
# CREDIT RISK MODELLING: DRIVING EFFICIENCY WITH AI

October 2025

A business of Marsh McLennan

# CONFIDENTIALITY

Our clients' industries are extremely competitive, and the maintenance of confidentiality with respect to our clients' plans and data is critical. Oliver Wyman rigorously applies internal confidentiality practices to protect the confidentiality of all client information.

Similarly, our industry is very competitive. We view our approaches and insights as proprietary and therefore look to our clients to protect our interests in our proposals, presentations, methodologies, and analytical techniques. Under no circumstances should this material be shared with any third party without the prior written consent of Oliver Wyman.

© Oliver Wyman

# TODAY, THE LARGEST CREDIT RISK MODELS ARE VERY EXPENSIVE WITH LONG TIME-TO-MARKET

**100-300 FTES**
working on largest credit risk models with €15-45 MM annual cost

**€100-150 MM**
cost of a large modelling programme

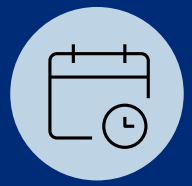

**3-4 years**
for roll-out of new models or material model changes…

**…with PMAs**
staying in place for a long time, incurring capital costs

# OUR HERCULES SOLUTION TRANSFORMS CREDIT RISK MODELLING WITH DRASTIC EFFICIENCY GAINS & ACCELERATION

Hercules is Oliver Wyman's **agentic AI solution for credit risk modelling**, built upon our **proven modelling codebase, methodology, and templates,** featuring a **no-code, natural language interface** for **steering and oversight by human modellers**

- The agentic AI orchestrator executes the end-to-end model development, refinement and/or re-calibration process according to the desired configuration
- It leverages our production-quality Python codebase, which has been deployed successfully across multiple modelling programmes and validated by both Model Risk functions and regulators
- Hercules produces a near-final modelling package with model documentation, all required regulatory artefacts and production code
- Its intuitive interface proactively presents results, options and recommendations for human modellers to review, discuss with experts and approve – ensuring effective 'human-in-the-loop' oversight

