Deloitte.





Platforming Domain Context for Agentic AI

Turning Data, Models, Controls, and Governance Artefacts into Reusable Building Blocks for Financial Modelling

> "Through creative destruction, **the accumulation of knowledge fuels endogenous growth**, relentlessly sweeping away obsolete technologies in waves of industrial innovation"
>
> *Philippe Aghion and Peter Howitt from A Model of Growth Through Creative Destruction*

Agentic AI is unlocking new possibilities in financial modelling — but only when grounded in deep, structured domain context. This post explores why data, models, controls, and governance artefacts must be thoughtfully platformed to scale responsibly. In regulated environments, context is not an input — it is a necessity. When designed for reuse, traceability, and control, it allows humans and machines to reason together at scale. This is how financial modelling capabilities evolve from experimentation to institutional strength.

LLM AS A NEW COMMODITY

- Every week seems to bring a new Large Language Model (LLM) announcement. That speed is a signal of progress.
- At this pace, it's fair to say LLMs are becoming a commodity: powerful, impressive — and increasingly interchangeable.
- LLMs are becoming a powerful tool to transform financial modelling, but being a general-purpose engine and force multiplier, they won't do it alone.
- The real transformation comes from domain aware context — and from how that context is platformed.
- Over the past few months, I deliberately went back to fundamentals: LLMs, agents, RAG, context windows, prompt engineering, and tool orchestration — with many hands-on hours spent deploying and experimenting on my home lab servers.
- Not because the concepts are inaccessible, but because they only truly click when grounded in real use cases — something that's increasingly hard in a world driven by hype and vibes (see the OpenClaw story...).

In financial modelling — particularly in banking and insurance — models are not abstract mathematical objects.

They are embedded in:

01	Regulatory constraints (OSFI, AMF, Basel, IFRS 9/17)
02	Strong governance and auditability requirements
03	Purpose driven, representative data selection
04	Lessons from major events (Covid, GFC 2008-09, geopolitical shocks)
05	Nuanced distinctions across risk types, practices, and vulnerabilities
06	Long time horizons, path dependency, and scenario driven decision-making
07	Explicit separation between judgment, assumptions, and computation
08	Assurance over implementation, controls, and execution

This is where agentic architectures become both powerful — and hard.



THIS IS WHERE PLATFORMING BECOMES CRITICAL

Agentic AI does not scale through prompts alone

It scales by making the right context available by design:

- Golden datasets
- Curated feature stores
- Approved modelling libraries & validation routines
- Governance artefacts

Platforming turns context into an asset.

It feeds agents and LLM tooling with trusted, reusable building blocks, while enforcing consistency, traceability, and control across use cases.

Without this foundation, agents remain custom-made experiments.

With it, they become part of an institutional modelling capability.

> **One insight crystallized quickly:**
> 👉 Agentic AI does not transform complex, regulated industries without deep, structured industry context. Vibe coding rarely survives the validation wall — or regulators' visits 😉
>
> *Nicolas Gemin, Risk Leader*
> *Deloitte Canada*

CONTEXT AND HUMAN AS ANCHOR

Because an agent is only as effective as the context it can reason over:

- The system prompt that encodes how decisions may be made
- The tools it can call (models, data, validation checks, reporting layers)
- The memory it retains (assumptions, scenarios, prior runs, artefacts)
- The constraints it must respect (regulatory, methodological, ethical)

When designed properly, this is where agentic AI starts to compound knowledge.

Not by replacing experts — but by persisting, reusing, and enriching institutional understanding over time.

Each run, stress test, and review feeds the next cycle, strengthening consistency, traceability, and insight.

In financial modelling — a domain that is both highly regulated and deeply dynamic — context is not a convenience. It is the fuel.

Without deep domain context:

1. Automation breaks governance
2. Speed creates risk instead of value
3. Agents hallucinate where judgment is required
4. Sparse data leads to unreliable outputs
5. Missed context erodes relevance and accuracy
6. Model uncertainty undermines trust
7. Regulatory expectations are not met

With domain expertise, thoughtfully platformed:

- Agents become decision accelerators — not decision makers
- Controls are embedded by design
- Modelling, validation, documentation, and reporting converge
- Humans remain accountable while machines handle complexity at scale
- Efficiency improves without loss of expertise

The real transformation is not “AI replacing models”.

