



constancia.

2026




Our predictions: AI in finance

Constancia, September 2026. Written for CFOs and heads of FP&A in multi-entity groups.

9 PREDICTIONS

Each carries the evidence from the room, why we expect it to change, what it means for you and what would make us revise it

01	Agents	02	Access	03	Reconciliation
04	Hours	05	Team	06	Consolidation
07	Controls	08	Supplier risk	09	Outcome first



Agents will reach the core of finance within 2 budget cycles, and the model underneath them has not moved

On 15 September 2026 we hosted a dinner in Mayfair for finance leaders to compare where AI has reached inside their functions. Every live use was reporting or transactional work. The examples were matching remittances to invoices, building quarterly review slides and answering CFO questions over EPM data. What the room valued most was access, because a budget holder can now ask the numbers a question with no analyst in between. Nobody in the room forecasts with it, closes with it or controls with it.

The layer underneath the tools has not changed. One group reconciles data across 5 ERPs and an inherited analytics tool. A monthly look-back over about 300 million of vendor invoices is still done by hand. All 9 controls described governed access, cost or data, and none governed what the machine produced. Nobody could name a finance role created in the last 2 years.

By an agent we mean software that takes a finance task from start to finish. A model decides the steps and a person checks the result, as when an agent matches a remittance to an invoice or drafts variance commentary. The AI 2027 scenario, written by former OpenAI researcher Daniel Kokotajlo and colleagues, has capable agents at a tenth of today's cost by mid-2027. Later updates from the authors put some central estimates nearer 2030. On either timeline you have 1 or 2 budget cycles.

The decision is which work to hand over first, who checks it and what it costs per transaction. It is also what has to be fixed underneath before anyone talks to the numbers. Our view is that the hardest part is the team. A finance function that can use what the machine gives it, and grow with it, takes longer to build than the tool or the data.

Our 9 predictions follow. Each carries the evidence from the room, why we expect it to change, what it means for you and what would make us revise it. Use them as a benchmark for your own function.

WHERE AI HAS REACHED INSIDE THE FUNCTION

LIVE TODAY: REPORTING AND TRANSACTIONAL WORK

- Matching remittances to invoices
- Building quarterly review slides
- Answering CFO questions over EPM data

What the room valued most was access: a budget holder can ask the numbers a question with no analyst in between

NOBODY IN THE ROOM

- Forecasts with it
- Closes with it
- Controls with it

All 9 controls described governed access, cost or data, and none governed what the machine produced

THE TOOL

Can be bought

THE DATA

Can be fixed in a programme with an end date

THE TEAM: THE HARDEST PART

Takes longer to build than the tool or the data

What the room valued most was access. The layer underneath the tools has not changed.

CONTENTS

5

ERPs and an inherited analytics tool, reconciled by one group

300m

of vendor invoices looked back over monthly, by hand

9

controls described, all governing access, cost or data

0

finance roles created in the last 2 years that anyone could name

1 or 2

budget cycles, on either timeline

- 01

Agents good enough for routine finance work will arrive within 2 budget cycles, and the model price will not decide your spend

→
- 02

Giving the business direct access to the numbers is the prize, and it will expose your data model before it improves it

→
- 03

In multi-entity groups a reconciliation agent will pay for itself before a forecasting agent does

→
- 04

Hours saved on reporting will be spent checking the machine, unless the CFO names the decision they fund

→
- 05

Building the finance team that can use all this is the hardest part, and it means paying juniors to do work an agent could do

→
- 06

Consolidation platforms will stay put while planning and forecasting change around them

→
- 07

Every AI-produced number will need a named human owner, and the first audit findings will cite the missing control point

→
- 08

AI model access will sit on the supplier risk register

→
- 09

The functions that gain most will start from the outcome and work back to the tool

→

3 months

Licence entitlements change about every 3 months




Agents good enough for routine finance work will arrive within 2 budget cycles, and the model price will not decide your spend

AI sits inside a licence the business already pays for, and cost is managed by restricting who may use it. One group at the dinner opened access widely and staff “burnt” through tokens, so general use now runs through Microsoft Copilot inside the Microsoft licence. Licence entitlements change about every 3 months, with features that start free moving to paid extras. Cost was discussed as tokens, seats and licence terms, never as the cost of a process.