## HERCULES WILL DELIVER TRANSFORMATIONAL IMPACT FOR CREDIT RISK MODELLING

**65%**
SAVING IN MODELLING FTES

**50%**
FASTER MODELLING

© Oliver Wyman

# HERCULES COMBINES OUR PROVEN METHODOLOGY AND CODEBASE WITH AN AGENTIC AI LAYER TO TRANSFORM CREDIT RISK MODELLING

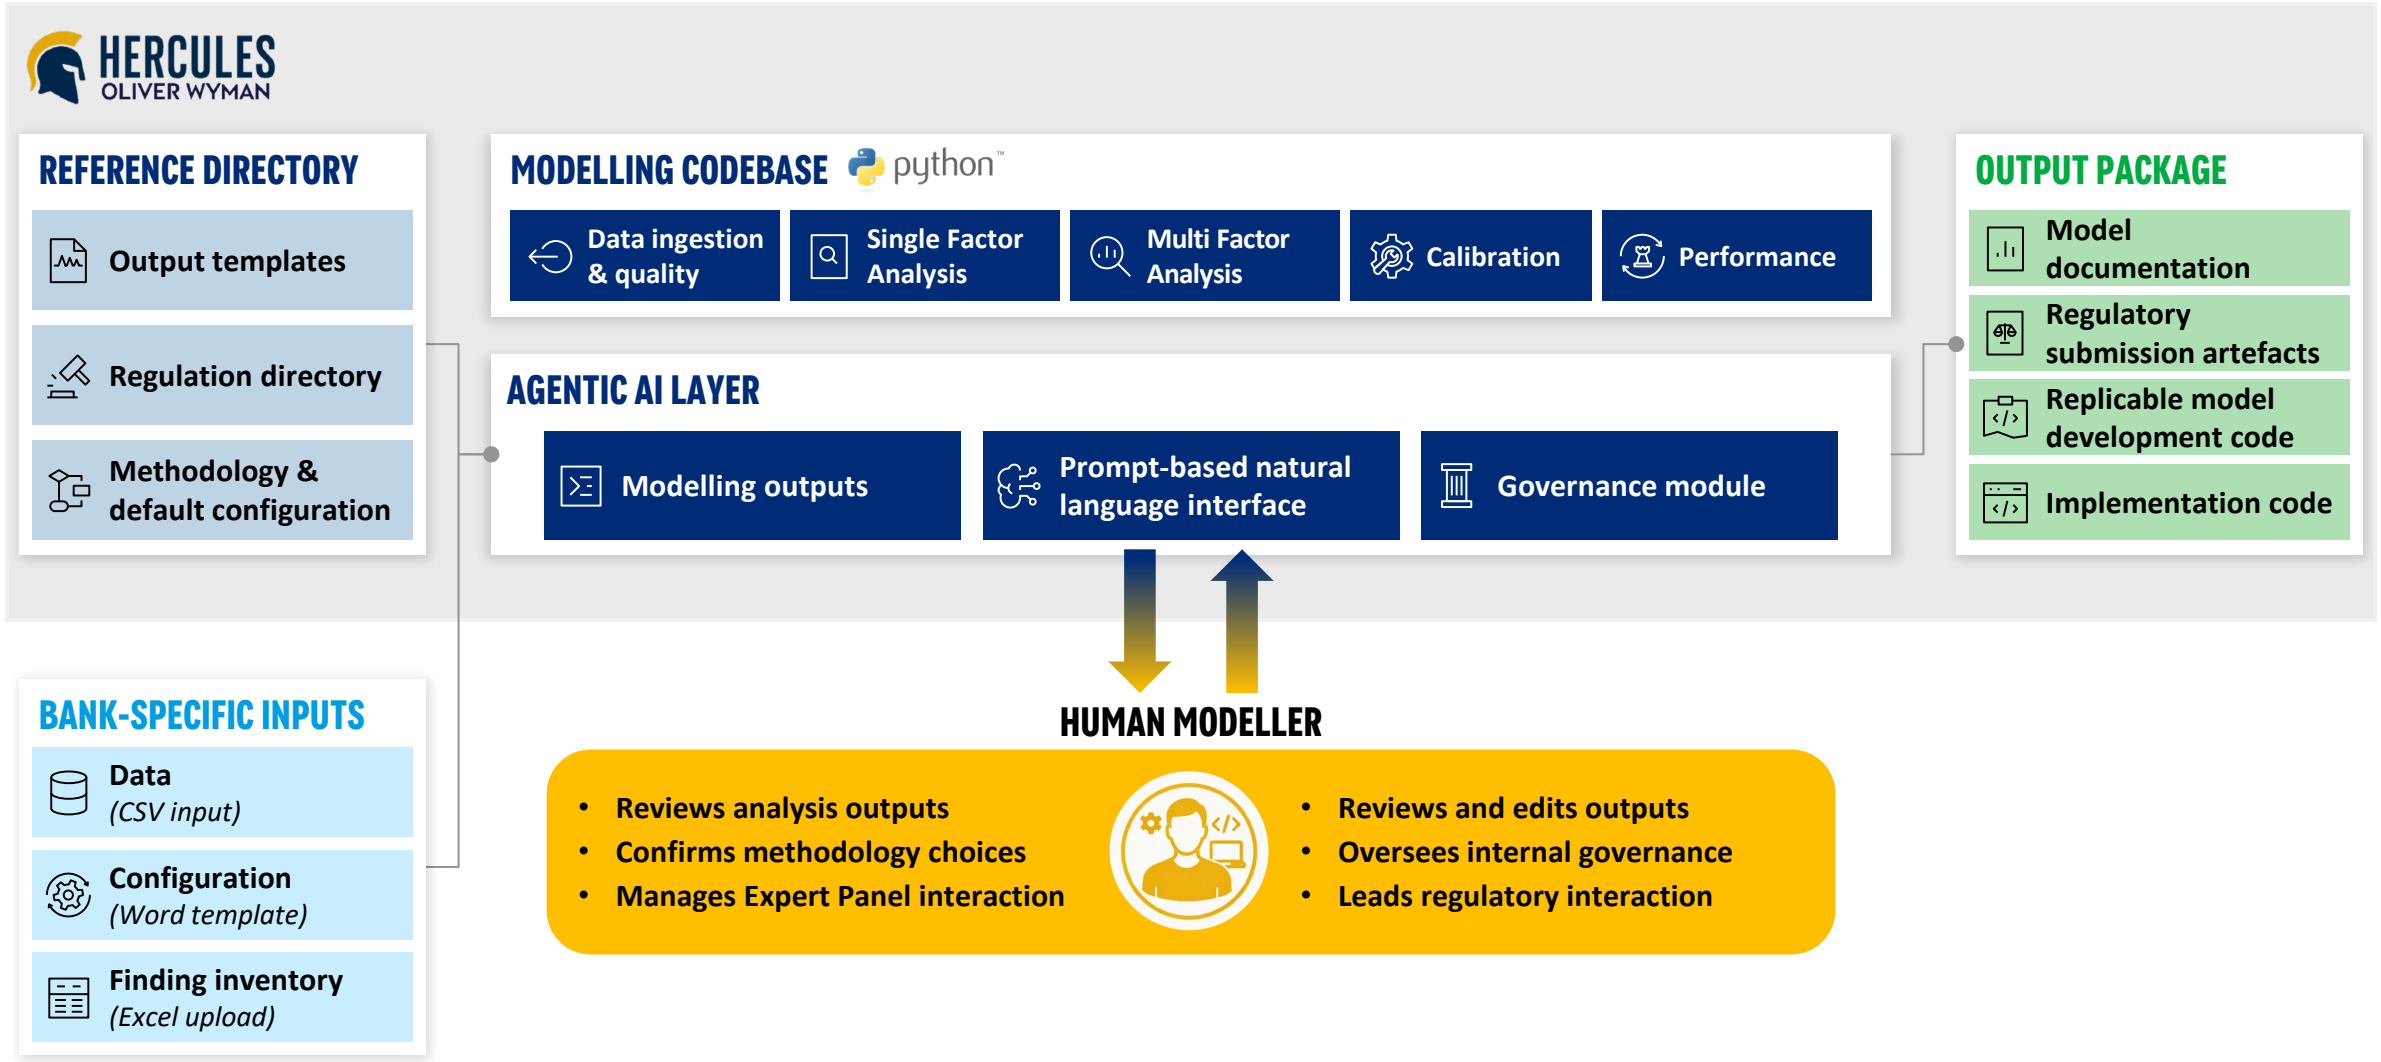


© Oliver Wyman

# HERCULES IS BUILT ON OUR FLEXIBLE AND VALIDATED MODELLING CODEBASE, WHICH HAS BEEN DEPLOYED IN MANY SUCCESSFUL CREDIT RISK MODELLING PROJECTS

