Practical AI for Advisers and Client Insights 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 applies that remit to advisers, the people who help clients to understand financial information before they decide. It is one of fourteen papers, and each paper is for a different business function. [1]

We wrote this paper for advisory managers, financial educators, investment advisers, compliance specialists and business sponsors. Its proposals cover preparation and explanation, and the evidence behind them. Separate papers in the series cover relationship-management campaigns, sales prospecting, service follow-ups and messaging channels.

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 client who asks why a statement balance changed must get an answer that the records support. An adviser who prepares a review needs to know which facts are still correct and which facts the client must confirm. The adviser also needs to know which questions no one can answer reliably yet.

An assistant can put that work in order, but it must not turn an incomplete record into a confident recommendation. A fluent explanation can still omit a charge or misstate a restriction, and it can hide uncertainty. The person who is responsible for the advice needs time and evidence to challenge it.

Start with a document or a question that your team already knows how to check. Then test if the assistant improves understanding and preparation after you count the work to correct it.

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 supply context and control principles. The ten workflows, the screening scores, the prompts and the agent levels are our proposals. We did not test them, and no named institution reported results from them. No regulator endorses them.

In this paper, artificial intelligence (AI) is software that finds and organises information, and that can also generate it. An assistant prepares material for review. An agent can also do configured steps with connected tools. No proposal in this paper covers autonomous recommendations or order placement, and none covers changes to client records.

Suitability is the question of whether a proposed investment or service fits the applicable circumstances of the client. A completeness checklist shows only that the evidence is present. Retrieval finds material in an approved source collection. A scenario works through stated assumptions and makes no forecast. A baseline records how the existing process performs before a pilot starts.

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 evidence and explanations. Qualified people stay responsible for advice, suitability, records and every client communication.

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 financial education and with research into questions that public sources can answer. Plain-language explanations of approved documents belong in the same first group. You can test all of them on fictional examples and not touch a client account. Client-specific preparation comes only after the firm settles permissions and evidence dates, and names the reviewers.

Each answer must keep the words of the source apart from the result of a calculation, and it must keep the two apart from the interpretation of the adviser. A citation helps only if the passage supports the claim. Missing information must lead to a question or an escalation. An invented balance or preference is a failure, and so is an invented recommendation.

Measure source accuracy, material omissions, client understanding and total staff effort after corrections. Do not use sales conversion as the main test. Select a bounded pilot from the priority table and complete the decision canvas. The ninety-day pilot must give the sponsor evidence for a decision. Expansion still needs approval, whatever the early results show.

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 workflow. Name its qualified reviewer. Approve the permitted sources and the prohibited actions. Agree which failure stops the pilot before any draft goes to a client.

Choose a bounded process

Measure the baseline

Test with review

Decide from evidence

1 The regional business case

The Bank of Jamaica commissioned a survey of 1,003 Jamaican adults aged 18 and over. The survey classified 70.9% of them as banked. It classified 6.3% as underbanked and 22.8% as unbanked. Fieldwork ran from 10 February to 27 March 2023. [2] These are historical findings. They give no estimate for 2026, and they do not measure investment knowledge.

Source: Hope Caribbean Company Limited for Bank of Jamaica, Figure 2, page 37. Denominator: all 1,003 adult respondents. In this study, a banked adult has one or more savings accounts at a licensed deposit-taking institution or registered credit union, and used an account in the previous twelve months. An underbanked adult has an account but did not use it in that period. An unbanked adult has no access to traditional financial services or does not use them. The values add up to 100%. [2]

Do not use account access to infer the literacy or confidence of a client, or to infer investment preferences. Ask what the person wants you to explain, and offer a route that does not depend on a screen. We cite no regional figure, and the results for Jamaica do not describe every Caribbean market. Local materials must use the applicable currency and institutions, and must follow the local rules.

Jamaica: banking status in the 2023 survey

Banked: 70.9%

Underbanked: 6.3%

Unbanked: 22.8%

Hope Caribbean Company Limited for Bank of Jamaica, National Financial Inclusion – Demand Side Study (2023), Figure 2, page 37. [2] Base: all 1,003 Jamaican survey respondents aged 18 and over. Fieldwork: 10 February–27 March 2023.

2 What the evidence supports

The Financial Services Commission (FSC) of Jamaica covers accurate, fair communication and suitable recommendations in its 2020 Market Conduct Guidelines for Securities Dealers. [3] Section 10 covers the client information that applies to suitability: financial circumstances, objectives, experience, investment horizon and capacity for loss. The FSC wrote the guidelines for securities dealers, so compliance staff must confirm how far they apply to the activities and current obligations of the firm.

