



BTRM

The Certificate
of Bank Treasury
Risk Management

Client Webinar Series

Webinar III: Artificial Intelligence (AI) and Bank Balance Sheet Management

30th July 2026

Professor Moorad Choudhry
Faculty BTRM



Agenda

- /// Webinar I: Updating your liquidity risk management framework
- /// Webinar II: Interest-rate risk in the banking book (IRRBB)
- /// **Webinar III: Artificial Intelligence (AI) and Bank Balance Sheet Management**
 - /// Opportunities and challenges in adopting AI
 - /// Where generative AI can deliver genuine productivity gains
 - /// When (and when not) to introduce agentic AI into processes
 - /// Why small, focused language models often outperform broad ones in ALM
 - /// Defining the outcome must come before building the model
 - /// Illustration and examples: AI in Bank Treasury

“AI is taking all (y)our jobs!”

- /// Attributed to some very wealthy, clever (*sic*) individuals
- /// Leaving aside at least 101 “jobs” that my wife and I compiled in about 10 minutes that we think AI will never, ever replace (see overleaf)....

We bite our thumb at AI.... *

* mis-quoting Sampson from "Romeo & Juliet"....

Professional sportsman/sportswoman	Make-up Artist
Sports coach	Fireman
Physiotherapist	Child Minder / Nanny
Chiropractor	Jockey
Electrician	Horse Trainer
Plumber	Dairy Farmer
Gas heating engineer	Electronic Warfare Specialist
Bricklayer	Actor
Podiatrist	Merchant Seaman
Midwife	Oil Rig Worker
Health visitor	Forest Ranger
Chef	Lifeguard
Sommelier	Dog Breeder
Caretaker	Milkman
Cleaner	Stonemason
Piling engineer	Movie Director
Carpenter	Entrepreneur
Tree Surgeon	Barista
Gardener	Harbour Pilot
Refuse collector	Window Cleaner
Vicar	Double-Glazing Salesman
Referee	Milliner
Paramedic	TV News Anchor
Ice Cream Seller	Wedding Photographer
Tailor	Antarctic Research Scientist
Cake Decorator	Tour Guide
Beautician	Mountain Rescue Guide
Masseuse	Security Guard / Doorman
Yoga Instructor	Children's Party Entertainer
Woodwind Instrument Teacher	Piano Tuner
Aircraft Carrier Flight Deck Crewman	Dry Dock Engineer
Hairdresser	Walter / Waitress
Mining engineer	After-Dinner Speaker / Raconteur
DEA Agent	Feng Shui Consultant
Customs Officer	Bank Board Chairman
Baggage handler	Abattoir Worker
North Sea Diver	Care Worker
Fisherman	Hot-Air Balloonist
Butcher	Gynaecologist
Optician	Police Dog Handler
Dentist	Chocolatier
Funeral Celebrant	Sports Groundsman
Funeral Bearer	Sewage Worker
Florist / Flower Arranger	Roadmen
Seamstress	Zoo Keeper
War Correspondent	Museum Curator
Stand-up Comedian	Archaeologist
Orchestra Musician	Palaeontologist
Vet	Antique Dealer
Fashion Model	Reality TV Contestant
Theme Park Attendant	The Pope






60%



“AI is taking all (y)our jobs!”....


- ///what roles / jobs can “AI” perform more effectively, and more efficiently, than humans (and processes), which it would be logical for bank Treasuries to adopt?
- /// This is a good question
- /// I’d like to kick off by offering an analogy:
 - /// You’re a 16-year old musician and have just discovered the electric guitar
 - /// You marvel at the endless possibilities offered by this magical new piece of kit, and instantly take it up
 - /// And then join an existing rock band as their new guitarist
 - /// How do you play it? What sort of guitarist do you end up as?
 - /// ***The possibilities are endless.....***

From the sublime....



To the ridiculous....

/// (Absolutely no disrespect intended! 😊)



Introduction

- /// Artificial intelligence (“AI”) presents almost limitless possibilities, for any kind of user
- /// Banks are already using AI in its different guises for different applications
- /// This discussion will limit itself to the potential impact of AI – advantages, challenges, risk management issues – on the Bank Treasury and B/S risk management function
- /// Our starting premise is, “Adopting Generative AI, and (with care) Agentic AI, brings benefits to a bank Treasury function such that the Asset-Liability Committee (ALCO) will be more effective in its search for the Holy Grail of ALM discipline, the ‘optimised balance sheet’.”
- /// But first, an observation from my friend Professor Alexandros Kontonikas (Essex University) who makes a very pertinent point, *viz.*,
 - /// Commentators, regulators and analysts always speak in terms of “Banks should do this...or banks are impacted by, etc etc” – as if banks are universal, identikit monolithic creatures with always one direction and one thought process
 - /// In fact banks come in many different shapes and sizes, and within each bank there are always competing interests and objectives, from the Board down
 - /// What may make sense for one bank to do – in any aspect of its operation – may be quite inappropriate for another bank
 - /// ***We must always remember this, whether we are discussing AI, or anything else in banking for that matter***

What is Artificial Intelligence?

(A) Definition:

AI: simulation of human intelligence processes by machines

/// Three key paradigms

/// **Narrow AI:** specific tasks — image recognition, language translation

/// **General AI:** human-level reasoning across all domains (*still theoretical*)

/// **Super AI:** exceeds human intelligence (*speculative, future horizon*)

/// In banking: **practical narrow AI** applied to financial workflows

/// Why narrow? ECB (FSR, May 2024): “broad consensus that all AI technologies are currently weak/narrow AIs”

/// Each model solves ONE task: a cash flow model cannot hedge FX or run a stress test

/// Narrow scope = auditable, back-testable, and aligned with regulatory model risk frameworks

/// AI is data-driven models that learn from historical patterns

/// Deployed to automate, augment and enhance treasury operations

AI is meant to amplify the Treasurer — not replace them



AI – Taxonomy & Overview

/// Succinct taxonomy:

/// Artificial Intelligence

/// | Machine Learning

/// | | Neural Networks

/// | | Deep Learning

/// | | Large Language Models (LLMs)

/// | | GenAI more generally, including images, sound

/// | Natural Language Processing (NLP)

/// | Natural Language Generation (NLG)

/// | Modern implementations use LLMs

/// | Where in this tree am I placing Agentic AI...?

/// Neural networks (deep learning etc) are black boxes. The lack of visibility of interpretable variables means overfitting is not a concern.

/// LLMs are black boxes and hallucinate “openly”.

/// *Small Language Models – where do they fit? – Moorad*

/// Machine learning includes automation of orthodox data science / statistical methods. Might increase risk of overfitting

AI vs ML vs Deep Learning vs GenAI

Traditional AI & Machine

- /// **Traditional AI** is Rule-based: explicit programmer-defined logic
- /// **ML**: learns from data without explicit rules
 - /// Supervised: labelled data (classification, regression)
 - /// Unsupervised: patterns without labels (clustering)
 - /// Reinforcement learning: reward-based trial and error
- /// **Deep Learning**: multi-layer neural networks
 - /// Excels at unstructured data: text, images, audio
 - /// Foundation for modern AI in banking

Generative AI & LLMs

- /// **LLMs** trained on vast text corporate via transformers
 - /// Few-shot learning: powerful from small examples
 - /// In-context learning: adapts via prompt engineering
- /// **GenAI**: creates new content — text, code, images
- /// **Key models in banking context**
 - /// **GPT-4o (OpenAI)**: strong coding and reasoning
 - /// **Claude (Anthropic)**: long context, safety-focused
 - /// **Gemini (Google)**: native multimodality, massive context window
 - /// **LLaMA 3 (Meta)**: open source, fine-tuneable
 - /// **CoPilot (Microsoft)**

How AI can be material in Bank Treasury

30%

**reduction in idle cash
buffers**

AI-powered cash flow forecasting improves accuracy, reducing the precautionary buffer held overnight. For a EUR 10bn balance sheet, this can free EUR 300M+ in deployable liquidity.