It is AI reshaping how modelling knowledge is captured, reused, stress tested, and governed over time.

Why this matters now :

- GenAI experimentation in financial modelling is already underway across institutions
- Early implementations often struggle to pass audit, validation, and portability thresholds - the learning curve may be steep
- The gap between perceived productivity gains and durable institutional capability continues to widen

The winners will combine deep domain expertise with platformed context — where data, models, controls, and knowledge are deliberately designed to be reusable by humans and machines.

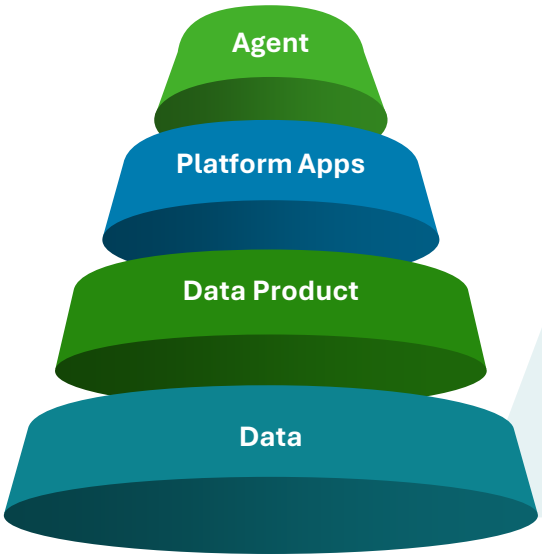

FROM DATA TO AGENTS | A new digital architecture.

Platforming establishes the stable foundations that allow back-testing to operate at scale.

FOUNDATION

Stable, governed assets

Data Assets

- Internal datasets (incl. SAS datasets)
- External / third-party data

Documentation & Standards

- Methodology notes
- Model documentation
- Validation guidelines

Nothing “agentic” happens without governed inputs.

USE CASE : MBS ACCELERATOR

- MBS Accelerator is a Deloitte developed application designed to orchestrate and streamline the end-to-end credit risk model back-testing lifecycle, from data ingestion to reporting and monitoring.
- MBS serves as a concrete illustration of how agent-enabled applications can be implemented in model governance and credit risk back-testing.
- The application automates the full back-testing lifecycle—from data import (Excel, CSV, SAS) and preparation to statistical testing, result tracking, decision statements, and audit-ready reporting—enabling continuous monitoring across portfolios and reporting periods through a single interface.

What platforming looks like ?

- Agent-enabled platforms are typically organized in layers (building blocks) with data product, apps and skills forming an embedded platform.
- Skills encode authoritative knowledge such as regulatory guidelines, internal policies, and standard operating procedures, defining the rules and constraints under which agents operate.
- Applications are the engines providing execution capabilities anchored in financial modeling.
- Other layers—such as documents, and data products provide context, evidence, and execution capabilities.
- Together, these layers enable agents to perform tasks in a governed, traceable, and repeatable manner.

LAYERED AND MANUAL BACKTESTING PRODUCTION EFFORT

The production of a single backtesting report requires a layered effort involving numerous recurrent and largely manual steps.