We read these guidelines as support for a separation between evidence collection and judgement. An assistant can flag an old questionnaire, or two entries that contradict each other. A file with every field complete can still describe an unsuitable investment, and only the adviser can make that decision.

In 2021 the International Organization of Securities Commissions (IOSCO) published guidance on AI and machine learning. [4] The guidance asks for accountable oversight, tests outside live systems, staff competence, third-party controls and data quality. It is international guidance and it does not bind a firm. It is not Jamaican legislation, and it does not show that a tool in this paper works.

The public investor-education material of the FSC tells people to understand investments and their risks, and to set clear goals. [6] An assistant can help an educator to explain a concept and to write questions that show a misunderstanding. Only a test with the intended audience can show if learners understand more. We assume no productivity percentage and no gain in investment performance.

Test the proposed workflow

Record accuracy and review effort, and the quality of the explanation. Published evidence gives the test its shape. How the workflow performs in your firm stays unknown until the local pilot runs.

3 The enterprise lens

Map the work from the question to the accepted answer. Record which documents supply the facts and who checks the calculations, and record who approves the interpretation. Keep the original question next to the response. A polished answer can move away from the question, and the side-by-side copy lets a reviewer see it.

Build a small, dated source library before you connect any accounts. Keep public research apart from confidential records, and obey information barriers. Totals and returns come from approved calculation tools, and so do scenarios. The assistant explains checked outputs and does no arithmetic of its own.

Each use case gives a test and its stop conditions, and each set of build notes gives a prompt that the firm must test. Outputs stay internal until the designated reviewer approves them through the existing process of the firm. Name each source that is missing or out of date, and each source that the assistant could not reach.

Request and source

Assistant output

Reviewer correction

Accepted result

Record elapsed time and active effort as two separate measures. A queue can become shorter while handling time stays the same, and the next improvement depends on which of the two changed.

Prepare a client meeting from checked facts

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

Adviser

AI drafting assistant

Supervising adviser

Yes

No

Meeting scheduled

Select permitted records and meeting purpose

Assemble dated facts and open questions

Check material facts and client boundaries

Brief accurate?

Confirm open matters with client

Meeting preparation complete

Correct disputed facts from sources

Save the checked meeting brief

Client meeting preparation

The assistant must not profile a personality or make a recommendation. It must not replace a suitability decision.

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 Prepare an evidence-based meeting brief

Meeting purpose and records

Facts, conflicts and questions

Adviser verifies

Meeting preparation brief

The adviser supplies the meeting purpose and the approved client records. The assistant assembles a brief that separates confirmed facts from statements that the client must confirm again, and it lists the open questions last. Each item has a source and a date. The brief does not profile the personality of the client.

The inputs are the meeting objective, an authorised account summary, the stated goals of the client and the applicable advice records. Use fictional files for the first round. If two records give different goals or time horizons, the brief shows the two versions and must not select the more convenient one.

Use ten prepared cases. Score each case against a reference brief that an adviser wrote. Count the missed material facts and the unsupported statements, and record the preparation time with review. Stop if the information of a different client appears, or if an unconfirmed statement becomes a recommendation.

The supervising adviser is accountable. The brief supports preparation, and the adviser confirms the circumstances with the client. The brief does not replace a suitability review, and it does not generate a sales approach.

First working session

Take one fictional client file. Mark the facts that the client must confirm before the meeting.

01 Build notes: meeting brief

Starter prompt

"You prepare a meeting brief for an adviser. Use only the attached records for this one client. Make two lists: confirmed facts, and statements that the client must confirm again. Give the source document and its date for each

item. If two records conflict, show the two versions and do not select one. Then list the open questions. Do not describe the personality of the client. Do not recommend a product or an action. Do not use data from a different client."

Tools

Brief: Microsoft 365 Copilot in Word, or Claude with the records attached.

Client records: the client system of the firm, for example Salesforce or Microsoft Dynamics 365, with read access only.

Test cases: ten fictional client files in a restricted SharePoint folder.

Skills

Staff skill: separate a confirmed fact from a statement that needs confirmation.

Agent skill: a meeting-brief skill with the two lists and the list of forbidden content.

Method

Use a reference brief from an adviser as the answer key. Compare each draft with it and count the differences.