3 days

↓ **2 hours**

**per ALCO reporting
cycle**

RAG-powered report generation pulls live data, drafts commentary, and formats the dashboard — replacing days of manual assembly.

1,000

**stress scenarios in
seconds**

AI runs full liquidity stress test scenarios in parallel, versus hours for a manual model. ECB and Basel IV demand granular, real-time analytics.

Your Treasury AI use cases

/// For you to complete after the webinar

/// ...

/// ...

/// ...

/// ...

/// (Not there are any prizes on offer – mainly as there are no wrong answers!)

/// But first....a bank needs an AI Policy



AI Policy case study: “neo” challenger bank

- /// A template for any bank to adopt...I quote from its AI Management Policy:
- /// “We use AI technologies to promote our customers' financial wellbeing”
- /// “Responsible AI”: Neo-bank is guided by the following Responsible AI Principles:
 - /// “Trustworthy and Transparent AI: we are transparent about our AI use, ensuring our AI Systems are aligned with our customers interests, building trust.
 - /// “Continuous innovation in AI: we are constantly improving AI capabilities to meet customers needs, improve efficiency and productivity, and lead innovation.
 - /// “Accountability and Governance: we take full ownership of our AI governance, with clear roles and responsibilities throughout our AI System's life-cycle.
 - /// “Robustness, security and safety: we adopt comprehensive protection measures to address AI-related risks.
 - /// “Bias and Awareness: we actively work to ensure data diversity where relevant, data accuracy, and reasonable impact of our AI Systems over people and our ecosystem.
 - /// “Democratising AI: we envision accessible AI for all, to empower our customers, and foster creativity and innovation.”
- /// [BTW this bank has also implemented an Agentic AI model to undertake model validation review and sign off on model approval...]

AI and ALCO decision making: Quoting BTRM Faculty **Dr. Tom Ho ***

- /// “A 500-page MI deck doesn’t help if it doesn’t help you make a decision”
- /// Amen.
- /// This is an excellent Use Case for any bank: Gen AI assisting decision making at ALCO level
- /// AI adds value when it is
 - /// Integrated into the decision-loop
- /// Limits of the traditional approach to bank ALM:
 - /// Which decisions should be made now?
 - /// Who is accountable for those decisions?
 - /// On what assumptions are decisions based?
 - /// How do decisions interact across the balance sheet?
 - /// How quickly can outcomes be evaluated and adjusted?

* Quoted from *Income Risk Intelligence* © Thomas Ho Company 2025

AI and ALCO decision making: Dr. Tom Ho...

/// When decisions are siloed by function and reviewed only through backward-looking reports, management is left reacting to outcomes rather than shaping them.


///centre ALM into the decision-making loop – AI facilitates this

/// And the execution element?

At any of these stages we can benefit from the integration of an AI model....

The training of that is key.

Figure 1: The duality of Decision Loop and Decision Collaboration



© Tom Ho 2026. Used with permission.

AI and ALCO decision-making...

- /// For example:
- /// “Claude” has been available in MS Excel since January this year
- /// One can produce Excel sheets quickly, or incorporate its output within existing Excel sheets...
 - /// Given how many banks are still using Excel for their (e.g.,) IRRBB calculations or inputs to regulatory reporting templates, this is an area for instant adoption
 - /// To quote Dr. Tom Ho, “*Claude is not using a quantitative ALM model yet. But the complexity is not in using Claude, but in quantifying the ALM key metrics....*”
 - /// We’ll consider this later
- /// Although, is this better as a training tool rather than for one’s actual output, which may demand a tailor-made system?



Case study: teaching and learning

- /// BTRM Alumnus Laurent Birade took this article from...
- /// *Journal of Risk Management in Financial Institutions*, Vol17 no3 (Summer 2024)
- /// ...and turned it into something that works better as a teaching and learning tool
- /// Quickly and simply by using Notebook LM to summarise it

The screenshot shows a LinkedIn feed. The main post is by Moarad Choudhry, Non-Executive Director of Loughborough Building Society. The post text reads: 'Possibly my last formal word on the subject, I'm pleased to post our article "Bank liquidity risk management through stable and volatile markets: The role of the asset-liability committee and lessons learned for the balance sheet governance operating model". This considers the ALCO governance framework in place at a number of failed banks in 2007-09 and 2023. The article is published this month in the Journal of Risk Management in Financial Institutions, vol. 17 no.3 (Summer 2024), see link to the journal itself in the comments below. Last word? I first published a chapter on ALCO back in 2001, and subsequently in books and articles in 2007, 2012, 2017 and 2023. Time to move on perhaps... :-). I am indebted to my co-authors Claire Trythall, (The Certificate of Bank Treasury Risk Management (BTRM)) and Diyama Awadallah (Birzeit University), it was my privilege and pleasure to work with them. Thank you also to the Editor at JRMFI, Julie Kerry, I hope you find the article of interest, or at the very least find the subject matter to be worthy of debate! Enjoy...'. Below the post, there is a section titled 'Practice Papers' with the title 'Bank liquidity risk management through stable and volatile markets: The role of the asset-liability committee and lessons learned for the balance sheet governance operating model'. The post also includes a 'Follow' button and a 'Messaging' button.


Case study: “Horizon scanning”

- /// Keeping up with the continuous stream of publications from regulatory authorities – guidance, policy statements, supervisory statements, Final Reports, consultation papers, etc etc – is a full-time job
- /// Senior execs have to be brought up to speed quickly on the latest pronouncements and exactly what implications they have for the bank in terms of regulatory compliance
- /// AI is easily employed to summarise the content of the latest PS...
 - /// How reliable is the summary? Is “97%” good enough? (True story, related to a rating agency review)
 - /// Will we still need someone to review the AI summary before issuing? If yes, does this defeat the object?
- /// But these issues are not insurmountable
 - /// I gave it a shot with a very recent UK PRA Consultation Paper, CP5/26 (March 2026), which is VERY relevant to bank ALM practitioners in the UK (and the principles it discusses to bank ALM desks everywhere)
 - /// And to keep it interesting, I opted for the “Kawaii” option in Notebook LM >>>

First cut...



Mustn't lose focus...the title is "PRA CP 5/26"!





The verdict?

