Management in Financial Services

Archive text edition of the revised manuscript · 11 November 2026

Prepared by the Harbour Think Tank under the leadership of Adrian Dunkley. The original evidence review is dated 26 October 2026.

## About this paper

The Harbour Think Tank of the Caribbean AI Association covers the economy, work and enterprise. This paper covers asset management and is one of fourteen in our financial-services collection. We wrote it for portfolio managers, investment researchers, fund operations teams and investment committees. [1]

The proposals prepare evidence and check records. They authorise no automated investment recommendations, portfolio changes, valuation decisions or trading. Those activities stay with authorised people. Research support and operational checks sit apart from the investment decision. Adviser communication and market execution each have their own paper, and so does compliance.

This revision adds a page of build notes to each use case. The notes give a starter prompt, example tools, skills, a method and the limits for an AI agent.

Foreword

A portfolio manager must not have to choose between a second read of every page and trust in a summary that no one can check. I want an assistant that points to the filing and says what changed, and that leaves the open question in plain sight.

I see the case most clearly on difficult days. A company changes its reporting currency, or a price stops its updates. A mandate hides an exception in an attachment. On those days, faster preparation helps only if the team can still find the original records and explain the judgement that followed.

My advice is to start with research records and reconciliations. Require evidence before you give the assistant a larger role in the investment process.

Adrian Dunkley, President, Caribbean AI Association

## Acknowledgements and terms

The Harbour Think Tank of the Caribbean AI Association prepared this paper, and President Adrian Dunkley led the work. The sources give professional and regional context. The workflows, the screening scores, the prompts and the agent levels are our proposals for local tests. We report no measured results from a named financial institution, and we claim no regulatory or institutional endorsement.

An artificial intelligence (AI) assistant finds and compares information, and it can draft text. An agent can also do configured steps with connected tools. In these use cases, neither can submit orders or change portfolio records. A mandate sets the investment objectives and constraints of the client. A benchmark is an agreed reference for the assessment of performance. A corporate action is an issuer event, such as a dividend or a share split, that can affect a security or its holder.

Reconciliation compares records that different parties keep independently, and explains the differences. A price observation needs its instrument, currency, date and source. A scenario tests stated assumptions about possible conditions, and it is not a forecast. The approved accounting and valuation policies of the firm govern the work, with its investment policy.

A prompt is the written instruction that a person gives to an assistant. The word skill has two meanings in this paper. A staff skill is an ability that a person learns. An agent skill is a saved set of instructions that an assistant loads for one task. [7]

### Approved source

Draft or retrieval

### Human review

Accepted output

### The operating boundary

The assistant prepares research and operational evidence. Investment decisions, valuations, portfolio changes and every trading instruction stay with authorised people.

## Contents

Section | Page

Executive summary | 5

1 The regional business case | 6

2 What the evidence supports | 7

3 The enterprise lens | 8

Process design: BPMN swimlanes | 9

4 Ten practical use cases: how to read each use case | 10

Use cases 01 to 10, each with build notes | 11-30

Priority and responsibility tables | 31-32

5 The foundations lens | 33

6 From pilot to practice | 34

Appendix A Pilot decision canvas | 35

Appendix B Measure value | 36

References | 37-38

Read the summary and the priority tables first. Then select one use case and define a pilot. Complete the decision canvas with the people who do the work.

## Executive summary

Start with filing comparisons and research evidence packs, then add meeting-action records. You can test each of them on approved document snapshots, with no connection to trading systems or client systems.

Operational checks come next, where identifiers and holdings are reliable and source dates are correct. Keep the calculation engine apart from the explanation. A model that drafts a clear account of an approved result does not show that it can calculate that result or make the investment decision.

Select a use case that has an accountable owner and reference examples, with a failure test that the assistant can plausibly fail. Count the unsupported statements and the missed exceptions, and record the effort that stays after review and correction. A more persuasive memo is not evidence of a better investment. A shorter research cycle is not proof of higher returns.

Each use case now has build notes with a starter prompt and example tools. The notes also set a limit for AI agents. In the pilots, no agent acts without the approval of a named person.

### Four decisions for the sponsor

Select the bounded research or operations task. Confirm the permitted documents and licences. Name the reviewers. Approve the error threshold and the person who can stop the work, before the pilot uses holdings or mandates.

**Choose a bounded process**

Measure the baseline

**Test with review**

Decide from evidence

## 1 The regional business case

The Financial Stability Report 2024 of the Bank of Jamaica shows what carefully bounded sector data looks like. [2] Table A.2 reports regulatory capital against risk-weighted assets for twelve securities dealers that cover 70.0% of the market. The table describes that group of dealers. It says nothing about the returns of any investment portfolio.

The six December observations for 2019 to 2024 are 21.4%, 22.3%, 23.0%, 21.2%, 20.8% and 20.2%. Regulatory capital is the numerator, and risk-weighted assets are the denominator. [2] The series is historical. Do not read it as an assessment for 2026 or as the capital position of one firm. It is not an estimate for the Caribbean.

If an assistant omits the coverage note of the table, the evidence changes its meaning. Research memos make the same error when they treat a group result as an issuer result. They also make it when they put amounts in different currencies side by side.