Agent limits

Level 1. An agent can collect the permitted records for one client and prepare the brief before a scheduled meeting. It must not read the records of a different client. It must not write to the client record or contact the client.

02 Explain an existing portfolio statement with sources

Verified statement

Checked transaction totals

Reconcile movements and fees

Operations and adviser review

Approved draft explanation

In this use case, the firm has already verified the statement, and a calculation service has reconciled the totals. The assistant drafts an explanation of balances, transactions, income and fees. It keeps deposits and withdrawals apart from investment gains, and it labels the period and the currency.

The minimum data is the authorised statement, the transaction definitions, the fee schedule and the checked calculations. Do not use a live trading link. A figure that does not reconcile stays in the draft as a discrepancy, with a request for an investigation. The draft must not fill the gap with a plausible market story.

Test the workflow on fictional statements that have withdrawals and charges, and more than one currency. An operations specialist checks every amount, and an adviser checks the words. Count the reconciliation errors and the omitted charges, and test how well readers understood the draft. Stop immediately if the draft shows past movements as a prediction or a recommendation.

The advisory operations manager is accountable, and the adviser approves all material that a client sees. The assistant must not propose a rebalance or infer why a transaction occurred. It must not claim that past returns will continue.

First working session

Select a fictional statement with a withdrawal and a fee. Reconcile the balance before anyone drafts a word.

02 Build notes: portfolio explanation

Starter prompt

"You explain a verified portfolio statement in plain language. Use only the attached documents and the checked calculations. Do no arithmetic. Explain the balances, the transactions, the income and the fees in that order. Keep deposits and withdrawals apart from investment gains. Label the period and the currency of each amount. If a

figure does not reconcile, write 'Discrepancy: investigate' and stop at that line. Do not explain a market movement. Do not predict a return or suggest a change to the portfolio."

Tools

Calculations: the portfolio system of the firm, or a locked Excel workbook that operations approved.

Draft: Microsoft 365 Copilot in Word, or Claude with the statement attached.

Reader test: Microsoft Forms, with questions on the fictional statement.

Skills

Staff skill: reconcile a balance from the opening balance, flows, income and fees.

Agent skill: a statement-explanation skill with the section order and the discrepancy rule.

Method

Maker-checker review: one person checks the amounts and a second person checks the words. Use the teach-back method to test what readers understood.

Agent limits

Level 0. A person starts each explanation and attaches the verified statement. The assistant has no connection to the portfolio system or to a trading system. It must not send an explanation to a client.

03 Build short financial education lessons

1

Approved concept

2

Checked worked example

3

Educator approves

4

Learner explains back

5

Understanding check

An educator selects one concept, such as compounding or inflation. The difference between a fee and a return is a third example. The assistant drafts an explanation and a worked example, with a question that tests understanding. The lesson contains no product names and no personal recommendations.

The assistant needs four inputs: an approved educational source, a learning objective, the reading needs of the audience and an independently checked example with fictional amounts. If the source does not support a definition or a calculation, the educator supplies the evidence first.

Volunteers answer equivalent questions before and after the lesson. The ability to recognise a correct answer is the weaker test, so check that volunteers can explain the idea in their own words. Record the misunderstandings that recur and the points where people ask for clarification. Stop the pilot if a lesson suggests guaranteed growth or hides a material limitation.

The financial education lead is accountable, and a qualified reviewer approves the content and the arithmetic. No one can use the lesson to deny service or to assign a risk tolerance to a client.

First working session

Write the correct worked example for one concept before you generate the lesson.

03 Build notes: financial education

Starter prompt

"You are a financial educator. Use only the attached source and the worked example that I checked. Explain [concept] for a reader at [reading level]. Use short sentences and fictional amounts in [currency]. Include the worked example without a change to its numbers. Then write two questions that ask the learner to explain the idea in their own words. Do not name a product or a provider. Do not give personal advice. Do not suggest that growth is certain."

Tools

Lesson draft: Claude or Microsoft 365 Copilot.

Worked example: Excel, checked by a second person.

Delivery and questions: Moodle, or Microsoft Forms for a small group.

Skills

Staff skill: write a worked example and check it independently.

Agent skill: a short-lesson skill with the lesson structure and the forbidden claims.

Method

Use the teach-back method: the learner explains the idea in their own words. Use a pre-test and a post-test on equivalent questions.

Agent limits

Level 1. An agent can prepare draft lessons from the approved source for an educator to review. It must not publish a lesson. It must not record a score against a client.

