

— ASIAN ACTUARIAL CONFERENCE


From Checklists to Conversations

Agentic AI in Actuarial Work

Daniel Tham

Head — Commercial, Actuarial & Strategy, Insurance

GrabInsure



— ACT ONE


Primer

What an agent actually is



You've all used this

It *advised* you. **You are still the one who opens the spreadsheet.**




What turns an advisor into a worker

BASELINE

Chat

Answers you. **You** open the spreadsheet, run the query, write the file.

+ WHAT AN AGENT ADDS



Tools

Runs the query. Writes the file. Posts the message.

INSTEAD OF

producing text that you then execute




State

Persistent memory, and a workspace it returns to.

INSTEAD OF

forgetting, while you re-paste the context



Loop

Plans, acts, observes, corrects — until it is done.

INSTEAD OF

one turn, one answer, stop

Chat + tools + state + loop = an agent

Same model. Different harness.

The models barely differ. **The harness does.**



Chat

 Claude  Gemini  ChatGPT

Answers you. You do the work.



Agentic

 Claude Code  Codex  Cursor

Given a terminal, your files, and a job.

What a loop actually looks like



PLAN — WRITTEN BY THE AGENT, NOT BY ME

- ✓ Pull the paid and incurred data for the period
- ✓ Rebuild the development triangles
- ✓ Run the reserving method and compare against last month
- ✓ Flag anything that moved more than expected
- ✓ Draft the note and stop for sign-off

It corrected itself twice on step 3 — and **stopped on its own** at step 5.

— ACT TWO

What this means for the work

None of this is actuarial-specific — but you didn't come here to watch an agent sort email.



Two checklists you already know

PRICING ACTUARY

- ☐ Assemble exposure & claims experience
- ☐ Set the assumptions
- ☐ Build the loss cost
- ☐ Draft the rate review
- ☐ Support the quotes
- ☒ **Sign off the basis — you**

RESERVING ACTUARY

- ☐ Pull the quarter's claims data
- ☐ Build the triangles
- ☐ Select factors, calculate IBNR
- ☐ Draft the pack
- ☐ Book & file
- ☒ **Sign off the reserves — you**

Different desks, same shape: **a sequence of specified steps, ending in your name.**

The checklist doesn't disappear. It changes hands.

THE RESERVING CHECKLIST FROM THE LAST SLIDE –
YOU EXECUTE IT

- ☐ Pull the quarter's claims data
- ☐ Build the triangles
- ☐ Select the factors, calculate IBNR
- ☐ Draft the pack
- ☐ Book the entries
- ☒ Review & sign off — you

WITH AN AGENT –
YOU SAY WHAT YOU WANT

“Run the June reserving on the locked basis. Flag anything outside tolerance.”

...the same steps run, by the agent

- ☒ Review & sign off — still you

“You want me to sign a number a **language model** produced?”

01

Your name is on the valuation.

02

The regulator is on the other end of it.

03

Professional standards have no AI exemption.

That instinct is not conservatism. **It is the correct professional reaction** — and nothing here asks you to trust the model.

The LLM doesn't give the number

It runs a **pre-approved program** — your methodology, your basis — and **the program** produces the number. No different from the numbers you already take from **Excel, Prophet, or your reserving software**.

You already sign numbers you didn't calculate by hand. The question was never *who pressed run*.

Where the actuary stays in the loop

01 Which program runs

The agent selects from your approved set — nothing else.

02 The shape of the program

The methodology is yours, written down as **code you can read**.

03 The context it works from

The basis, the definitions, the documented method — **you author what the agent knows before it starts**. It reads that basis; it cannot change it.

Context engineering is actuarial work. **Not the vendor's, not IT's — yours.**

One cycle, end to end

Monthly reserving, as it runs today. The only thing that changed hands is the keyboard.

CONTEXT ENGINEERING — The basis, the definitions, the documented method — written by you before the cycle starts. Every step runs against it.

STEP 1 · AGENT

Data audit

pulls & validates the quarter

STEP 2 · AGENT

Triangles

builds from locked scripts

STEP 3 · AGENT

IBNR

runs the approved methods

▲ HARD GATE · YOU

Refine & review

Stops. Waits for your call on the numbers.

▲ HARD GATE · YOU

Lock & post

Stops. Nothing books without your sign-off.

CONTINUOUS OVERSIGHT — Every step is visible as it runs — the plan, every tool call, every diff. You can interrupt at any point.

The context is yours. Oversight is continuous. The gates are absolute — **the cycle cannot pass them without you.**

— ACT THREE • EVIDENCE

The Ark

Eight agents, in production, at a live insurer.