For a Caribbean portfolio, confirm the local issuer identifiers, the filing periods, the currencies and the access rights before you compare documents. Where no authoritative price or disclosure is available, keep the exception open. No one knows if an international data feed covers every local instrument until someone checks.

**Jamaica: reported securities-dealer capital ratio**

Dec 2019: 21.4%

Dec 2020: 22.3%

Dec 2021: 23.0%

Dec 2022: 21.2%

Dec 2023: 20.8%

Dec 2024: 20.2%

Bank of Jamaica, Financial Stability Report 2024, Table A.2, page 68. [2] Ratio of regulatory capital to risk-weighted assets. Twelve securities dealers that represent 70.0% of the market.

## 2 What the evidence supports

In March 2025 the International Organization of Securities Commissions (IOSCO) published a consultation on the use of AI in capital markets, which covers investment research. [3] It names risks from unreliable outputs and from model behaviour that changes, and risks from poor data. It is a consultation report. It sets no final rule, and it gives no evidence that pilots such as ours improve returns.

In November 2025 CFA Institute, the body for the investment profession, announced a publication on AI in asset management. [4] The announcement stresses professional judgement and governance, and evaluation. We cite the announcement for context only. We did not reproduce or validate the technical methods of the book.

In 2020 the Financial Services Commission of Jamaica issued market conduct guidelines for securities dealers. [5] The guidelines cover confidentiality and communications, and suitability among other matters. A workflow must fit the licence and the approved policies of the firm, and the requirements that apply. Those duties stay with the firm, whatever the assistant drafts.

These sources support controlled tests in research and operations. None of them promises a percentage saving or an excess return, and we have no regional estimate. Evaluate the workflow on the documents and authorised data of your own firm before you widen access.

**Test the proposed workflow**

Test source accuracy and missed exceptions, and record the review effort. Whether better research records lead to better returns is an open question. No attractive memo and no document pilot will settle it.

## 3 The enterprise lens

Decide where research ends and a decision starts. The assistant collects evidence for an investment thesis, and the analyst must judge if the evidence is credible and applicable. A committee needs the contrary evidence and the open assumptions. A sharper pitch for its existing view helps no one.

Use approved calculations for returns and exposures, and for reconciliations. Record the data version and the method before you ask for an explanation. When the figures change, generate the narrative again and review it again. Do not put one new number into an explanation that the assistant wrote for the old data.

Keep public research apart from confidential holdings and client information. A document that the firm has a licence to read can have no licence for external model processing, so confirm with the data owner and the supplier. No assistant can hold trading credentials. That rule includes an assistant that only prepares the portfolio report. The prompts in the build notes are start points that each firm must test.

**Request and source**

Assistant output

**Reviewer correction**

Accepted result

Research preparation and numerical calculation are separate steps, and investment approval is a third step. Each step needs its own recorded source and reviewer before the firm accepts the final output.

## Review changes against both issuer filings

Business Process Model and Notation (BPMN) swimlanes. Proposed workflow for use case 01.

**Research analyst**

AI drafting assistant

**Investment research lead**

Yes

No

**New filing available**

Supply complete versions and reporting basis

**Build a source-linked change table**

Trace material changes to both filings

Comparison sound?

**Accept reviewed changes and limitations**

Comparison available

**Correct mismatches and missing evidence**

Save the reviewed change table

**Issuer filing comparison**

The assistant must not change a rating or make an investment recommendation. It must not invent performance that the issuer did not disclose.

How to read the model. Rounded rectangle: task. X diamond: one decision path. Thin circle: start. Thick circle: end. Solid arrows show the sequence in one process. Lanes show who is responsible.

Notation: Object Management Group, BPMN 2.0.2. https://www.omg.org/spec/BPMN/2.0.2

# 4 Ten practical use cases

How to read each use case

Each use case has two pages. The first page gives the workflow, the minimum data, the pilot test and the accountable role. The second page gives build notes in five parts.

Starter prompt. Copy the prompt into the approved assistant, and replace the text in square brackets. Each prompt gives the assistant a role, its sources, one task, the rules and the output format.

Tools. The notes name example products. The names are examples and not endorsements. Vendors change product names frequently. Use only a service that the firm approved.

Skills. A staff skill is an ability that the team must build. An agent skill is an instruction file that the assistant loads for the task. Agent Skills is an open format, and many assistants can read it. [7]

Method. The notes name an established method that gives the work a structure.

Agent limits. The notes say what an AI agent can do, and where it must stop.

### Four agent levels

Level 0, assist. The assistant answers one request. A person does every step.

Level 1, prepare. The agent collects approved inputs and prepares a draft. It changes no system and sends no message.

Level 2, act with approval. The agent proposes an action. A named person approves each action before it runs.

Level 3, act alone. The agent acts without approval inside fixed limits.

Each pilot in this paper starts at Level 0 or Level 1. A move to a higher level needs a new approval after the pilot. No pilot is at Level 3.

# 01 Compare successive issuer filings

Earlier issuer filing

### Later issuer filing

Trace changes and restatements

### Research lead checks

Source-linked change table

An analyst supplies two authorised versions of the financial report of an issuer and gets a change table. Each entry has a link to its page. An entry covers the reporting scope, the accounting notes, the management discussion or the disclosed risks. The table shows if a fact changed, or only the description of the fact.