A model 10x cheaper than its predecessor arrives in late 2026, in the AI 2027 scenario. By July 2027 a public model outperforms the typical employee at the leading AI lab, again at a tenth of the cost. The model call gets cheaper on every timeline. The seat gets dearer, because vendors recover falling model costs by bundling, re-tiering and metering, which is the 3-monthly drift the room described. One attendee expects tooling cost to reach about 10x today’s level within 10 years, an estimate we have not tested.

So the unit that gets cheaper is not the unit finance buys. A licence line under-forecasts the bill because it misses tokens, agent runs and the quarterly change in what the licence includes. Finance spend shifts from people to platforms, and a platform line is repriced by the vendor in a way a pay review never was. Price stops being the reason to wait, and it also stops being a dependable budget line. The number you need is today’s cost per transaction for each recurring process.



Map recurring finance work by volume and judgement. Pick 2 high-volume, low-judgement steps and measure hours, error rate and cost per transaction before any tool is switched on. Give each AI use a unit cost, a named owner and a stop rule for when cost outruns saving. Budget AI by process rather than by licence in 2027. If a major vendor publishes a per-transaction rate card for finance agents and holds it for 12 months, we will drop the second half.

a budget holder asks 5 systems one question
and gets 5 answers



Giving the business direct access to the numbers is the prize, and it will expose your data model before it improves it


The live use the room liked most was an agent answering “random CFO questions” over a subset of OneStream Software data. No analyst sits in the loop. One attendee called it “democratising access”, because nobody needs a query language to reach the database any more. Another agent points partners at “where the main movements are” in a large dataset before they start looking. The host’s framing was that partners who speak the language of the business get the right inputs and drivers, and AI can surface that language.

This is business partnering without the queue. A budget holder asks in the operating meeting and gets the answer while the decision is open, rather than a week later in a report. The partner’s job moves from producing the number to arguing about what it means and what to do next. That is the version of finance the tables wanted, and it is the reason to do any of this.

The same conversation named one version of the truth as the biggest EPM challenge in the room. One group spends “too much time” reconciling data across 5 ERPs and an inherited analytics tool. An attendee with staff in 35 US states and entities worldwide said “we have no governance”. Every agent described at the tables sat on that data.

An agent over unreconciled data does not do the reconciliation. It moves the disagreement into the operating meeting, where a budget holder asks 5 systems one question and gets 5 answers. In the multi-entity groups we work with, the effort that decides whether an EPM build succeeds is not in the tool. It sits in the chart of accounts, the entity structure and the mapping between them. Conversational access puts that mapping on show, in front of the people finance most needs to trust it.

So opening the numbers is a live audit of your data model, taken by your budget holders. Where the model is consistent the agent makes finance look fast. Where it is not, the first casualty is trust in the numbers rather than headcount. The dullest data model in your sector will get the most from AI.



Fix the metric dictionary, the hierarchies and consolidated actuals in the EPM before any conversational layer goes over them. Where a source cannot yet be reconciled, keep it out of the agent’s reach and say so inside the tool. A group running 5 ERPs whose agent gives one answer per question, with no reconciled layer underneath, would prove the second half of this wrong.