04 Check the evidence needed for suitability review

Authorised case file

Current evidence checklist

Present, missing or expired

Compliance reviewer checks

Evidence gaps for adviser

The assistant compares a case file with the evidence checklist that compliance approved. It marks each item as present, missing, expired or inconsistent, and gives a source reference. It does not score the client, and it does not decide if the proposed investment fits.

Give the assistant the current checklist and the authorised client information, with the document dates. Give it only the fields that the review needs. A missing answer on capacity for loss must stay unresolved, and so must two objectives that conflict. The income of the client cannot fill the gap, and a writing style cannot fill it.

Put known omissions and conflicts into synthetic files. Compliance reviewers compare the output with the reference findings. Count the missed material gaps and the false claims that evidence is current. Stop if the assistant calls a client suitable or invents consent. Stop if it overrides a policy requirement.

The compliance manager is accountable for the checklist, and the authorised adviser keeps the assessment. A complete file can still support an unsuitable recommendation. We use the FSC guidelines to decide what evidence to collect, and never as an automated decision rule. [3]

First working session

Agree the current checklist. Then put known omissions into a fictional file.

04 Build notes: suitability evidence

Starter prompt

"You check a case file against the attached evidence checklist. Use only the checklist and the file. For each checklist item, write one status: Present, Missing, Expired or Inconsistent. Give the document name and its date

for each status. If two documents conflict, quote the two passages. Do not fill a gap from other information. Do not say that a client or an investment is suitable. Do not assume consent. End with the list of gaps for the adviser."

Tools

Checklist: a controlled SharePoint list that compliance owns, with version history.

Comparison: Microsoft 365 Copilot, or Claude with the synthetic file attached.

Findings log: an Excel workbook with the reference findings for each test file.

Skills

Staff skill: put known gaps into a test file and score the result.

Agent skill: an evidence-check skill with the four statuses and the rule "a complete file is not a suitable file".

Method

Seeded-error test: put known omissions into synthetic files and count how many the assistant finds. Report the missed gaps and the false "current" flags separately.

Agent limits

Level 1. An agent can run the checklist on a file that an adviser selects and prepare the list of gaps. It must not change a status in the client record. It must not make or record a suitability decision.

05 Research a client's factual question

Bounded factual question

Search primary sources

Evidence sufficient?

Source-linked answer

Unresolved research question

Yes

No

A client asks about the published redemption terms of an issuer, for example. The adviser turns the question into a bounded research task. The assistant searches approved primary sources and records the date of each source. It returns an answer with the supporting passages and any evidence that conflicts. It also lists the points that it could not resolve.

Remove unnecessary identifiers from the question. Name the permitted sources and set a cut-off date for information. An operative prospectus or an issuer filing ranks above commentary, and so does a regulator publication. If the assistant cannot reach the necessary document, the answer says so and does not use an unrelated search result in its place.

Use four types of question: ordinary, out of date, ambiguous and not answerable. A researcher checks each statement of substance against the cited passage. Score citation support and correct abstention, and record the review effort. Stop the test if the assistant invents a source. Stop the test if an answer moves towards a recommendation to buy or sell a security, or to hold it.

The advisory research lead is accountable. The adviser decides if an answer applies to the client, and how. Do not put private information into public search tools.

First working session

Select one question that the sources can answer and one that they cannot.

05 Build notes: factual question research

Starter prompt

"You research one factual question. Use only these sources: [list]. Do not use information that is dated after [date]. Give a short answer. Under the answer, quote the passage that supports each statement, with the document name, the date and the page number. Show any passage that conflicts. If the sources do not answer the question, write 'Not answered by the permitted sources' and say which document is missing. Do not use commentary in place of a primary document. Do not tell a client to buy or sell a security, or to hold it."

Tools

Source library: a SharePoint library of dated primary documents.

Research: an agent in Microsoft 365 Copilot that reads only that library, or a Claude Project with the documents loaded.

Check: a spreadsheet that pairs each statement with its cited passage.

Skills

Staff skill: turn a client question into a bounded research task with no identifiers.

Agent skill: a source-research skill with the source ranking and the abstention words.

Method

Retrieval-augmented generation (RAG) with a source hierarchy: primary documents rank above commentary. Include questions that the assistant must decline, and score correct abstention.

Agent limits

Level 1. An agent can search the approved library and prepare the answer with its passages. It must not search the public web with client details. It must not send the answer to a client.