Flag each restatement and compare on the restated basis, or explain why a comparison is not possible. The work needs complete filings, with the publication dates, the reporting periods, the currency, the group boundaries and the page references.

Pilot on sections of different types. One section has no change, and one has a material amendment. In a third section, a table moved. A second analyst traces every material entry to the two versions. The analyst counts the missed changes and the incorrect matches, and the interpretations that the filings do not support. Stop the test if the table reports a currency change as growth, or if the assistant rebuilds a missing paragraph and does not say so.

The investment research lead is accountable. Analysts judge materiality and its consequences. The assistant cannot upgrade a recommendation or infer undisclosed performance from a change of wording. Its table is an aid to the reader, and it reaches no investment conclusion.

### First working session

Select two issuer filings. Mark one material change that you already know about.

## 01 Build notes: filing changes

Starter prompt

"You compare two versions of an issuer filing. Use only the two attached documents [issuer, periods]. Make a change table with five columns: section, earlier text or figure, later text or figure, page references and change type. Use these change types only: Fact changed, Description changed, Restated, Moved and Removed. State the currency and the reporting basis of each figure. If a passage is missing from one version, write 'Not found' and do not rebuild it. Do not say what a change means for the investment. Do not change a rating."

Tools

Filings: complete documents from the issuer or the exchange, saved as dated snapshots in SharePoint.

Comparison: Claude or Microsoft 365 Copilot with the two versions attached, and the Compare function in Word as a cross-check.

Trace record: an Excel sheet where the second analyst records each traced entry.

Skills

Staff skill: read a restatement note and compare on the restated basis.

Agent skill: a filing-comparison skill with the change types and the rule "not found is not rebuilt".

Method

A two-analyst trace: the second analyst checks every material entry against the two source pages. Report missed changes and incorrect matches separately.

### Agent limits

Level 1. An agent can prepare the change table when a new filing arrives in the approved folder. It must not send the table outside the research team. It must not change a research rating or a model input.

## 02 Assemble a balanced research evidence pack

Defined investment question

### Approved filings and notes

Supporting and contrary facts

### Research manager challenges

Balanced evidence pack

Start from a defined investment question. The assistant sorts approved filings and research notes into four groups: supporting evidence, contrary evidence, unresolved assumptions and questions for the issuer. Every factual entry has its source and date.

If the collection contains no contrary view, say that this is a limit of the collection. Do not claim that no contrary evidence exists. Management claims keep the label "claim" until someone verifies them independently. The inputs are the question, the authorised collection, the instrument identity and the evidence cut-off date.

Pilot on a historical question where the team already knows that a disclosure contradicts the thesis. Ask if a second analyst can rebuild the evidence with no trust in the summary. Count the unsupported claims and the omitted contradictions, and record the retrieval time. Stop the pilot if the pack states an assumption as a fact or hides a conflict.

The portfolio research manager is accountable. The investment professional forms the thesis, and committee members challenge it. The assistant must not estimate a target price or rank securities for purchase, and it must not issue a recommendation. Confidential research must stay inside its approved access boundary.

### First working session

Write one investment question. Include a source that challenges the first view.

## 02 Build notes: research evidence

Starter prompt

"You assemble an evidence pack for this investment question: [question]. Use only the attached collection, with evidence dated before [date]. Make four lists: Supporting evidence, Contrary evidence, Unresolved assumptions and Questions for the issuer. Give the document and the page for each entry, with its date. Label each statement from management as 'Management claim'. If the collection has no contrary evidence, say that the collection has none. Do not form a thesis. Do not estimate a price or a return. Do not rank or recommend a security."

Tools

Collection: approved filings and research notes in a restricted SharePoint library or a Claude Project.

Pack: Claude or Microsoft 365 Copilot, with the collection attached.

Market data: the licensed terminal of the firm, for example Bloomberg, used by the analyst and not by the assistant.

Skills

Staff skill: write a bounded investment question with an evidence cut-off date.

Agent skill: an evidence-pack skill with the four lists and the "Management claim" label.

Method

An evidence table with a mandatory column for contrary evidence, and a devil's advocate review by the research manager. A second analyst rebuilds a sample of entries from the sources.

#### Agent limits

Level 1. An agent can collect entries from the approved collection for an analyst. It must not read confidential holdings. It must not write to a research model or send a pack to a committee.

## 03 Translate mandates into a review checklist

Current mandate and amendments

#### Extract clauses and exceptions

Meaning unambiguous?

#### Checklist for specialist review

Interpretation question

Yes

No

Operations staff supply an approved investment mandate with its amendments. The assistant drafts a checklist of permitted instruments, concentration limits, reporting obligations and exceptions. Each line cites its clause. Ambiguous wording becomes a question.

Use only signed, current documents, with their effective dates and approved interpretations, and the identity of the portfolio. A draft amendment is never effective. When an exception seems to be missing or a definition is unclear, the question must go to the authorised specialist. The assistant must not select an interpretation.

Test the workflow on a mandate that has a superseded limit and a conditional exception. Legal and compliance staff compare the checklist with their own reference, together with portfolio staff. They record the omitted conditions and the incorrect effective dates, and the review time. Stop the pilot if the assistant invents a restriction, or treats an ambiguous permission as settled.

The investment operations manager is accountable. Authorised specialists approve the interpretation, and existing controls enforce the restrictions. The checklist prepares for human review and never replaces the operative agreement. The assistant has no power to change a mandate or give a waiver, and it has no power to release a

trade.