**Oliver Wyman's fully modularised Python codebase covers the end-to-end model lifecycle:**

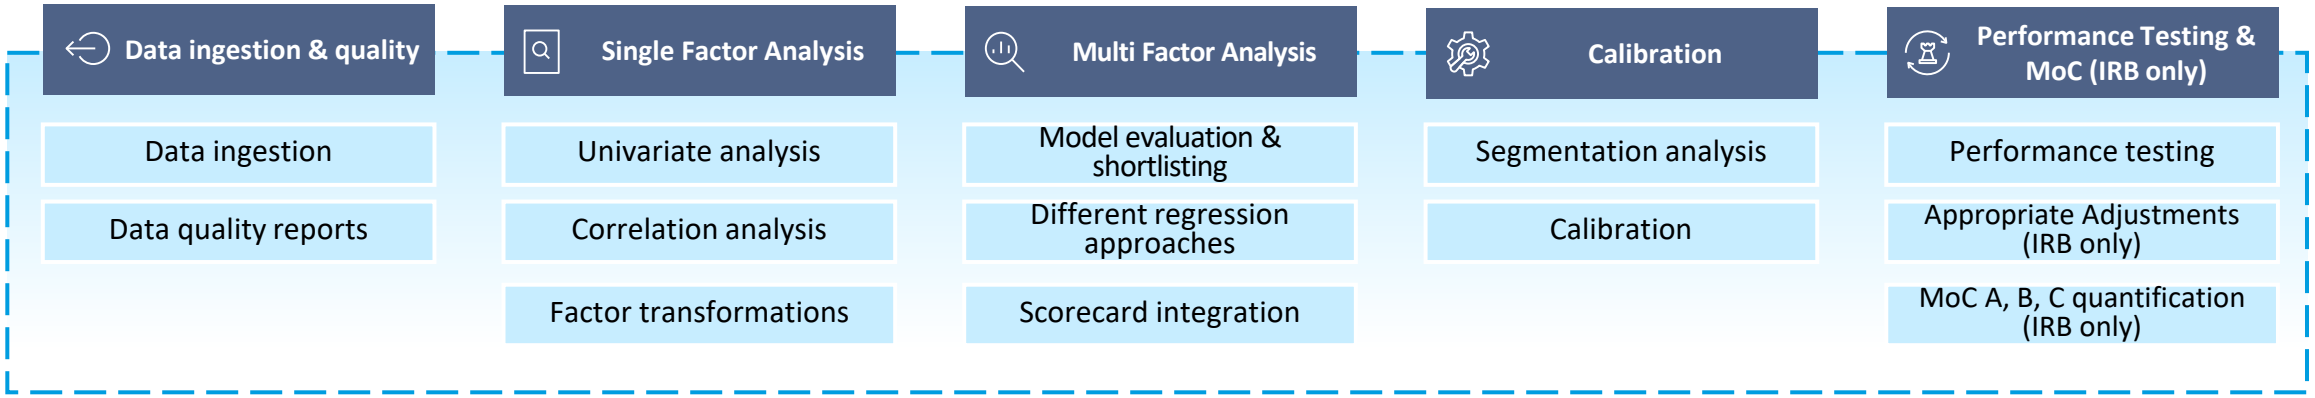


- Our codebase has been developed and validated through years of **extensive credit risk modelling experience**, resulting in numerous successful regulatory submissions
- This codebase encapsulates our proven modelling methodologies while also supporting a range of the most common alternative approaches, such as various transformation techniques and calibration methods
- Designed for horizontal scalability, the codebase is **regularly updated with new features** to address **evolving regulatory requirements**.
- It offers **flexible configuration options** to accommodate bank-specific thresholds (such as for discriminatory power, multicollinearity, etc)

**THE HERCULES CODEBASE IS RING-FENCED FROM THE AGENTIC LAYER TO ENSURE RELIABILITY AND TRACEABILITY**

© Oliver Wyman

# THE HERCULES' AGENTIC AI LAYER ACCELERATES CREDIT RISK MODEL DEVELOPMENT VIA END-TO-END ORCHESTRATION, PROACTIVE INSIGHTS AND AUTOMATED DOCUMENTATION

## End-to-end orchestration

Hercules **executes the end-to-end model lifecycle**, with **human modeller 'in-the-loop'** via an intuitive, prompt-based interface with **key pre-defined review checkpoints**

**Cycle times cut drastically** with automatic orchestration & error-free re-runs in hours

- **Models updated** based on recent data **much faster and more frequently**
- **Findings fixed much faster**, add-ons stay in place much shorter

**No-code interface** with **human modellers focused on value-adding steer** with reduced need for deep coding skills

**Ensures alignment with pre-selected methodology,** with **AI agent acting within well-defined guardrails** and embedded constraints

## Proactive insights

Hercules provides a **first interpretation of results**, **proactively asks for human review** & input, and makes **recommendations for methodology choices** and further analyses

**Key insights directly reflected** in documentation

**AI detects anomalies and potential findings**, leveraging regulation directory and findings inventory

- **Proactively flagging** to human modeler **with quantified impact**
- **Suggesting mitigants** where possible

Ability to **compare results across comparable portfolios and models** via access to full model inventory and data

## Automated documentation

Hercules drafts **documentation** per template specifications, **updated 'at the push of a button'** – with ability to **review & edit AI outputs iteratively**, preserving feedback across versions