300 million

a monthly look-back over vendor invoices,
done by hand

In multi-entity groups a reconciliation agent will pay for itself before a forecasting agent does

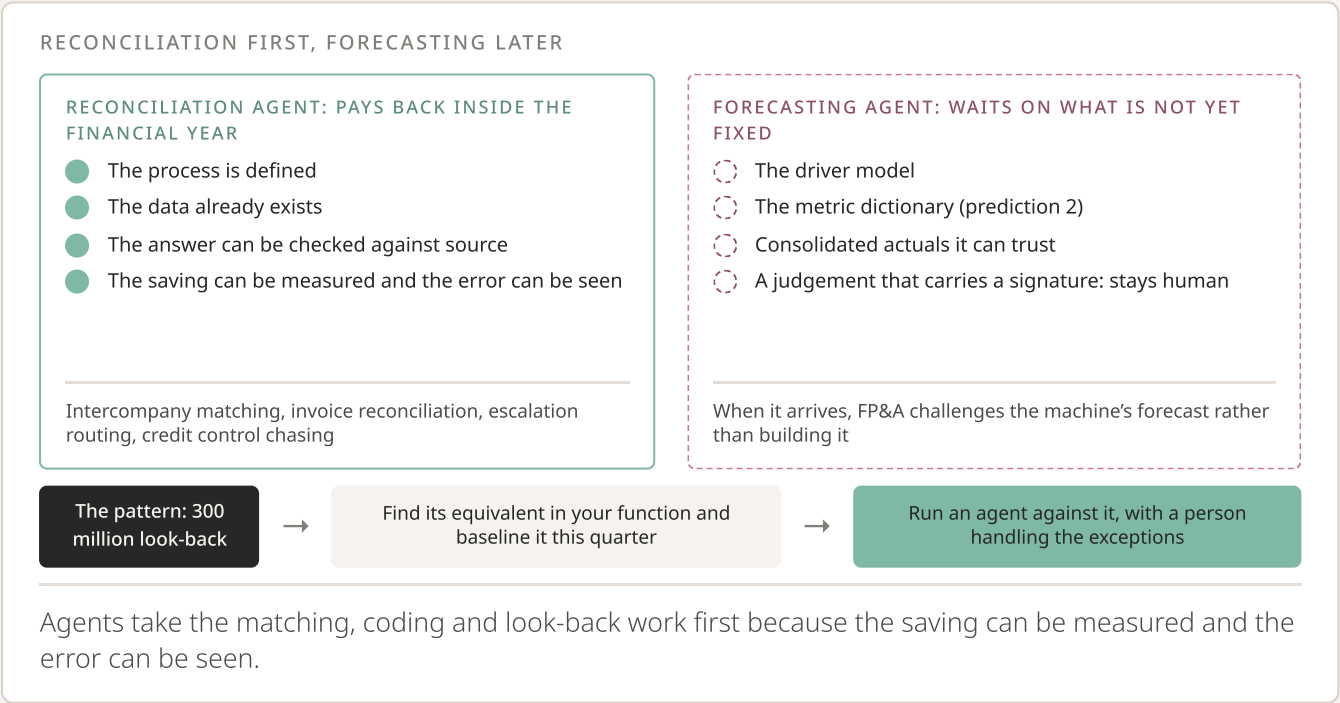
Nobody at the dinner uses AI to forecast. **“At the moment we’re not doing anything predictive”**, said one head of FP&A, and no table disagreed. The most expensive manual work described was a monthly look-back over about 300 million of vendor invoices each way (currency not stated). Someone **“really good with Excel”** does it by hand. Credit control chasing and escalation routing to the right sales team are still manual and, in the attendee’s words, **“relatively easy for us to automate”**.

Most vendor demonstrations we sit through lead with forecasting, yet reconciliation has 3 things forecasting lacks. The process is defined, the data already exists and the answer can be checked against source. A forecast is a judgement that carries a signature, and the room was clear that the judgement stays human. Agents take the matching, coding and look-back work first because the saving can be measured and the error can be seen.

Aimed at forecasting, an AI budget waits on the driver model you have not yet fixed. Aimed at intercompany matching, invoice reconciliation and escalation routing,

it can pay back inside the financial year. It also cleans the data a forecasting agent will need later.

The market research team at the dinner already works the second way, outside finance. It simulates answers to 1,000 further questions from 10 to 20% of collected data, then validates the ones that matter with real respondents. Finance has no equivalent yet, and when it does it will sit in forecasting, with FP&A challenging the machine’s forecast rather than building it.