- /// Not bad at all....
- /// Strangely, the reiteration to include the CP reference omitted the most important part of the paper...the requirement to assess immediate short-term stressed outflow risk (circled in the first cut)
 - /// If a human has to review the AI summary before issuing to the Senior Execs (“97%” rule again!), does that defeat the object? I won't know if the AI output is a useful – and complete – summary without reading the regulator document myself, will I..?!
- /// As John Lennon sang in Nobody Told Me, “....strange days indeed!”
- /// Maybe I should have opted for the serious non-Kawaii option? 😊

Sidebar: PRA CP 5/26

- /// This requirement from the CP was amongst the items in the Notebook LM summary:
 - /// ***Design "Initial Day" Stress Scenarios: Develop a new stress scenario focused on sudden and severe liquidity outflows that peak within the first week of a stress event***
- /// That's a result! Yours truly advocated this in June 2023, suggesting a 3-, 5- and 7-day "LCR" (my "LCR" is actually a "Liquid Cash Ratio" as I include G-SIFI bank deposits in the numerator 😊)
- /// BTRM liked the recommendation so much we made it into a YouTube video:
- /// <https://www.youtube.com/watch?v=CbyiFhtBOYA>



Case study: the 1-pager ALCO dashboard

- /// But as Jermain Jackson once sang, “Let’s Get Serious...”
- /// Board and ALCO MI packs are – in most banks – not as effective and as efficient as they can be, in part due to plethora of key risk indicators included
 - /// Horror story from recent personal experience: a 400-slide ALCO pack (*sic*)
- /// Should we ask Notebook LM to suggest a 1-page summary KRI dashboard for Board and (for the start of) the ALCO pack?
- /// Yes – and here, ideally one would build a Small Language Model trained on the bank’s own balance sheet data, and existing ALCO pack design, and use prompts to target the senior exec readership
- /// But for now, we’ll just use Notebook LM (and eschew the Kawaii design sadly ☹)
- /// And modify this to suit our specific bank ALCO requirements

1-page ALCO MI dashboard

- /// Asked to summarise the most important KRIs from an ALCO pack that contained a LARGE number of metrics....see **overleaf**
- /// Its intriguing the KRIs the AI picked – again, more human review needed.
- /// ***My takeaway: flexibility in reiteration will allow the bank to tailor MI summaries for specific readership, and also modify quickly if/when required***
- /// ***But I still need to approve the final output myself if I'm ALCO Chair***
- /// True story:
 - /// Confronted with a Board pack over 400 slides long, where the “CRO Report” appeared well after slide 200, and which contained a “red” KRI value, and which didn't appear until then, I asked the CRO to produce a 1-slide MI dashboard of Tier 1 KRIs, to appear at the front of the Board pack.
 - /// That way the Board would be directed immediately to any risk appetite indicator outside “green”
 - /// Three months later I was still waiting for it...
- /// That CRO can now go to Notebook LM and deliver the NED's ask in an hour!

1. Regulatory Requirements

Liquidity Coverage
Ratio (LCR)

> 180%

Rationale: Statutory regulatory requirement ensuring sufficient high-quality liquid assets to survive a 30-day stress scenario.

Net Stable Funding
Ratio (NSFR)

> 125%

Rationale: Statutory regulatory requirement ensuring the bank maintains a stable funding profile in relation to its on- and off-balance sheet activities.

2. Core Appetite Metrics

Customer Loan-Deposit
Ratio (LDR)

< 85%

Rationale: Defines the bank's appetite for the extent of customer surplus funding of the balance sheet (excludes central bank facilities).

Survival Days

> 180 Days

Rationale: Defines how long the bank wishes to remain liquid in a severe, stressed "market lockout" scenario without relying on central bank funding.

3. Early Warning Indicators (EWIs)

3m LIBOR-OIS Spread

< 75 bps

Rationale: A vital macroeconomic leading indicator signalling systemic money market stress or widespread illiquidity.

Notice of Withdrawal
Balances

< 7.5% (Total
Balance)

Rationale: A crucial behavioural metric acting as a firm-specific early warning of excessive fund withdrawals and potential deposit runs.

4. Primary IRRBB Metrics

IRRBB Shock Loss Exposure

< £3m

Rationale: Measures the maximum capital loss exposure to Net Interest Income / Economic Value of Equity in a severe 200 bps interest rate shock scenario.

The KRI Ecosystem: From Early Warning to Regulatory Protection

**Components:**

EWIs (LIBOR-OIS Spread & Notice of Withdrawals).

Role: **The Tripwire.**

Insight: Detects systemic liquidity freezes and idiosyncratic deposit flight weeks before they manifest on the balance sheet.

Components:

Core Appetite & IRRBB (LDR, Survival Days, Shock Loss Exposure).

Role: **The Shock Absorbers.**

Insight: Ensures structural surplus funding and restricts capital loss exposure, buying the Board a guaranteed 180-day runway to execute Recovery Plans if the tripwire is breached.

Components:

Regulatory Thresholds (LCR & NSFR).

Role: **The Absolute Floor.**

Insight: By targeting internal thresholds drastically above statutory minimums (>180% LCR vs standard 100%), the bank entirely insulates its regulatory standing from intraday volatility and shock absorption impacts.




Board / ALCO 1-pager 21-Jun-2024

Key observations

- reduced survival days due to higher loan pipeline
- pressure on NIM as peer rates raising funding cost



Surplus Cash



BoE Reserve Account

£0.91bn

HQLA £1.01bn

Base rate 5.25%

Balance sheet £12.2


Non-equity liabilities:

Retail £6.4bn

SME £ 4.1bn

Corp £1.09 bn

3m SONIA



© Moorad Choudhry 2014, 2024

CET 1 ratio



23%



>21% <21% <18%

Excess over TCR



88bps

>75 <75 <50

IFRS9 Stage 3 of loan book



1.1%

<1.5% >1.5% >2.0%

EVE Delta T1 Shock



4.10%

<7.5% >7.5% >10%

NII Delta 12M T1 Shock



3.90%

<5% >5% >8.5%

LCR



218%

>160% <160% <140%

Survival Days



182

>180 >120 <90

LDR



78%

<85% >85% >95%

WA Contractual Tenor Funding



93

>90 <70 <60

Liquidity Risk Factor



7.2

>5.5 <5.5 <3.5

NIM

380bps

Weighted Average Rates

Lending Excl. Credit Cards

7.66%

Lending Credit Cards

23.10%

Liabilities Customer


3.84%

Liabilities Wholesale

5.01%

Loan Book £10.28bn

Pipeline 90d £398mIn




Notice Savings Deposits

On-notice outflow
£335mIn



NIM Series





BTRM

The Certificate
of Bank Treasury
Risk Management

Issues, challenges, ...

...and sundry other matters

/// Prompts in AI are all-important...LinkedIn post 5/3/2026

Inbox - mooradchoudhry@gmail.com x Google Calendar - March 2026 x (2) Feed | LinkedIn x +

linkedin.com/feed/

Star Bookmarks Elsevier Royalty Link Nubank - Sign In E*TRADE Fidelity Login John Lewis Credit C... COURTTS login NatWest QuickBooks Nationwide Login Lloyds Bank Google Santander ac All Bookmarks

in Search

Home My Network Jobs Messaging Notifications Me For Business Learning



Moorad Choudhry 
Chair, Board Risk Committee, SBUK Limited
Guildford, England
Sonali Bangladesh UK (SBUK) Limited

Profile viewers 1,995
Post impressions 20,153

Your Premium features

Moorad Choudhry Financ...

Activity 0
Page visitors 6

Grow your business faster
Try Premium Page for £0
Advertise on LinkedIn

See visitor analytics

Saved items
Groups
Newsletters
Events



Syed Arsalan Mushtaq, CFA  • 1st
VP - Head of Treasury Business Management | CFA Charterholder | Performance A...
2h • 

Most people prompt AI the way they used to Google things — one line, vague, hoping for the best.

[Ruben Hassid](#) (worth following) broke down what a well-structured Claude prompt actually looks like, and it's a useful framework for anyone using AI seriously at work.

The short version:

- (1) define the task with success criteria, not just "do this."
- (2) Upload context files with your expertise, rules, and preferences so the model doesn't start blind.
- (3) Give it a reference — show it what good looks like.
- (4) Write a brief specifying output format, length, and what it definitely shouldn't sound like.
- (5) Set your rules.
- (6) Then have a conversation before it executes, not after.