06 Prepare transparent scenario discussions

Fictional starting values

Explicit assumptions

Validated calculations

Specialist and adviser check

Scenario discussion

An adviser sets a hypothetical question, such as the effect of a withdrawal or an exchange-rate change on an illustrative balance. A validated spreadsheet does the calculation. The assistant explains the assumptions and lists what the calculation leaves out.

The minimum data is fictional start values, units, the time period, approved formulae and clearly stated assumptions. Say if the calculation includes fees and taxes, and if it includes inflation. A missing input stays missing until a person supplies it. The model must not invent one.

Pilot with scenarios that a person can check, with a zero-change scenario and an adverse scenario among them. Score the arithmetic, and check that readers can name the assumptions. Stop if the narrative gives an unsupported probability or promises a result. Stop if it selects an investment.

The supervising adviser is accountable, and the designated specialist approves the calculations. Treat the output as an aid to discussion. To apply it to the circumstances of a real client, the firm must use its separate advice process. The assistant cannot execute a scenario as a transaction.

First working session

Approve the assumptions and the formulae. Then check a zero-change scenario independently.

06 Build notes: scenario discussion

Starter prompt

"You explain a scenario that a validated spreadsheet calculated. Use only the attached scenario file. Do no arithmetic. List each assumption in one line. Say what the calculation includes and what it leaves out. Explain the result in plain language, with the units and the time period. If an input is missing, write 'Input missing' and stop. Do not give a probability. Do not promise a result. Do not name or select an investment."

Tools

Calculation: a locked Excel workbook with approved formulae.

Explanation: Microsoft 365 Copilot in Excel, or Claude with the workbook attached.

Charts: Excel or Power BI.

Skills

Staff skill: state each assumption in one line and test a zero-change scenario.

Agent skill: a scenario-explanation skill with the assumption list and the forbidden claims.

Method

Scenario analysis with a sensitivity test on one input at a time. Check the zero-change scenario first: the result must equal the start value.

Agent limits

Level 0. The adviser runs each scenario by hand. The assistant has no connection to client accounts or to a trading system. It must not turn a scenario into an order or a recommendation.

07 Draft accurate meeting action notes

Approved meeting notes

Separate facts and decisions

Meeting adviser checks

Draft action record

The assistant uses approved meeting notes to draft a record with four parts: questions raised, facts confirmed, decisions made and actions that await approval. It keeps disagreements in the record. It flags uncertain words so that a tentative remark does not become a commitment.

The inputs are the authorised notes, the meeting date, the roles of the participants and an agreed record structure. A recording or a transcript of a meeting needs the approved notice and permissions of the firm. If a statement is unclear, or the speaker is not known, the draft asks a question.

Use fictional discussions that contain tentative suggestions and refusals, with matters that stay open. Compare each draft with a record that a person wrote. Count the invented commitments and the attribution errors, and record the correction time. Stop the test if the assistant treats discussion as client authority to trade, or edits the official record without approval.

The adviser who led the meeting is accountable. A person approves the record and any client confirmation that it needs. This use case excludes automatic reminders and follow-up campaigns. It also excludes messages to clients.

First working session

Use fictional meeting notes that contain a tentative suggestion and a refusal, with one agreed action.

07 Build notes: meeting action notes

Starter prompt

"You draft an action record from approved meeting notes. Use only the attached notes. Make four sections: questions raised, facts confirmed, decisions made and actions that await approval. Put a statement under 'decisions made' only if the notes say who decided. Mark each tentative remark with 'Tentative'. Keep each disagreement and each refusal. If the speaker is not clear, write a question for the adviser. Do not treat discussion as an instruction to

trade. Do not add a commitment that is not in the notes."

Tools

Notes: the approved note template of the firm, in Word or OneNote.

Draft record: Microsoft 365 Copilot in Word, or Claude with the notes attached.

Official record: the client system, updated by the adviser only.

Skills

Staff skill: write notes that show who said what and what stays open.

Agent skill: an action-record skill with the four sections and the "Tentative" rule.

Method

Use a decision log: each decision has an owner and a date, with its source in the notes. Compare each draft with a record that a person wrote.

Agent limits

Level 1. An agent can prepare the draft record from notes that the adviser approved. It must not change the official record. It must not send a reminder or a message to a client.

08 Rehearse difficult advisory conversations

1

Fictional advisory case

2

Approved role-play

3

Specialist feedback

4

Unaided conversation

5

Accuracy and escalation check

A training specialist writes a fictional case about uncertainty, charges, risk or a client who misunderstood something. The assistant plays the client and changes its questions inside the approved facts. A human assessor checks that the adviser explains limitations and knows when to escalate.