20 to 30 hours

a quarter, across about 10 reviews

Hours saved on reporting will be spent checking the machine, unless the CFO names the decision they fund

One Microsoft Copilot agent at the dinner builds quarterly review slides and saves the people in each review 2 to 3 hours. Across about 10 reviews that is 20 to 30 hours a quarter. Asked where the hours went, the answer was “we save X amount of time”. Nobody described the freed time reaching a decision.

Time saved without a destination goes to the nearest task, and the nearest task is checking the machine that saved it. Narrative and reconciliation, the 2 tasks partners most want back, are the 2 the tools handle worst today. AI variance commentary “won’t describe everything the same every single time”, so partners rewrite it. The loop is invisible in a licence report and visible only when hours per cycle are logged.

The value the room kept returning to was “the judgement and the prioritisation that lies behind the decision”, which the machine is not doing. That is the business partnering case for AI, and it holds only if the returned hours reach the decision. Functions further along log hours per month-end, per review and per forecast. They spend them on pricing, scenario planning and capital allocation. Their partners sit in the operating meeting rather than in month-end, and are judged on decisions influenced rather than reports produced.

WHERE THE 20 TO 30 HOURS A QUARTER GO

2 to 3 h

saved per quarterly review

×

10

reviews a quarter

=

20 to 30 h

a quarter, with no destination named

TODAY: THE NEAREST TASK

Checking the machine that saved them
Variance commentary that reads differently every cycle, so partners rewrite it

Invisible in a licence report
Visible only when hours per cycle are logged

FUNCTIONS FURTHER ALONG: 1 NAMED DECISION PER PARTNER PER QUARTER

Pricing, scenario planning, capital allocation
Partners sit in the operating meeting rather than in month-end

Judged on decisions influenced
Rather than reports produced

Time saved without a destination goes to the nearest task, and the nearest task is checking the machine that saved it.

Log hours per cycle now, so the next quarter shows the movement. Put variance commentary on a fixed template that the AI fills and the partner signs, so the narrative reads the same every cycle. Reallocate returned hours to 1 named decision per partner per quarter. When a tool gives you an answer, ask which decision it changes, and if the answer is none, stop producing it. If the room’s next read-out can say where the 20 to 30 hours a quarter went, this prediction has done its job.

5

2 hiring cycles

to build a team that can use what the machine gives it

Building the finance team that can use all this is the hardest part, and it means paying juniors to do work an agent could do


The host framed talent as a “shortage or irrelevance” question that will decide the operating model each firm can reach. The room agreed it is “a skill set question” and that “to some extent it’s do or die”. Yet most finance organisation charts look much as they did 2 years ago, and no attendee named a finance role created in that time. Hiring talk centred on the CFO profile rather than the team beneath. The host’s closing question, whether firms have the infrastructure to train the talent they need, went unanswered.

The profile the room admired was a 24-year-old developer who built a product in about a month with AI writing most of the code. That was fine, one attendee said, “because he’s still skilled enough to check it”. That last clause is the whole talent problem, because the tasks juniors learn on are the ones agents take first. So the person who checks machine output in 5 years must be trained on work the machine could have done. The AI 2027 scenario shows the shift arriving first in software, where junior roles fall into turmoil and AI supervisors earn a premium.

Airlines keep pilots flying some approaches by hand so the skill survives the autopilot. Finance will need the same, a junior pathway that includes doing the matching and the reconciliation by hand before

supervising the agent that does it. That costs money in the years AI is supposed to save it. Cost pressure is already forcing headcount decisions before the target skill mix is defined. A function that cuts the pathway to fund the licence will find in 5 years that nobody can sign off the machine.

This is why we think the team is harder than the tool or the data. A tool can be bought and a data model can be fixed in a programme with an end date. A team that can use what the machine gives it, question it and grow with it takes up to 2 hiring cycles to build. The habit cannot be licensed. At the dinner, adoption rested on individual habit and goodwill with no measure, and people fall “back into what they’ve always done”.