### First working session

Take a signed mandate with its current amendments. Find one ambiguous clause.

## 03 Build notes: mandate checklist

Starter prompt

"You draft a review checklist from an investment mandate. Use only the attached signed mandate and its signed amendments. Make one line for each permitted instrument, concentration limit, reporting obligation and exception. Give the clause number and the effective date for each line. If an amendment replaces a clause, use the amendment and name the superseded clause. If the wording is ambiguous, write a question for the specialist and give no interpretation. Do not use a draft amendment. Do not add a restriction that is not in the documents."

Tools

Documents: the signed mandate and amendments, in a restricted SharePoint library.

Checklist draft: Claude or Microsoft 365 Copilot, with the documents attached.

Enforcement: the existing compliance rules in the order or portfolio system, changed only by authorised staff.

Skills

Staff skill: read an amendment against the clause that it replaces.

Agent skill: a mandate-checklist skill with the four line types and the rule "ambiguous wording becomes a question".

Method

A traceability matrix: each checklist line maps to one clause and one effective date. Specialist reviewers compare it with their own reference.

### Agent limits

Level 1. An agent can draft the checklist when a signed amendment arrives. It must not change a compliance rule in any system. It must not give a waiver or release a trade.

## 04 Reconcile holdings across approved records

Controlled holdings extract

### Independent custodian record

Fixed-rule position matching

### Fund operations investigates

Reviewed reconciliation issues

Fixed rules do the match. They compare a controlled holdings extract with an independent custodian statement. The assistant then explains each difference in an identifier or a quantity, and each difference in a date. Its exception list connects each issue to the two source records.

The minimum data is the security identifier, the account scope, the position date, the quantity, the currency and the settlement treatment. Keep cash movements apart from security positions. Where one system has an unverified local identifier, ask for a mapping. Never match on a similar issuer name.

Use fictional records that contain an unsettled trade and a share split, with a missing position. Operations staff calculate every difference again, independently. Record the discrepancies that the workflow did not find and the false matches, and the time to resolve each question about a record. Stop immediately if the assistant overwrites a holding. Stop immediately if it decides, without evidence, which system is correct.

The fund operations lead is accountable. People investigate and approve corrections through the established process. The assistant sees only the fields that the comparison needs. It cannot book adjustments or instruct the

custodian. It cannot declare the reconciliation complete while an exception is open.

### First working session

Prepare fictional holdings with a settlement difference and a share split.

## 04 Build notes: holdings reconciliation

Starter prompt

"You explain reconciliation exceptions. A fixed-rule match already compared the holdings extract with the custodian statement. Use only the attached exception list and the two source records. For each exception, show the two values and name the field that is different. Suggest a possible cause, and label it 'To verify'. Keep cash and securities in separate sections. If an identifier has no verified mapping, write 'Mapping needed'. Do not match on an issuer name. Do not say which system is correct. Do not propose an adjustment entry."

Tools

Match: fixed rules in Excel Power Query, or the reconciliation module of the fund accounting system.

Explanation: Claude or Microsoft 365 Copilot, with the exception list only.

Break log: a SharePoint list where operations staff record the cause and the approval for each correction.

Skills

Staff skill: calculate a difference independently and find its cause.

Agent skill: a reconciliation-notes skill with the exception format and the rule "no match on names".

Method

Two-way reconciliation with fixed matching rules and a break log. Classify each break by its cause, and count the breaks that the rules did not find.

### Agent limits

Level 1. An agent can prepare the exception notes after each scheduled match. It must not book an adjustment or message the custodian. It must not mark a reconciliation as complete.

## 05 Prepare a corporate-action verification queue

Official issuer notices

### Authorised security list

Match events and dates

### Operations verifies

Corporate-action review queue

The assistant compares authorised issuer notices with the security list of the portfolio. It drafts a queue of events that can be applicable. Each event has its announcement source, its event type, its key dates and questions for operations staff.

The assistant needs the official notice, the instrument identifier, the holdings date and the approved event conventions of the firm. The announcement date, the record date, the ex-dividend date and the payment date stay distinct where they apply. When an issuer amends a notice, keep the version history and mark the changed instruction.

Historical notices make a fair test if they include a correction and an event that needs an election. Include also an instrument that the portfolio does not hold. Count the missed notices and the incorrect date labels, and count the false portfolio matches. Stop the pilot if the assistant invents an election deadline, or treats an announcement as proof that cash arrived.

Coverage of local markets can be incomplete, and we did not measure how complete it is. Where coverage is missing, record a monitoring gap. An empty queue for those securities can mean only that no one looked.

The corporate-actions operations lead is accountable. Staff verify eligibility, and authorised people decide and submit elections. The assistant never replies to an issuer or changes a cash forecast.

#### First working session

Select an original issuer notice and its correction. Confirm the instrument identifier.

## 05 Build notes: corporate actions

Starter prompt

"You prepare a review queue for corporate actions. Use only the attached issuer notices and the security list dated [date]. For each notice, give the instrument identifier, the event type and the source. Show each date with its label: announcement, record, ex-dividend or payment. If a notice amends an earlier notice, show the two versions and mark the change. Say if the portfolio holds the instrument on the holdings date. Write the questions that operations staff must answer. Do not invent a deadline. Do not say that the portfolio is eligible or that cash arrived."