What strikes me is how different this is from how most of us actually prompt. We expect the model to read our minds, then blame the model when it doesn't. Better inputs produce better outputs. That's true in financial models. It's equally true in AI prompts. (Credit: Ruben Hassid)

[#AI](#) [#Banking](#) [#Treasury](#) [#Productivity](#) [#Claude](#)

£22m fine for South East Water

6h ago • 1,968 readers

Firms call for shipping climate deal

6h ago • 641 readers

How tech can aid staff with dementia

6h ago • 336 readers

Show more

Puzzles by LinkedIn

 **Zip - a quick brain teaser**
Solve in 60s or less!
97 connections played

Promoted ...

 **wire 2026 trade fair**
The world's leading expo for the wire industry - 13th-17th, April 2026.
 **Dr. Vassilia I.** & 3 other connections also follow wire and...

 **University Transcriptions**
Fast & Accurate Audio & Video Transcription Services. UK GDPR, CE+ ISO27001

About Accessibility Help Center
Privacy & Terms Ad Choices
Advertising Business Services
Get the LinkedIn app More

LinkedIn LinkedIn Corporation © 2026

The Anatomy of a Claude prompt

I want to [TASK] so that [SUCCESS CRITERIA].
First, read these files completely before responding:

```
[filename.md] — [what it contains]
[filename.md] — [what it contains]
[filename.md] — [what it contains]
```

Here is a reference to what I want to achieve:

[Upload reference file as markdown, or paste it here]
Here's what makes this reference work:

[Paste your reverse-engineered blueprint - the patterns, tone, structure, and rules you extracted from the reference. Format each one as a rule starting with "Always" or "Never."]

Here's what I need for my version:

SUCCESS BRIEF
Type of output + length:

Task

Context Files

Reference

11°C Sunny

Search

17:56 05/03/2026

Messaging 2



The Anatomy of a Claude prompt

I want to [TASK] so that [SUCCESS CRITERIA].
First, read these files completely before responding:

[filename.md] — [what it contains]
[filename.md] — [what it contains]
[filename.md] — [what it contains]

Here is a reference to what I want to achieve:

[Upload reference file as markdown, or paste it here]
Here's what makes this reference work:

[Paste your reverse-engineered blueprint - the patterns, tone, structure, and rules you extracted from the reference. Format each one as a rule starting with "Always" or "Never."]

Here's what I need for my version:

SUCCESS BRIEF
Type of output + length:
[Contract, memo, report, proposal, landing page, post?]
Recipient's reaction:
[What should they think/feel/do after reading?]
Does NOT sound like:
[What to avoid - generic AI, too casual, formal, jargon-heavy?]
Success means:
[They sign? They approve? They reply? They take action?]

My context file contains my standards, constraints, landmines, and audience. Read it fully before starting. If you're about to break one of my rules, stop and tell me.

DO NOT start executing yet. Instead, ask me clarifying questions (use 'AskUserQuestion' tool) so we can refine the approach together step by step.

Before you write anything, list the 3 rules from my context file that matter most for this task.

Then give me your execution plan (5 steps maximum).
Only begin work once we've aligned.

- Task
- Context Files
- Reference
- Success Brief
- Rules
- Conversation
- Plan
- Alignment

Opus 4.6



Syed Arsalan Mushtaq, CFA · 1st
VP - Head of Treasury Business Management | ...
2h ·

Most people prompt AI the way they used to Google things — one line, vague, hoping for the best.

Ruben Hassid (worth following) broke down what a well-structured Claude prompt actually looks like, and it's a useful framework for anyone using AI seriously at work.

The short version:

- (1) define the task with success criteria, not just "do this."
- (2) Upload context files with your expertise, rules, and preferences so the model doesn't start blind.
- (3) Give it a reference — show it what good looks like.
- (4) Write a brief specifying output format, length, and what it definitely shouldn't sound like.
- (5) Set your rules. ...more

2

- Like
- Comment
- Repost
- Send



Add a comment...

Case study: AI prompts for Treasury learning

Re: FPC submission - mooradch | Moorad Choudhry AI | Google Calendar - March 2026 | Feed | LinkedIn | The Beatles - I Saw Her Star

agmi.co/go/choudhry

Bookmarks: Elsevier Royalty Link, Nubank - Sign In, E*TRADE, Fidelity Login, John Lewis Credit C..., COUTTS login, NatWest, QuickBooks, Nationwide Login, Lloyds Bank, Google, Santander ac, All Bookmarks

**Moorad Choudhry AI**

When the pressure's on, Moorad Choudhry AI thinks with you

Whether you're prepping for the CFA, switching roles, just been promoted or facing that Board or Grad-level interview, this AI-powered chatbot helps you work smarter, faster, and with confidence.

"Moorad Choudhry AI has been instrumental in my finance studies. The interactive mode that walks me through concepts from Moorad's books is incredibly effective. It feels like having a personal tutor guiding me through complex banking principles anytime I need." —Rod Pienaar, UBS

[Start Your Free Trial](#)



0:15 / 1:21

Moorad Choudhry AI is trained solely on the content of every book written by Professor Moorad Choudhry, and built for how you actually work:

The anatomy of a [Moorad Choudhry] AI prompt

/// Crafting [MC] AI Prompts

- /// **Set the frame clearly:** Begin by defining your role or scenario (e.g., “Act as the Group Treasurer designing a future-state ALM architecture...”). Clear context helps MCAI calibrate both depth and tone.
- /// **Layer the scope:** Break complex requests into structured components (FTP, LRM, IRRBB, data architecture, etc.). This encourages comprehensive yet concise outputs.
- /// **State your preferred style:** For example, “Summarise key design choices at executive level” or “Provide a technical blueprint with regulatory references.”
- /// **Iterate through refinement:** After the initial response, guide MCAI using follow-up prompts such as “Expand section 3 with illustrative data flows” or “Add governance implications.” This incremental approach produces richer, more tailored material.

/// Approaches to Adopt / Avoid

/// Adopt:

- /// Use explicit labels or numbering (e.g., “JCOQ2 - Data Architecture”) - this helps maintain clarity and continuity within structured, multi-part dialogues.
- /// Encourage comparative framing (“Contrast centralised vs federated ALM models”) to surface strategic trade-offs.
- /// Reference regulatory context directly (e.g., “Assume EBA/ECB jurisdiction”) for greater precision.

/// Avoid:

- /// Overloading a single prompt with too many nested instructions; instead, divide it into thematic sections.
- /// Using vague objectives (“make this better”), specify what “better” means: shorter, more technical, more strategic, and so on.

/// Maximising MCAI for ALM, Business, and Data Work

- /// **Architecture Design:** Use MCAI to draft target-state blueprints, then iterate toward technical depth with scenario-specific follow-ups.
- /// **Data Strategy:** Ask MCAI to map “data lineage from source systems to regulatory reports,” highlighting governance checkpoints.
- /// **Business Decision Support:** Use it to generate executive briefing notes linking risk metrics (LCR/NSFR/IRRBB) to P&L or capital steering narratives.
- /// **AI Integration:** Request “practical use cases for generative AI in ALM analytics” - MCAI can outline feasibility, control considerations, and explainability layers.

Our checklist items

/// Important questions for bank Treasury:

- /// Using generative AI in ALM: are the benefits only incremental
- /// What are the obstacles to progress?
- /// What should Treasury be doing to make the most of the potential?

/// Is “productivity gain” a sufficient use case?