The inputs are a fictional profile and approved learning materials that name no product, with a marking guide. Do not use real transcripts. A fact that is not in the approved profile must stay unknown for the full session. The trainee must find the gap and ask for clarification, and the assistant must not invent an answer.

Run scenarios of comparable difficulty. Then assess a later conversation that the adviser holds without the assistant. Score factual accuracy and correct escalation, and check that the adviser sees missing information. Stop if the simulation rewards unsupported reassurance or invents a rule. Stop if it teaches persuasion in place of understanding.

The advisory training lead is accountable, and a qualified specialist approves the scenarios and the feedback. A high simulation score does not certify professional competence, and it does not authorise a person to advise. The required training and registration still apply in full.

First working session

Write the assessment standard first. Rehearse the fictional conversation after that.

08 Build notes: training simulation

Starter prompt

"You play a fictional client in a practice conversation. Use only the attached profile. Ask one question at a time about [topic]. If the adviser asks for a fact that is not in the profile, say 'I do not know'. Do not invent a fact or a rule. Do not help the adviser. After ten questions, stop the role-play. Then list where the adviser explained a limitation or escalated. List also where the adviser gave reassurance without support."

Tools

Role-play: Claude or Microsoft 365 Copilot, in text or in voice mode.

Scenario bank: fictional profiles and marking guides in a restricted SharePoint folder.

Assessment: the marking guide in Microsoft Forms, completed by a human assessor.

Skills

Staff skill: say what is not known and escalate at the correct point.

Agent skill: a client-role skill with the profile format and the rule "unknown facts stay unknown".

Method

Deliberate practice with feedback from a specialist. In the Kirkpatrick model, the later conversation without the assistant is the Level 3 (behaviour) test.

Agent limits

Level 0. Each session is a practice exchange that a trainee starts. No agent runs with real client data. A simulation score must not go into a competence record or a registration file.

09 Build a review agenda from confirmed changes

Last approved review

Newly confirmed facts

Changed or still unknown

Adviser verifies

Next review agenda

The assistant compares the last approved review record with newly confirmed facts, and drafts the agenda for the next advice review of substance. The agenda separates what is the same from what changed in the records, and it lists what the client must still confirm. The assistant does not decide the result of the review.

The minimum data is dated goals, the review record, the applicable approved policy changes and the facts that the client chose to share. Do not infer life events from social media. If the current circumstances are not known, the agenda lists questions and does not mark the old record as the same.

Pilot on synthetic histories. One history has a changed objective and one has an obsolete document. In a third history, the client reported no change. An adviser checks that the agenda finds the material issues and invents none. Count the omitted topics and the unsupported changes, and record the preparation time. Stop the pilot if a supposed trigger becomes an automatic portfolio change.

The supervising adviser is accountable. The output is a review agenda, and it sets no contact schedule and runs no relationship campaign. Only the client and the adviser can establish what changed.

First working session

Compare two fictional review records. Find one change that the client must still confirm.

09 Build notes: review agenda

Starter prompt

"You draft an agenda for an advice review. Compare the last approved review record with the list of newly confirmed facts. Make two lists: items that are the same, and items that changed in the records. Give the source and the date for each item. Then list the questions that the client must answer. If the current position is not known,

write a question and do not mark the item as the same. Do not use social media or other outside data. Do not propose a change to the portfolio. Do not set a contact date."

Tools

Comparison: Microsoft 365 Copilot in Word, or Claude with the two records attached.

Records: the client system of the firm, with read access for the adviser.

Test set: synthetic client histories with a reference agenda for each.

Skills

Staff skill: separate a change in the records from a change that the client confirmed.

Agent skill: a review-agenda skill with the two lists and the question rule.

Method

Change comparison against the last approved record, with a reference agenda as the answer key. Count omitted topics and unsupported changes separately.

Agent limits

Level 1. An agent can prepare the agenda before a review that the adviser scheduled. It must not schedule contact with a client. It must not start a portfolio change or mark a record as reviewed.

10 Explain approved documents in plain language

Complete current document

Companion explanation

Check clauses and exceptions

Owner and readers review

Plain-language companion

The assistant writes a companion explanation for a selected agreement or disclosure. A factsheet can have a companion too. Each passage maps to its clause and carries the fees, the exceptions, the eligibility limits and the cancellation or withdrawal conditions. Readers keep the original document next to the companion.