Tools

Notices: official issuer and exchange notices, for example from the Jamaica Stock Exchange, saved with their dates.

Queue draft: Claude or Microsoft 365 Copilot, with the notices and the security list attached.

Queue: a SharePoint list with an owner and a status for each event.

Skills

Staff skill: tell the four event dates apart and verify eligibility.

Agent skill: a corporate-action skill with the date labels and the amendment rule.

Method

An event calendar with a version history for each notice. Record a monitoring gap for each market or security that no source covers.

#### Agent limits

Level 1. An agent can check the notice sources each day and add events to the queue. It must not submit an election or reply to an issuer. It must not change a cash forecast.

## 06 Flag price observations that need review

Price and source timestamp

#### Apply approved freshness rules

Observation needs review?

#### Valuation review packet

No exception; retain source

Yes

No

An approved rule set flags valuation inputs that have a missing date or an inconsistent currency. It also flags observations that are older than the permitted threshold of the firm. The assistant describes each exception, and it collects the authorised source evidence that the valuation team needs.

Each observation needs the instrument identity, the price, the time, the source and the currency. Read each one against the approved freshness rules. Keep a last-traded price apart from an independently assessed valuation. Thin trading or a market holiday needs context. A generated price is never an acceptable replacement.

Put a currency error and a stale source into fictional data, with a quote that is the same as before but valid. Review the false alerts and the missed material exceptions. Measure the time to recover each source, and count the issues that the workflow classified correctly. Stop the pilot if the assistant invents a price or changes a valuation. Stop the

pilot if it reads an absence of trades as proof that economic value stayed the same.

The valuation control lead is accountable. The authorised valuation process decides which input and which method the firm accepts. The job of the assistant ends at evidence. Its report keeps the uncertainty visible and puts no estimate in its place.

### First working session

Prepare one stale price and one currency mismatch, with one quote that is the same as before but valid.

## 06 Build notes: price exceptions

Starter prompt

"You describe price exceptions for a valuation team. A rule set already flagged the observations. Use only the attached flag list and source records. For each flag, show the instrument, the price, the time, the source and the currency. Name the rule that the observation failed. Say if the price is a last trade or an assessed valuation. Add market context only from the attached calendar. Do not suggest a replacement price. Do not change a valuation. Do not say that value is the same because no trade occurred."

Tools

Rules: approved freshness and tolerance rules in Excel, or in the pricing system of the firm.

Description: Claude or Microsoft 365 Copilot, with the flag list and the source records.

Evidence packet: a folder for each exception, with the source record and the rule.

Skills

Staff skill: tell a stale price from a valid price that did not change.

Agent skill: a price-exception skill with the rule names and the forbidden estimates.

Method

Exception-based review with staleness and tolerance rules. Test with seeded data, and report the false alerts and the missed exceptions separately.

### Agent limits

Level 1. An agent can prepare the evidence packet for each flagged observation. It must not write a price to any system. It must not close an exception or approve a valuation.

## 07 Explain approved performance calculations

Checked performance outputs

### Reconcile result and attribution

Figures reconcile?

### Draft commentary for approval

Return calculation exception

Yes

No

The assistant drafts commentary from checked portfolio and benchmark calculations. The draft keeps the measured result apart from the verified contributors. It keeps the two apart from interpretation, which an analyst must approve. Every material figure keeps its period and its currency, with its calculation basis.

The minimum data is the approved returns, the benchmark version, the cash-flow treatment, the fee basis and the authorised attribution output. Keep gross figures and net figures separate, and compare periods that match. If the attribution does not reconcile with the reported result, hold the explanation and send the exception back.

Pilot on a completed reporting period that had a fee change or a benchmark change. Check that the narrative states the comparison correctly, and that its causal claims stay inside what the calculations show. Count the transcription errors and the omitted basis changes, and count the reviewer corrections. Stop the pilot if the assistant calculates returns informally. Stop the pilot if a sentence turns past performance into a promise.

The performance measurement lead is accountable. Analysts verify the interpretation, and approved reviewers approve publication. The assistant must not select a flattering benchmark or hide a negative contributor. It must not distribute the report.

#### First working session

Select an approved reporting period. Confirm the fee basis and the currency, and the benchmark version.

## 07 Build notes: performance explanation

Starter prompt

"You draft performance commentary from checked calculations. Use only the attached performance and attribution outputs for [period]. Do no arithmetic. First check that the attribution total equals the reported result. If it does not, write 'Does not reconcile' and stop. State the result with its period, its currency, gross or net, and the benchmark version. List the verified contributors, positive and negative. Put interpretation in a separate section with the label 'For analyst approval'. Do not explain a result that the attribution does not show. Do not suggest that past performance will continue."

Tools

Calculations: the performance system of the firm, with attribution output.

Commentary draft: Claude or Microsoft 365 Copilot, with the checked outputs attached.

Publication: the existing approval workflow for client reports.

Skills

Staff skill: read an attribution report and state the calculation basis.

Agent skill: a performance-commentary skill with the reconcile check and the section order.

Method

Performance attribution from the approved system, for example the Brinson model. Use the Global Investment Performance Standards (GIPS) as the reference for a fair presentation basis.