- /// *“In December [2025], OpenAI reported that ChatGPT Enterprise was saving users an average of 40-60 minutes on each day they used it.” (Source: The Economist, 31/1/2026)*
- /// Back to my musician who just discovered the electric guitar...or a better analogy: company book-keepers who have just discovered Excel (*ibid.*)

/// The use case for bank Treasury is in “Vertical AI models”

- /// “Horizontal” capabilities: generic skills that everyone uses / needs: writing summaries, producing PowerPoint slides
- /// “Vertical” capabilities: specific skills, like producing an ALCO MI deck for a bank

/// BUT vertical capability AI model building requires specialist knowledge, specialist training...

///which means more than the software engineer just sitting next to the ALM desk

/// This is where the specific ask from the bank must be articulated precisely, and the specific skills to build the model identified



Checklist items....

- /// Bank / Treasury: it's not a case of vendors selling models, but a case of banks acquiring models (including agentic models) as solutions to problems
- /// Hence we must identify the “problem” statement first...
- /// ...we define the outcome before engaging the model build
- /// This list is (i) potentially vast (ii) specific to each bank Treasury desk
- /// What is yours? [I asked already on slide 13!]
- /// “*In other words, companies* [read: Treasury desks, ALCO] are still having to make sense of the technology for themselves.” (*Ibid.*)
- /// This is a good thing, because....
- ///not every Treasury necessarily needs agentic AI



Checklist items....

/// A not-uncommon horror story in corporate world:

- /// A “cutting edge” AI tool (for Sales people, for example) is introduced as an efficiency gain but is not user-friendly enough, summaries riddled with errors, etc
- /// There may not be any problem with rejecting a new customer credit card application but there will be with generating errors in administration or handling of existing customers

/// Build – and train – the model to the specific use case

- /// Training it on the right data (banks should not have a problem with this. Er...)
 - /// “...*Historical data inconsistencies — originating in system migrations, portfolio acquisitions, and shifting regulatory interpretation — have been a known, unresolved challenge for financial risk models for years.*” – Jason Rodrigues, author of *AI and Banking: Understanding the AI-Era in Financial Risk Management*, DeGruyter, Global Banking Series, forthcoming...
 - /// “*Banks need a clear business case for investing in a proper historical data setup — and while AI itself can be that business case, it would need more substantiation for practical implementation.*” Jason Rodrigues again....being addressed in his book.
- /// Fine tuning the model thru reiteration and usage
- /// **Small Language Models** – trained on subsets of data and focused on specific tasks – may be a solution for Treasury (depending on the use case of course – the problem we are trying to solve)
- /// “I don’t need the model that you’re ordering a pizza from to be able to tell you about the most famous economist in the world” (*Ibid.*)
- /// If my problem is a 400-page ALCO deck, I want the Treasury AI model to fix that – and nothing beyond my bank’s Treasury needs

Checklist items....

- /// *“Problems become much more tractable when the tasks given to agents are narrow...” (Ibid.)*
- /// Set a very specific task for AI to undertake (or solve)
 - /// IRRBB hedging is one that will occupy you for ages 😊
 - /// ***A problem or task with well-defined inputs and outputs, and a finite universe of possibilities, is well-suited for an AI solution***
 - /// The more complex the task, the more permutations in outcome that are possible, the higher the probability of model errors
 - /// Which Treasury activities fall into our desired category?!
- /// The list of tasks we wish to tackle and solve for requires cross-disciplinary collaboration from the start
 - /// That means Treasury f/o and m/o, other technical experts in ALM, the business analysts, and the software engineers addressing the issue together

Checklist items...final word

- /// *“Specify a business outcome you’re trying to drive more than a technical outcome” (ibid.)*
 - /// Makes sense. But in Treasury they can be one and the same:
 - /// “Getting the dynamic hedge re-balancing more accurate over time saves me on hedge costs.”
 - /// That sounds both technical and business alike!
- /// Ultimately, for generative AI (the PRA CP summary) or agentic AI (undertaking the hedge transaction execution automatically), if a human being has to get involved (review output; check the transaction impact) then we’ve not exactly arrived at the optimum solution...
- /// *“To make AI work...a set of management problems has to be solved.”*
 - /// Well-designed incentives for adoption
 - /// Guardrails to mitigate against problems, errors
 - /// Process for choosing, implementing and measuring applications
- /// ***Which is why defining the outcome must come before building the model....!***



BTRM

The Certificate
of Bank Treasury
Risk Management

Illustration

AI and Bank Treasury

AI vs GenAI/LLM: What Each Does in Treasury

AI / ML

Numerical / pattern-based models — learns from data, predicts & optimises

Predictive Forecasting

- /// Cash flows, liquidity demand, NII under rate scenarios

Anomaly Detection

- /// Unusual payments, data outliers, early warning signals

Behavioural Modelling

- /// NMD stability, repricing lags, core deposit decay curves

Optimisation Algorithms

- /// Hedge ratios, payment sequencing, funding cost minimisation

GenAI / Large Language Models

Language / text-based models — reads, writes, reasons over documents

Information Summarisation

- /// Regulatory updates (Basel IV, EU Digital Operational Resilience Act¹), internal policies

Data Retrieval & Integration

- /// Cross-system aggregation, Natural Language Programming (NLP) queries on data

Code Generation

- /// Python/SQL extraction scripts, Excel automation, dashboards

Document Drafting

- /// ALCO packs, hedge documentation, IFRS 9 narratives, policy drafts

Decision Commentary (RAG – Retrieval Augmented Generation)

- /// Scenario commentary grounded in verified internal data via RAG

Example: AI-Enhanced Funding Plan

BUSINESS NEEDS

- /// Produce an optimised funding plan
 - /// as key output and pillar of the bank's budget/strategic plan
 - /// updated monthly (with latest b/s), aligned to bank and market conditions
 - /// covering all instruments integrated with LCR, NSFR projections within ILAAP framework

BUSINESS GOALS

- /// Avoid funding gaps and maturity concentration risk
- /// Comply with regulatory liquidity ratios at all times
 - /// Minimise cost of funds within the above constraints
- /// Multicurrency management

OUTPUT

- /// Optimised funding plan with instrument-by-instrument as well as currency allocation
 - /// cost of funds breakdown by tenor, instrument and currency
 - /// LCR and NSFR projected daily over the 12-month horizon
 - /// scenario comparison: base case vs stress scenarios (*at least 3*)
 - /// LLM-generated ALCO commentary for each scenario



AI-Enhanced Funding Plan...

AI Contribution by Technology

AI / ML

Learns from data — predicts, detects and optimises

Cost optimisation

- Minimises funding cost subject to LCR, NSFR and maturity ladder constraints
- Identifies cheapest instrument mix across wholesale, retail and central bank facilities

Scenario simulation

- Runs base, stress and reverse-stress cases in seconds; ML behavioural deposit model

GenAI / LLM

Reads, writes and reasons over documents

ALCO commentary

- Full narrative generated for each funding scenario — rationale, constraints, trade-offs
- Policy compliance check: does the plan respect treasury policy and board limits?