**Drastically reduced human effort in document writing** & production

**Automated re-writes** to reflect model and/or data changes **within hours** with version tracking

**Internal consistency** from enforcing standardised templates, terminology, and formatting

**No version control errors**

© Oliver Wyman

# A NATURAL LANGUAGE INTERFACE FACILITATES DECISION-MAKING BY PROACTIVELY ASKING FOR INPUT AND GIVING SUGGESTIONS

**Hercules prompts the user to make decisions and provide input as it works through the credit risk modelling process – It can pick up the analysis at any time from any point in the process, allowing for governance steps such as Expert Panel discussions**

- Database/previous runs available via UI
- Interface allows natural language input
- Agent runs relevant code module with ability to bring up results visually
- Ability to assess results against pre-configured methodology settings
- Seamless transition between code modules

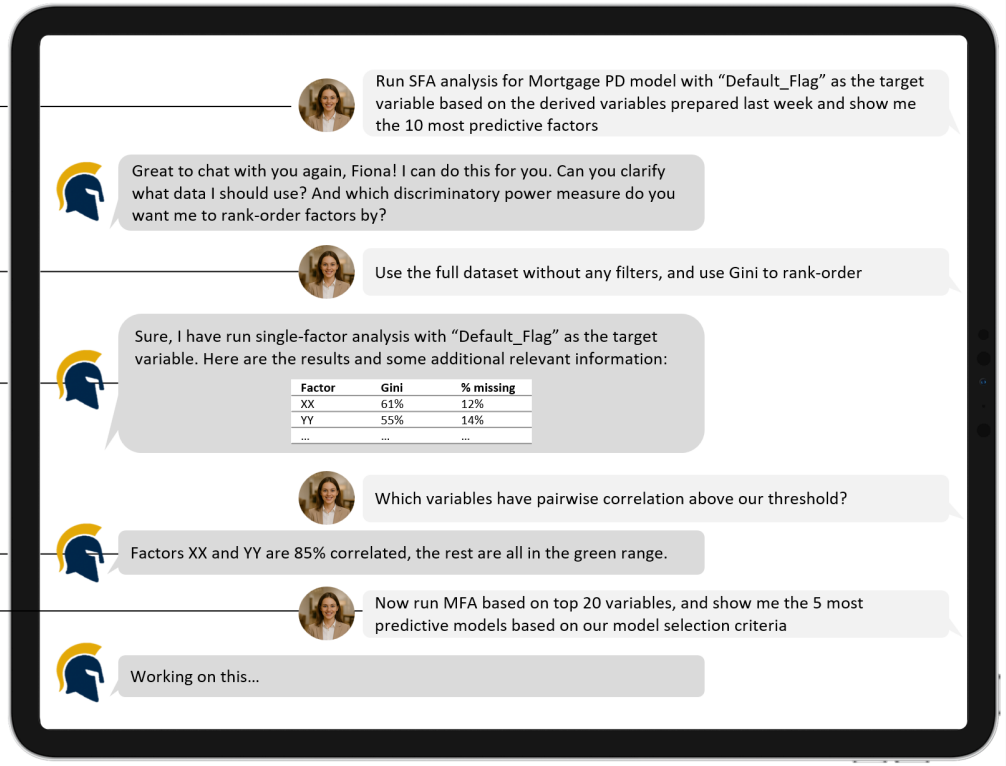


## HERCULES interface provides…

- ✓ Display of key results and recommends next steps
- ✓ Iterative engagement with the modeler via targeted questions
- ✓ Repeatable runs when the modeler approves changes
- ✓ Decision logs that allow inputs and feedback to be carried over across sessions

## To safeguard reliability, replicability & traceability, it does not…

- × Modify the Hercules codebase
- × Operate autonomously without modeller approval
- × Bypass or override system-level prompts and/or configuration to overcome issues it may encounter when performing tasks

© Oliver Wyman

# ROBUST GUARDRAILS & SOUND ARCHITECTURAL DESIGN MITIGATE THE INHERENT AI RISK IN HERCULES – FOCUSING ON KEY CONCERNS AROUND REPLICABILITY AND TRACEABILITY