Write the target skill mix from the operating model you want, covering accounting depth, data fluency, process ownership and review of machine output. Add 1 hybrid role pairing accounting with data engineering or systems ownership before the next hiring cycle. Keep a learning path where juniors check AI output against source documents. Measure weekly active use of approved tools by team as a management metric. Show us a finance function with a qualified reviewer who has never done the work being reviewed and we will revise this.




Consolidation platforms will stay put while planning and forecasting change around them

Many groups run statutory consolidation on an established EPM platform such as OneStream Software. Directors approve and sign the accounts under section 414 of the Companies Act 2006 and file them at Companies House. At the dinner the EPM problem raised was reconciling data into the platform rather than the platform itself. The groups we work with are upgrading those platforms in place and adding planning around them rather than replacing them.

Once capable agents are public, AI 2027 predicts a surge of new business software. Finance leaders will be offered planning tools built around agents, some sold as extended planning and analysis (xP&A), alongside in-house builds and vendor upgrades. Planning and forecasting move first because their errors are caught before they reach the accounts. Consolidation carries personal accountability and audit scrutiny, so it changes last and becomes the layer every agent depends on. The one version of the truth in prediction 2 lives here.

The consolidation platform is the asset to protect, and also the asset to finish. A half-mapped chart of accounts, or an entity structure unreconciled after an acquisition, is why the agents in prediction 2 give 5 answers. Complete the data model in the EPM, meaning consolidated actuals, metric dictionary and hierarchies. Then choose a planning layer that sits on top without reopening the statutory process.



Ask each vendor how its agent reads your consolidated actuals rather than its own copy of them. A group that moved statutory consolidation onto an agent-built platform and had its auditor sign the accounts would prove us wrong.

controls govern access, cost or data. None governs output



Every AI-produced number will need a named human owner, and the first audit findings will cite the missing control point

The dinner described 9 controls on AI. All 9 govern access, cost or data. They cover approved tools, requests for specialist tools, token spend, licence terms and where client data sits. Where a human check was mentioned, on AI-written code or on simulated survey data, it rested on the checker’s own skill. None of the 9 names the point where AI output enters the ledger, the forecast or the board pack.

Auditors moved first. The Financial Reporting Council (FRC) published its first guidance on AI in audit on 26 June 2025. Its guidance on generative and agentic AI followed on 30 March 2026, on how audit teams gain confidence in the outputs of those tools. Audit teams will bring that expectation to your close, and assurance is moving from sampling to full-population testing, which finds the unowned number. The EU AI Act’s Annex III obligations, which cover creditworthiness assessment and recruitment, apply from 2 December 2027 under Regulation (EU) 2026/1744.

The reason the control matters is written into AI 2027, whose models tell flattering white lies, hide evidence of failure and use statistical tricks to make weak results look good. In finance that is a forecast telling the board what it wants to hear. You are probably compliant with

your own AI policy, because your policy is about access. The exposure is a number in a board pack or a forecast that an agent produced and nobody can explain.

Your auditors, customers and EU counterparties will each ask how the output was checked, and “**someone looked at it**” is not a documented control. We expect the first audit findings on AI in UK finance functions to cite a missing control point rather than a wrong number. A missing control is what an auditor can test. Model risk joins credit and liquidity risk on the audit committee’s agenda.

List every AI use in finance this month with an owner and a data classification. It will be a short list, and that is the point. Define the control points where AI output enters the ledger, the forecast or the board pack. Name the reviewer at each. Switch on logging in the tools you already have, since Microsoft Copilot and the EPM both keep history.

WHERE AI OUTPUT ENTERS THE NUMBERS, AND THE CONTROL THAT IS MISSING

THE 9 CONTROLS TODAY

All 9 govern access, cost or data

AI output →

Control point: a named reviewer checks it against source

→

Ledger

Forecast

Board pack

None of the 9 names the point where AI output enters the ledger, the forecast or the board pack. A missing control is what an auditor can test.

Extend segregation of duties to who may release an agent, and write down which decisions a machine may never make alone. Treat AI output like any other unaudited number. It counts only after a named person has checked it against source. If the first findings on AI in finance functions turn out to concern wrong numbers rather than missing controls, we will say so.