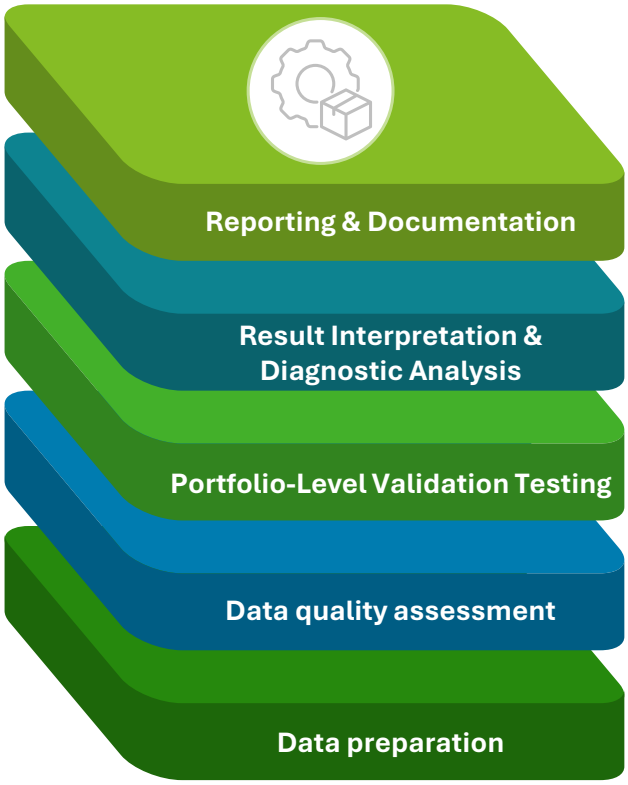

COMPOUNDING EFFORT ACROSS PORTFOLIOS AND REPORTING CYCLES

This layering effect scales across both portfolios and revision cycles, with limited persistence of historical results, resulting in duplicated effort instead of progressive build-up over time.

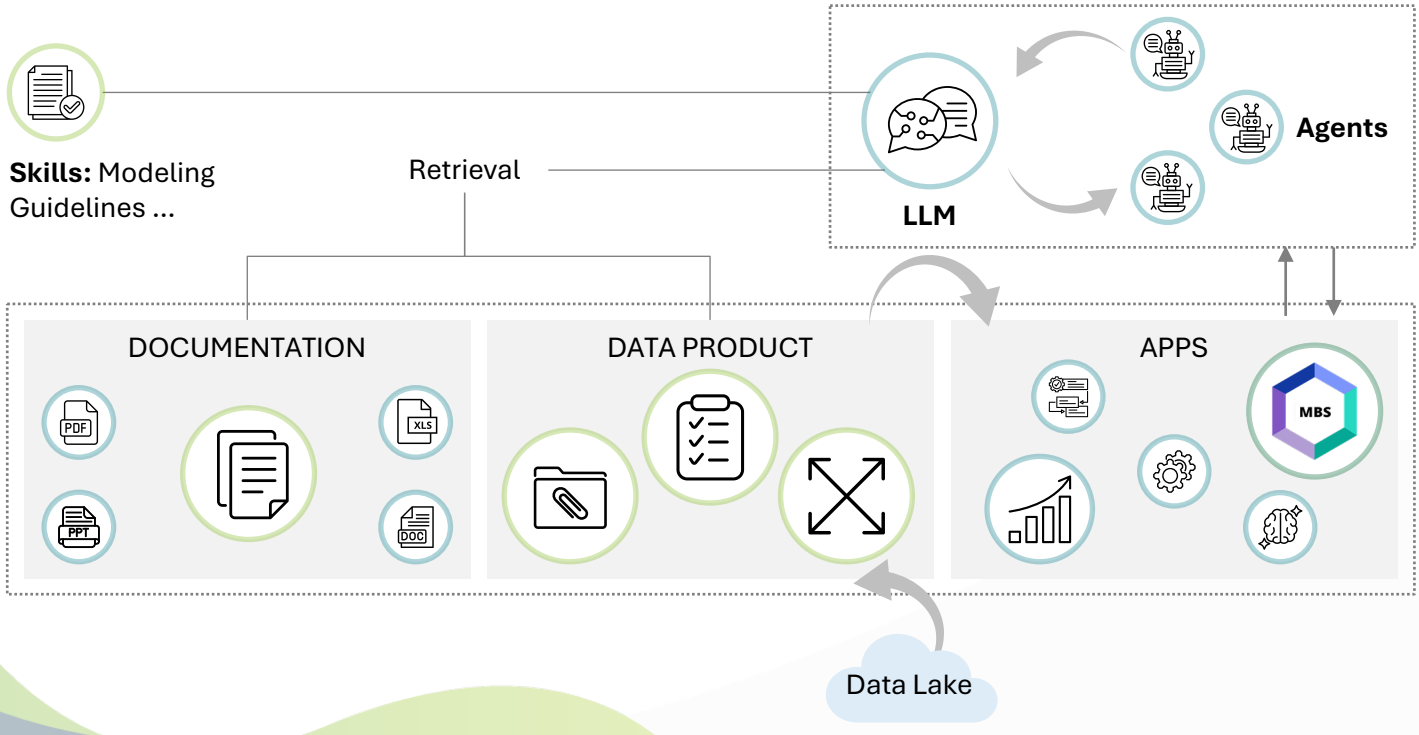

WHERE AGENTS CREATE VALUE ?

