



SYG
CONSULTING

AGENTIC SIMULATION · A WORKING SERVICE

Compare policy options across thousands of simulated futures.

Each option runs hundreds to thousands of times against matched draws of disruptive events. You receive a distribution of futures per option, a traceable record, and a method note.


● BUILT · PILOTED · IN PRODUCTION USE

THE PLATFORM


Three surfaces run every engagement. Only the design pack changes.



Authoring and silent review



The browsable simulation record



Any descriptor, any option, every run

Dozens to hundreds

agents, one per faction

Hundreds to thousands

runs per option, Monte Carlo scale when needed

Multi-year

horizon, quarterly steps, two decision rounds each

Frequencies and conditional probabilities

read across all runs

Every claim traceable

to run, decision, persona, source

Four expert gates

before any finding is read

Findings are distributions across simulated futures, read under the design assumptions. No run is a forecast.

HOW IT WORKS



Build does not start until the design is frozen; production does not start until calibration clears.

Humans decide at every gate.

HOW BIAS IS TACKLED

Bias is measured before results are read: a pre-registered battery, fixed before the data is opened, re-run after every correction.


SOURCE OF BIAS	DESIGN RESPONSE
● Model priors on the domain	— Grounding in documented sources and historical anchors
● Escalation propensity of language models	— Restraint written into every brief as a costed option; propensity measured and disclosed
● Persona drafting	— Expert review with a diverse panel; disagreement register
● Prompt framing and format	— One shared template; every rewrite measured on fixed seeds
● Model-strength asymmetry between agents	— Uniform-model test on pre-registered metrics
● State-transition rules	— Referee with inertia, decay and expert-set bounds
● Cultural and linguistic normalization	— Disclosed; heavier expert calibration

PRE-REGISTERED BIAS BATTERY

CROSS-VENDOR ADJUDICATION ENSEMBLE AS A DIAGNOSTIC


RESIDUAL BIAS DISCLOSED IN THE METHOD NOTE

MODELS, SELECTED AND PINNED



- Versions pinned per batch, recorded in every call
- Provider allow-list, fallbacks disabled: a call reaches the pinned model or fails visibly
- Contextual judgments adjudicated by two or three vendor families; disagreement routed to experts
- Regionally hosted endpoints where data handling requires

HUMANS REVIEW AND DECIDE



- Your experts design the world: actors, factions, options, events, sources
- Silent review before the room: confirm, not my area, revise, flag
- Expert splits are recorded, never averaged; they define the sensitivity sweep
- Outcome questions, plausibility bounds and sign-off are yours

What you receive

- **Trajectory database**, one self-contained file
- **Browsable simulation record**
- **Method note**: claims, limits, residual bias
- **Replication kit**

Who it is for

- Government, defense and policy institutions comparing options under shocks
- Corporate strategy stress-testing plans against disruption
- Macro and financial analysts running what-if scenarios

What we do not claim

- No run is a forecast.
- Scale improves reliability, not validity.
- Bias is measured, never assumed away.