#### Agent limits

Level 1. An agent can draft commentary when the checked outputs for a period are final. It must not calculate or change a return. It must not select a benchmark or send a report.

## 08 Organise liquidity-scenario assumptions

Portfolio exercise scope

#### Liquidity assumptions

Validated scenario outputs

#### Specialists challenge

Owned assumption register

Portfolio specialists supply an approved liquidity exercise and its data requirements. The assistant builds an assumption register from them. The register covers redemption requests, available cash, settlement timing and constraints, and each entry has an owner and evidence.

Sale timing and market depth keep an explicit label as assumptions. This use case has lower scores for feasibility and data readiness in our screening. A missing local trading observation must stay a visible gap. Never fill it from an unrelated market. The inputs are the exercise scope, the approved scenario, the portfolio snapshot and the

validated outputs.

Use a historical or fictional exercise that has a settlement delay and a disputed assumption. The register must show the issue and the person who decides it. Count the omitted assumptions and the stale sources, and count the disagreements that stay open. Stop if the assistant invents a liquidation price. Stop if it says that the portfolio can meet redemptions and no validated analysis supports that statement.

The portfolio liquidity lead is accountable. Specialists approve the scenarios, and validated tools do the calculations. The assistant starts no sales and changes no redemption terms. It certifies nothing about liquidity. A register makes the exercise easier to challenge, and the uncertainty is still there after the exercise.

#### First working session

List the assumptions in one approved liquidity exercise. Name the owner of each assumption.

## 08 Build notes: liquidity assumptions

Starter prompt

"You build an assumption register for a liquidity exercise. Use only the attached exercise documents and validated outputs. Make one row for each assumption: the assumption, its value, its source, its date and its owner. Label sale timing and market depth as 'Assumption' in every case. If a local observation is missing, write 'Gap' and do not use data from a different market. If specialists disagree, record the two positions and the person who decides. Do no calculations. Do not give a liquidation price. Do not say that the portfolio can meet redemptions."

Tools

Calculations: the validated liquidity model of the firm, in Excel or in the risk system.

Register: an Excel workbook or a SharePoint list, with an owner for each row.

Draft: Claude or Microsoft 365 Copilot, with the exercise documents attached.

Skills

Staff skill: state an assumption with its source and challenge it.

Agent skill: an assumption-register skill with the row format and the gap rule.

Method

A liquidity stress test with an assumption register and a sensitivity test on the disputed assumptions. Specialists challenge each assumption before they accept the outputs.

#### Agent limits

Level 0. Specialists run each exercise by hand, and the assistant has no connection to a risk system or a trading system. It must not start a sale or change a redemption term.

## 09 Track investment-committee evidence and actions

Approved papers and minutes

#### Current action register

Separate approval and execution

#### Committee secretary verifies

Evidence-linked action list

The assistant uses the approved papers and minutes, with the action register. It prepares a list of decisions, conditions, open questions and promised follow-ups, and each item has a link to its source. It keeps an item that the committee discussed apart from an item that the committee approved. It keeps the two apart from an item that someone executed.

The work needs the final minutes, the paper versions, the meeting date, the named action owners and the agreed due dates. A decision that is not in the record stays unconfirmed. Keep dissent and conditions as the minutes

record them. Never compress them into one favourable conclusion.

Take one completed committee cycle that had a deferred proposal and a conditional approval. Check that you can trace every action to its source and that evidence supports each current status. Count the invented approvals and the omitted conditions, and count each incorrect closure status. Stop the test if a draft minute becomes an executed instruction.

The investment committee secretary is accountable. Authorised people confirm the records. An action closes when the owner gives evidence of completion, and never on a statement alone. The assistant cannot send trading instructions or change committee decisions. It administers evidence and has no investment authority.

#### First working session

Reconstruct one committee cycle with a deferred decision and a conditional approval.

## 09 Build notes: committee actions

Starter prompt

"You prepare an action list for an investment committee secretary. Use only the attached committee documents and the action register. List each decision, condition, open question and follow-up, with its source page. Give each item one status: Discussed, Approved, Executed or Deferred. Mark an item 'Approved' only if the final minutes record the approval. Mark an item 'Executed' only if the register has evidence of completion. Keep each condition and each dissent in the words of the minutes. If a decision is not in the record, write 'Unconfirmed'. Do not write an instruction to trade."

Tools

Records: final minutes and papers in a restricted SharePoint library.

Action register: a SharePoint list or Microsoft Planner, with an owner and a due date for each action.

Draft: Microsoft 365 Copilot, or Claude with the committee documents attached.

Skills

Staff skill: tell an approval from a discussion, and close an action on evidence.

Agent skill: a committee-actions skill with the four statuses and the "Unconfirmed" rule.

Method

A decision log that records the conditions and the owner of each action, with evidence of completion. Trace a sample of actions to the minutes.

#### Agent limits

Level 1. An agent can prepare the action list after the secretary approves the final minutes. It must not close an action. It must not send an instruction to a trading desk or change a recorded decision.

## 10 Check an investment reporting pack for consistency

Draft reporting pack

#### Approved source tables

Flag figures and wording

#### Content owners approve fixes

Consistent reporting draft

In this use case, the assistant reads an approved reporting pack against its source tables and the previous version. It flags dates that conflict, portfolio names that do not match, inconsistent figures and narrative that does not agree with the checked results.