3

The Ark



Six specialists — scoping, research, pricing, drafting, monitoring. In Slack,
where the team already works.

THE-ARK-PROMO-24S.MP4 **24S** 1920×1080 · SILENT CUT · LIVE NARRATION

Eight agents, in production



Ginger

Assistant



Biscuit

Scoping



Hoot

Reserving & regs



Bao Bao

Pricing



Twiga

General



Puffy

Ticket intel



Warden

Compliance



Track

PMO

It shows its working

GrabInsure

Channels

the-ark

reserving

Apps

Hoot the Owl

Puffy the Seal

Twiga the Giraffe

Track the Otter

Hoot the Owl

Actuarial analysis · reserving & pricing data

RL

Risk lead

10:21

@Hoot the Owl Can you tell me what is the average TPA cost per GPML claim? Also how many claims per quarter in Q1 and Q2 of 2026?



Hoot the Owl

APP 10:21

✓ Worked through 4 steps

From the lake ([redacted] , latest valuation snapshot 202607):

TPA COST flat [redacted] — a fixed per-claim fee in the reserving pipeline, not a variable cost

Q1 / Q2 2026 [redacted] and [redacted] reported claims by accident quarter

One caveat for your exposure view: **Q2 is still developing**, so its count will drift up as late reports land — the reserving IBNR covers that tail.

SOURCES [redacted] (valuation_yyyymm = '202607' · product = 'GrabGroupPML_SG')

✓ 1

Message Hoot the Owl

B I ●●●

Real thread, re-set to be legible. Figures and the source table withheld, colleague names replaced with roles — external talk. The caveat and the sources line are verbatim.

Public filings, read for us

GrabInsure

Channels

the-ark

reserving

Apps

☐ Puffy the Seal

☒ Hoot the Owl

☐ Twiga the Giraffe

☐ Track the Otter

Hoot the Owl

Market research · MAS returns

DT

daniel.tham

20:23

@Hoot the Owl estimate for me the rough loss ratios of these reinsurance on motor since you should have the ceded premium and the ceded loss as well



Hoot the Owl

APP

20:24

✓ Worked through 1 step

Motor reinsurers ran a ~65% ceded loss ratio in 2022–2025 (2021: 51%) — derived from Form G1 as reinsurers' share of claims incurred ÷ ceded earned premium.

BY YEAR 2021 51% · 2022 66% · 2023 72% · 2024 64% · 2025 65%

Caveat: this is **claims-only**. Add RI commission paid to cedants and the reinsurers' motor combined ratio is likely ~80–90% — profitable but not fat.

SOURCES — Form G1, SIF, Motor class · MAS Insurance Company Returns

✓

1

Message Hoot the Owl

B I ⋮

Every figure here is from MAS's published insurance returns — public data. Per-insurer attribution and the source table withheld.

— NOBODY ASKED FOR THIS

It posts before anyone asks

GrabInsure

Channels

the-ark

reserving

Apps

Puffy the Seal

Hoot the Owl

Twiga the Giraffe

Track the Otter

Puffy the Seal Insurance operations · daily digest, 08:00 SGT

 **Puffy the Seal** APP Fri 31 Jul · 08:02

2 spikes flagged — insurance tickets running above the 56-day baseline:
FLAGGED [Thailand] Enquiry regarding a claim · [Singapore] Accident Reporting
BY MARKET Thailand · Singapore · Philippines · Malaysia · Indonesia

✓ 2

👁 1

 **Puffy the Seal** APP Sat 1 Aug · 08:01

Calm waters — nothing breaching the spike threshold.
THEMES Enquiry regarding a claim · Enquiry on coverage details · Delivery-guarantee claim

 **Puffy the Seal** APP Sun 2 Aug · 08:00


Calm waters — nothing breaching the spike threshold.
THEMES Enquiry regarding a claim · User enquires on coverage · Cancellation or refund

Message Puffy the Seal

B I ⋮


Three consecutive mornings, including the weekend. Ticket volumes and per-market counts removed — external talk.

Not everything is actuarial



Colleague names replaced with roles. Track runs on a different harness, built by another team — the pattern spread without central planning.

They have a voice, and opinions



Verbatim. Nothing withheld on this one — it is a conversation about the agent’s own writing style.

— ACT FOUR

What we got wrong

Three lessons, honestly



What we got wrong

01 A wrong fact propagates at fleet speed.

Capture is not canon. Promotion needs a human.

02 The material step must stop and ask.

Every time — not by exception.

03 We under-rated the context engineering.

Of tools, state and loop — only **state** was ever ours to build. We spent the first months on everything else.

The actuary stops executing the procedure and starts specifying and reviewing it.

Which is what the profession already trains for.

QUESTIONS