Board narrative

- One-page summary of recommended funding strategy with risk/cost trade-off framing

Agentic AI

Acts autonomously within approved boundaries

Plan execution

- Approved funding transactions routed for execution within pre-approved counterparty/notional limits
- Intraday monitoring: alert if LCR/NSFR approaches regulatory minimum

Rollover automation

- Refinancing instruction triggered automatically as a maturing instrument approaches rollover date

Example: FX Risk: *tactical fx hedging (short term)*

BUSINESS NEED

- /// Understand exposure evolution
- /// Gauge market conditions

Part of Treasury
operations linked to
daily management

BUSINESS GOAL

- /// **reduce risk of P&L losses P&L**
- /// Avoid missing payment deadlines

ACTIONS

/// **Tactical FX hedging (short-term)**

- /// AI identifies optimal execution windows from intraday patterns
- /// NLP processes news and central bank communications
- /// Sentiment scoring feeds into hedge timing recommendations

/// **The AI co-pilot concept**

- /// AI generates recommendation with rationale and confidence score
- /// Treasury reviews, modifies if needed, then approves execution



BTRM

The Certificate
of Bank Treasury
Risk Management

FX Risk: Tactical Hedging...

AI Contribution by Technology

AI / ML

*Learns from data — predicts,
detects and optimises*

Sentiment analysis

- Real-time NLP on central bank communications, macro releases and news flow
- Short-term FX direction prediction with confidence intervals and scenario ranking

Execution timing

- ML model identifies optimal execution windows to minimise

GenAI / LLM

*Reads, writes and reasons over
documents*

Market commentary

- Interprets central bank language and quantifies impact on treasury FX positions
- Trade rationale documented automatically for compliance and audit trail

Post-trade analysis

- LLM reviews: "Was this the right hedging decision given subsequent market

Agentic AI

*Acts autonomously within
approved boundaries*

Recommendation engine

- Generates hedge recommendation with rationale, confidence score and cost estimate
- Treasurer approves with one click — full audit trail captured automatically

Intraday monitoring

- Flags deteriorating tactical positions before end-of-day; proposes corrective action

Example: Treasury Reporting & Documentation

BUSINESS NEED/GOAL

- /// More efficient process for documentation (eliminate 3–5 day manual ALCO pack production)
- /// Improve quality: increase number of scenarios, reduce drafting errors, ensure source citation

ACTIONS

- /// **Prompt engineering for treasury documentation**
 - /// Specify role: ‘You are a senior bank treasurer at a European bank...’
 - /// Provide context: paste the relevant data table or metric
 - /// Define format: ‘Three sections: fact, driver, management action’
- /// **ALCO pack generation (current state: 3–5 days manual)**
 - /// AI automates data extraction, chart generation, and narrative commentary
 - /// Human review and approval before distribution
- /// **RAG (Retrieval-Augmented Generation) for accuracy**
 - /// Ground LLM on verified internal documents to prevent hallucinations
 - /// LLM cites source for every numerical claim it makes



Treasury Reporting...

AI Contribution by Technology

AI / ML

*Learns from data — predicts,
detects and optimises*

Data extraction

- Automated KPI extraction and trend identification from all source systems
- Anomaly detection: metrics that deviate from prior period or budget flagged before publication

Data completeness

- Every number verified against originating system before report generation begins

GenAI / LLM

*Reads, writes and reasons over
documents*

ALCO pack drafting

- Full narrative generated with RAG grounding — every claim cited back to source document
- Regulatory narratives: ICAAP, ILAAP and IRRBB sections drafted from verified data

Policy updates

- Treasury policy document updated automatically to reflect latest balance sheet conditions

Agentic AI

*Acts autonomously within
approved boundaries*

End-to-end pipeline

- Data pull → compute KPIs → draft narrative → human review → distribute — fully orchestrated
- Source verification: agent traces every published number to its originating system record

Distribution workflow

- Approved report sent to correct recipients on schedule — with version control and audit log

Key Takeaways

AI transforms every layer of Treasury — from data infrastructure to executive reporting

01 — Data is the foundation

Integration, cleaning and completion are prerequisites. No AI model can outperform its input data.

02 — Two AI paradigms, distinct roles

ML models predict and optimise over numbers. LLMs read, write and reason over documents. Both are needed.

03 — The co-pilot model

AI recommends with rationale. The treasurer decides. Human-in-the-loop is not optional — it is the design.

04 — Integration is organisational

Process mapping, ALCO buy-in and change management matter more than the technology choice.



BTRM

The Certificate
of Bank Treasury
Risk Management

Agentic AI and Bank Treasury

What is an AI Agent?

Definition: An AI Agent is a system that uses an LLM as its reasoning engine to **pursue a goal autonomously** through iterative planning and tool use — it acts, not just generates

How an Agent Works: ReAct Loop

1. **Reason:** receives goal, breaks into sub-tasks
2. **Act:** calls tools (SQL, APIs, Python, dashboards)
3. **Observe:** reads tool output, updates context
4. **Repeat:** iterates until goal is achieved

Key difference: an LLM answers a question. An agent completes a workflow.

Treasury Agent Examples

- ▶ **Hedging co-pilot:** monitors exposure, proposes hedge, drafts confirmation
- ▶ **Middle office:** checks trade confirmation, flags discrepancies
- ▶ **ALCO reporting:** pulls data, drafts commentary, formats dashboard
- ▶ **Margin call agent:** detects trigger, calculates requirement, sends notice

Governance rule: the treasurer approves; the agent executes

What makes a system "agentic"? — six defining properties

// 1. Autonomy

// Acts on its own initiative **further to a trigger** without a user prompt; pursues a defined goal

// 2. Perception

// Reads and interprets diverse inputs: market data, emails, documents, system alerts

// 3. Planning

// Breaks a high-level goal into a structured, ordered sequence of sub-tasks

// 4. Tool use

// Calls external APIs, databases, calculators, and execution systems

// 5. Memory

// Retains context across multi-step tasks; accesses policy documents via RAG

// 6. Multi-step execution

// Executes a chain of actions, adapting at each step based on the result

An LLM assistant has properties 1 and 2 only — agentic AI has all six

Agentic AI

- /// Do you REALLY want AI to execute your Treasury decisions?
- /// Use cases are more by the day...
- /// A common one:
 - /// Gen AI suggests the optimum hedge for your fixed-rate residential loan origination volume this month (actual or forecast)
 - /// Remember this isn't a straightforward "Buy bond-sell future using DV01 ratios" use case...
 - /// There is a timing issue ("pipeline risk"), an early repayment issue ("pre-payment risk"), an organic hedge option issue (in full? Part? Not at all?) and an ongoing dynamic rebalancing issue
 - /// Agent executes the AI solution for you, direct automated with the dealing counterparty
- /// There are unsurprisingly different opinions on this!

Deploying Agentic AI for risk management

- /// Is this the answer?
- /// Peter Simon, Archon Labs:
- /// “Tabular Foundation Models”
 - /// Trained on tabular structured data not layered
 - /// Not a half-way house between LLM and SLM...
 - /// ...rather, half way between ML and SLM
- /// “It isn’t cutting edge AI that is the key to success, it’s a data edge...”
- /// **Moorad’s response:** But, if we set up an Agent to enforce approved limits, will this cure the perennial problem (endemic in firms with poor risk culture) where risk limits are in place but just ignored and exceeded when deemed acceptable?
 - /// Examples of THAT are legion....!
- /// Set up your AI so that the bank cannot exceed a limit
- /// So, NOT necessarily asset allocation (or putting on the hedge transaction), rather adhering to “risk management” limit itself

Why Use Agentic AI in Treasury Operations?