Use the reporting cut-off date, the final calculation outputs, the approved disclosures and the version history. Check the source permissions before any client-specific information goes in. A figure that still awaits approval

gets a dependency label. The value of the previous period never goes in its place.

Put an old benchmark and an incorrect currency label into a draft, with an unsupported positive claim. A reviewer compares the findings with these planted errors. Count the missed material inconsistencies and the false flags, and record the final review effort. Stop if a required disclosure disappears. Stop if the assistant invents a replacement result.

The investment reporting manager is accountable. Content owners approve corrections, and release goes through the existing publication process. The assistant sends no client report. A clean layout can still contain an unsupported investment claim, so the source evidence stays attached.

#### First working session

Put an old benchmark and an incorrect currency label into a draft report, with an unsupported claim.

## 10 Build notes: report consistency

Starter prompt

"You check a draft reporting pack for consistency. Use only the attached draft and its reference documents. Make a table with four columns: page, item in the draft, value in the source and problem type. Use these problem types only: Date conflict, Name mismatch, Figure differs from source, and Narrative differs from result. If a figure awaits approval, write 'Dependency' and do not use the value of the previous period. List each required disclosure and say if it is present. Do not correct the draft. Do not invent a figure. Do not send the pack."

Tools

Sources: final calculation outputs and approved disclosures, with version history in SharePoint.

Check: Claude or Microsoft 365 Copilot for the narrative, and Excel for the tie-out of figures.

Release: the existing publication workflow, with approval by the content owners.

Skills

Staff skill: tie each figure in a pack to its source table.

Agent skill: a report-consistency skill with the four problem types and the dependency rule.

Method

Tie-out review: every figure and date in the pack maps to one source cell. Test with planted errors and count the misses.

#### Agent limits

Level 1. An agent can check each draft pack and prepare the findings table. It must not edit the pack or release a report. Content owners approve each correction.

## Choose the first pilot

These screening scores assume an investment firm that has authorised document snapshots and qualified reviewers. The firm permits no autonomous investment decisions or execution. The scores are our first assessments, and this paper reports no measured benefits from these workflows. Score them again with the process owner.

Case | Value | Feasible | Data | Effort | Risk

01 Filing changes | 4 | 4 | 4 | Small | Moderate

02 Research evidence | 4 | 4 | 4 | Medium | Higher

03 Mandate checklist | 4 | 3 | 3 | Medium | Higher

04 Holdings reconciliation | 4 | 3 | 3 | Medium | Higher

05 Corporate actions | 4 | 3 | 3 | Medium | Higher

06 Price exceptions | 4 | 3 | 3 | Medium | Higher

07 Performance explanation | 4 | 3 | 3 | Medium | Higher

08 Liquidity assumptions | 4 | 2 | 2 | Larger | Higher

09 Committee actions | 4 | 4 | 4 | Small | Moderate

10 Report consistency | 4 | 4 | 3 | Medium | Higher

### How to read the table

Value is the possible gain in evidence quality or the possible decrease in rework. Feasible shows how well the firm can run a controlled test. Data shows how ready the minimum inputs are. Each score is from 1 (weak) to 5 (strong).

Small effort uses a bounded document set. Medium effort needs structured records or specialist review. Larger effort needs more than one source system or validated models. Risk rates the consequence of an error, and it does not rate the probability.

Use cases 01 and 02 are the first choices, together with use case 09. They use approved document snapshots. Portfolio records come in only when the identifiers and the permissions are ready, and the reviewers are in place.

## Name an owner and a measure

Case | Accountable role | Pilot measures | Agent level

01 | Investment research lead | Missed changes; incorrect comparisons | 1

02 | Portfolio research manager | Unsupported claims; omitted contradictions | 1

03 | Investment operations manager | Omitted conditions; wrong effective dates | 1

04 | Fund operations lead | Missed differences; false matches | 1

05 | Corporate-actions lead | Missed notices; wrong date labels | 1

06 | Valuation control lead | Missed exceptions; false alerts | 1

07 | Performance measurement lead | Wrong basis; unsupported attribution | 1

08 | Portfolio liquidity lead | Missing assumptions; stale evidence | 0

09 | Committee secretary | Invented approvals; omitted conditions | 1

10 | Investment reporting manager | Missed inconsistencies; review effort | 1

Each accountable role is a proposed responsibility and assigns nothing to a named employee. The firm approves ownership, thresholds, authority and agent levels before the test.

## 5 The foundations lens

Keep decisions and calculations under accountable control

Write down where research, calculation, interpretation and execution each start and end. Arithmetic belongs to tested calculation tools. Keep a link that a reviewer can follow from their inputs to the reported outputs. During the pilot, give assistants read-only access to approved snapshots.

An authorised reviewer must be free to reject a convincing narrative. Put contradictory documents and stale prices in the test set, with missing data. Check again after any change to the model, the prompt, the data source or the audience. Keep a manual process for the times when the assistant cannot trace its sources.

### Give each agent its own identity

An AI agent must have its own account. Do not let an agent use the account of an employee. Give the account read access to the approved sources only. Keep a log of each tool call, with the time, the tool, the source and the result. Set a limit on the number of steps and on the run time.