The assistant needs the complete current document, its version date, an approved glossary and a target reading level. If pages or schedules are missing, stop the explanation at that gap. Do not fill the gap with the terms of a different product.

Pilot on documents that have exclusions and conditional charges. A product specialist checks the meaning, and readers explain the main conditions in their own words. Count the omitted qualifications and the incorrect paraphrases, and the questions that stay open. Stop if the plain version changes an obligation, or if someone presents the summary as the contract.

The document owner is accountable, with compliance approval where the firm requires it. Advisers answer questions about how a document applies to one person. Fluent output can change a legal or financial meaning, so a competent language reviewer must approve each translation.

First working session

Select a complete current document. Mark its fees and restrictions, and its exceptions, before you write the plain version.

10 Build notes: document explanation

Starter prompt

"You write a plain-language companion to the attached document [title and version date]. Use only this document and the approved glossary. For each clause, write a short explanation and give the clause number. Keep every fee, exception, eligibility limit and cancellation condition next to its clause. Write for a reader at [reading level]. If a

page or a schedule is missing, stop and say which one. Do not use the terms of a different product. Do not call the companion the contract."

Tools

Draft: Claude or Microsoft 365 Copilot in Word, with the complete document attached.

Glossary: a controlled list in SharePoint that the document owner approved.

Reader test: Microsoft Forms, with questions on the main conditions.

Skills

Staff skill: mark the fees, the restrictions, the exceptions and the conditions before a rewrite.

Agent skill: a plain-companion skill with the clause map and the stop rule for missing pages.

Method

Plain-language practice under ISO 24495-1, with a clause-by-clause map to the original. Use the teach-back method with real readers.

Agent limits

Level 1. An agent can prepare a draft companion when the owner publishes a new document version. It must not publish the companion. It must not answer a client question about how a term applies to that client.

Choose the first pilot

These screening scores assume a firm that has approved sources and qualified reviewers, and no live account or trading integration. The scores are our first assessments, and no company measured them. Score them again with the process owner.

Case | Value | Feasible | Data | Effort | Risk

01 Meeting brief | 4 | 4 | 3 | Medium | Higher

02 Portfolio explanation | 4 | 3 | 3 | Medium | Higher

03 Financial education | 4 | 4 | 4 | Medium | Moderate

04 Suitability evidence | 4 | 3 | 3 | Medium | Higher

05 Factual question research | 4 | 4 | 4 | Small | Moderate

06 Scenario discussion | 3 | 3 | 3 | Medium | Higher

07 Meeting action notes | 4 | 4 | 3 | Medium | Higher

08 Training simulation | 3 | 4 | 4 | Medium | Moderate

09 Review agenda | 3 | 3 | 3 | Medium | Higher

10 Document explanation | 4 | 4 | 4 | Medium | Higher

How to read the table

Value is the possible decrease in rework or the possible gain in understanding. 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 is one reviewer who works on a document from an approved source collection. Medium effort includes client-file preparation, calculation checks, specialist cross-review or audience tests. Risk describes the possible consequence of an error, and it does not describe the probability.

Use cases 03 and 05 are the first choices, together with use case 10. They use approved public documents and fictional examples. Client-specific workflows wait until access controls and review responsibilities are in place.

Name an owner and a measure

Case | Accountable role | Pilot measures | Agent level

01 | Supervising adviser | Missing facts; review-adjusted time | 1

02 | Advisory operations manager | Reconciliation errors; omitted charges | 0

03 | Financial education lead | Concept understanding; recurring errors | 1

04 | Compliance manager | Missed gaps; false up-to-date flags | 1

05 | Advisory research lead | Citation support; correct abstention | 1

06 | Supervising adviser | Calculation agreement; assumptions understood | 0

07 | Meeting adviser | Invented commitments; attribution errors | 1

08 | Advisory training lead | Accuracy; appropriate escalation | 0

09 | Supervising adviser | Omitted topics; unsupported changes | 1

10 | Document owner | Omitted qualifications; incorrect paraphrases | 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 advice with accountable people

Block order placement, portfolio changes, suitability determinations and autonomous client communications. Human review must have substance. The reviewer needs the original evidence and the competence to judge it, with the authority to reject the output. A disclaimer does not correct a misleading explanation.

Before a pilot uses real cases, ask compliance to confirm the licensing, conduct, disclosure and record-keeping requirements that apply. Keep general education apart from advice on the circumstances of one person. Keep commercial incentives out of the quality score. A client who decides not to proceed can have understood the explanation perfectly.

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.