Two distinct payoffs: doing the work faster and cheaper, making better decisions

- /// The problem with today's manual operations
- /// Skilled staff absorbed by repetitive monitoring, reconciliation and pack assembly
- /// Month-end and event-driven gaps: breaches could happen to surface late, after the window to act
- /// LLM copilots help draft, but wait to be asked — nothing initiates

EFFICIENCY — time and cost

- /// Initiative: monitors limits 24/7 and triggers a response unprompted
- /// Throughput: routine middle-office tasks handled end-to-end behind a human gate
- /// Scale: many scenarios at no extra marginal effort, no added headcount

EFFECTIVENESS — decision quality and P&L

- /// **Earlier hedging:** event-driven replanning closes IRRBB / FX gaps weeks sooner — less time running unhedged
- /// **Lower carry and slippage:** timely rebalancing cuts the cost of carrying open positions and reduces execution slippage
- /// **Funding cost:** continuous Linear Programming re-optimisation lowers WACF vs. monthly plan
- /// **Fewer breaches and incidents:** real-time limit surveillance avoids breach penalties and late settlement costs
- /// **Consistency:** deterministic pipeline removes transcription errors; every figure cross-checked vs source

Autonomy is earned use case by use case — prove both efficiency and effectiveness on the bank's own data before widening it

Treasury operations map: *where agents can act*

/// Front office

- /// **Use Case 1** • Interest rate risk hedging: monitor DV01/NII → propose hedge → human approves
- /// **UC2** • FX hedging (structural and tactical): exposure aggregation → hedge strategy → approval
- /// **UC3** • Intraday liquidity management: monitor nostro / LCR in real time → rank rebalancing options → human executes

/// Middle office

- /// **UC4** • Contract confirmation matching: extract terms → match blotter → STP or escalate
- /// **UC5** • Margin call management: reconcile MTM → instruct collateral or draft dispute

/// Risk Management

- /// **UC6** • ALCO pack assembly and commentary: pull data → compute metrics → draft text → review
- /// **UC7** • Stress test and scenario execution: parameterise → run model → draft narrative
- /// **UC8** • Limit monitoring and breach alert: real-time exposure vs limits → automated alert
- /// **UC9** • Regulatory report drafting: RAG-backed drafting with human sign-off

Example: Interest rate risk hedging agent

TRIGGER

- /// **DV01 exceeds board-approved limit:** Daily limit breach detected in TMS position
- /// **NII-at-Risk projection outside policy band:** Yield curve shift triggers stress threshold
- /// **Repricing gap outside ALCO-approved range**

AGENT LOGIC

- /// 1. Read DV01 and NII position from TMS limit report
- /// 2. Calculate NII impact under +/- 100bps AND +/- 200bps parallel shift
 - /// Define to which currency to be applied
- /// 3. Identify eligible hedging instruments (IRS, caps, floors) from approved ctp
- /// 4. Optimise hedge notional and tenor to restore compliance at minimum cost
- /// 5. Draft hedge memo with rationale, instrument choice, and P&L impact
- /// 6. Route to Treasurer for approval — no execution without sign-off

OUTPUT

- /// **Hedge recommendation memo (draft)**
 - /// Instrument, notional, tenor, expected post-hedge DV01
- /// **P&L and NII impact estimate**
- /// **ALCO-ready rationale**
- /// **Human approval gate before any execution**

**AI technology: LLM (reasoning + memo drafting) · ML (yield curve model) ·
Agentic (monitor → plan → route to human)**

ALCO report and dashboard agent

TRIGGER

- /// Scheduled ALCO meeting T-3 business days
- /// Regular pack assembly triggered automatically by calendar
- /// Ad hoc board request for liquidity update
- /// Regulatory review pack due

AGENT LOGIC

- /// 1. Pull latest data from TMS, risk system, and market data feed
- /// 2. Compute standardised metrics: LCR, NSFR, NII, EVE, DV01, FX exposure
- /// 3. Compare metrics against limits and prior period using RAG over policy documents
- /// 4. Flag limit breaches and trend deviations automatically
- /// 5. Draft executive commentary with trend analysis and policy rationale
- /// 6. Assemble charts and tables into ALCO pack template; route to ALCO secretary

OUTPUT

- /// Fully assembled ALCO pack (draft)
- /// Executive commentary with trend analysis
- /// Limit breaches flagged with policy reference
- /// RAG-validated against board-approved limits document
- /// Auto-generated metric charts and tables
- /// ALCO secretary receives for final review and approval before distribution

AI technology: LLM (commentary + RAG policy check) · ML (metric forecasts) · Agentic (schedule → collect → draft → route)

**BTRM**The Certificate
of Bank Treasury
Risk Management

AI Agents: Advantages and Risks

Advantages

Efficiency & speed

- /// End-to-end funding plan cycle automated from data ingestion to ALCO pack delivery; monthly cycle compressed from 3–4 weeks manual to 2–3 hours — analyst time freed for judgment and challenge

Consistency auditability

- /// Deterministic pipeline eliminates manual transcription errors, missed steps and version inconsistencies; every action logged with timestamp — full trail for model risk review and regulatory inspection

Scalability and monitoring continuity

- /// Same agent runs 5+ stress scenarios and sensitivity analysis without added effort; background agent detects RAF threshold proximity and triggers replanning without waiting for the monthly cycle

Risks & Limitations

Prompt injection and tool misuse

- /// Adversarial content in market data feeds or regulatory emails may redirect agent actions; agent may call the correct tool (LP solver, LLM) with incorrect parameters — silent errors with real funding consequences

Accountability gap and over-automation

- /// When the agent acts autonomously, ownership of errors is unclear without explicit governance; human judgement may be bypassed when anomalies fall just below escalation thresholds — ALCO pack distributed without adequate

LLM hallucination and regulatory uncertainty

- /// Commentary may contain plausible but factually incorrect figures — cross-check against solver output is mandatory before ALCO distribution; ECB/EBA model risk frameworks do not yet fully address autonomous agentic systems in treasury

Agents amplify both benefit and risk — governance is the bridge

AI Hallucinations in Treasury

- /// Definition: an LLM hallucination occurs when the model generates plausible-sounding output that is factually incorrect and not grounded in the data it received
- /// The LLM does not know when it is wrong — it has no access to 'truth'; only the grounding check does
- /// Hallucinated output looks identical in style and confidence to correct output — there is no visible warning signal

Hallucinated LLM output

“The forecast is driven by a EURIBOR change of -EUR 25M and month-end settlement of -EUR 110M, partially offset by wholesale inflows of +EUR 20M”

Actual SHAP values (source of truth)

- /// month_end: -EUR 85M
- /// euribor_ch5: -EUR 18M
- /// wholesale_cf_lag1: -EUR 4M

The error is invisible until cross-checked against the SHAP table

Why hallucinations are especially dangerous in treasury

- /// ALCO decisions are made on the cited numbers — a hallucinated EUR 25M EURIBOR driver may redirect hedging strategy
- /// Errors look confident and well-formatted — there is no visible difference between correct and hallucinated commentary
- /// The LLM does not verify its output against the input data — **cross-check layer needed**




BTRM

The Certificate
of Bank Treasury
Risk Management

Final Words

Sidebar: one AI model that won't hallucinate 😊



**Moorad
Choudhry AI**
agmi.co/go/choudhry

/// Because it's trained on my books, and only on my books!

/// AI Bot: <https://agmi.co/go/choudhry>

A governance starter kit: what to put in place first

1

Inventory

Register every model and agent

• DAY ONE

2

Validation gate

2nd-line sign-off before any model goes live

• DAY ONE

3

Owner & risk tier

Assign a named owner and a risk tier to every model

4

Limits & floors

Set confidence floors, exposure limits, human-approval thresholds

5

Prompt repository

Centralise prompts in a version-controlled shared repository

6

Incident wiring

Route AI incidents through the operational-risk process

7

Non-AI fallback

Keep a manual path for every critical Treasury function

Start with steps 1 and 2 — the inventory and the validation gate must come before the first model goes live.