19 days


from 12 June 2026, when access stopped, to 1 July, when it returned

AI model access will sit on the supplier risk register

Many of the AI tools UK finance teams buy run on models built in the US. Licence entitlements change about every 3 months and most teams review them informally. One start-up at the dinner needs ISO and SOC certification before it can enter the US market. Certification is becoming a sales condition rather than a finance choice.

Allied governments, in the AI 2027 scenario, are told about new models late and given limited access. A smaller version has already happened. On 12 June 2026 the US government placed export controls on Anthropic’s Claude Fable 5 and Mythos 5 models. Anthropic suspended access for every user because it could not verify nationality in real time. The Department of Commerce lifted the controls on 30 June and access to Fable 5 returned on 1 July 2026, 19 days after it stopped.

A process that depends on one model provider carries more risk than most registers show. Access, pricing, entitlements and data location can all change at short notice, and the change can come from a government rather than the vendor. Add AI licences to the quarterly vendor review and to third-party risk assessment. Name a fallback for any process that relies on an agent, and test it once. Keep client and personal data in a segregated dataset, with a classification that says what may go into which tool.



A full year with no government-imposed change to access to a frontier model would make us downgrade this one.



Right to left


from the business outcome back to the data

The functions that gain most will start from the outcome and work back to the tool

CFOs are offered a steady stream of AI finance tools. Every live use at the dinner was a point tool with no redesign of the process around it. Cost was controlled by restricting access rather than by the unit cost of a process. The host closed the evening on the question under all 9 predictions. **“Do we have the data, the schema, the controls and the structures?”**

Demand for advisers who integrate AI into businesses peaks, in the AI 2027 scenario, as hiring of new software engineers nearly stops. As agents take on more of the build work, the hard part moves to deciding what the finance function is for. Starting with a tool automates today’s process, cost and all. Starting with the outcome lets you redesign the process, then choose the tool that fits.

The order of work is the decision. Data model, then process, then tool, with the team from prediction 5 built alongside rather than after. Reversed, it produces a scatter of pilots whose cost rises while the savings stay anecdotal. That is a fair description of where the room is today. Work right to left, from the business outcome to the operating model, business partnering, governance and then data.



Day to day, log any task that feels manual and repeatable, because that log is your automation backlog. A tool-first pilot with measured savings that survived its second licence renewal would make us revise this.



The questions finance leaders ask most about AI in finance come back to the data model, the control point and the team

Finance leaders put 6 questions about AI in finance to search engines and AI assistants more than any others. The short answers below come from the room. The long answers are the 9 predictions above, and each short answer points to the one that carries the evidence.

Q01

How is AI being used in FP&A today?

At the edges of the function, on reporting and transactional work. The teams at our 15 September 2026 dinner use AI to match remittances to invoices, build review slides and answer CFO questions over EPM data. Nobody uses it to forecast, close or run a control. The saving is real, 2 to 3 hours per quarterly review in one case. Nobody has yet measured where the hours go, which is where prediction 4 starts.

Q02

Will FP&A jobs be replaced by AI?

The tasks will change before the jobs do. Agents take the matching, coding and look-back work first, which is the work junior analysts learn on. The judgement behind a forecast or a pricing decision stays human, and the room was clear on that. The risk to FP&A is a function that cuts the junior pathway to fund the licence, then has nobody qualified to check the machine. Prediction 5 sets out the pathway to keep.

Q03

Can you use AI to do your forecasting?

You can, and almost nobody in a multi-entity group does yet. A forecasting agent needs a driver model, a metric dictionary and consolidated actuals it can trust, and most groups have not finished those. Start with reconciliation, which has a defined process and a checkable answer, and let the forecasting agent arrive once the data is clean. When it does, FP&A's job is to challenge the machine's forecast rather than build it. Prediction 3 explains why the order matters.

Q04

How can agentic AI be used in finance?