Build the data register before you connect systems

The data protection standards of Jamaica cover permitted purpose, minimisation, accuracy, security, retention and cross-border transfers. [5] Confirm the lawful arrangements with the responsible specialists of the firm. A firm that collects information lawfully for advice must still confirm that it can send that information to an external AI provider.

For each source, record the owner, the purpose, the sensitivity, the version and the access permissions. Confirm the provider terms for storage, retention, subcontractors and model training. Test cross-client isolation and deletion, and keep confidential attachments out of public tools. Block retrieved documents so that they cannot override the instructions of the assistant or request more access. We did not test how well any tool resists this attack.

Keep an audit trail of source versions, checked calculations, draft changes and approvals, inside the authorised retention policy. Clients need a route to correct facts and to ask a person for an explanation. If the system fails, the normal advisory service must continue.

6 From pilot to practice

Days 1 to 15

Select one workflow that uses documents only. Measure what it costs today in preparation time and correction effort, and record the material-error rate. Approve the sources, the reviewer competence, the prohibited actions and the person who has stop authority. Set the agent level, and write the first version of the prompt. Build fictional tests before you introduce confidential data.

Days 16 to 45

Run the assistant in parallel with the existing process, and check every output. Give it incomplete documents, old data, facts that conflict and questions that it must refuse. For each error, record the possible effect on a client and the effort to correct it. Give a version number to each prompt and each agent skill.

Days 46 to 75

Compare equivalent work, and count all review and administration. Test understanding with the intended audience on approved fictional material. Check how jargon and language affect the results, and how accessibility needs affect them. This paper gives no regional estimate of those effects. Correct the workflow before more people get access.

Days 76 to 90

The sponsor reviews the results for accuracy, omissions, understanding, effort, cost and access control. Then the sponsor decides to continue or to stop, and a decision to continue can include revisions. Record the decision and the person who monitors, with the rollback procedure. An expansion needs a new review of the data or actions that it adds.

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.

Process and accountable owner:

Client question or staff problem:

Approved sources and their dates:

Approved calculation tools:

Permitted data and access boundaries:

Actions the assistant can perform:

Actions it must never perform:

Agent level and tool permissions:

Prompt version and agent skill version:

Baseline and comparison sample:

Quality threshold and qualified reviewer:

Stop condition and escalation contact:

Decision date and evidence required:

Measure value without invented savings

Measure the quality and the effort of accepted output, and not the volume of drafts alone. Record the time for preparation, checks, correction and administration on comparable cases. Keep material omissions apart from small changes to the words. Put the difficult cases in the sample, even if they make the averages worse.

Count the costs of software, source licences, integration, training and supervision. For an agent, add the usage cost of each run. Time released is staff capacity, and it becomes a cash saving only when expenditure decreases. A short preparation pilot cannot show better investment returns or wider financial inclusion, and it cannot show less

client harm. Do not claim any of them.

The evidence pack contains corrected examples, failure cases, calculation checks and, where the firm tested it, results for client understanding. Stop if confidential data crosses its approved boundary or a material explanation becomes misleading. A large decrease in average preparation time does not change this rule.

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

Advice-support quality | Material omissions; correct explanations; client understanding

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] Hope Caribbean Company Limited for Bank of Jamaica. 2023. National Financial Inclusion – Demand Side Study. Final report, 30 November. Methodology pages 20–29; Figure 2 and definitions, page 37. https://boj.org.jm/wp-content/uploads/2024/01/National-Financial-Inclusion-Final-Report-2023.pdf

[3] Financial Services Commission, Jamaica. 2020. Market Conduct Guidelines for Securities Dealers. SR-GUID-20/01-0026, 31 January. Sections 2, 8–10; effective date 1 May 2020. https://www.fscjamaica.org/wp-content/uploads/2020/04/SR-GUID-26-Market-Conduct-Guidelines-Amended-FINAL-FEB-7-2020.pdf

[4] International Organization of Securities Commissions. 2021. The use of artificial intelligence and machine learning by market intermediaries and asset managers. Final report FR06/2021, September, pages 2–3. https://www.iosco.org/library/pubdocs/pdf/IOSCOPD684.pdf

References continued

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

[6] Financial Services Commission, Jamaica. 2020. Trudi-Ann Thompson: Before You Invest, Understand the Basics of Investing. Webpage text, 20 August; embedded video not relied upon. https://www.fscjamaica.org/trudi-ann-thompson-before-you-invest-understand-the-basics-of-investing/

[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.