## Frequently Asked Questions

| | | |
|---|---|---|
|   | **Is this 'vibe-coding'?** | • No, the AI agent only creates the 'config' to call out different inputs or modules from the ring-fenced Hercules codebase, where the prescribed methodology is embedded – the AI agent does not create or change modelling code |
|  | **What guardrails are in place?** | • We enforce guardrails via system-level prompts, data filters and pre-defined action limits to prevent unsafe or unauthorised behaviour, whether it stems from the AI agent or human user |
|  | **How to you guarantee traceability?** | • Every accepted decision is recorded with timestamps, user IDs, and artefact versions to create an auditable decision trail<br>• All sessions are logged with the data and reference materials used, configuration followed, and user inputs provided |
|  | **How do you ensure explainability?** | • Each recommendation includes a short rationale and key drivers; explanation metadata is saved with the run for review<br>• Hercules also provides human-readable decision logs for each suggestion |
|  | **Are results replicable?** | • Yes – Hercules records environment, dependency versions, seeds, and full configuration to ensure runs are replicable<br>• Artefacts and scripts are stored in a versioned repository, which can be easily accessed and executed outside the Hercules environment for validation of outputs |
|  | **What if the human makes a mistake?** | • As with any other model development, human error is a factor, which is controlled by banks' governance processes<br>• Within Hercules, any decision requires explicit confirmation, can be postponed until governance has been followed or made dependent on sign-off from authorised user – and can be rolled back |
|  | **How do you monitor Hercules performance?** | • We monitor Hercules using a combination of efficiency and risk metrics against established benchmarks<br>• For efficiency, we track the number of required 'hard code' fixes and the time spent on each step of the analysis<br>• For risk, we measure output reproducibility, pipeline error rates, SME acceptance rates, and the extent of documentation rework needed |

© Oliver Wyman

# HERCULES DELIVERS TRANSFORMATIONAL IMPACT FOR CREDIT RISK MODELLING FROM DAY 1 – WITH GREAT SCALABILITY POTENTIAL AND EXPANDING USE CASES

## Day 1 benefits

**65%**
SAVING IN MODELLING FTES

AVERAGE MODEL DEVELOPMENT TIME DOWN FROM
**FROM 6 TO 2 MONTHS**

**FASTER VALIDATION & REMEDIATION OF FINDINGS**

## Future

- Continuous model monitoring based on monthly data feeds
- Model validation module with challenger AI agent
- Automating scenario generation and stress-testing orchestration at scale
- Enabling secure cross-bank collaboration for data pooling and model benchmarking

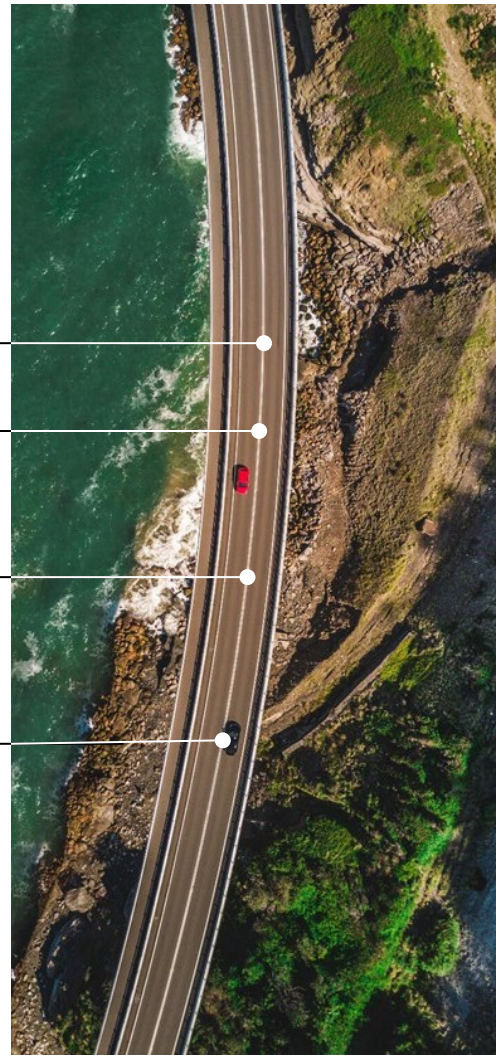

© Oliver Wyman

# HOW WE PROPOSE TO PARTNER WITH YOU

## WE CAN IMPLEMENT HERCULES IN 12 WEEKS

**Our Python codebase is fully developed and tested**, having undergone multiple reviews by banks' Model Validation functions