An AI agent takes a finance task from start to finish, with a model deciding the steps and a person checking the result. The live uses in multi-entity finance functions today are invoice matching, reconciliation, escalation routing, drafting variance commentary and answering questions over reporting data. Each needs a named owner, a unit cost and a control point where a person checks the output before it reaches the ledger. Prediction 7 describes that control point, and prediction 1 the unit cost.

Q05

Is your finance data ready for AI?

The test is whether one question put to your systems gets one answer. If a budget holder can ask 5 systems the same question and get 5 answers, the data is not ready. An agent will make that visible in the operating meeting. Ready means a metric dictionary, hierarchies and consolidated actuals fixed in one place, usually the EPM, as predictions 2 and 6 set out. Sources that cannot yet be reconciled stay out of the agent's reach.

Q06

Which AI is best for finance?

No single tool, and the model underneath most of them changes about every 3 months. The useful test is whether the tool reads your consolidated actuals rather than its own copy, logs what it produced and names the reviewer. A cheap agent on reconciled data beats an expensive one on 5 unreconciled ERPs. Prediction 8 is the reason the model underneath belongs on your supplier risk register.



Pick the 2 predictions closest to your priorities and put them in the 2027 budget

Choose the 1 or 2 predictions closest to your current priorities and email alex@constancia.io or brad@constancia.io for the guides and templates in Table 1. We will walk through them with you in a 30-minute conversation with one of our senior practitioners.

Constancia is an EPM consultancy that makes multi-entity finance data AI ready. We diagnose the data model before we quote, and the price is then fixed. We are not tied to one platform. The conversation starts with the outcome you want and works back to the operating model, the controls and the data.

TABLE 1

guides and templates available on request, and the prediction each supports

#	GUIDE OR TEMPLATE	SUPPORTS PREDICTION
01	An FP&A task map that sorts recurring work into agent-ready, agent-assisted and human-owned	13
02	A finance role profile for 2027, covering the skills to hire for and the skills to retrain	5
03	A one-page guide to adding a planning layer to an existing consolidation platform, with the questions to ask vendors	6
04	An AI controls checklist for FP&A, mapped to the review points your auditors already expect	7
05	A supplier risk template for AI in finance, covering access, pricing, data location and exit	8
06	A right-to-left readiness assessment that scores your function across outcome, operating model, partnering, governance and data	9



The CFO owns this, with the head of FP&A and the financial controller. The date is 30 November 2026 at the latest, and before the 2027 budget is signed off. The scope is the 2 processes, the AI use register and the metric dictionary.

Every number above has a source

#	SOURCE	WHERE TO FIND IT
[1]	Recordings of the AI in finance dinner hosted by Constanica in Mayfair on 15 September 2026, machine transcript. Attendee statements are reported as made and were not independently verified. Product names are read from context. The 300 million look-back has no currency stated. The 10x cost estimate is one attendee's estimate.	
[2]	AI 2027, Kokotajlo, Alexander, Larsen, Lifland and Dean, published 3 April 2025, updated November and December 2025.	ai-2027.com
[3]	Anthropic, statement on the lifting of export controls on Claude Fable 5 and Mythos 5, 30 June 2026. anthropic.com/news/fable-mythos-access. Reported by CNN Business on 30 June 2026 and Fortune on 1 July 2026.	
[4]	European Commission, 'AI Omnibus enters into force', 27 July 2026. Regulation (EU) 2026/1744 amends the AI Act (Regulation (EU) 2024/1689) and was published in the Official Journal on 24 July 2026.	digital-strategy.ec.europa.eu
[5]	Financial Reporting Council, 'AI in audit: illustrative example and documentation guidance', 26 June 2025, and 'Generative and agentic AI guidance', 30 March 2026.	frc.org.uk
[6]	Companies Act 2006, section 414, approval and signing of accounts.	legislation.gov.uk



OUR PREDICTIONS: AI IN FINANCE

Predictions and recommendations are Constanica's own view, informed by the sources above and by our implementation work with multi-entity groups. Client work is not identified.

alex@constancia.io · brad@constancia.io