Banks and AI

- /// I know one product design I'd love a bank to offer....can AI help a bank to offer this:
 - /// Tailored Loan / deposit product, relevant to demands and profile to individual customer
 - /// Customer interaction that moves with customer life cycle
- /// But that's for another time...



Final word

- /// We restricted this conversation to banks and AI, and specifically bank ALCO and Treasury and AI
- /// The ask seems to be:
 - /// What is it you want to do more efficiently / differently / more cheaply / more effectively etc etc in your ALCO and Treasury process
 - /// Straightforward items like the ALCO pack or Reg Reporting
 - /// Arcane subtleties associated with decision making on how best to manage Pillar 2 liquidity risk decision-making
 - /// Timing-related items such as pipeline hedging for IRR
 - /// A combination of these and others....or something else entirely?
 - /// Or do you want to redesign it altogether? Imagine whole new outputs and processes as part of achieving a new, higher goal?
 - /// Articulate this FIRST, and then start work with a combined project team of the sponsor / users, subject matter experts, business analysts and software programmers and work together deliver it
- /// I finished this slide deck last Friday....
- /// Then I remembered this in *The Economist* on from my BTRM webinar back in April ☺ >>>

BY INVITATION

Ethan Mollick

Companies should treat AI as what it is: profoundly weird

A SYSTEM DESIGNED to predict the most likely next word in a sentence can also write good computer code, offer strategic advice and respond with remarkable empathy to human problems. We don't fully understand why. And yet the dominant instinct across the corporate world is to treat artificial intelligence as if it were just another piece of enterprise software: slot it into existing processes, assign it Key Performance Indicators and hand it to the IT department for management.

This is a profound strategic mistake. Companies are racing to de-weird AI, and in doing so they are squandering what makes it transformative, turning it into just the latest wave of office automation. To be clear, making AI easy to use is important, and building it into familiar tools and workflows is just smart engineering. The mistake is letting those smooth interfaces flatten your understanding of the technology's possibilities.

The urge to de-weird is understandable. Executives are trained to normalise new technologies and fit them into familiar categories. So AI becomes a fuzzy-logic processor built into a workflow or a tool that shaves minutes off a task. Normal technology gets normal rollout plans, like requiring that 90% of employees must use a new software package each week.

But when you set that target for AI, what actually happens? Employees use AI to transcribe meetings, or produce an endless stream of "workslop" in the form of dozens of extra memos or PowerPoints. Treating this technology as another software deployment is like receiving a mysterious alien artefact and immediately using it as a paperweight.

The de-weirding impulse produces a second, deeper failure: it leads companies to default towards automation rather than augmentation. When leaders see studies showing productivity gains of 30% from AI, their instinct is to cut 30% of the workforce. That arithmetic is simple. What is hard, and requires genuine imagination, is asking a different question: what does it mean to rebuild an organisation around the fact that a single programmer can now write a hundred times more code? What new products become possible? What new markets open up? No vendor can answer those questions for you. No consultant has a playbook (much as

they might claim they do). The hard strategic work of reimagining what your organisation could become is precisely the work that de-weirding AI allows companies to avoid.

And there is a natural place where de-weirding AI goes to die: the IT department. This is not a criticism of IT professionals, who do essential work. But in most firms their mandate is to minimise risk. If they could take away your keyboards, they would sleep better at night. Every creative thing an employee does on a computer leaves the firm potentially vulnerable. AI, by contrast, demands that organisations embrace risk by experimenting wildly, tolerating failure and accepting that nobody yet knows the right way to use these tools. Handing sole control over AI to a department whose core mission is risk elimination is a category error.

So what should firms do instead? In my work in this area I have come to advocate a three-part model: Leadership, Crowd and Lab.

Leadership means that direction must come from the top. The CEO and other senior managers cannot delegate AI strategy to middle management or IT. They must articulate a vision for how AI changes what the organisation is, not merely how it operates, and they must create incentives that make experimentation safe. And they have to lead by using these systems themselves.

If they do, they can inspire the Crowd, a company's employees, who, when given access to AI tools and genuine permission to experiment, will figure out use cases not even the AI firms expected. Since AI is most effective in the hands of experts, the Crowd is where the best ideas come from.

Those ideas then go to the Lab, where a team of technical and non-technical employees work on generative AI full-time. These are people whose job is to push boundaries, develop new workflows and feed discoveries back into the organisation. I am shocked by how many large companies still lack even this. Without it, they have no mechanism for learning what AI can actually do for them. They are flying blind, relying on vendor demos and conference keynotes instead of building institutional knowledge.

There is one more problem that de-weirding creates, and it may be the most consequential. When companies fail to create the right incentives, employees respond rationally: they hide their AI use. Some fear punishment. Some do not trust that productivity gains will be shared with them rather than captured by the firm. Some quietly work 90% less and see no reason to volunteer that information. The result is an enormous information gap. Managers cannot see the true impact AI is already having inside their own organisations, which makes it even harder to develop a real strategy.

Weird AI, the stage is yours


Resisting the de-weirding trend does not guarantee AI will go well. There is no default good outcome. But a failure to see AI for what it is—a profoundly odd, risky and powerful technology—will guarantee bad ones. Firms that sand down AI's strange edges will veer towards automation and layoffs, as that is all they can see. Those willing to confront the technology head-on can find something far more interesting, including ways to help make their people, and their organisations, capable of things that were impossible a year ago and will be impossible to predict a year from now. Nobody knows exactly where this is going. But you don't navigate strange territory by pretending your old maps will work. ■

Ethan Mollick is an associate professor at the Wharton School and the author of "Co-Intelligence: Living and Working with AI".


BTRM

The Certificate
of Bank Treasury
Risk Management

Further reading on bank ALM



The Principles of Banking, 2nd Edition, Singapore:
John Wiley & Sons Ltd 2022, chapters 11-15



*Moorad Choudhry Anthology: Past, Present and Future
Principles of Banking and Finance*, Singapore: John
Wiley & Sons Ltd 2018, chapters 8, 10-14

Email: mooradchoudhry@gmail.com

DISCLAIMER

The material in this presentation is based on information that we consider reliable, but we do not warrant that it is accurate or complete, and it should not be relied on as such. Opinions expressed are current opinions only. We are not soliciting any action based upon this material. Neither the author, his/her employers, any operating arm of his/her employers nor any affiliated body can be held liable or responsible for any outcomes resulting from actions arising as a result of delivering this presentation. This presentation does not constitute investment advice nor should it be considered as such.

The views expressed in this presentation represent those of the lecturer in his or her individual private capacity and should not be taken to be the views of any employer or any affiliated body, including any bank that employs any member of the BTRM Faculty, or of the lecturer as an employee of any institution or affiliated body, including The BTRM or of FHNW. Either he/she or his/her employers may or may not hold, or have recently held, a position in any security identified in this document.

No AI training: Without limitation of BTRM's or the Author's exclusive rights under copyright, it is expressly forbidden to use any part of this publication to "train" Artificial Intelligence (AI) systems to generate text. BTRM Limited and the Author reserve all rights to license use of this work for all purposes, including for generative AI training and development of Machine Learning language models.

This presentation is © BTRM / Moorad Choudhry / Andrea Cremonino 2014, 2026. No part of this presentation may be copied, reproduced, distributed or stored in any form, including electronically, without express written permission in advance from BTRM Ltd and/or the author.