Purpose-built agents execute well-defined back-testing and governance tasks as part of the normal workflow, powered by applications, such as MBS, acting as execution engines anchored in financial modeling.

- DATA PREPARATION AGENT: Controls data admissibility.
- BACK-TESTING AGENT: Executes approved methods at scale.
- MONITORING AGENT: Surfaces signals requiring human judgement.
- DOCUMENTATION AGENT: Drafts artefacts and supports traceability mapping.
- QA / CONTROL AGENT: Enforces institutional boundaries.

These agents are enabled by large language models with retrieval-augmented generation (RAG), internal skills and rules, and secure access to third-party tools, ensuring automation remains governed, explainable, and aligned with model-risk standards.

↓ DEMO AVAILABLE ↓

Compatible with any type of data set including SAS

Treatment of voluminous datasets across multiple time periods and portfolios, powered by 25+ automated import, cleaning, and harmonization tools.

30+

integrated statistical, regulatory, and analytics functions across PD, LGD, and EAD dimensions.

Less than 60 min

from data upload to full statistical and visual-rich report generation spanning multi-period credit portfolios.

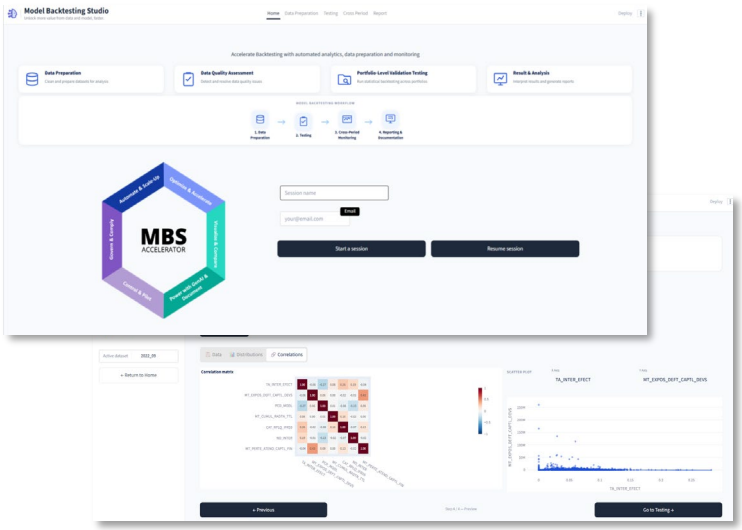

WHO WE ARE:

Our Services on Modelling for Financial Institutions:

- Platforming of domain-aware financial modelling capabilities, making data, models, assumptions, and controls reusable by humans and AI
- Integration of AI Risk and Trustworthy AI principles, embedding transparency, traceability, explainability, and model risk governance
- Development and review of models and scenarios using approved libraries, golden datasets, and auditable workflows
- Agent-enabled scenario analysis, stress testing, treasury, hedging, cash-flow forecasting, and pricing impact modelling—at scale and under control

In financial modelling, AI advantage will not come from adopting tools faster —but from institutionalizing knowledge better.

CONTACT US:



Nicolas Gemin
Risk leader
ngemin@deloitte.ca



Daniel Eklove
Risk Expert
deklove@deloitte.ca

About Deloitte

Deloitte provides audit and assurance, consulting, financial advisory, risk advisory, tax, and related services to public and private clients spanning multiple industries. Deloitte serves four out of five Fortune Global 500® companies through a globally connected network of member firms in more than 150 countries and territories bringing world-class capabilities, insights, and service to address clients' most complex business challenges. Deloitte LLP, an Ontario limited liability partnership, is the Canadian member firm of Deloitte Touche Tohmatsu Limited. Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee, and its network of member firms, each of which is a legally separate and independent entity. Please seewww.deloitte.com/about for a detailed description of the legal structure of Deloitte Touche Tohmatsu Limited and its member firms.

To learn more about how Deloitte's approximately 330,000 professionals, over 11,000 of whom are part of the Canadian firm, please connect with us on LinkedIn, Twitter, Instagram, or Facebook.

© Deloitte LLP and affiliated entities.