**Agentic AI layer is proven to work** with successful execution of **modelling steps and automated documentation**

We would implement Hercules **in your tech and AI environment**, working with your AI partner if necessary

As **Hercules** ingests modelling data in a simple CSV or similar format, there is **no lead time for data pipelines to be created**

## COST SAVINGS ARE IMMEDIATE WITH LIMITED UPFRONT INVESTMENT & LOW RISK

- We are eager to partner with you to make this a reality – and happy to devise a setup that **limits any investment by you without proven cost savings**
- Regulatory risks are minimal as Hercules is designed to be a hyper-competent AI assistant to an experienced modeller, with 'human-in-the-loop' principle for all decisions and configuration
- We can work with your Model Risk function to explore and address any concerns they may raise

## WE SUGGEST TO SCHEDULE A DEMO SESSION TO SHOWCASE HERCULES' CAPABILITIES – AND BUILD A MODEL TOGETHER IN AN HOUR

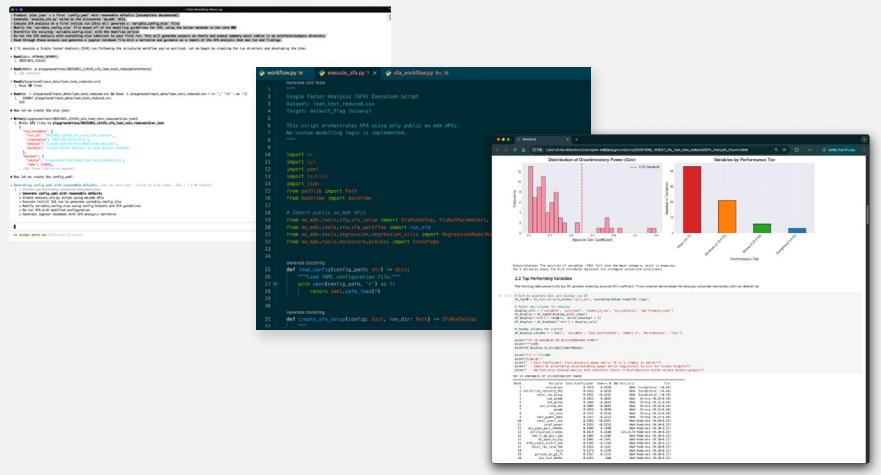

© Oliver Wyman

# QUALIFICATIONS, ASSUMPTIONS, AND LIMITING CONDITIONS

This report is for the exclusive use of the Oliver Wyman client named herein. This report is not intended for general circulation or publication, nor is it to be reproduced, quoted, or distributed for any purpose without the prior written permission of Oliver Wyman. There are no third-party beneficiaries with respect to this report, and Oliver Wyman does not accept any liability to any third party.

Information furnished by others, upon which all or portions of this report are based, is believed to be reliable but has not been independently verified, unless otherwise expressly indicated. Public information and industry and statistical data are from sources we deem to be reliable; however, we make no representation as to the accuracy or completeness of such information. The findings contained in this report may contain predictions based on current data and historical trends. Any such predictions are subject to inherent risks and uncertainties. Oliver Wyman accepts no responsibility for actual results or future events.

The opinions expressed in this report are valid only for the purpose stated herein and as of the date of this report. No obligation is assumed to revise this report to reflect changes, events, or conditions, which occur subsequent to the date hereof.

All decisions in connection with the implementation or use of advice or recommendations contained in this report are the sole responsibility of the client. This report does not represent investment advice nor does it provide an opinion regarding the fairness of any transaction to any and all parties. In addition, this report does not represent legal, medical, accounting, safety, or other specialized advice. For any such advice, Oliver Wyman recommends seeking and obtaining advice from a qualified professional.


A business of Marsh McLennan