Treat text inside a document as data. An agent must not obey an instruction that it finds in a document or an email. Review each agent skill before use. A skill can contain scripts, so the information technology team must check it as it checks other software.

### Protect client information and licensed material

The Office of the Information Commissioner of Jamaica describes standards on purpose, minimisation, accuracy, retention, security and overseas transfers. [6] The firm must assess each information flow and each provider arrangement separately, because investment work gives no general authority to disclose.

Before you connect a source collection, check the research licences and the permitted processing, and the provider terms for retention and model training. Then separate the portfolios and the client permissions. Test that a request about one client cannot return the mandate or the holdings of a different client.

Instructions inside filings and research documents are content, and they have no authority. A document must never be able to give access or change a tool setting. It must never be able to send information outside the firm. We did not test any provider against this risk, and the pilot must do so. Keep sufficient evidence, under the approved retention policy, to reconstruct material outputs. A material access failure stops the pilot.

## 6 From pilot to practice

Days 1 to 15

Select one document-based use case and an accountable reviewer. Measure the preparation and correction time of the current process. Agree the sources, the licences, the definition of a material error and the person who can stop the work. Set the agent level, and write the first version of the prompt. The reference examples must include an adverse fact and a missing source.

### Days 16 to 45

Run the pilot in parallel with the current process, and review every output. Log errors of source, retrieval, calculation and interpretation as four separate types. Keep trading systems and live portfolio systems disconnected for the full period. Confirm that the reviewer can recover the evidence behind each material statement. Give a version number to each prompt and each agent skill.

### Days 46 to 75

Compare equivalent work, and include the time for checks and corrections. A second specialist reviews a sample and does not see the first review. Then test a changed filing and a data feed that is not available. Where the two reviewers disagree, settle the point before you extend the pilot.

### Days 76 to 90

Assess quality, confidentiality, cost and the capacity of reviewers. Then the sponsor decides to continue or to stop, and a decision to continue can include revisions. A new portfolio or a confidential source needs a new approval, and so does a new action. A good document pilot gives no one authority to automate investment decisions.

## Pilot decision canvas

Complete this canvas with the process owner before you start the assistant. Attach the approved source list and the test results to the decision record.

### Research or operations task:

Accountable owner and reviewer:

### Approved sources and licences:

Evidence cut-off date:

### Portfolio and client access boundary:

Permitted assistant actions:

### Decisions and systems excluded:

Agent level and tool permissions:

**Prompt version and agent skill version:**

Reference examples and baseline:

**Material-error threshold:**

Stop contact and manual fallback:

**Review date and next decision:**

## Measure value without invented savings

Measure accepted outputs, and not page counts alone. On comparable work, record the time for preparation, source retrieval, review, correction and administration. Count material errors apart from revisions of style.

Include the costs of software, integration, data licences, training and supervision. For an agent, add the usage cost of each run. Time released is capacity until spending changes. A short pilot cannot show excess returns or lower portfolio risk, and it cannot show better results for clients. A report that claims them from a short pilot goes beyond its evidence.

Report the sample size and the task mix, and the examples that failed. None of our proposals has results from a named firm yet. The sponsor must see which cases the assistant could not complete safely, and how much human review was still necessary. Where polish hides unreliable evidence, keep the existing process.

Record | What to enter

Current process | Work volume; handling time; review time; correction rate

Pilot process | Same measures on comparable work; all review and correction effort

Costs | Licence; integration; training; continuing supervision; agent usage

Investment support | Missed exceptions; unsupported statements; source traceability

Decision | Continue (with or without revisions) or stop; owner; review date

## References

[1] Caribbean AI Association. Think Tanks. Harbour: economy, work and enterprise. Accessed 26 October 2026. https://caribbeanaiassociation.com/think-tanks

[2] Bank of Jamaica. Financial Stability Report 2024. Table A.2, page 68; securities-dealer coverage note. https://boj.org.jm/wp-content/uploads/2025/03/Financial-Stability-Report-2024.pdf

[3] International Organization of Securities Commissions. 2025. Artificial Intelligence in Capital Markets: Use Cases, Risks, and Challenges. Consultation Report CR/01/2025, March; asset-management uses, pages 26–27; model and data limitations, pages 38–40. https://www.iosco.org/library/pubdocs/pdf/IOSCOPD788.pdf

[4] CFA Institute. 2025. CFA Institute Equips the Investment Sector to Navigate AI Developments. Official publication announcement, 18 November. https://www.cfainstitute.org/about/press-room/2025/ai-in-asset-management-report-2025

## References continued

[5] Financial Services Commission, Jamaica. 2020. Market Conduct Guidelines for Securities Dealers. Effective 1 May; confidentiality, communications and suitability. https://www.fscjamaica.org/wp-content/uploads/2020/04/SR-GUID-26-Market-Conduct-Guidelines-Amended-FINAL-FEB-7-2020.pdf

[6] Office of the Information Commissioner, Jamaica. Data Protection Standards. Accessed 26 October 2026. https://oic.gov.jm/page/data-protection-standards

[7] Agent Skills. Agent Skills Overview. Accessed 10 November 2026. https://agentskills.io

### Evidence and version

Evidence reviewed on 26 October 2026. Text revised and build notes added on 11 November 2026. Prepared by the Harbour Think Tank, Caribbean AI Association. Proposed workflows and prompts need local validation before deployment.